Can intergroup contact foster more continuous, fluid, and inclusive social identities?



Nils Karl Reimer
Brasenose College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Trinity 2018

Can intergroup contact foster more continuous, fluid, and inclusive social identities?

Nils Karl Reimer
Brasenose College, University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Trinity 2018

Abstract

This thesis examines whether intergroup contact can foster more continuous, fluid, and inclusive social identities—and, if so, how more continuous, fluid, and inclusive social identities relate to intergroup bias and support for social change. Chapter 1 introduces relevant research on social identification along a single category and across multiple categories, and develops the hypotheses tested throughout this thesis. Chapter 2 gives an overview of the empirical research, and introduces the statistical methods used to analyse findings from that research. Chapter 3 reports research developing and validating a novel measure of social identity continuity and fluidity, and examines whether contact with sexual and gender minorities fosters more continuous and fluid conceptions of sexuality and gender. Studies 1–3 found that intergroup contact was associated with more continuous and fluid social identities, but that more continuous and fluid social identities were not, in turn, associated with less intergroup bias. Chapter 4 reports research testing whether contact with caste or religious outgroups fosters more inclusive social identities in South India. Study 4 found that, as hypothesised, cross-group friendship was associated with more inclusive identities, while more inclusive identities were associated with less intergroup bias. Chapter 5 reports research testing how intergroup contact relates to support for social change in advantaged and disadvantaged groups. Studies 1 and 3 found that for disadvantaged-group members, negative contact, but not positive contact, was associated with more collective action. For advantaged-group members, positive contact predicted more solidarity-based collective action. More continuous and fluid identities did not mediate any of these relationships. Study 4 found neither intergroup contact nor more inclusive identities to be associated with support for or opposition to social change. Chapter 6 summarises and discusses the research presented in this thesis, highlighting its implications for intergroup relations research.

Acknowledgements

Foremost, I thank my supervisors, Katharina Schmid and Miles Hewstone, for opening so many of the doors through which I have passed to get where I am now. I am indebted to them for making me a better writer and scientist.

I thank the institutions who, through their generous support, have made this thesis possible. I am grateful to the Economic and Social Research Council, for supporting my doctoral studies with an Advanced Quantitative Methods studentship and for funding my research visit to Yale University; to Brasenose College, for supporting me throughout two graduate degrees; and to the Society of Personality and Social Psychology, for giving me the funds to present my research at its annual meeting.

I thank my collaborators, without whom the present research would not have been possible: Julia Becker, Oliver Christ, Kristof Dhont, Ulrich Klocke, Angelika Love, Sybille Neji, and Magdalena Rychlowska (Study 3). I am grateful to Shanmukh Kamble (Study 4) for placing his faith in my research idea, and for his patience awaiting its eventual publication. I am indebted to 2953 participants who, together, have donated almost 100,000 hours of their time to this research.

I thank past and present members of the Oxford Centre for the Study of Intergroup Conflict, as well as the wonderful people in its orbit, for welcoming, supporting, and inspiring me over the years. In particular, I would like to thank Ben Fell for being there for me since the beginning, and for sharing a great many zoological facts. I am also grateful to Jack Dovidio, Jenn Richeson, and others for welcoming and mentoring me while visiting Yale University. I would like to express my appreciation for the Massachusetts Institute of Technology libraries who gave me a place to write while living abroad.

I wish to thank the people I love for being there for me throughout this journey. I consider myself fortunate to have a great many friends in my life—I appreciate all of you. Among them, I would like to thank two who have made my time in Oxford such a delight: Huma, for reminding me that there is more to life than studying (and for being a great lead and producer), and Todd, for being my best friend. At this time, I am reminded of all I owe to my family: I thank my parents, as well as my siblings, for their love and support over all those years. To Ananya, I owe you more than would fit this page. I would still be writing had it not been for your love and support. Thank you for keeping my life exciting.

On a final note, I would like to thank my examiners, Nadira Faber and Nicole Tausch, for being perhaps the only people to continue reading past this page.

Abstract

This thesis examines whether intergroup contact can foster more continuous, fluid, and inclusive social identities—and, if so, how more continuous, fluid, and inclusive social identities relate to intergroup bias and support for social change. Chapter 1 introduces relevant research on social identification along a single category and across multiple categories, and develops the hypotheses tested throughout this thesis. Chapter 2 gives an overview of the empirical research, and introduces the statistical methods used to analyse findings from that research. Chapter 3 reports research developing and validating a novel measure of social identity continuity and fluidity, and examines whether contact with sexual and gender minorities fosters more continuous and fluid conceptions of sexuality and gender. Studies 1–3 found that intergroup contact was associated with more continuous and fluid social identities, but that more continuous and fluid social identities were not, in turn, associated with less intergroup bias. Chapter 4 reports research testing whether contact with caste or religious outgroups fosters more inclusive social identities in South India. Study 4 found that, as hypothesised, cross-group friendship was associated with more inclusive identities, while more inclusive identities were associated with less intergroup bias. Chapter 5 reports research testing how intergroup contact relates to support for social change in advantaged and disadvantaged groups. Studies 1 and 3 found that for disadvantaged-group members, negative contact, but not positive contact, was associated with more collective action. For advantaged-group members, positive contact predicted more solidarity-based collective action. More continuous and fluid identities did not mediate any of these relationships. Study 4 found neither intergroup contact nor more inclusive identities to be associated with support for or opposition to social change. Chapter 6 summarises and discusses the research presented in this thesis, highlighting its implications for intergroup relations research.

Contents

Acknowledgements i

Abstract iii

1 Introduction 1
Social identification and intergroup relations . . . 2
Social categorisation and social stereotypes . . . 2
Social categorisation and intergroup bias . . . 3
Interpersonal–intergroup continuum . . . 5
Social categorisation and intergroup contact . . . 5
Social identification along a single category . . . 7
Psychological essentialism . . . 8
Bridging binaries . . . 11
Social identification across multiple categories . . . 14
Crossed categorisation . . . 15
Bridging divides . . . 17
Social identification, intergroup contact, and social change . . . 19
A darker side of 'we'? . . . 20
A path to solidarity? . . . 22
Overview of thesis structure . . . 24

2 Method 25
Overview of studies . . . 25
Bayesian data analysis . . . 26
Likelihood . . . 27
Estimation . . . 30
Comparison . . . 36
Extensions . . . 41
Generalised linear models . . . 41
Multilevel models . . . 43
Multivariate models . . . 45
Structural equation models . . . 46
Software . . . 47

3 *Bridging binaries*: Can intergroup contact foster more fluid and continuous conceptions of sexuality and gender? 49
Sexuality as case study . . . 50
Gender as case study . . . 52
Present research . . . 55

Study 1 . . . 58
Method . . . 58
Results . . . 60
Study 2 . . . 64
Method . . . 65
Results . . . 68
Study 3 . . . 76
Method . . . 78
Results . . . 79
Discussion . . . 85
Summary of findings . . . 85
Strengths and limitations . . . 87
Implications . . . 88

4 *Bridging divides*: Can intergroup contact foster more inclusive national and religious identities? 91
Caste, religion, and nation in South India . . . 92
Study 4 . . . 95
Pilot study . . . 96
Method . . . 98
Results . . . 101
Discussion . . . 115
Summary of findings . . . 115
Strengths and limitations . . . 116
Implications . . . 118

5 *Intergroup contact and social change*: A darker side of 'we' or a path to solidarity? 121
Present research . . . 122
Study 1 . . . 123
Method . . . 124
Results . . . 125
Study 3 . . . 129
Method . . . 130
Results . . . 131
Study 4 . . . 140
Method . . . 141
Results . . . 142
Discussion . . . 149
A darker side of 'we'? . . . 149
A path to solidarity? . . . 151
Strengths and limitations . . . 153

6 Conclusion 155
When does intergroup contact reduce prejudice? . . . 156
Positive contact *sometimes* decreases prejudice . . . 156
Negative contact *sometimes* increases prejudice . . . 159
How does intergroup contact reduce prejudice? . . . 160
Bridging binaries . . . 160

Bridging divides 162
Does intergroup contact help or hinder social change? 164
Doing research inside the reproducibility crisis 166
Conclusions 168

Appendices **170**

A Materials for Chapter 3 **173**

B Materials for Chapter 4 **179**
Instructions for crossed-categorisation task 179
Descriptive statistics 180
Correlates of categorisation 182
Measures 182
Results 182

C Materials for Chapter 5 **185**

D Materials for Chapter 6 **189**

References **191**

List of Figures

1.1 Results from Haslam, Rothschild and Ernst (2000, 2002) 10
1.2 Schematic representations of different conceptions of sexuality (for a man) 13
1.3 Schematic representations of social identity structures 17

2.1 Binomial likelihood for 10 coin flips 28
2.2 Binomial likelihood for 100 coin flips 28
2.3 Normal likelihood for continuous observations 29
2.4 Bayesian inference for 10 coin flips 31
2.5 Prior and posterior probabilities for 10 coin flips 33
2.6 Prior and posterior probabilities for 100 coin flips 35
2.7 R^2 as a function of the number of predictors in linear regression 38
2.8 Prior and posterior probabilities for generalised linear model 42
2.9 Posterior probabilities with complete, partial, and no pooling 44
2.10 Posterior probabilities with two varying intercepts 44

3.1 Results from a representative survey of British adults (YouGov, 2015) . . . 51
3.2 Schematic representations of different conceptions of sexuality (for a woman) 53
3.3 Results from a representative survey of British adults (YouGov, 2016) . . . 55
3.4 Schematic representations of different conceptions of sexuality (for a woman) 56
3.5 Inter-item correlations for social identity continuity and fluidity (Study 1) 61
3.6 Correlations of scale factors with variables related to construct validity and variables hypothesised as antecedent and consequent (Study 1) 63
3.7 ESEM for 1-factor and 2-factor solutions (Study 1) 64
3.8 Inter-item correlations for social identity continuity and fluidity (Study 2a) 69
3.9 Factor structures in exploratory factor analyses (Study 2a) 69
3.10 Correlations of scale factors with variables variables related to construct validity and variables hypothesised as antecedent and consequent (Study 2a) 72
3.11 ESEM for sexuality (Study 2a) 72
3.12 Inter-item correlations for social identity continuity and fluidity (Study 2b) 73
3.13 Factor structures in exploratory factor analyses (Study 2b) 74
3.14 Correlations of scale factors with variables related to construct validity and variables hypothesised as antecedent and consequent (Study 2b) . . . 76
3.15 ESEM for gender (Study 2b) 77
3.16 Longitudinal structural equation model for sexuality (Study 3a) 83
3.17 Longitudinal structural equation model for sexuality (Study 3b) 85

4.1 Map of South India . . . 93
4.2 Proportion of SC/ST vs GM classifications per target stimulus (Pilot study) 97
4.3 Foreign targets selected based on participants' average typicality ratings (Pilot study) . . . 98
4.4 Examples of stimuli used in Study 4 . . . 99
4.5 Edit distances between social identity structures and participants' responses 103
4.6 Estimated probabilities of participants' categorising a target as "us" versus "not us" . . . 107
4.7 Estimated probabilities of target categorisations as a function of inter-group contact . . . 109
4.8 Estimated social distance ratings . . . 113
4.9 Estimated outgroup attitudes ratings . . . 114

5.1 Structural equation model for sexual-minority participants (Study 1a) . . . 127
5.2 Effect sizes for indirect effects in Study 1a . . . 127
5.3 Alternative model for sexual-minority participants (Study 1a) . . . 128
5.4 Structural equation model for sexual-majority participants (Study 1b) . . . 129
5.5 Panel model for sexual-minority participants (Study 3a) . . . 133
5.6 Effect sizes for indirect effects in Study 3a . . . 134
5.7 Panel model for sexual-majority participants (Study 3b) . . . 135
5.8 Panel model (with continuity/fluidity) for sexual-majority participants (Study 3b) . . . 136
5.9 Panel model for sexual-majority participants (Study 3c) . . . 137
5.10 Panel model for sexual-minority participants (Study 3c) . . . 138
5.11 Panel model (with continuity/fluidity) for sexual-minority participants (Study 3c) . . . 139
5.12 Panel model (with continuity/fluidity) for sexual-majority participants (Study 3c) . . . 140
5.13 Posterior probabilities of mean perceived life difficulty ratings (Study 4) . . . 144
5.14 Posterior probabilities of mean policy support ratings (Study 4) . . . 145

B.1 Correlations between social identity variables and target categorisations . . . 183

C.1 Three-wave panel model for sexual-minority participants (Study 3a) . . . 188

List of Tables

2.1 Model comparison using leave-one-out cross validation 40

3.1 Geomin-rotated factor loadings for the 1-factor and 2-factor solutions in Study 1 62
3.2 Geomin-rotated factor loadings for the 4-factor solution in Study 2a . . . 71
3.3 Geomin-rotated factor loadings for the 5-factor solution in Study 2b . . . 75
3.4 Model fit indices for measurement invariance comparisons in Study 3a . 80
3.5 Estimates of internal consistency in Study 3a 81
3.6 Model fit indices for measurement invariance comparisons in Study 3b . 83
3.7 Estimates of internal consistency in Study 3b 84

4.1 Participants by gender, age, nationality, religion, and caste 100
4.2 Mean intra and inter-category correlations for target categorisations . . . 102
4.3 Model comparison for group and individual differences 106
4.4 Model comparison for consequences 112
4.5 Model comparison for intergroup threat 114

5.1 Model fit indices for measurement invariance comparisons in Study 1 . . 125
5.2 Model fit indices for measurement invariance comparisons in Study 3a . 132
5.3 Estimated residual correlations in Models 0 and 1 for Study 4 145
5.4 Model comparison for individual differences in perceived life difficulty for Study 4 148
5.5 Model comparison for individual differences in policy support for Study 4 149

6.1 Example items for the social identity continuity and fluidity scales 161

A.1 Complete item pool for scale development in Study 2a 174
A.2 Complete item pool for scale development in Study 2b 175
A.3 German translations for Study 3a 176
A.4 German translations for Study 3b 177

B.1 Correlations and descriptive statistics for Study 4 181

C.1 Correlations for LGBT+ respondents within time points (Study 3a) 186
C.2 Correlations for LGBT+ respondents across time points (Study 3a) 186
C.3 Correlations for heterosexual respondents across time points (Study 3a) . 187
C.4 Model fit indices for measurement invariance comparisons for sexual-minority participants in Study 3c 187
C.5 Model fit indices for measurement invariance comparisons for sexual-majority participants in Study 3c 187

D.1 Items for a generalised social identity continuity and fluidity scale 190

1

Introduction

Past research has found that contact between groups reduces prejudice and social conflict (Al Ramiah & Hewstone, 2013; Pettigrew & Tropp, 2006, 2011). This research has shown that interacting with diverse others affects how we think and feel about them. In this thesis, I examine whether intergroup contact can change not only how we see *others*, but also how we see *ourselves*. That is, I study whether contact changes how we understand our social identities and group memberships. Below, I outline how answering this question contributes to ongoing debates in intergroup relations research.

Researchers have debated *how* or *why* intergroup contact reduces prejudice. Pettigrew and Tropp (2008) compared three such mechanisms in a meta-analysis. Fifty-four studies tested whether contact reduces prejudice by decreasing anxiety about interacting with outgroup members, by fostering empathy and perspective-taking, and by increasing knowledge about the outgroup. Pettigrew and Tropp (2008) found evidence for all three mediators, though the two affective mediators (anxiety, empathy) explained more variance than the cognitive mediator (knowledge). Swart, Hewstone, Christ and Voci (2011) found longitudinal evidence for the mediating roles of anxiety and empathy. All three mechanisms entail contact changing how we think and feel about others. Other researchers have proposed that contact could reduce prejudice by changing how we see ourselves (Pettigrew, 1997, 1998). For example, contact could change attitudes by fostering more inclusive (Gaertner, Dovidio, Guerra, Hehman & Saguy, 2016) or more complex (Schmid, Hewstone, Tausch, Cairns & Hughes, 2009) social identities. Adding to this literature, I propose two novel mechanisms that could explain how contact reduces prejudice.

In this chapter, I argue that social identification is central to understanding intergroup relations. In our feelings, thoughts, and actions, we tend to prefer those whom we consider "us" to those whom we consider "them". I propose two ways in which intergroup contact could change how we understand our social identities and, thereby, reduce intergroup bias. First, I consider social identification along a single category (e.g., sexual orientation). I hypothesise that contact encourages people to perceive a social identity as their unique position on a continuum, rather than as belonging to one of several discrete categories; and as flexible and changing, rather than as fixed and stable throughout their lifetime. By blurring boundaries between "us" and "them", intergroup contact could reduce intergroup bias. Second, I consider social identification across multiple categories (e.g., religion and nationality). I hypothesise that contact fosters inclusive social identities that include a wider range of groups in one's ingroup. By extending ingroup preferences to (former) outgroup members, intergroup contact could reduce intergroup bias. Overall, I

expect that intergroup contact encourages more continuous, fluid, and inclusive identities, and that these changes explain the prejudice-reducing effects of intergroup contact.

Beyond prejudice, researchers have debated whether intergroup contact *helps* or *hinders* social change. Above, I have discussed ways in which contact could change attitudes and beliefs. Changing attitudes and beliefs, however, is seldom enough to overcome structural inequalities. Instead, social change often requires collective action by the disadvantaged, and sometimes the advantaged. On the one hand, researchers have argued that, by fostering more inclusive identities, contact could discourage disadvantaged-group members from engaging in collective action (Dovidio, Saguy, Gaertner & Thomas, 2012). On the other hand, more inclusive identities could break down the boundaries between the advantaged ingroup and the disadvantaged outgroup—making injustice faced by "them" into "our" problem. Contact could thus encourage advantaged-group members to support the struggle for social justice. In this chapter, I discuss how intergroup contact could—by fostering more continuous, fluid, and inclusive identities—affect support for social change in disadvantaged and advantaged groups.

In the next section, I draw on social identity theory (Tajfel, 1981) and self-categorization theory (J. C. Turner, Hogg, Oakes, Reicher & Wetherell, 1987) to review the role of social identification in intergroup relations. I discuss social categorisation as a *necessary* condition for social stereotypes. I argue that social identification makes social categorisation a *necessary* and, under some circumstances, *sufficient* condition for intergroup bias—when contrasting "us" and "not us" (ingroup favouritism), "me" and "them" (outgroup derogation), and "us" and "them" (intergroup-directed bias). I then review *when* social categorisation shapes our thoughts and behaviour. Finally, I discuss how various perspectives have applied these insights to intergroup contact.

Social identification and intergroup relations

> *The situation, it seems, can best be stated as follows: although we could not perceive our own in-groups excepting as they contrast to out-groups, still the in-groups are psychologically primary. We live in them, by them, and, sometimes, for them.*
>
> – Gordon W. Allport (1954/1979, p. 42)

As humans, we readily categorise ourselves and others into social groups. That is, we spontaneously divide people according to *ascribed* categories (e.g., ethnicity or gender), *achieved* states (e.g., occupation or political affiliation), and other relevant distinctions (Deaux, 1996)—or, more generally, based on perceptions of similarity, proximity, and common fate (Campbell, 1958). On the one hand, this process of *social categorisation* helps us navigate a social world marked by group distinctions and hierarchies. On the other hand, social categorisation has perceptual and motivational consequences that underlie harmful social stereotypes and intergroup bias.

Social categorisation and social stereotypes

Social categorisation changes how we process information about members of relevant groups. Specifically, social categorisation leads to the accentuation of *within*-category similarities and *between*-category differences (Tajfel, 1969). That is, we tend to perceive members of the same group as more similar to each other, and members of different groups

as more different from each other, than they actually are, once we have categorised them into their respective groups. This process reduces continuous human traits to discrete category stereotypes. Social categorisation thus shapes our beliefs about members of the relevant categories.

Social categorisation explains how *individuals* form stereotypes—but not what the stereotype's content is. Category stereotypes become *social* stereotypes when their contents are widely shared between members of a group, or even in a society. The contents of social stereotypes arise from the psychological functions they serve for a group (Tajfel, 1981, Chapter 7; see also, Fiske, 2015; Fiske, Cuddy, Glick & Xu, 2002). What these functions are, in turn, depends on competition and power relations between groups. Social categorisation is thus *necessary*, but not *sufficient*, for the formation of social stereotypes.

Social stereotypes give rise to *stereotype-based discrimination*, which occurs "when persons are treated differentially as a function of their category membership based on beliefs about the category as a whole" (Brewer, 2007, p. 696). As the example of gender stereotypes shows, stereotypes need not be exclusively negative to elicit discrimination (Glick & Fiske, 2001), are sometimes shared by members of the target group, and can justify differential treatment and the resulting inequality (Jost & Kay, 2005). Social categorisation thus relates to some forms of discrimination, without recourse to the perceiver's own group membership and its motivational consequences.

Social categorisation and intergroup bias

Social categorisation, however, acquires further significance when applied to ourselves. As humans, we depend on groups for survival (Brewer, 2007), and have developed a fundamental need to belong to social groups (Baumeister & Leary, 1995). This means that we are motivated to categorise ourselves *into* some groups, and *out of* others. As we do so, we begin to define ourselves in terms of our group memberships, and attach value and meaning to them. Social categorisation, when applied to ourselves, thus begets social identification (Tajfel, 1981, Chapter 11). This process results in the formation of *social identities*, defined as "that *part* of an individual's self-concept which derives from his knowledge of his membership of a social group (or groups) together with the value and emotional significance attached to that membership" (p. 255).

Social identification, in turn, changes the nature of social categorisation. We divide others not just into groups, but into groups we are also part of (*ingroups*) and those we are not part of (*outgroups*). Social identities determine whom we categorise as "us" and "them". Dividing others into "us" and "them" changes how we treat members of these groups. *Intergroup bias* refers to "the systematic tendency to evaluate one's [ingroup] or its members more favourably than [an outgroup] or its members" (Hewstone, Rubin & Willis, 2002, p. 576). In this definition, intergroup bias arises as differential evaluation—that is, as the extent to which individuals favour ingroup members in their attitudes. Social categorisation into ingroups and outgroups is, by definition, *necessary* for intergroup bias to occur. But is it also *sufficient*?

To answer this question, Tajfel, Flament, Billig and Bundy (1971) studied social categorisation in 'minimal group' experiments. That is, researchers assigned participants to groups using trivial criteria (e.g., preference for different abstract painters). This artificial setting removed the usual trappings of social groups (e.g., social ties, social norms, and historical narratives), and thus reduced group membership to *mere* categorisation. To

measure intergroup bias, researchers then asked participants to allocate rewards between (unnamed) ingroup and outgroup members, excluding themselves. Participants' choices were anonymous and not linked to their own rewards. Still, most participants (e.g., 72% in Experiment 2) favoured ingroup members in their reward allocations. As the experimental setting avoided alternative explanations for intergroup bias (e.g., past hostilities, conflicts of interest, and self-interest), Tajfel et al. concluded that social categorisation was *sufficient* to elicit discrimination. Tajfel et al. defined intergroup bias as differential treatment—that is, as the extent to which individuals favour ingroup members in their behaviour.

Researchers have debated various processes that could explain why social categorisation leads to discrimination in minimal group studies (for a review, see Hewstone et al., 2002). Most likely, participants' reactions to social categorisation are multiply determined (Spears & Otten, 2012). Whatever the specific aetiology, research has consistently found social categorisation to elicit intergroup bias. Beyond reward allocations, participants spontaneously evaluated ingroup members more positively than outgroup members (Otten & Moskowitz, 2000), and tended to trust ingroup members more than outgroup members (Insko & Schopler, 1998). Likewise, participants favoured ingroup over outgroup members in causal attributions (Hewstone, 1990). Participants, for example, made more internal attributions for positive acts, and more external attributions for negative acts, when rating behaviour by ingroup versus outgroup members. Finally, participants were more likely to help ingroup strangers than outgroup strangers (Cuddy, Rock & Norton, 2007; Levine, Prosser, Evans & Reicher, 2005).

In these examples, participants show *ingroup-serving* preferences that resemble well-documented *self-serving* biases and behaviours (e.g., Mezulis, Abramson, Hyde & Hankin, 2004). In other words, people think about ingroup and outgroup in "me"–"not me" terms, and favour ingroup members ("us") over others ("not us").[1] Such *ingroup favouritism* is directed toward the ingroup, and not against any specific outgroup. Discrimination, in this case, does not stem from direct hostilities against outgroup members (Brewer, 1999). Instead, discrimination results from reserving favours and benefits for ingroup members. Ingroup favouritism arguably explains contemporary discrimination better than outgroup-directed hostilities (Greenwald & Pettigrew, 2014).

In other cases, intergroup bias is directed against an outgroup and its members. Mummendey et al. (1992), however, found that participants did not differentiate between (minimal) ingroup and outgroup members when distributing negative outcomes. This suggests that *mere* categorisation into "us" and "not us" is not sufficient to elicit outgroup-directed discrimination. Sexual prejudice (Herek & McLemore, 2013), for example, does not seem to be motivated by the desire to advantage one's ingroup. Instead, sexual prejudice is often motivated by the desire to distance *oneself* from the outgroup: "I'm not gay…I'm a real man!" (Falomir-Pichastor & Mugny, 2009). Discrimination thus arises from individuals harming outgroup members in the effort of distancing themselves ("me") from a stigmatised outgroup ("them").

As discussed above, ingroup favouritism does not usually necessitate outgroup derogation. The two are, however, intrinsically linked in some situations (Brewer, 1999, 2007). When comparing the ingroup to an outgroup, any ingroup favouritism entails outgroup derogation (*intergroup comparison*). Likewise, people sometimes perceive an outgroup's interests at odds with their ingroup's—especially when they see the outgroup as a threat to the ingroup (Sherif, Harvey, White, Hood & Sherif, 1954/1961; Stephan &

[1] I follow Brewer's (2007) typology, which distinguishes three forms of prejudice as consequences of social categorisation: ingroup (us–not us), outgroup (me–them), and intergroup (us–them) prejudice.

Stephan, 2000). In *intergroup conflict*, harming the outgroup can benefit the ingroup. Discrimination thus arises from protecting the ingroup ("us") against perceived threat from an outgroup ("them").

The three types of intergroup bias—ingroup-directed ("us"–"not us"), outgroup--directed ("me"–"them"), and intergroup-directed ("us"–"them") discrimination—are not mutually exclusive. Indeed, they often co-occur. All are linked to social categorisation. When we divide others into ingroup and outgroup members, we tend to favour ingroup members. When we want to distance ourselves from the outgroup, or when we see the outgroup as a threat to our ingroup, we further tend to derogate outgroup members. Social categorisation is thus *sufficient* to elicit ingroup favouritism and, in some circumstances, outgroup derogation.

Interpersonal–intergroup continuum

This conclusion raises the question of *when* we think of others in terms of their social categories. According to Tajfel (1981, p. 240), situations vary along an interpersonal–intergroup continuum. At the *interpersonal* extreme, a social encounter is "determined by the personal relationships between the individuals and by their respective individual characteristics". At the *intergroup* extreme, "*all* behaviour of two or more individuals towards each other is determined by their membership of different social groups or categories". As a situation moves closer to the intergroup extreme, we increasingly think of others in terms of their category memberships. Rather than act toward them as individual persons, we treat others as generic exemplars of their social category—a condition Tajfel (1969) calls depersonalisation.

Group salience—"the extent to which group memberships are psychologically 'present' during contact"—determines whether we experience an interaction as interpersonal or intergroup (R. Brown & Hewstone, 2005, p. 257). Whether a specific ingroup–outgroup categorisation is salient, in turn, depends on an interaction between its relative accessibility and fit in a given situation (Oakes, 1987). A social categorisation fits a situation if observed similarities and differences between individuals correlate with the relevant category distinctions (*structural fit*), and if these similarities and differences correspond to the content of relevant category stereotypes (*normative fit*). A social categorisation is *situationally accessible* if features of a situation make that categorisation clearly relevant. An ingroup–outgroup distinction becomes salient when it is both accessible and fits a given situation. On the one hand, group salience is thus situational.

On the other hand, group salience is also personal. That is, people differ in how often they use specific social categories. For some people, a social categorisation is more often on the forefront of their mind—it is *chronically accessible*. This means that they should find this social categorisation salient across more situations (Oakes, 1987). In social identification, a social category acquires emotional significance and becomes central to an individual's self-concept. Social identification should thus make a social categorisation more accessible.

Social categorisation and intergroup contact

If social categorisation is *necessary*, and often *sufficient*, for intergroup bias, then blurring the boundaries between groups should reduce such bias. Whether we categorise others into "us" and "them" depends on the situational and individual determinants of social

categorisation, as reviewed above. Below, I discuss how various perspectives draw on these insights to define the ideal conditions for intergroup contact, to conceptualise the psychological mechanisms through which contact reduces intergroup bias, and to understand how and when the effects of contact generalise beyond the immediate situation.

Brewer and Miller (1984; Brewer, 2007) argued that the contact situation should be structured to resemble the interpersonal end of Tajfel's (1981) continuum. That is, the contact situation should minimise the salience of group boundaries, and thus discourage social categorisation (*decategorisation model*). Such interpersonal contact should reduce intergroup bias through two processes. By learning unique information about individual outgroup members, interaction partners begin to draw distinctions among them. This process of *differentiation* erodes the perceptual foundation of social stereotypes: the accentuation of within-category similarities (see Tajfel, 1969). By interacting with outgroup members in terms of their personal relationships to oneself, interaction partners find that these relationships supersede the usefulness of category distinctions. This process of *personalisation* undermines the motivational foundation of intergroup bias: dividing others into "us" and "not us". Brewer and Miller assumed that repeated interpersonal contact with outgroup members should generalise to new situations by permanently eroding the perceived usefulness of social categorisation.

Doubts persist, however, about the feasibility of Brewer and Miller's proposal. Tajfel (1981, Chapter 11), for example, questioned whether an interpersonal interaction is possible, given how central social categorisation is to human cognition.[2] Instead, Gaertner and Dovidio (2000; Gaertner et al., 2016) argued that the goal of contact should be *re*categorisation, not *de*categorisation. That is, the contact situation should be structured to encourage current ingroup and outgroup members to adopt a more inclusive, superordinate identity (*common ingroup identity model*). This could mean, for example, that interaction partners of different ethnicities espouse their shared national identity over their respective ethnic identities. Rather than reduce ingroup favouritism, contact thus expands the ingroup, and extends the benefits of ingroup favouritism to former outgroup members. As the former ingroup–outgroup distinction loses psychological significance, outgroup-directed and intergroup-directed bias similarly lose their motivational foundation. Instead of deemphasising the salience of *all* group boundaries, Gaertner and Dovidio thus proposed emphasising superordinate over less inclusive category distinctions during contact.

Hewstone and Brown (1986; R. Brown & Hewstone, 2005) criticised both approaches. If *interpersonal* contact achieves decategorisation (Brewer & Miller, 1984), participants should associate their positive contact experiences with those immediately encountered—but not with other members of the same outgroup. If *intragroup* contact results in recategorisation (Gaertner & Dovidio, 2000), contact should reduce bias toward outgroup members *inside* the new superordinate category—but not toward those *outside* of it. Instead, Hewstone and Brown argued that category distinctions should remain salient in the contact situation (*intergroup contact model*). If interaction partners see each other as typical or representative of their respective groups, then any positive contact experiences should transfer to other outgroup members. This way, the effects of contact generalise beyond the immediate contact situation. By reducing anxiety, increasing empathy, and deepening knowledge about the outgroup, intergroup contact thus reduces negative evaluations of the entire outgroup (Pettigrew & Tropp, 2006, 2008). Beyond gen-

[2] Tajfel maintained that the intergroup extreme *is* possible, citing extreme forms of dehumanisation as an example.

eralisation, R. Brown and Hewstone questioned whether contact participants would give up valued identities outside laboratory settings (as required for decategorisation and recategorisation). *Intergroup* contact—that is, contact that maintains group salience—does not threaten participants' identities, and is thus more practical.

While addressing some limitations of previous models, Hewstone and Brown's model faces its own challenges. Without disrupting social categorisation, it is unclear how, if at all, intergroup contact could diminish harmful social stereotypes (see Tropp & Pettigrew, 2005). As intergroup contact leaves the "us"–"not us" distinction intact, it is also unclear how, if at all, contact could reduce ingroup favouritism. Various researchers have attempted to integrate the three perspectives to address the limitations of individual models. Pettigrew (1998) proposed a staged model, in which interpersonal contact is followed by intergroup contact, culminating in intragroup contact and recategorisation. R. Brown and Hewstone (2005) argued that interpersonal and intergroup comparisons need not be mutually exclusive, and could both be salient in a contact situation. In both approaches, the ideal contact situation could reap the benefits of personalisation *and* generalisation. Gaertner et al. (2016; Gaertner, Dovidio, Anastasio, Bachman & Rust, 1993) proposed an alternative dual-identity strategy in which the contact situation maintains the salience of the original subordinate *and* the novel superordinate categorisations. Thereby, contact could reap the benefits of recategorisation without threatening valued identities.

Social categorisation aside, the perspectives reviewed above mostly agree on Allport's (1954/1979) optimal conditions for contact (see Brewer, 2007; R. Brown & Hewstone, 2005; Gaertner et al., 2016). In this thesis, however, I focus less on the ideal conditions for intergroup contact—and more on the mechanisms by which contact reduces intergroup bias. Still, the mechanisms proposed by Brewer and Miller (decategorisation) and Gaertner and Dovidio (recategorisation), as well as the concerns raised by R. Brown and Hewstone (generalisation), inform the research in this thesis.

In the remainder of this chapter, I discuss omissions and limitations of previous explanations of how intergroup contact reduces intergroup bias. Specifically, I argue that contact can reduce bias by fostering more fluid and continuous identities along a single dimension of categorisation, and that contact can reduce bias by fostering more inclusive identities across multiple, cross-cutting categories.

Social identification along a single category

In this section, I reconsider decategorisation as a seemingly intuitive explanation of how contact reduces prejudice. I highlight practical and conceptual limitations of this perspective, and introduce a novel approach to conceptualising the effects of contact on social identification along a single axis of categorisation.

As social categorisation is necessary, and often sufficient, for intergroup bias, decategorisation seems an intuitive solution to group-based discrimination. If we do not think of someone as an outgroup member, we cannot discriminate against them based on their group membership. This process can explain how cross-group friendship can persist in an otherwise antagonistic climate. By interacting with an outgroup member on an interpersonal level, we can begin to see that person in terms of their relationship and function to ourself. According to this view, we may learn to see them as a friend—that is, as an individual person rather than a generic exemplar of their social category. Cross-group contact would then have personalised the outgroup friend, and would have

removed them from the outgroup category and its attendant biases.

This example, however, also shows the limits of Brewer and Miller's (1984) approach. As detailed above, interpersonal contact can decategorise outgroup friends, and thus prevent discrimination against *them*. Brewer and Miller proposed that repeated interpersonal contact could altogether erode a social categorisation, and thus reduce discrimination against *all* outgroup members. This proposition about generalisation is less convincing. First, social categorisation is a fast and automatic perceptual process (Macrae & Bodenhausen, 2000). People, for example, perceive another person's race (Richeson & Trawalter, 2005) and gender (Martin & Macrae, 2007) within less than a second of encountering them. We have only limited control over this process—it seems unlikely that we could permanently 'unsee' prominent social categories (see Tajfel, 1981). Second, social categorisation can persist even if individual friends are excluded from the outgroup category. This is reflected in the following sentiment: "I don't like [outgroup]. But my friend [name] is OK—I don't even think of him as [outgroup], he's just [name]!" It seems that people can selectively decategorise individuals, without altogether abandoning a category distinction. While intuitive, decategorisation thus seems an unrealistic solution to intergroup bias.

Brewer and Miller (1984) argued that under the right conditions, cross-group contact changes how we see *others*—how we categorise others in terms of their category memberships. Gaertner and Dovidio (2000) instead argued that under the right conditions, contact changes how we see *ourselves*—how we categorise ourselves in terms of multiple category memberships available to us. This recategorisation process relies on hierarchical group memberships wherein superordinate group identities replace subordinate ones. Superordinate categories, however, can provoke resistance (when they threaten valued identities) and may not be available in all situations. Below, I review research on psychological essentialism, and argue that research has previously overlooked how social identification could change along a single axis of categorisation. I then propose that intergroup contact could reduce bias by fostering more continuous and fluid conceptions of social identification.

Psychological essentialism

> *There is an inherent "Jewishness" in every Jew. The "soul of the Oriental," "Negro blood," Hitler's "Aryanism," "the peculiar genius of America," "the logical Frenchman," "the passionate Latin"—all represent a belief in essence.*
>
> – Gordon W. Allport (1954/1979, p. 174)

Thus far, I have treated social categorisation as an all-or-nothing process which results in *binary* categorisations. We categorise ourselves and others as *either* belonging *or* not belonging to a specific group. We categorise others as *either* outgroup *or* ingroup member—as either "us" or "them"—without any grey area between the two choices. For some categories, this dichotomy seems to accurately represent lay beliefs. Many people would agree, for example, that a person can either be male or female, black or white, but rarely both.[3] For other categories, this conceptualisation seems less accurate. Most people

[3] I do not suggest that this is an accurate representation of the world. Indeed, genderfluid and other non-binary people challenge the idea of gender as a binary category (see Chapter 3), while multiracial people challenge dichotomised conceptions of race (see Ho, Sidanius, Levin & Banaji, 2011). The remainder of this subsection shows that many people nonetheless think in binary terms.

would concede, for example, that a person can be both an amateur actor and a gym goer, both a doctoral student and a teaching assistant, without either precluding the other.

Research on psychological essentialism seeks to understand what differentiates the former categories (gender, race) from the latter (interest, occupation). Medin and Ortony (1989) recognised that superficial similarities are not enough to explain how humans divide their environment into categories. When asked to define a bird, for example, we would likely mention that birds can fly and have feathers. Yet, we accept that penguins and ostriches are birds even though they cannot fly—and would probably concede that a rare featherless bird is still a bird. What makes a bird is not their superficial similarities, but an underlying, unchangeable, and unobserved nature or essence they have in common with other birds. This essence, in turn, gives rise to bird-like features and other similarities within the category. In other words, "twins are not twins *because* they are similar; they *are* similar because they are twins" (Medin & Ortony, 1989, p. 186). This way of thinking about the natural world—essentialism—is disputed in science and philosophy (Bird & Tobin, 2018). *Psychological essentialism* is the proposition that in everyday thinking, humans nonetheless rely on deep, underlying essences to determine category membership (S. A. Gelman, 2003; Medin & Ortony, 1989).

Not all categories are essentialised, at least not to the same extent. S. A. Gelman (2003) found that people readily essentialise categories of 'natural kinds', but less so categories that are 'human artefacts'. *Natural kinds* are categories that people believe to be natural, to be discovered rather than invented, and to allow rich inferences about exemplars of the category. *Human artefacts* are categories that people believe to be artificial, to be invented rather than discovered, and to not allow any inferences beyond the superficial similarities that define the category. What distinguishes natural kinds—such as "bird" or "water"—from human artefacts—such as "chair" or "shirt"—is that people tend to ascribe essences to the former but not the latter. This also means that people believe natural-kind distinctions to be stable (birds have always existed), discrete (a bird is never also a mammal), and unchangeable (a bird will always be a bird).

Social categories—such as race (Richeson & Sommers, 2016) and gender (Butler, 1990/1999)—are social constructs. In other words, they are human artefacts rather than natural kinds. Rothbart and Taylor (1992) hypothesised that people largely misunderstand social groups as natural kinds and ascribe essential qualities to them. To test this hypothesis, Haslam, Rothschild and Ernst (2000) measured lay people's essentialist beliefs about various social categories. Categories varied widely in the extent to which participants characterised them as natural kinds (Figure 1.1, on the left). That is, categories varied in the extent to which participants rated them as *natural*, *stable*, *discrete*, *immutable*, and as defined by *necessary characteristics*.[4] Race and gender were among the most essentialised category domains. Psychological essentialism thus explains why lay people tend to see the differences between male and female, black and white, as discrete and immutable—and why they recognise interests and occupations as continuous and malleable.

Essentialist beliefs shape how we perceive members of relevant groups. Rothbart and Taylor (1992) argued that when we ascribe essences to social categories, we endow

[4]Haslam et al. (2000) found that participants' beliefs formed two factors. Above, I only discuss the 'natural kinds' factor as it most closely reflects the concept of psychological essentialism. I omitted the 'reification' factor—the extent to which categories were *informative*, *inherent*, and *exclusive*—as it aligns more closely with the concept of entitativity (Campbell, 1958). Figure 1.1 only shows results for the 'natural kinds' factor.

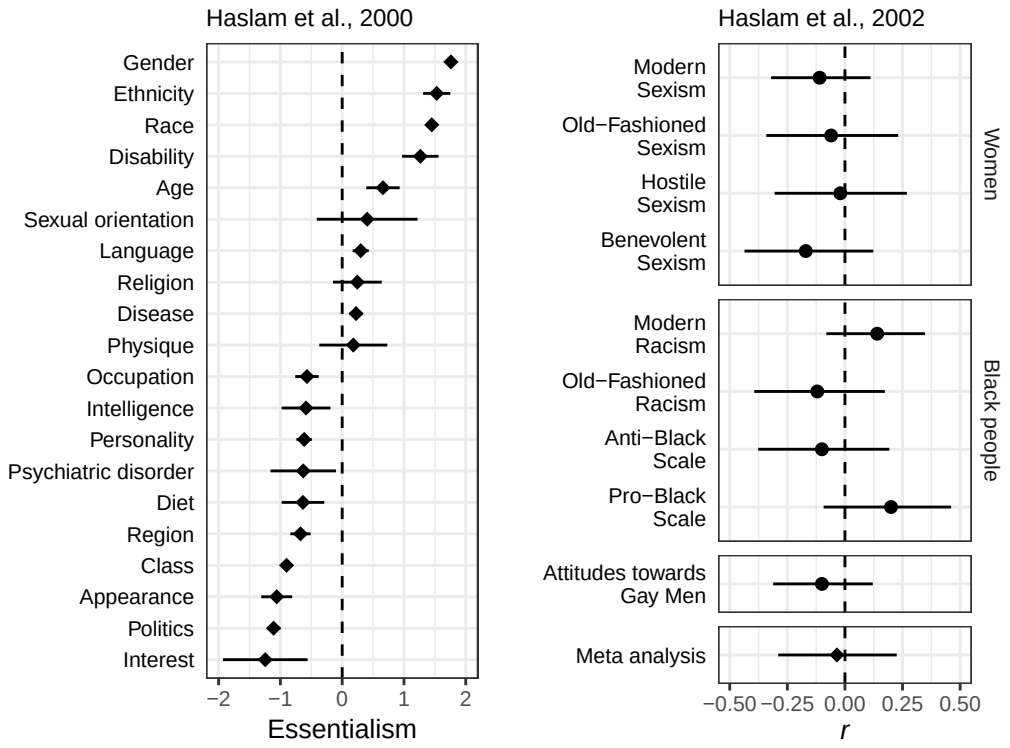

Figure 1.1. **On the left:** Essentialism ratings for each category domain surveyed by Haslam et al. (2000). Participants rated two categories for each domain (e.g., "male" and "female" for gender). Diamonds show the mean ratings across categories; lines show the range between the mean ratings for each category. **On the right:** Correlations between essentialist beliefs (about women, black people, and gay men) and prejudice in Haslam et al. (2002). Lines show confidence intervals; the last row shows a meta-analytic summary (Goh et al., 2016).

these categories with *inductive potential*. That is, we believe that knowing someone's category memberships allows us to infer other characteristics of that person. Likewise, essentialist beliefs define category boundaries as discrete and stable. Rothbart and Taylor argued that essentialist beliefs should thus reinforce the perceptual consequences of social categorisation. That is, essentialist beliefs should *accentuate* within-category similarities and between-category differences (see also Prentice & Miller, 2007). Supporting this hypothesis, Yzerbyt, Corneille and Estrada (2001) found that participants showed greater accentuation when comparing two groups that were defined by different genetic markers, rather than by waiting for a different bus. Together, Rothbart and Taylor's arguments suggest that essentialist beliefs about social categories strengthen social stereotypes.

Thus far, I have discussed *category* differences in psychological essentialism—participants essentialise some categories more than others. Research on the consequences of essentialist beliefs highlights *individual* differences—some participants endorse essentialist beliefs more than others. Bastian and Haslam (2006) found individual differences in essentialist beliefs which, as predicted by Rothbart and Taylor (1992), correlated with differences in stereotype endorsement. Participants who believed that a person's attributes were discrete, informative, and based in biology were more likely to hold stereotypic beliefs about various social categories. Brescoll and LaFrance (2004) exposed participants to biological and sociocultural explanations of gender differences. Participants who had

read about biological explanations endorsed gender stereotypes more than participants who had read about sociocultural explanations. Reimer, Hospers and Kok (2014) exposed participants to scientific evidence for genetic explanations of sexual orientation, against genetic explanations, or a summary of unrelated research. Participants who had read evidence for genetic explanations expected gay men to show more 'feminine' traits (in line with social stereotypes) than participants who had read evidence against genetic explanations. Correlational and experimental research thus support Rothbart and Taylor's prediction that essentialist beliefs increase stereotype endorsement.

As discussed, intergroup bias differs from stereotype endorsement in that the latter relates to the *perceptual* consequences of social categorisation while the former relates to its *motivational* consequences. Haslam et al. (2002) measured essentialist beliefs about various social categories (women, black people, gay men) to test whether participants' beliefs were associated with prejudice against these categories. Other than expected, essentialist beliefs did not correlate with any of the outcome measures (Figure 1.1, on the right). Other than Haslam et al. (2002), Haslam and Levy (2006) found that essentialist beliefs about homosexuality correlated with participants' attitudes toward gay people. The direction of this relationship, however, varied across three dimensions of essentialist beliefs. Participants who believed that homosexuality constituted a "discrete, entitative type with defining features" reported, on average, *more* negative attitudes. Participants who believed that homosexuality was "biologically based, immutable, and fixed early in life" or who believed that it was "cross-culturally and historically universal" reported *less* negative attitudes (p. 471). Haslam and Levy argued that essentialist beliefs reduce sexual prejudice to the extent that they frame sexual orientation as uncontrollable. Attributing homosexuality to stable, uncontrollable causes should elicit pity and a willingness to help rather than derogation (Weiner, Perry & Magnusson, 1988).

Other research, however, contradicts Haslam and Levy's (2006) findings. Hegarty (2002) showed that immutability beliefs correlated with more tolerant attitudes *only* among participants who thought that tolerant people would express such beliefs. Reimer et al. (2014) found that manipulating participants' beliefs about genetic explanations of sexual orientation did not reduce or increase sexual prejudice. These findings align with research on social identification. Reducing (or increasing) essentialist beliefs changes how participants *perceive* social categories. As social identification *motivates* intergroup bias, changing perceptions should not be enough to reduce bias. Reducing essentialist thinking is thus not a promising strategy to reduce intergroup bias—unless it was able to altogether eliminate a social categorisation (see *decategorisation*). Below, I argue that contact could reduce intergroup bias by changing how we understand *our own* category memberships—if contact results in *less discrete* and *immutable* conceptions of our social identities.

Bridging binaries

> *the problem [is] the relation between a set of attributes which vary on continuous dimensions and classifications which are discontinuous*
>
> – Henri Tajfel (1969, p. 177)

Thus far, I have argued that *decategorisation* is an unlikely explanation of how contact reduces bias. Repeated positive interactions with outgroup members will not make us 'unsee' the social categories of others. Further, I have argued that research on *psychological*

essentialism shows that individuals differ in the extent to which they believe various social categories to be discrete and immutable. Essentialist beliefs usually concern perceptions of *abstract* category distinctions rather than *concrete* ingroup–outgroup distinctions. Researchers, for example, measured beliefs about 'heterosexuals' and 'homosexuals' (Haslam et al., 2000, 2002) instead of asking how discrete and immutable participants' *own* sexual orientation was.[5] In other examples, researchers studied category domains in which both categories were (presumably) outgroups for most participants (e.g., 'blind' and 'paraplegic'). Research on decategorisation and psychological essentialism thus relates to how we perceive the social categories *of others*.

I propose that intergroup contact could reduce intergroup bias by changing how we perceive the social categories we ascribe to *ourselves*, our social identities. Past research on *social identification* has shown that people differ in how strongly they identify with a given social category (for a review, see Ashmore, Deaux & McLaughlin-Volpe, 2004). That is, how central a category membership is to their self concept, how much they value that category membership, how committed they are to an ingroup, and how much they define themselves in terms of a group membership (Leach et al., 2008). This research, however, has mostly overlooked how participants define the boundaries of the social categories they identify with, and how malleable they consider their identities to be.[6] Prior research on *recategorisation* has studied whether contact changes how we categorise ourselves in terms of *multiple* subordinate and superordinate categories. Again, this research has mostly considered participants' social identities as binary (either they consider themselves part of a group, or not) or as a matter of strength of identification (see above).

In this thesis, I investigate whether intergroup contact can change how participants understand their social identity along a *single* axis of categorisation. Based on the literature on psychological essentialism, I propose that these changes could fall along two dimensions. *Social identity continuity* is the extent to which individuals perceive an identity as their unique position on a continuum, rather than as belonging to one of several discrete categories. *Social identity fluidity* is the degree to which individuals perceive an identity as flexible and changing, rather than as fixed and stable throughout their lifetime. Someone high in both dimensions would, for example, understand their sexual orientation as a unique combination of same-gender and other-gender attraction, that varies across situations, has changed in the past, and may change in the future (see Figure 1.2). Below, I discuss why intergroup contact could foster more fluid and continuous conceptions of social identity, and why more fluid and continuous conceptions could lead to less intergroup bias.

Intergroup contact could foster fluid and continuous social identities in two ways. First, interacting with diverse others should trigger a process of *cognitive differentiation* (see Brewer & Miller, 1984; Schmid & Hewstone, 2011). Interacting with persons of other sexual orientations, for example, should make heterosexual people more aware of the various ways in which people can experience their sexuality, and in which they choose to identify. This awareness, in turn, might prompt heterosexual people to adopt a more nuanced conception of their own sexual orientation. Second, Page-Gould, Mendoza-Denton, Alegre and Siy (2010) demonstrated that intergroup interactions can drive people to include the outgroup of an interaction partner in their self-concept. As far as contact

[5] Some studies measured essentialist beliefs in even more abstract terms, asking about 'personal attributes' in general (Bastian & Haslam, 2006).

[6] Questions about the nature of participants' self-categorisations usually do not go beyond binary yes–no choices, or selecting one of several categories in a multiple-choice question (Ashmore et al., 2004).

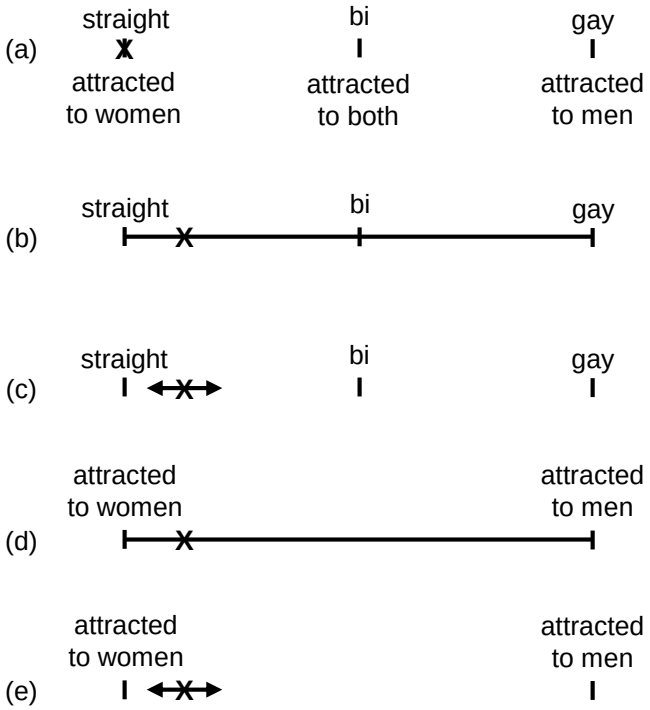

Figure 1.2. Schematic representations of different conceptions of sexuality (for a man). Row (a) shows a binary, stable understanding of sexuality; (b) and (d) show more continuous conceptions; (c) and (e) show more fluid conceptions.

motivates people to *include the outgroup in the self*, it should encourage less discrete and definite boundaries between the self and the outgroup. Together, these mechanisms suggest that intergroup contact should foster more fluid and continuous social identities.

In turn, more fluid and continuous social identities could reduce intergroup bias. As reviewed earlier, dividing others into "us" and "not us" is a necessary and sufficient condition for ingroup favouritism. If contact weakens the boundaries between groups, it should lead to a wider sense of who is "us", and thus extend ingroup favouritism to a wider group of people. Dividing others into "us" and "them" leads to outgroup derogation when an outgroup ("they") threatens the ingroup ("us"). More fluid and continuous social identities should encourage people to see outgroup members as less distinct and more familiar—and thus as less threatening. In other cases, intergroup bias arises from efforts to distance oneself ("me") from a stigmatised outgroup ("them"). More fluid and continuous identities reduce the psychological distance between self and outgroup, and should thus reduce the need to distance oneself from that outgroup.

In Chapter 3, I test these hypotheses in the realms of sexuality and gender. In Studies 1 to 3, I develop and validate a novel measure of social identity fluidity/continuity. I examine whether intergroup contact increases social identity fluidity/continuity, and whether social identity fluidity/continuity reduces intergroup bias (*bridging-binaries hypothesis*). If intergroup contact fosters more fluid and continuous identities, and if more fluid and continuous identities reduce intergroup bias, then social identity fluidity/continuity will be a plausible mediator of the relationship between contact and bias.

Social identification across multiple categories

Thus far, I have discussed social identification within a *single* category domain. I have argued that individuals differ in how discrete and immutable they understand their category membership within that domain to be. I have hypothesised that intergroup contact should foster more continuous and fluid identities, and thus reduce intergroup bias.

As humans, however, we live in complex societies in which we think of ourselves and others in terms of *multiple* category domains. Depending on the situation, we categorise ourselves and others in terms of gender, race, religion, nationality, occupation, interest, or any other salient category domain. Each of us holds a unique combination of overlapping group memberships. Theoretical perspectives (Tajfel, 1981; J. C. Turner et al., 1987) recognise this multiplicity of social identities. Empirical studies, however, often focus on one social identity at a time, ignoring other identities participants simultaneously ascribe to themselves (see Ashmore et al., 2004). This approach is sufficient in situations where only one category domain is relevant or salient.

In other situations, multiple identities shape participants' thoughts and behaviours. Immigration, for example, confronts both immigrants and natives with multiple national, cultural, and religious identities. Berry (1997) recognised that immigrants have to navigate two cultures—that from which they have originated, and that in which they have settled. Berry described four adaptation strategies that map how immigrants negotiate their ethnic and national identities (see Dovidio, Gaertner & Saguy, 2009): Immigrants maintain their ethnic identity but reject the national identity (separation); immigrants abandon their ethnic identity but adopt the national identity (assimilation); immigrants reject both identities (marginalisation); or immigrants maintain their cultural identity *and* adopt the national identity (integration). Berry's model thus describes the relationship between two social identities as *orthogonal* (see van Dommelen, Schmid, Hewstone, Gonsalkorale & Brewer, 2015). That is, both identity dimensions vary independently across individuals.

Like Berry (1997), the common ingroup identity model describes the relationship between two social identities (Gaertner et al., 2016). This model, however, describes relationships between *hierarchical* identities where a superordinate identity encompasses subordinate identities. Both immigrants and natives can identify with their ethnic group (e.g., White or Asian) and/or with the national group (e.g., British). The latter is more inclusive as it encompasses more residents of the country (e.g., all British citizens). Gaertner et al. (2016) proposed two ways in which a superordinate identity can relate to someone's subordinate identity. In the *one-group* scenario, individuals adopt the superordinate identity and abandon the subordinate identity. In the *dual-identity* scenario, individuals adopt the superordinate identity but also maintain the subordinate identity. Gaertner et al.'s model thus describes the relationship between two social identity dimensions as either *negative* or *orthogonal*.

Both Berry (1997) and Gaertner et al. (2016) studied social identification along *two* dimensions. Both models reflect today's multicultural societies better than earlier *one*-dimensional approaches. Yet, both fall short in other ways. First, two dimensions are often not enough to understand a complex social situation. Political allegiances in India, for example, often fall along the fault lines of caste, religion, *and* nationality. Second, social identities vary in inclusiveness—but are seldom hierarchical. Indian Muslims, for example, form the intersection of two more inclusive groups, Indians and Muslims. Each group encompasses more than one billion people, so neither is strictly superordinate.

Third, relationships between social identities are often neither orthogonal nor negative. Instead, individuals (and groups) construe these relationships in more complex ways. A Hindu nationalist, for example, may see being Hindu and being Indian as *converging*, and see being Muslim and being Indian as *conflicting*. Someone else may see being Hindu/Muslim and being Indian as entirely independent.

Berry's and Gaertner et al.'s models thus fail to represent situations in which *three or more* identities are relevant, in which identities are *cross-cutting* rather than hierarchical, and in which identities *converge* and *conflict* in complex patters. Below, I review *crossed categorisation* and *social identity complexity* as approaches that seek to address some of these limitations.

Crossed categorisation

Research on *crossed categorisation* examines how perceiving the social categories *of others* influences intergroup bias (for reviews, see Crisp, 2010; Crisp & Hewstone, 2007). This research recognises that every person we meet holds multiple cross-cutting identities. This means that we often meet people with whom we share some group memberships (e.g., gender) but not others (e.g., nationality). This is more likely when category dimensions are uncorrelated (like gender and nationality). When two such category dimensions are salient, we can meet people who share *both* category memberships with us (double ingroup), share *either* of the category memberships with us (mixed groups), or share *neither* category membership with us (double outgroup). Deschamps and Doise (1978) studied how crossed categorisation affects evaluations of others. To that end, participants rated targets from each of four combinations of gender and two minimal groups (Experiment 2). In this and similar studies, participants tended to evaluate double ingroups more positively than mixed groups, whom participants evaluated more positively than double outgroups (Migdal, Hewstone & Mullen, 1998). This corresponds to an additive pattern of evaluation where more shared group memberships lead to more positive evaluations.

Researchers have, however, observed other patterns of evaluation (Crisp & Hewstone, 2007). In some situations, a single category dimension dominates target evaluations. That is, all targets that share the dominant category membership are evaluated more positively than all targets that lack this category membership. In other situations, observers may evaluate all targets that share at least one category membership equally positively (social inclusion). Alternatively, observers may evaluate all targets that do not share both category memberships as equally negative (social exclusion). Other situations lead to patters of evaluation in which the effects of one category dimension depend on whether targets share another, more important category membership with the observer. For example, a target's gender could affect that target's evaluation only if that target shares the participant's nationality (hierarchical acceptance), or only if that target does not share the participant's nationality (hierarchical rejection). Finally, observers may evaluate all target groups as equally positive (equivalence). Crisp and Hewstone (1999) found that although the additive pattern was the most common result of crossed-categorisation studies, all the other patterns occurred in particular social circumstances.

Research on crossed categorisation offers a paradigm for studying multiple, cross-cutting categorisations. This perspective, however, confines itself to studying the social categorisation *of others*. That is, researchers assume that participants' categorisations *of themselves* match their objective group memberships, and that participants perceive mixed-group members as partial ingroup members. Roccas and Brewer (2002), however, argued

that whether we categorise others as ingroup members depends on how we represent our own group memberships. The researchers proposed that individuals vary in *social identity complexity*—that is, the extent to which they represent their group memberships as largely overlapping and highly similar (low complexity) or as non-overlapping and dissimilar (high complexity). Low complexity means that individuals perceive their group memberships as converging into a single ingroup—and that they regard only people who share most or all their group memberships as fellow ingroup members. High complexity means that individuals recognise their group memberships as forming various smaller intersections—and that they consider a wider range of people as ingroup members. Social identity complexity could reduce intergroup bias by extending ingroup favouritism to a wider group of people.

Roccas and Brewer (2002) operationalised social identity complexity in two ways. *Overlap complexity* measures the perceived numerical overlap between two of the participant's ingroups (e.g., "How many Indians do you think are Hindus?", 0–100%, reverse scored). *Similarity complexity* measures the perceived similarity in meaning and prototypes of the participant's ingroups (e.g., "How similar or different are the typical Hindu and the typical Indian to each other?", 1 = *very different*, 5 = *very similar*, reverse scored). As hypothesised, these measures correlated with less intergroup bias in various settings (Brewer & Pierce, 2005; Schmid, Hewstone & Al Ramiah, 2013; Schmid et al., 2009). These measures indicate to what extent participants are *aware* of some, but not all, of their multiple group memberships. They do not, however, measure whether participants also *identify* with these partial ingroups.

Van Dommelen (2014; van Dommelen et al., 2015) developed the *triple crossed-categorisation task* (TCCT) to address the shortcomings of prior research on crossed categorisation and social identity complexity. In this task, participants view stimuli that each resemble an identity card and communicate a fictional person's group memberships. Group memberships cross three category dimensions so that each target person shares none, one, two, or all three of the relevant group memberships with a participant.[7] For each target, participants then decide whether they consider this person "us" (ingroup) or "not us" (outgroup). Responses indicate how participants construe their ingroup from multiple, cross-cutting group memberships. This task thus combines the theoretical insights of social identity complexity with the methodological insights of the crossed categorisation paradigm. That is, the TCCT studies subjective representations of multiple social identities by using fictional targets with cross-cutting group memberships; it incorporates more than two category dimensions, addressing a limitation shared by all previous approaches.[8]

Van Dommelen et al. (2015) analysed participants' responses quantitatively, as *Social Identity Inclusiveness* (SII), and qualitatively, as *Social Identity Structure* (SIS). They operationalised SII as the total number of targets a participant had categorised as "us" in the TCCT. They coded SIS according to structures proposed by Roccas and Brewer (2002). Figure 1.3 shows schematic representations of various social identity structures. An Indian Hindu, for example, may consider *only* people who share their nationality, religion, *and* caste as ingroup members (*intersection*). Someone else may consider *all* their fellow Indians, whatever their religion or caste, as ingroup members (*dominance*).

[7] Van Dommelen et al. (2015) included targets for each of the $2 \times 2 \times 2 = 8$ ingroup–outgroup combinations. They tested 24 targets: 6 for each the triple-ingroup and triple-outgroup combinations and 2 for each of the 3 single-ingroup and 3 double-ingroup combinations.

[8] Van Dommelen et al. 2014 crossed three categories. Researchers can extend the TCCT to more dimensions, though the number of targets required grows exponentially.

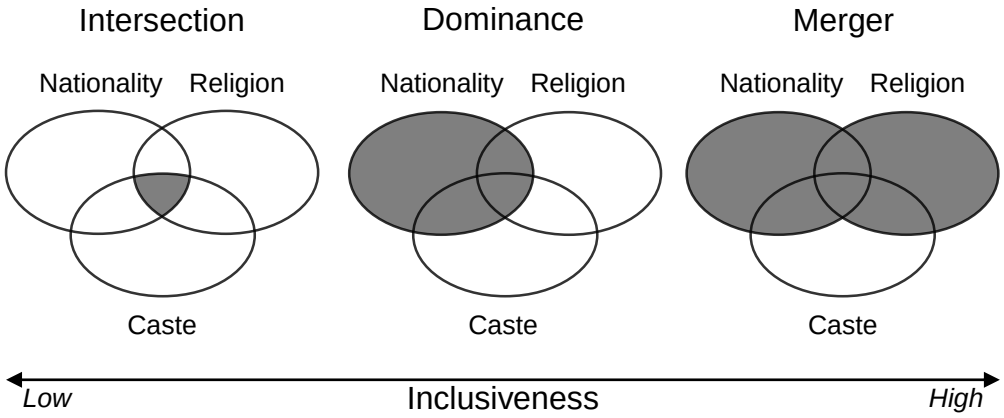

Figure 1.3. Schematic representations of social identity structures. Shaded regions represent groups considered "us" in each social identity structure.

Another person may consider *anyone* who shares their nationality *or* religion as ingroup members (*merger*). These structures vary in inclusiveness. A person who represents their social identity as a merger of nationality and religion is higher in SII than someone who represents their social identity as the exclusive intersection between their nationality, religion, and caste.

In two studies, van Dommelen et al. (2015) studied SII and SIC among Turkish-Belgian Muslims (Study 1) and Turkish-Australian Muslims (Study 2). Participants viewed 24 targets, representing each of $2 \times 2 \times 2 = 8$ combinations of nationality, religion, and ethnicity. On average, participants categorised $M_{S_1} = 14.06$ (59%) and $M_{S_1} = 15.61$ (65%) of 24 targets as "us". Participants endorsed a wide range of social identity structures, including intersection, dominance, merger, and egalitarian structures. Accordingly, the researchers observed substantial *individual differences* in social identity inclusiveness ($SD_{S_1} = 2.97$, $SD_{S_2} = 4.64$). As hypothesised, these differences correlated with differences in intergroup bias, so that more inclusive identities correlated with more favourable attitudes toward triple, double, and single outgroup members. In turn, contact with ethnic and religious outgroup members was associated with more social identity inclusiveness (in Study 2, but not Study 1). These findings suggest social identity inclusiveness as a potential *mediator* of the relationship between contact and intergroup bias. Below, I describe how I adapt the TCCT to study how intergroup contact relates to *individual differences* and *group differences* in social identity inclusiveness.

Bridging divides

Thus far, I have discussed past approaches that studied how contact relates to social identification across multiple category dimensions. I have argued that Gaertner et al.'s (2016) common ingroup identity model fails to represent situations in which three or more identities are relevant, in which identities are cross-cutting rather than hierarchical, and in which identities converge and conflict in complex patters. I have then introduced van Dommelen et al.'s (2015) triple crossed-categorisation task as a paradigm to study how individuals construct their social identities from multiple, cross-cutting group

memberships.

As described above, van Dommelen et al. (2015) studied *individual differences* in how many targets participants included in their ingroup. The researchers recruited homogeneous samples in which all participants shared the same group memberships. Participants' group memberships, however, should influence whom they consider "us" and "not us". For a Turkish-Belgian Muslim, for example, social identity structures could reflect how they negotiate their ethnic and national identities (see Berry, 1997). For a White-Belgian Christian, social identity structures could instead reflect whether they hold nativist or multicultural views. These *group differences* arise from differences in social status and power. As individuals within groups share these experiences, I expect that group differences are at least as great as, if not greater than, individual differences in social identity structure—especially when groups differ in status and power.

Shared experiences, however, make only part of an individual's day-to-day experiences. Of these, contact experiences could overcome group differences in social identity structure. In particular, intergroup contact could foster more inclusive social identity structures in two ways. First, encountering diverse others should start a process of *cognitive differentiation* (Schmid & Hewstone, 2011) whereby interaction partners become aware of the complex interrelations of their cross-cutting group memberships (see Roccas & Brewer, 2002). This awareness, in turn, should foster more inclusive representations of social identity structure. Second, as noted previously, Page-Gould et al. (2010) found that intergroup interactions can drive people to include the outgroup of an interaction partner in their self-concept. As far as contact motivates people to *include (partial) outgroups in the self*, it should encourage ingroup construals that include the interaction partner's combination of group memberships. Together, these mechanisms suggest that intergroup contact could overcome group differences, and foster more inclusive social identities.

In turn, more inclusive social identities should reduce intergroup bias. As argued before, dividing others into "us" and "not us" is a necessary and sufficient condition for ingroup favouritism. If contact leads people to consider more people as "us", it should also extend ingroup favouritism to a wider range of people, and thus reduce discrimination (see Gaertner et al., 2016). Ingroup favouritism can turn into outgroup derogation (Brewer, 1999) when an outgroup ("they") threatens the ingroup ("us"). If contact leads people to regard fewer people as "not us", it should reduce the number of groups ("they") which could be perceived to threaten the ingroup. Together, these mechanisms suggest that intergroup contact could reduce intergroup bias by fostering more inclusive social identities.

In Chapter 4, I test whether contact with ethnoreligious and caste outgroups fosters more inclusive social identities among South Indian students, and whether more inclusive identities reduce intergroup bias (*bridging-divides hypothesis*). I adapt the TCCT to the novel setting, and address some methodological limitations of van Dommelen et al.'s paradigm. First, I recruit participants from multiple caste backgrounds. This means that I can study *group differences* as well as *individual differences* in participants' ingroup construals. Second, I do not aggregate participants' "us"–"not us" categorisations into a single SII score. Instead, I analyse response probabilities for each target using *multilevel models*. Finally, I do not aggregate participants' contact experiences and outgroup attitudes. Instead, I test how contact with a *specific* outgroup relates to participants' categorisations of and attitudes toward this outgroup. Together, these changes allow more fine-grained analyses of how contact relates to social identity inclusiveness and structure, and how these outcomes, in turn, relate to intergroup bias.

Social identification, intergroup contact, and social change

Thus far, I have discussed social identification as a necessary, and sometimes sufficient, condition for intergroup bias. I have proposed that intergroup contact can reduce intergroup bias by fostering more fluid, continuous, and inclusive social identities. This hypothesis focuses on intergroup contact and social identification as means to reduce intergroup bias. I have defined intergroup bias as *differential evaluation* (see Hewstone et al., 2002) and as *differential treatment* (see Tajfel et al., 1971)—that is, as the extent to which individuals favour ingroup members in their attitudes (prejudice) and their behaviour (discrimination).[9] The proposed research thus examines how, if at all, contact with outgroup members improves individuals' attitudes and behaviour toward that outgroup.

Most research on intergroup contact studies relations between groups of unequal status and power. That is, some groups (*advantaged groups*) have more power and resources at their disposal than others (*disadvantaged groups*). As such, advantaged-group members are often in positions of power that allow them to enact their prejudices. If they do, this discrimination hurts disadvantaged-group members, and thus maintains or exacerbates existing inequalities.[10] Reducing prejudice undermines the motivation of advantaged-group members to discriminate against disadvantaged-group members. Reducing discrimination in intergroup interactions should, in the long run, reduce intergroup inequality. Proponents of *prejudice reduction*—such as intergroup contact—thus seek to redress social inequalities by changing how the advantaged think and feel about the disadvantaged group.

Rather than focus on the "hearts and minds" of the advantaged group, collective action researchers study the conditions under which disadvantaged-group members take action to improve their group's position in society (Wright, 2010). Central to this perspective is that discrimination is the product of historical and structural inequalities rather than individual minds (Blumer, 1958). That is, discrepancies result from historical injustices against disadvantaged groups that have compounded into cumulative inequalities.[11] Discrepancies are maintained by institutions that systematically disadvantage members of already disadvantaged groups (Pratto, Sidanius & Levin, 2006).[12] Further, this perspective notes that advantaged groups are motivated to protect rather than challenge social inequalities (Blumer, 1958). Instead, social change often begins with disadvantaged-group members challenging the status quo through mass mobilization and political action.[13] Proponents of *collective action* thus seek social change by mobilizing the disadvantaged against structural injustice.

---

[9] Prejudice typically correlates with discrimination, though this association can be small (e.g., r = .26, [.23, .30]; Talaska, Fiske & Chaiken, 2008).

[10] Operario and Fiske (1998) argue that oppression results from power *plus* prejudice. I would argue that oppression results from prejudice *times* power—prejudice is not sufficient for discrimination as power is necessary for prejudice to result in oppression.

[11] In the United States, racialised slavery left former slaves with little or no wealth to pass on to their descendants. As wealth compounds over time, Black Americans are still a lot less wealthy than White Americans (M. K. Brown et al., 2003). Indeed, disparities in wealth and income have changed little in the last 40–60 years (Kraus, Rucker & Richeson, 2017).

[12] In the United States, the "War on Drugs" has disproportionately targeted people of colour, leaving many Black Americans incarcerated and disenfranchised (Alexander, 2012).

[13] "In all three examples [the end of apartheid in South Africa, civil rights reforms in the United States, and the abolition of New World slavery], then, equality was won rather than given away. In all three, change was the result of sustained collective resistance rather than some kind of general improvement, whether incremental or dramatic, in intergroup attitudes." (Dixon, Levine, Reicher & Durrheim, 2012, pp. 418–419).

Thus far, I have proposed research concerning how intergroup contact and social identification relate to intergroup bias in attitudes and behaviour (*prejudice reduction*). Beyond bias, however, social identification is a also necessary condition for *collective action* (Tajfel, 1981; van Zomeren, Postmes & Spears, 2008). If contact fosters more fluid, continuous, and inclusive social identities, it could influence whether individuals engage in collective action. Below, I discuss possible implications of intergroup contact for collective action in *advantaged* and *disadvantaged* groups.

A darker side of 'we'?

Demobilising effects of positive contact

In principle, prejudice reduction and collective action could coexist as *complementary* routes to social change. On the one hand, contact with disadvantaged-group members could reduce prejudice among advantaged-group members. Improving attitudes of those who hold social power (e.g., over employment or law) should lessen discrimination, both by influencing behaviour (e.g., hiring decisions) and by weakening resistance to legislation (e.g., affirmative action). On the other hand, disadvantaged-group members could engage in collective action to advance legislation that would reduce structural inequalities. Prejudice reduction could weaken advantaged-group members' resistance to collective action by disadvantaged-group members, and could thus complement the disadvantaged group's struggle against social injustice.

Some researchers argued that, in practice, the two routes to social change are *conflicting* rather than complementary. Researchers argued that in historically unequal societies, intergroup contact could entrench, rather than reduce, social inequality by stifling disadvantaged-group members' support for social change (Dixon et al., 2012; Wright & Baray, 2012). Specifically, researchers proposed three unintended consequences of intergroup contact First, friendly interactions with advantaged-group members could prompt disadvantaged-group members to doubt the structural *discrimination* they, in fact, face. This could undermine the perceived need for social change. Dixon et al. (2010) found Black South Africans' perceptions of personal and group discrimination to be negatively correlated with positive contact. Second, researchers found contact with advantaged-group members to correlate with less *support for policies* that would benefit the disadvantaged ingroup (Cakal, Hewstone, Schwär & Heath, 2011; Dixon, Durrheim & Tredoux, 2007; Sengupta & Sibley, 2013). Third, intergroup contact could reduce disadvantaged-group members' willingness to engage in *collective action*. Several studies found contact with advantaged-group members to correlate with less collective action in disadvantaged groups (Cakal et al., 2011; Tropp, Hawi, Van Laar & Levin, 2012). Intergroup contact could thus maintain an unjust status quo by reducing disadvantaged-group members' perceptions of discrimination, support for redistributive policies, and participation in collective action.

Researchers proposed three mechanisms that could explain *how* or *why* intergroup contact weakens disadvantaged-group members' support for social change. Contact with advantaged-group members could change how disadvantaged-group members think and feel *about others*. That is, if contact fosters more favourable attitudes toward the advantaged (Tropp & Pettigrew, 2005), it could make discrimination seem less plausible (Dixon et al., 2010) and reduce support for actions and policies that stand to antagonise the advantaged (Dixon et al., 2012). Contact with advantaged-group members could also change how disadvantaged-group members think and feel *about the system*. That is, intergroup contact

could promote system-legitimising ideologies, and thus quell support for redistributive policies (Sengupta & Sibley, 2013). Contact with advantaged-group members' could also change how disadvantaged-group members think and feel *about themselves*. Below, I further discuss this mechanism—that is, how contact could reduce disadvantaged-group members' support for social change by changing how they understand their social identities.

In previous sections, I have proposed that contact could reduce intergroup bias by fostering more fluid, continuous, and inclusive social identities. Social identification motivates individuals to "protect, maintain, or enhance their *group* interests [rather than] their *personal* interests" (van Zomeren, 2013, p. 380, italics in the original). Under specific circumstances, disadvantaged-group members pursue *collective* rather than *individual* strategies against perceived disadvantages (Tajfel, 1981). Individuals engage in collective action, rather than hope for individual social mobility, when they understand group boundaries as stable and impermeable, regard group differences as illegitimate, and believe that group differences can be changed. If the hypothesised effects of intergroup contact apply to disadvantaged groups, they could have the unintended consequence of undermining the conditions under which disadvantaged-group members support social change. I consider two reasons why these unintended consequences may occur.

First, I have hypothesised that intergroup contact could "blur the boundaries" of binary ingroup–outgroup distinctions (bridging-binaries hypothesis). If stable group boundaries are a necessary condition for collective action, contact could undermine collective action by fostering more fluid and continuous conceptions of social identity. Second, I have hypothesised that intergroup contact could extend individuals' ingroup construals to include former outgroup members (bridging-divides hypothesis). Dovidio et al. (2012) argued that more inclusive identities can have a "darker side" for disadvantaged groups. Some situations encourage individuals to abandon a subordinate identity in favour of a superordinate identity (*one-group identity*). Espousing a superordinate identity, which includes the advantaged outgroup, should reduce how central the disadvantaged ingroup is to disadvantaged-group members' self-concept, and how invested they are in the fate of their disadvantaged ingroup. As a result, disadvantaged-group members should be less outraged about the relative disadvantage of their group (van Zomeren, 2013), and thus pursue individual mobility instead of collective action. Further, a *one-group* representation deemphasises distinctions between the disadvantaged ingroup and the advantaged outgroup. As a result, more inclusive identities could distract disadvantaged-group members from group differences in status and power, and thus remove the perceived need for social change.

Research supported these predictions. Ufkes, Calcagno, Glasford and Dovidio (2016) found that emphasising a *common ingroup* decreased group-based anger and collective action intentions (relative to *separate-group* and *dual-identity* conditions) for disadvantaged participants. Saguy, Tausch, Dovidio and Pratto (2009) showed that *commonality*-focused contact (relative to *difference*-focused contact) reduced expectations of unfair treatment. More broadly, van Zomeren et al.'s (2008) meta-analysis found that strength of social identification (that is, how central a group membership was to a participant's self concept, and how invested participants were in a group membership) correlated with more collective action. Strength of social identification also correlated with greater perceived injustice which, in turn, correlated with more collective action. *If*, as hypothesised, intergroup contact fosters more fluid, continuous, and inclusive social identities, *and if* more fluid, continuous, and inclusive social identities reduce identific-

ation with the disadvantaged ingroup, *then* contact could reduce disadvantaged-group members' support for social change (*demobilisation hypothesis*).

Mobilising effects of negative contact

Most research on intergroup contact, including on its demobilizing effect, has studied the effects of positive contact. Recent research, however, has highlighted the importance of also studying the effects of negative contact to understand the overall impact of cross-group interactions. Stephan et al. (2002) found negative interactions, such as being insulted by an outgroup member, to be, not surprisingly, associated with more prejudice. Although less frequent than positive contact, some researchers found the prejudice-increasing effect of negative contact to be stronger than the prejudice-reducing effect of positive contact (Barlow et al., 2012; Graf, Paolini & Rubin, 2014). Importantly, while both dimensions of contact uniquely predicted attitudes, they are, if at all, weakly or moderately correlated ($r = -.26$, Barlow et al., 2012; $r = -.06$, Dhont & Van Hiel, 2009; $r = -.16$, Heitmeyer, 2005 in Pettigrew & Tropp, 2011). Overall, recent research has underlined the importance of studying negative and positive contact as related but separate dimensions of intergroup contact.

As disadvantaged-group members have been found to experience more negative contact than advantaged-group members (Stephan et al., 2002), the potential mobilising effect of negative contact on collective action should be considered alongside the demobilising effect of positive contact. Negative contact raises the salience of group membership (Paolini, Harwood & Rubin, 2010) and could thus encourage disadvantaged-group members to think and feel as group members rather than individuals. Similarly, negative contact could, by increasing group salience, evoke intergroup comparisons and facilitate the perception of discrimination (see Wright, 2010), especially if these experiences are interpreted as manifestations of discrimination. Relatedly, aversive contact experiences could threaten one's self-worth. Disadvantaged-group members could be motivated to attribute negative contact to discrimination to protect their self-esteem (Major, Kaiser & McCoy, 2003). Unfair treatment, in turn, should incite group-based anger (van Zomeren, Spears, Fischer & Leach, 2004). Negative contact could hence encourage collective action by fostering identification, highlighting discrimination, and inciting anger.

Finally, as reviewed above, negative contact is associated with less favourable outgroup attitudes. If negative characterizations of advantaged outgroups benefit collective action, negative contact could further mobilise disadvantaged-group members. Although positive contact could, as reviewed above, prevent collective action by reducing social identification, negative contact could instead foster collective action (*mobilisation hypothesis*). Thus far, I have discussed the implications of positive and negative contact in *disadvantaged groups*. Below, I discuss possible consequences of intergroup contact for members of *advantaged groups*.

A path to solidarity?

Collective action approaches focus on disadvantaged-group members as agents of social change. Advantaged-group members, however, sometimes join the struggle of the disadvantaged, for example, as heterosexual allies in the LGBT (lesbian, gay, bisexual, and transgender) movement. Both qualitative (e.g., Russell, 2011) and quantitative (e.g., Fingerhut, 2011) evidence suggests that positive contact often motivates heterosexual

allies' involvement in the LGBT movement, but prior research has not examined *how* positive contact increases advantaged-group members' participation in collective action.

Intergroup contact provides an opportunity for social comparison for both disadvantaged and advantaged-group members. In this comparison, advantaged-group members should experience a sense of relative advantage rather than relative deprivation (Leach, Snider & Iyer, 2007). Relatedly, positive contact promotes perspective-taking and self-disclosure (Pettigrew & Tropp, 2008) and could, thus, sensitize advantaged-group members to structural discrimination. Furthermore, positive contact improves outgroup attitudes (Pettigrew & Tropp, 2006), and these positive feelings toward a disadvantaged outgroup are, in turn, associated with support for egalitarian policies (Pittinsky & Montoya, 2009). Positive contact could hence encourage advantaged-group members to engage in collective action by highlighting discrimination and improving attitudes.

As discussed before, social identification motivates actions aimed at advancing the interests of one's ingroup. Blumer (1958) argued that identification with an advantaged ingroup implies *identification* with one's privileged group position and *resistance* to any challenge to that group position. Cakal et al. (2011) indeed found that White South Africans' strength of ethnic identification correlated positively with their motivation to improve conditions for their advantaged ingroup, and negatively with their support for egalitarian policies that would help Black South Africans. In historically unequal societies, advantaged-group members have to reckon with their relative advantage over less privileged groups. Recognising their advantage as unfair could threaten advantaged-group members group esteem (Wohl, Branscombe & Klar, 2006). If advantaged-group members want to maintain the status quo *and* protect their group image from the threat of unjust privilege, they are motivated to deny the existence (or the immorality) of the ingroup's advantage (Knowles, Lowery, Chow & Unzueta, 2014). This strategy, however, is difficult to maintain in the face of obvious injustices.

Another strategy of avoiding the threat of unjust privilege is to exclude disadvantaged outgroups from the ingroup's circle of moral concern. Singer (1981/2011) argued that humans have a *moral circle* that centres around the self and encompasses all beings whom we consider worthy of moral regard. Individuals differ with respect to whom they include in their moral circle. Some are only concerned for their immediate family, while others regard all human beings, or even all living things, as worthy of moral concern. Psychological factors shape whether individuals extend their moral circle to include distant others, or narrow their moral circle to include only close relatives (Graham, Waytz, Meindl, Iyer & Young, 2017). Whom individuals consider "us" and "them" often shapes whom they include in their moral circle. If, as hypothesised, positive contact fosters more fluid, continuous, and inclusive identities, it breaks down the boundaries between the advantaged ingroup and the disadvantaged outgroup. Positive contact could thus force advantaged-group members to include the disadvantaged group in their moral circle—making injustice faced by "them" into "our" problem. Intergroup contact could thus foster solidarity with the disadvantaged group's struggle for social justice.

In Chapter 5, I study the implications of positive and negative contact for social change in advantaged and disadvantaged groups. In Studies 1–4, I recruited members of advantaged majorities and of disadvantaged minorities. For *advantaged-group* participants, I test whether positive contact with disadvantaged groups fosters solidarity-based collective action and support for egalitarian policies (*solidarity hypothesis*). For *disadvantaged-group* participants, I test whether positive contact decreases collective action intentions and support for egalitarian policies (*demobilisation hypothesis*), and whether negative contact

increases collective action intentions and support for egalitarian policies (*mobilisation hypothesis*). I test to what extent more continuous, fluid, and inclusive identities mediate the hypothesised effects.

Overview of thesis structure

In this chapter, I have reviewed relevant literature, and have introduced research questions and hypotheses. First, I have discussed whether intergroup contact could reduce intergroup bias by changing how we understand our social identities. Specifically, I have hypothesised that intergroup contact reduces bias by fostering more continuous and fluid identities along a single category (*bridging-binaries hypothesis*), and by fostering more inclusive identities across multiple categories (*bridging-divides hypothesis*). Second, I have discussed whether intergroup contact benefits or impedes struggles for social justice. For disadvantaged groups, I have hypothesised that positive contact discourages support for social change (*demobilisation hypothesis*), while negative contact encourages support for social change (*mobilisation hypothesis*). For advantaged groups, I have hypothesised that positive contact encourages support for social change (*solidarity hypothesis*). In Chapters 3–5, I describe and discuss findings from four empirical studies to test the hypotheses proposed above.

In Chapter 3, I test the bridging-binaries hypothesis in the realms of sexuality and gender. In Studies 1–3, I develop and validate a novel measure of social identity continuity and fluidity. I examine whether intergroup contact fosters more fluid and continuous social identities and whether, in turn, more fluid and continuous social identities reduce intergroup bias. In Chapter 4, I test the bridging-divides hypothesis in the context of caste, religion, and nation in South India. In Study 4, I adapt the triple crossed-categorisation task (van Dommelen et al., 2015) to study whom participants consider "us" and "not us". I examine whether contact with ethnoreligions and caste outgroups fosters more inclusive social identities, and whether, in turn, more inclusive identities reduce intergroup bias.

In Chapter 5, I test the demobilisation, mobilisation, and solidarity hypotheses. In Studies 1–4, I examine the implications of positive and negative contact for social change in advantaged and disadvantaged groups. For advantaged-group participants, I test whether positive contact with disadvantaged groups fosters solidarity-based collective action and support for egalitarian policies. For disadvantaged-group participants, I test whether positive contact decreases support for social change, and whether negative contact increases support for social change. For both groups, I test to what extent more continuous, fluid, and inclusive identities mediate the effects of intergroup contact on social change.

In the next chapter, Chapter 2, I introduce and discuss methods necessary to understand the empirical chapters thereafter. I focus on Bayesian data analysis as these techniques are not yet common practice in psychological research. After describing and discussing the empirical studies, I then summarise the main findings in Chapter 6. I discuss limitations and implications of the current research, and highlight avenues for future research.

2

Method

In Chapter 1, I have introduced five hypotheses. I hypothesise that *i*. intergroup contact reduces bias by fostering more continuous and fluid identities along a single category (*bridging-binaries hypothesis*), and that *ii*. intergroup contact reduces bias by fostering more inclusive identities across multiple categories (*bridging-divides hypothesis*), that *iii*. positive contact discourages support for social change among disadvantaged-group members (*demobilisation hypothesis*), while *iv.* negative contact encourages support for social change among disadvantaged-group members (*mobilisation hypothesis*), and that *v*. positive contact encourages support for social change among advantaged-group members (*solidarity hypothesis*). In this chapter, I give a brief overview of the empirical research used to test these hypotheses, before introducing the statistical techniques used to analyse findings from the empirical research.

Overview of studies

In Chapters 3–5, I test the hypotheses described above using data from four original studies. In total, I report results from 2953 participants across two cross-sectional surveys (Studies 1, 2), one longitudinal survey (Study 3), and one quasi-experimental task (Study 4). Research spans advantaged and disadvantaged groups in four countries (England, Wales, Germany, India) and multiple intergroup contexts (sexuality, gender, caste, religion, and nationality). Below, I give a more detailed overview of the four studies, highlighting how each study relates to the hypotheses described above.

Studies 1 and 2 examined the relationships between intergroup contact, social identification, outgroup attitudes, and support for social change with a cross-sectional survey design. Studies 1 and 2 included 426 and 375 students, respectively. Participants were recruited at University of Oxford in May/June 2014 (Study 1) and in May/June 2015 (Study 2). Studies measured contact with and attitudes toward heterosexual, non-heterosexual, and/or transgender people. Both studies include heterosexual and non-heterosexual participants. In Chapter 3, I report results from heterosexual (Studies 1, 2a) and cisgender (Study 2b) respondents to develop novel measures of social identity continuity and fluidity, and to test the bridging-binaries hypothesis cross-sectionally. In Chapter 5, I report results from sexual-minority (Study 1a) and sexual-majority (Study 1b) respondents to test, respectively, the mobilisation and demobilisation hypotheses, and the solidarity hypothesis.

Study 3 examined the relationships between intergroup contact, social identification, outgroup attitudes, and support for social change with a longitudinal survey design.

Study 3 included 1801 students, recruited at Humboldt University of Berlin (n = 799), University of Hagen (n = 250), University of Oxford (n = 206), University of Kent (n = 204), Cardiff University (n = 177), and University of Osnabrück (n = 165). Participants completed three surveys, at the beginning, toward the middle, and toward the end of the 2015/2016 academic year. Surveys measured contact with and attitudes toward heterosexual, non-heterosexual, cisgender, and transgender people. Participants included students of different sexual orientations and gender identities. In Chapter 3, I report results from heterosexual (Study 3a) and cisgender (Study 3b) respondents to validate novel measures of social identity continuity and fluidity, and to test the bridging-binaries hypothesis longitudinally. In Chapter 3, I report results from sexual-minority (Study 3a), sexual-majority (Study 3b), and cisgender (Study 3c) respondents to test the mobilisation and demobilisation hypotheses (Study 3a), and the solidarity hypothesis (Studies 3b, 3c).

Study 4 adapted the triple crossed-categorisation task to examine the relationships between intergroup contact, social identification across multiple categories, outgroup attitudes, and support for social change. Study 4 included 351 students, recruited at Karnatak University, South India in Spring 2017. Participants came from advantaged and disadvantaged caste backgrounds. Study 4 thus tested group differences in responses to the triple crossed-categorisation task with a quasi-experimental design. Participants responded to a questionnaire on contact with and attitudes toward caste and religious outgroups. Study 4 thus tested antecedents and consequences of participants' categorisations with a correlational design. In Chapter 4, I report results from the triple crossed-categorisation task and the associated questionnaire to test the bridging-divides hypothesis. In Chapter 5, I report results from additional measures included in the questionnaire to test the mobilisation, demobilisation, and solidarity hypotheses in the South Indian context. Studies 1–4 feature complex research designs and data structures. In the next sections, I introduce the statistical techniques that I use to analyse findings from the studies presented above.

Bayesian data analysis

Most research in social and behavioural sciences relies on null hypothesis significance testing to derive inferences from observations. This approach, however, rests on stringent assumptions that are rarely met in research practice (A. Gelman & Loken, 2013; McShane, Gal, A. Gelman & Tackett, 2018), lends itself to misinterpretation (Gigerenzer, 2004; Morey, Hoekstra, Rouder, Lee & Wagenmakers, 2016), and biases scientific practice toward erroneous conclusions (Ioannidis, 2005; Smaldino & McElreath, 2016). In view of these shortcomings, the American Statistical Association concluded that "by itself, a p-value does not provide a good measure of evidence regarding a model or hypothesis", and that "scientific conclusions…should not be based only on whether a p-value passes a specific threshold" (Wasserstein & Lazar, 2016, p. 131–132).[1]

Instead, I use Bayesian data analysis (A. Gelman et al., 2014; Lambert, 2018; McElreath, 2016) as a flexible and powerful approach to statistical inference. As this approach is not (yet) standard practice in psychological research, I use the remainder of this chapter to introduce Bayesian data analysis, as used throughout this thesis. Below, I explain the central components of Bayesian inference—that is, the likelihood, prior, and

[1]Some researchers (e.g., Cumming, 2014) proposed estimation based on confidence intervals as a solution to the shortcomings of p-values. Confidence intervals, however, are closely related to null hypothesis significance testing, and share many of the same problems (Kruschke & Liddell, 2018; Morey et al., 2016).

posterior distributions—and introduce model comparison and selection as a strategy to reach accurate conclusions about a phenomenon of interest. Throughout, I rely on simplified examples that nonetheless illustrate the kinds of analysis used in this thesis. In the next section, I discuss the likelihood as a central concept in statistical inference, and introduce notation used throughout this thesis.

Likelihood

Imagine two people, Miles and Nils, settling a dispute by flipping a coin. Miles wins if the coin comes up tails, while Nils wins if the coin comes up heads. The coin comes up heads. Miles disputes the result, claiming that the coin is unfair. Miles and Nils agree to put the coin to the test. They flip the coin ten more times, and find that the coin comes up heads six times. Miles concludes that the coin is indeed unfair, as observing six heads in ten coin flips is most likely when the coin favours heads six to four. Nils disagrees, arguing that this outcome is quite likely even if the coin is fair. Who is right?

In this case—if we assume that every coin flip is independent of the other coin flips, and that every coin flip is equally likely to come up heads—we can use the binomial distribution to decide how plausible Miles' and Nils' conclusions are. This distribution gives the probability of observing a certain number of successes (in this case, heads) in a certain number of trials (in this case, coin flips) given a certain probability of success:

$$\Pr(k|n,p) = \frac{n!}{k!(n-k)!}p^k(1-p)^{n-k}$$

where k is the observed number of successes, n is the number of trials, p is the probability of success at each trial, and $\Pr(k|n,p)$ is the probability of k given n and p. I call $\Pr(k|n,p)$ the likelihood of k given n and p. Figure 2.1 shows the likelihood of all possible outcomes for various probabilities of success. In line with Nils' conclusion, observing six heads in ten coin flips is quite likely when the coin is fair, $\Pr(6|10,0.5) \approx 0.21$. In line with Miles' conclusion, this outcome is more likely when the coin favours heads six to four, $\Pr(6|10,0.6) \approx 0.25$, than when the coin is fair. In this case, both are right.

Miles and Nils could decide to settle their discussion by flipping the coin 100 more times (see Figure 2.2). If they again observed every sixth coin flip to come up heads—that is, 60 out of 100 coin flips—Miles would have a much better case. This outcome would be a lot more likely when the coin favours heads six to four, $\Pr(60|100,0.6) \approx 0.08$, than when the coin is fair, $\Pr(60|100,0.5) \approx 0.01$.[2] This scenario illustrates that the likelihood is sensitive to the available evidence: 100 trials provide stronger evidence than 10 trials. I return to this example in later sections.

The binomial distribution is only one of many distributions that researchers may use to derive the likelihood of a set of observations. In this thesis, I use the following notation to describe the likelihood in any given model:

$$y \sim \mathit{Distribution}(\theta)$$

where y is a set of observations, *Distribution* stands in for the specific distribution used, θ ("theta") is a set of parameters that determine the shape of that distribution, and $\sim$ expresses that y is distributed according to $\mathit{Distribution}(\theta)$. For continuous variables,

[2] Indeed, this outcome would be about eight times more likely in the former case than in the latter. This is an example of a likelihood ratio.

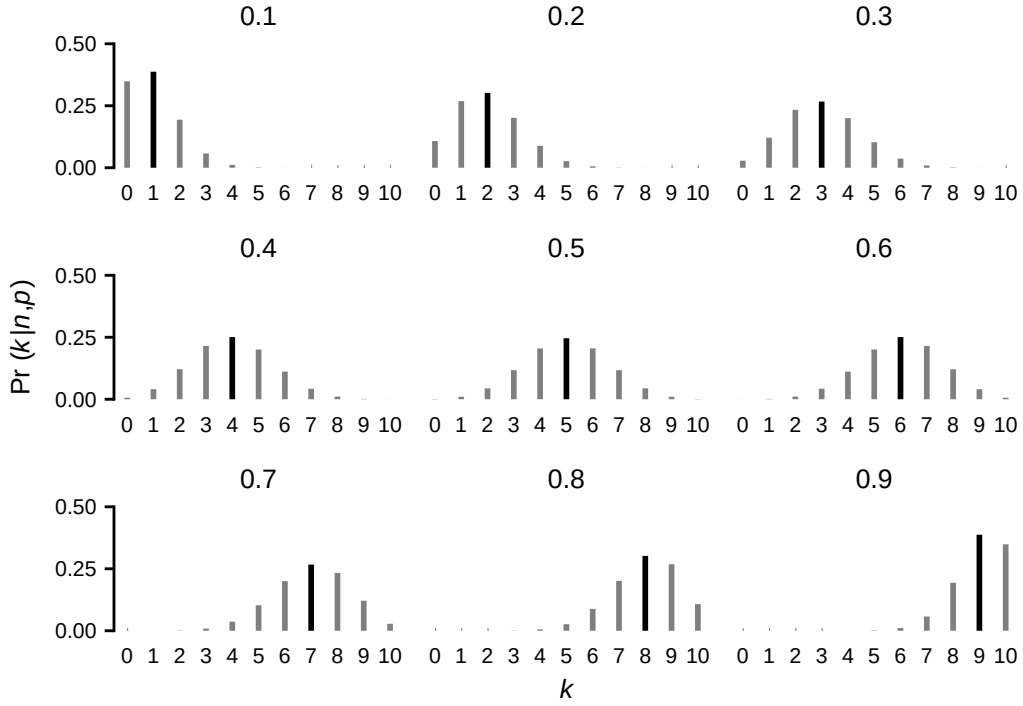

Figure 2.1. Likelihood of observing k successes in $n = 10$ trials for different probabilities p, highlighting the maximum likelihood estimate for each probability p. The central facet, for example, corresponds to the probabilities of observing 0–10 heads when flipping a fair coin 10 times.

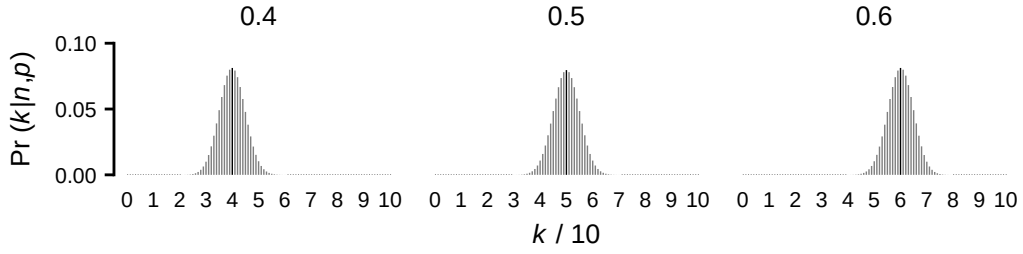

Figure 2.2. Likelihood of observing k successes in $n = 100$ trials for different probabilities p, highlighting the maximum likelihood estimate for each probability p. The x-axis shows the number of successes ($0 \leq k \leq 100$) divided by 10 to aid comparison with Figure 2.1.

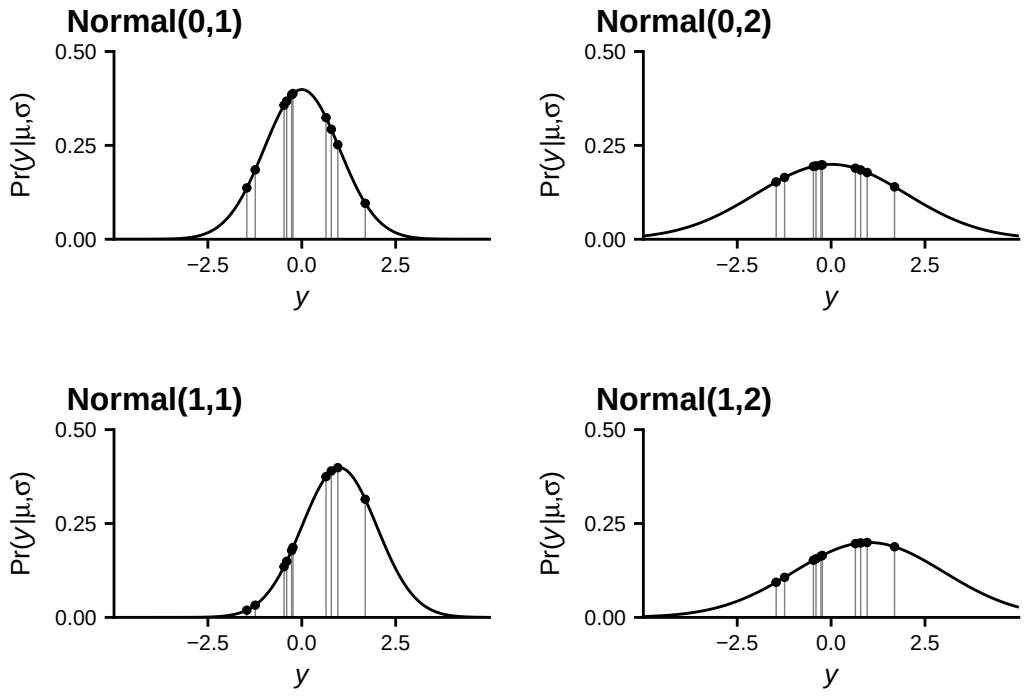

Figure 2.3. Likelihood of ten draws from a continuous random variable under four different normal distributions.

the normal distribution is often a useful way to describe the data (for a justification, see McElreath, 2016, Chapter 4). Any normal distribution is exhaustively described by two parameters, its mean and its standard deviation. I use the following notation:

$$y \sim \text{Normal}(\mu, \sigma)$$

where μ ("mu") is the mean of the normal distribution, and σ ("sigma") is its standard deviation. As an example, take a researcher who has made ten observations of a continuous random variable. Figure 2.3 shows the likelihood of each observation under normal distributions with different parameter values. Assuming that observations are independent and identically distributed, the joint likelihood of ten observations is the product of their individual likelihoods. As such, the researcher determines that her observations are more likely under a Normal$(0, 1)$ distribution than under Normal$(1, 1)$, Normal$(0, 2)$, and Normal$(1, 2)$ distributions (see Figure 2.3).

Often, researchers are interested in more than one variable—specifically, how one or more variables relate to another. To answer this question, researchers use linear regression. That is, they assign a distinct mean to each observation that corresponds to a linear combination of one or more other variables:

$$y_i \sim \text{Normal}(\mu_i, \sigma) \qquad \text{(likelihood)}$$
$$\mu_i = \beta_0 + \beta_1 x_i \qquad \text{(regression)}$$

where the index i marks the ith of n observations, and mu_i, the estimated mean for observation i, equals β_0, the fixed intercept across observations, plus β_1 times x_i, the observed value of x for observation i. In this case, the coefficient β_1 quantifies the linear

relationship between variables x and y, whereas σ quantifies the residual variance after accounting for the relationship between x and y. In principle, a researcher could again use the likelihood function to determine how likely her observations are given specific values of the relevant parameters (e.g., $\beta_0 = \bar{y}$ and $\beta_1 = 0$). In practice, a researcher is usually more interested in how plausible various values of the relevant parameters are given her observations—and not the reverse. In the next section, I introduce maximum likelihood estimation and Bayesian inference as ways to move from the likelihood to estimating parameters.

Estimation

Returning to the earlier example, imagine that, instead of arguing over *whether* their coin is fair or biased, Miles and Nils want to find out *how* biased, if at all, their coin is. Again, they base their analysis on observing $k = 6$ heads (success) in $n = 10$ coin flips (trials). One way to estimate p, the probability of success, is to find the value $\hat{p}$ that maximises the likelihood function. For a binomial distribution, the maximum likelihood estimate is $\hat{p} = k/n$, that is, the observed proportion of successes (see Figures 2.1 and 2.2). For a normal distribution, the maximum likelihood estimates are $\hat{\mu} = \bar{y}$ and $\hat{\sigma} = s$, the arithmetic mean and standard deviation of the observed values (see Figure 2.3, upper left corner). Using maximum likelihood estimation, Miles and Nils would conclude that their coin comes up heads three in five times ($\hat{p} = 0.6$). Still, this conclusion is unsatisfying. Researchers make observations to find out how plausible various beliefs about the world are given their observations—not how plausible their observations are given various beliefs about the world. Below, I introduce Bayesian inference as a way to do just this.

Expressed formally, Miles and Nils want to find out $\Pr(p|k)$, the probability of different parameter values given the observed number of successes.[3] This is the *posterior distribution* or, more simply, the posterior. Above, I have introduced $\Pr(k|p)$, the probability of the observed number of successes given different parameter values. This is the *likelihood distribution* or, simply, the likelihood. Both distributions are conditional probabilities. Bayes' theorem states how to move from one conditional probability to the other:

$$\Pr(p|k) = \frac{\Pr(k|p)\Pr(p)}{\Pr(k)}$$

where $\Pr(p)$ is the *prior distribution* or, simply, the prior, and $\Pr(k)$ is the *average likelihood* or marginal likelihood.[4] $\Pr(p|k)$, $\Pr(k|p)$, and $\Pr(p)$ are distributions. That is, they return a distinct probability for each possible value of parameter p. $\Pr(k)$ is constant across parameter values.

Specifically, $\Pr(k)$ is the likelihood averaged over all possible values of the prior distribution. In our example:

$$\Pr(k) = \int_0^1 \Pr(k|p)\Pr(p)\mathrm{d}p$$

As such, $\Pr(k)$ ensures that $\Pr(p|k)$ is a probability distribution (that is, that it sums to one). As the denominator is thus a constant, Bayes' theorem implies:

$$\Pr(p|k) \propto \Pr(k|p)\Pr(p)$$

[3] I omit n, the number of trials, in this section.

[4] I borrow this terminology from McElreath (2016), to whom this chapter is heavily indebted.

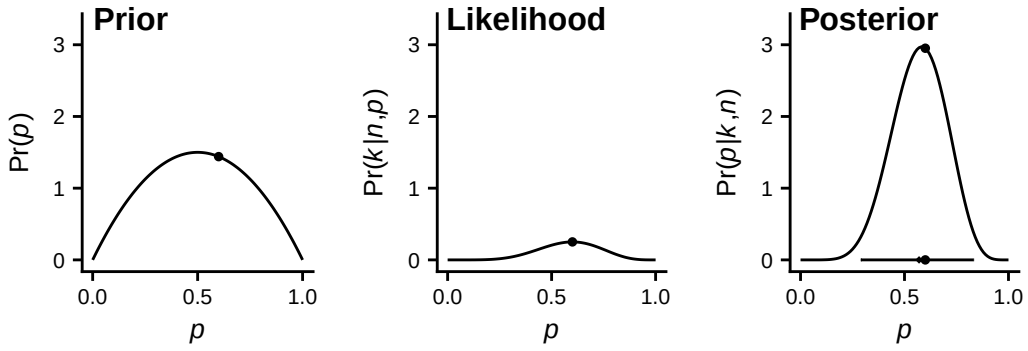

Figure 2.4. Prior, likelihood, and posterior (with 97% uncertainty interval) distributions for parameter p (the probability of success) after observing $k = 6$ successes in $n = 10$ trials.

which says that the posterior is proportional to the likelihood times the prior. Below, I discuss each of these three elements of Bayesian inference.

The likelihood distribution, $\Pr(k|p)$, defines the probability of the data (in the example, observing $k = 6$ heads in ten coin flips) given different values of one or more parameters (in the example, the probability of success p). In Bayesian inference, the likelihood fulfils three roles. First, the likelihood defines the parameters to which Bayes' theorem is applied. While the binomial likelihood has only one unobserved parameter, other likelihood distributions have hundreds or thousands of estimated parameters. In most cases, the likelihood reduces a set of observations to a smaller set of interpretable parameters.[5] Second, the likelihood relates these parameters to the observed data. That is, the likelihood quantifies how plausible a set of observations are under all parameter values under consideration. As such, it rules out parameter values under which the observations are impossible or implausible. In the current example, observing six heads in ten coin flips is clearly impossible if $p = 0$ (the coin never comes up heads) or $p = 1$ (the coin always comes up heads). Similarly, observing six heads in ten coin flips is clearly implausible if $p = 0.01$ (the coin comes up heads in 1 of 100 flips), $\Pr(6|10, 0.01) \approx 2 \times 10^{-10}$. As discussed previously, observing six heads in ten coin flips is most likely when $p = 0.6$, $\Pr(6|10, 0.6) \approx 0.25$. For all other parameter values, the likelihood falls somewhere in between these extremes (see Figure 2.4, centre).[6] Third, the likelihood encapsulates a set of assumptions about the phenomenon under study. In the current example, using a binomial likelihood assumes that coin flips can have only one of two outcomes (heads or tails), that the probability of heads is the same for every flip of a coin, and that coin flips are independent of each other. If these assumptions are mistaken, the likelihood misrepresents the phenomenon under study. Below, I discuss the role of prior distributions in Bayesian inference.

The prior distribution, $\Pr(p)$, assigns an a priori probability to each possible value of the parameters that define the likelihood distribution. That is, the prior defines how plausible different values of the parameters are before considering the observed data.

[5] In multilevel models, the number of parameters sometimes exceeds the number of observations. Still, the varying parameters are represented by a smaller set of hyperparameters (as discussed later).

[6] As a conditional probability distribution, $\Pr(k|p)$ can be solved either for k or for p. Figures 2.1 and 2.2 do the former, showing the likelihood of all possible observations for a given parameter value. Figure 2.4 (centre) does the latter, showing the likelihood of a given set of observations as a function of different parameter values. This reflects the role of the likelihood in Bayesian inference.

Often, the prior distribution reflects some knowledge about the phenomenon under study. In the example, Miles and Nils know that the probability of a coin coming up heads falls between zero and one. They also know that it is unlikely that a regular looking coin will always, or nearly always, come up heads (or that it will always, or nearly always, come up tails). They choose to reflect these considerations in the prior distribution. I use the following notation:

$$p \sim \text{Beta}(2, 2)$$

which defines the prior probability of parameter p as a Beta distribution with parameters $\alpha = 2$ and $\beta = 2$ (see Figure 2.4, on the left). Bayesian inference thus involves choosing a likelihood distribution, and assigning a prior probability to all possible values of that likelihood distribution's parameters. Bayes' theorem then weights the prior probability of each parameter value by the likelihood of observing the data given that parameter value. In short, Bayesian inference updates prior probabilities in light of the observed data to produce posterior probabilities (McElreath, 2016, p. 29). Below, I discuss how different prior distributions influence the shape of the posterior distribution.

As discussed, the posterior distribution equals the prior distribution weighted by the likelihood distribution, divided by the average likelihood. Thus, the posterior distribution reflects both prior probabilities and the observed data (that is, the likelihood distribution). In the example, the likelihood distribution is centred around the maximum likelihood estimate, $\hat{p} = 0.6$, and concentrates 97% of the likelihood between $p = 0.28$ and $p = 0.85$. The prior distribution is centred around $p = 0.5$, and concentrates 97% of the prior probability between $p = 0.07$ and $p = 0.93$. The likelihood and prior distribution combine to form a posterior distribution that is centred around $p \approx 0.58$, and concentrates 97% of the posterior probability between $p = 0.29$ and $p = 0.83$ (see Figure 2.4). Once the data has been taken into account, the posterior's shape depends solely on the choice of prior. Figure 2.5 illustrates how various prior distributions affect the posterior distribution. In the example, the Beta$(2, 2)$ prior has gently moved the posterior closer to $p = 0.5$ relative to the likelihood (second row). If I instead use a uniform Beta$(1, 1)$ prior, the posterior is proportional to the likelihood (first row). If I use distributions that concentrate prior probability more closely around $p = 0.5$, the posterior increasingly resembles the prior (third to fifth row). For a Beta$(100, 100)$ prior distribution, for example, the posterior distribution is centred around $p \approx 0.50$, with 97% of the posterior probability falling between $p = 0.43$ and $p = 0.58$ (fifth row). This comparison shows that, especially with few observations, the choice of prior influences the posterior.

Researchers might worry about the influence of prior probabilities on their conclusions, as incorporating prior probabilities seems to contradict the ideal that researchers should be objective and open to all results before looking at the data. To curb these concerns, researchers might advocate uniform priors that consider all possible parameter values as equally plausible. I present three arguments against this approach. First, researchers almost always have some knowledge about the parameters under study. A researcher might, for example, analyse a set of ten responses to a questionnaire that uses Likert scales ranging from 1 = *strongly disagree* to 7 = *strongly agree*. Without looking at the responses, the researcher knows that the parameters of a normal likelihood are necessarily constrained, so that $1 \leq \mu \leq 7$ and $0 < \sigma < 3.08$. If the researcher standardised participants' responses, they would know that, without any predictors, the parameters of the normal likelihood should be close to $\mu = 0$ and $\sigma = 1$. Not using this knowledge makes for worse inferences—especially with few observations. Second, prior distributions

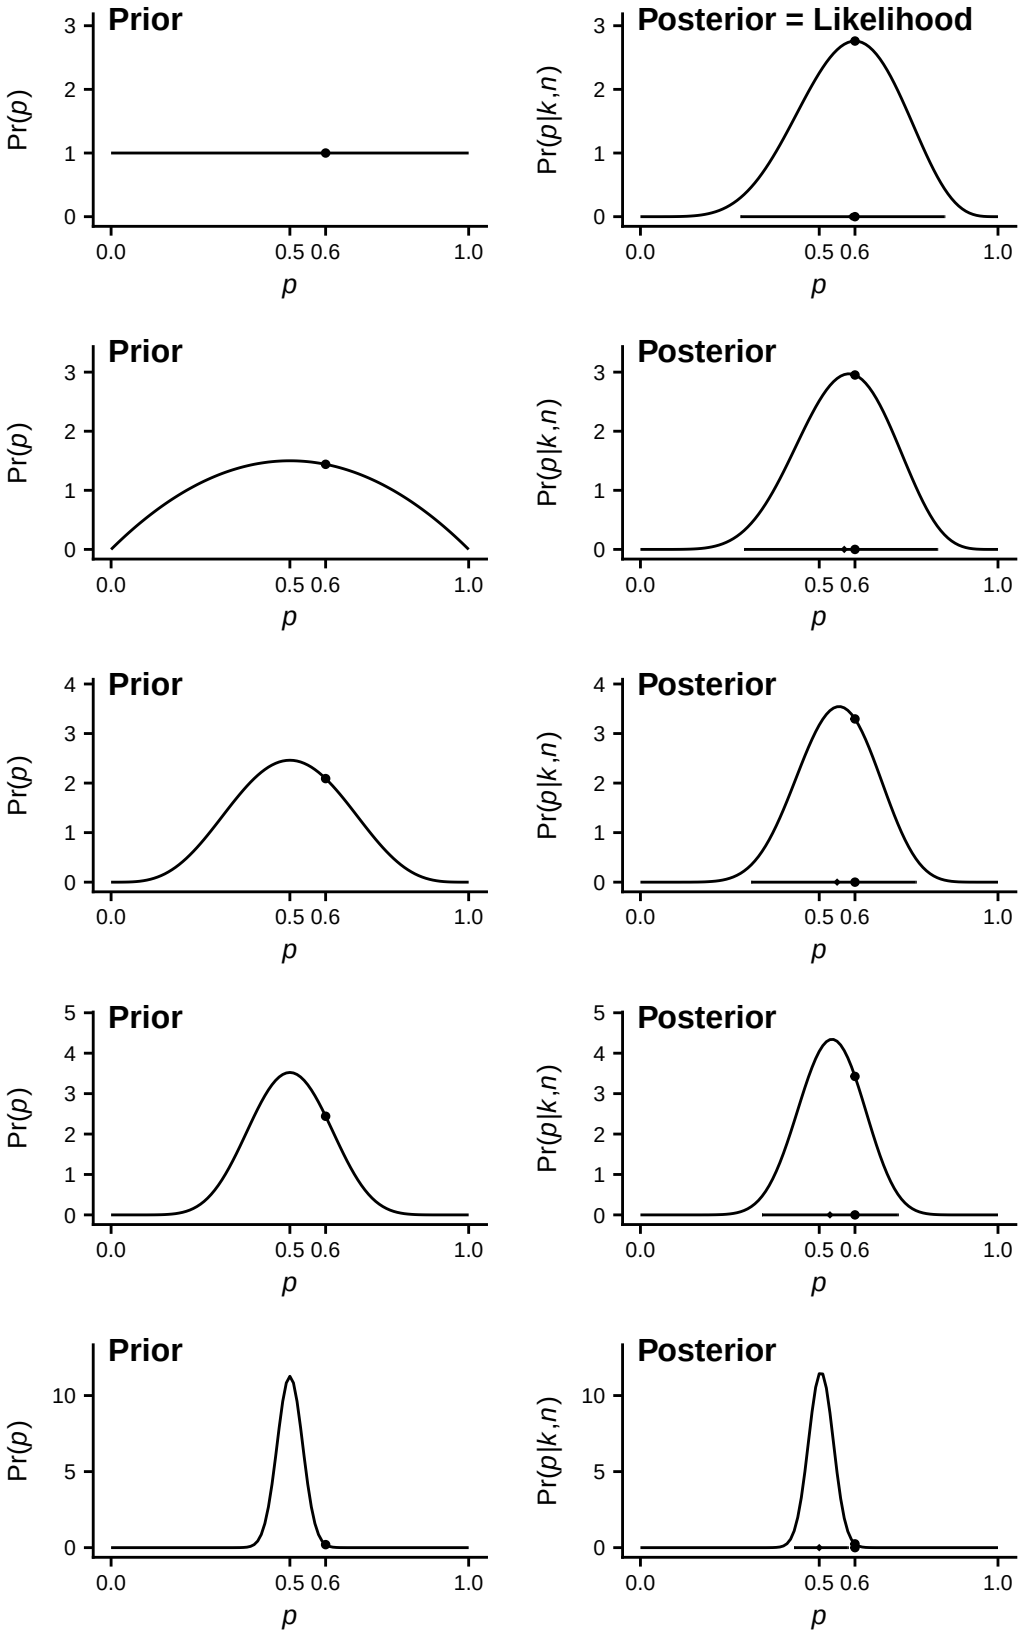

Figure 2.5. Prior probabilities (on the left) and posterior probabilities (on the right) for parameter p (the probability of success) after observing $k = 6$ successes in $n = 10$ trials. Rows (from top to bottom) use increasingly informative prior distributions.

need not represent the researcher's beliefs about a phenomenon.[7] Instead, researchers might choose prior distributions that, in the long run, yield more accurate and stable inferences. A regularising prior is usually conservative, penalising extreme values. As such, a well-designed prior prevents overfitting a statistical model to the observed data.[8] Third, the prior reflects the researcher's assumptions about the phenomenon under study. Above, I have shown that the likelihood encapsulates a set of assumptions about how the observed data were generated. Similarly, the prior interacts with the likelihood to constrain the parameter values under consideration. As such, the prior is best understood as a testable assumption about the data-generating processes (A. Gelman, Simpson & Betancourt, 2017). As any other assumption, researchers need to justify their choice of prior. Far from noninformative, uniform priors can bias inference towards highly implausible combinations of extreme parameter values (A. Gelman et al., 2017), and might not be a justifiable choice in many situations.

In this thesis, I use weakly informative priors "that are explicitly designed to encode information that applies to a general class of problems without taking full advantage of problem-specific knowledge" (A. Gelman et al., 2017, p. 13). Specifically, I rely on conservative priors—that is, distributions centred around null effects—to regularise statistical inferences. What makes for a weakly informative prior depends on the likelihood. Figure 2.6 illustrates how the priors in Figure 2.5 would influence the posterior if the likelihood was $\Pr(60|100, p)$ instead of $\Pr(6|10, p)$. Relative to the likelihood, the $\text{Beta}(10, 10)$ prior (fourth row) is strongly informative in Figure 2.5, but weakly informative in Figure 2.6. Reflecting these considerations, I use the following prior distributions for intercepts and coefficients in, respectively, logistic and linear regression models:

$$\beta \sim \text{Student}(2.5, 0, 1)$$
$$\beta \sim \text{Normal}(0, 1)$$

Below, I discuss how researchers can use the posterior distribution to answer questions about the parameters under study.

The posterior distribution, $\Pr(p|k)$, assigns an a posteriori probability to each possible value of the parameters that define the likelihood distribution. That is, the posterior defines how plausible different values of the parameters are after considering the observed data. Bayesian inference thus results in a distribution of plausible parameter estimates, rather than in discrete point estimates and intervals. Researchers can interrogate this distribution to answer questions about model parameters. Miles and Nils, for example, want to find out whether their coin is biased or not. They note that the 97% most plausible estimates include parameter values which imply that the coin favours heads, favours tails, or is fair ($p = 0.58$, $[0.29, 0.83]$). Alternatively, they could define a range of parameter values that they, for all practical purposes, consider equivalent to a fair coin, and test how much of the posterior distribution falls in this region of practical equivalence (Kruschke & Liddell, 2018). They note that similar proportions of the posterior distribution fall inside, $\Pr(p < 0.4, p > 0.6|k = 6) = 0.57$, and outside, $\Pr(0.4 \leq p \leq 0.6|k = 6) = 0.43$, the specified region. Based on either approach, Miles and Nils would conclude that their observations have been inconclusive. In both cases, however, researchers could have made different choices (e.g., a narrower interval or a wider region of practical equivalence) that would have been just as valid. Wherever feasible, I graphically present the full posterior

[7] Indeed, I do not think that I could accurately express my beliefs with a mathematical distribution.
[8] A similar idea underlies penalised likelihood estimation.

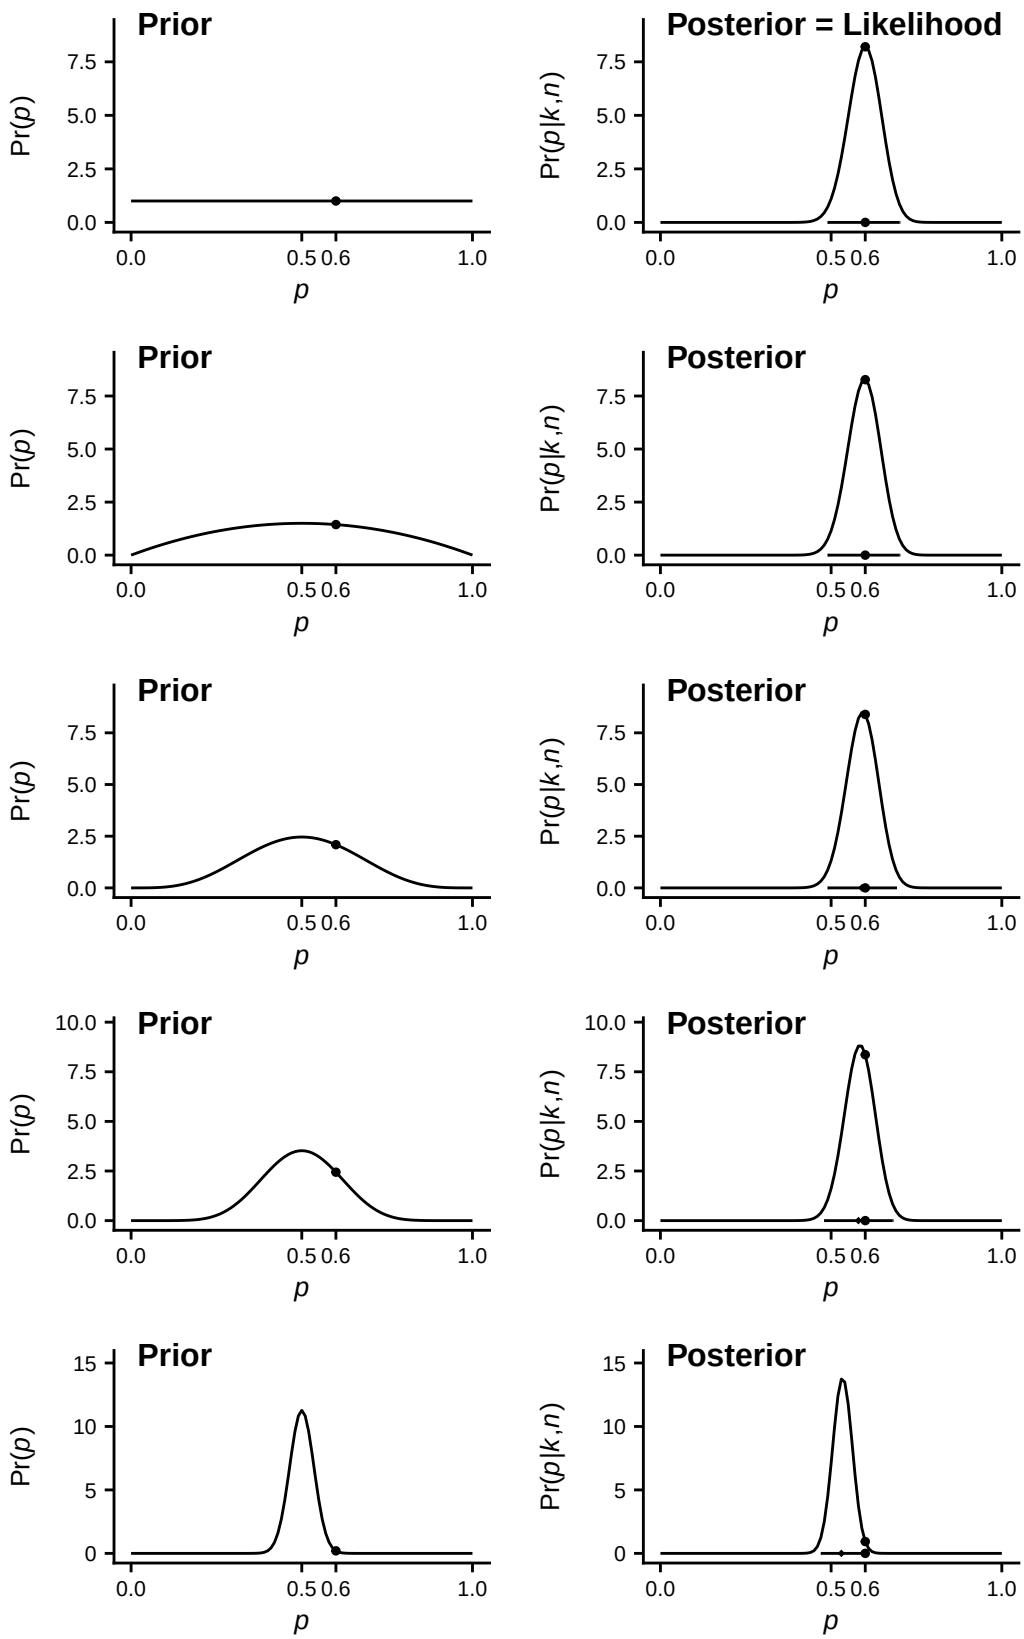

Figure 2.6. Prior probabilities (on the left) and posterior probabilities (on the right) for parameter p (the probability of success) after observing $k = 60$ successes in $n = 100$ trials. Compared to Figure 2.5, the prior distributions have less influence on posterior probabilities.

distribution for parameters of interest (for an example, see Figure 2.4). In addition, I report point estimates, based on the median of the posterior distribution, as well as intervals spanning the 97% most plausible estimates, based on the highest posterior density (Plummer, Best, Cowles & Vines, 2006).

Thus far, I have introduced Bayesian inference as a way to combine a set of assumptions, encapsulated in the likelihood function and prior probabilities, with a set of observations to estimate posterior probabilities for a set of parameters. I have shown that, for the same observations, different assumptions produce different posterior probabilities. I have not considered that, for the same assumptions, different observations would have produced different posterior probabilities. If, for example, Miles and Nils had observed five heads in ten coin flips, they would have been more confident that the coin is fair. If they had observed eight heads in ten coin flips, they would have been more confident that the coin is biased. Both cases are plausible when observations are sampled from a random process. Every sample has both regular features—that is, features that generalise across samples and answer a question of interest—and irregular features—that is, features that do not generalise and provide misleading answers to a question of interest (McElreath, 2016, p. 168). In any data analysis, the challenge is to avoid learning from the irregular features of a sample (overfitting), while not overlooking the regular features of that sample (underfitting). I have already introduced regularising priors as one way to prevent overfitting. In the next section, I introduce posterior predictive checking as a strategy for balancing overfitting and underfitting in Bayesian data analysis.

Comparison

Suppose a researcher has collected $n = 150$ observations from four continuous random variables, and wants to test whether three predictor variables (x_1, x_2, x_3) relate to one outcome variable (y). To that end, she regresses the outcome variable on the three predictor variables in a linear regression model:

$$
\begin{aligned}
y_i &\sim \text{Normal}(\mu_i, \sigma) && \text{(likelihood)} \\
\mu_i &= \beta_0 + \beta_1 x_{1,i} + \beta_2 x_{2,i} + \beta_3 x_{3,i} && \text{(regression)} \\
\beta_0 &\sim \text{Normal}(0, 1) && \text{(prior for intercept)} \\
\{\beta_1, \beta_2, \beta_3\} &\sim \text{Normal}(0, 1) && \text{(prior for coefficients)} \\
\sigma &\sim \text{Cauchy}(0, 1) && \text{(prior for residual variance)}
\end{aligned}
$$

where y_i is the observed value of outcome variable y for observation i, μ_i is the estimated value of outcome variable y for observation i, $x_{1,i}$ to $x_{3,i}$ are the observed values of the predictor variables for observation i, and β_1 to β_3 are the corresponding coefficients. She uses Bayes' theorem to derive posterior probabilities for all coefficients:

$$\Pr(\beta|y) \propto \Pr(y|\beta) \Pr(\beta)$$

where $\Pr(\beta|y)$, the posterior distribution for coefficient β, is proportional to $\Pr(y|\beta)$, the likelihood of observations y, times $\Pr(\beta)$, the prior distribution for coefficient β. She notes that for coefficients β_1 and β_2, the model assigns most posterior probability to parameter values above zero: $\Pr(\beta_1 > 0|y) > 0.99$ and $\Pr(\beta_2 > 0|y) > 0.99$. For coefficient β_3, the model assigns similar posterior probabilities to parameter values above and below zero: $\Pr(\beta_3 > 0|y) = 0.68$ and $\Pr(\beta_3 < 0|y) = 0.32$. Based on posterior probabilities,

the researcher would thus conclude that x_1 ($\beta_1 = 0.43$, $[0.26, 0.59]$) and x_2 ($\beta_2 = 0.22$, $[0.06, 0.37]$), but not x_3 ($\beta_3 = 0.03$, $[-0.11, 0.19]$), are associated with higher values of the outcome variable. Posterior probabilities alone, however, cannot differentiate between regular and irregular features of a particular sample of observations. Below, I introduce posterior predictive checking as another way to evaluate a statistical model.

In one view, Bayesian data analysis begins with a researcher's initial beliefs about a phenomenon, and ends with that researcher updating her beliefs in light of their observations. In this view, prior and posterior probabilities reflect the researcher's subjective beliefs about a phenomenon. By collecting more observations, researchers improve their estimates of a relevant model parameter, resulting in more precise posterior probabilities (Kruschke & Liddell, 2018).[9] In another view, Bayesian data analysis begins with a set of assumptions about a phenomenon, and then tests these assumptions against a set of observations. In this view, prior probabilities (alongside the likelihood) are testable assumptions about the data-generating process underlying the observations. Posterior probabilities are the logical implications of these assumptions applied to observations. As such, they serve as testable predictions about the data-generating processes. By pitting these predictions against observations, researchers evaluate and revise their models of a phenomenon of interest (A. Gelman & Shalizi, 2013). I agree with the latter perspective, which informs data analyses in this thesis. Below, I discuss ways to apply this perspective to balance overfitting and underfitting in research practice.

Broadly, this perspective can be summarised as posterior predictive checking. For a given posterior distribution and likelihood function, there exists a posterior predictive distribution:

$$\Pr(\tilde{y}|y) = \int \Pr(\tilde{y}|\theta) \Pr(\theta|y) \mathrm{d}\theta$$

where y are observed values, $\tilde{y}$ are predicted values, $\Pr(\theta|y)$ is the posterior distribution of parameters θ given observations y, $\Pr(\tilde{y}|\theta)$ is the likelihood of predicted values $\tilde{y}$ given parameters θ, and $\Pr(\tilde{y}|y)$ is the implied posterior distribution of predicted values $\tilde{y}$. This distribution incorporates two uncertainties. First, a fitted model implies a distribution of plausible parameter values, that is, the posterior distribution. Second, any set of parameter values implies a distinct distribution of plausible observations, for example, $\tilde{y} \sim \text{Normal}(0, 1)$ if $\mu = 0$ and $\sigma = 1$. Averaging over both distributions, the posterior predictive distribution indicates how plausible various observations would be given a fitted model's implied posterior probabilities about parameter values and sampling uncertainty. From this distribution, researchers can derive how plausible their observations would be given the data-generating process implied by a fitted model. In the example, the researcher finds that 146 of 150 (97.3%) observations fall within the interval around the 97% most plausible posterior predictions. She thus concludes that the fitted model plausibly represents the data-generating process. In other cases, posterior predictive checks point to ways in which a particular model misrepresents the data-generating process (underfitting). Still, posterior probabilities can tell a researcher which parameter values (and predicted values) are plausible given a specific model of the data, but cannot tell a researcher whether a different model might be a more plausible approximation of the underlying data-generating process. Below, I discuss how model comparison can answer the latter question.

As discussed, posterior probabilities alone cannot discriminate between paramet-

[9] Cumming (2014) argued for a similar cumulative research programme within the frequentist framework.

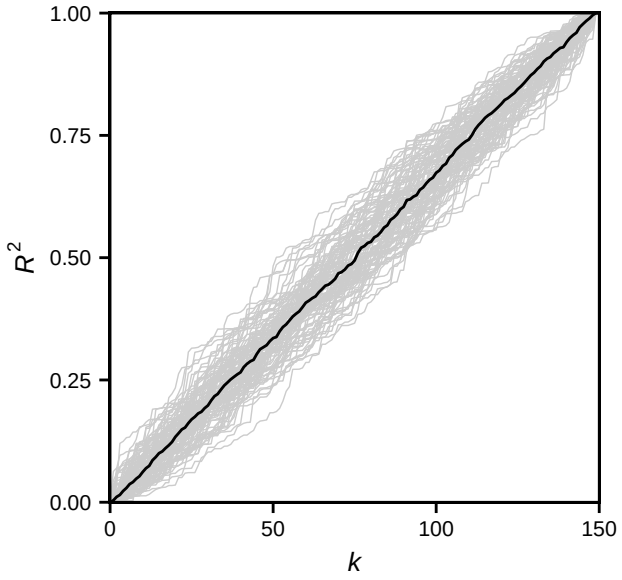

Figure 2.7. Explained variance (R^2) as a function of the number of predictor variables ($0 \leq k \leq 150$) in linear regression models, based on 100 simulated data sets with 150 observations of 151 variables (1 outcome variable, 150 predictor variables). All observations are random draws from a Normal$(0, 1)$ distribution.

ers that are necessary to account for the regular features of the observed data, and parameters that are superfluous. To tell one from the other, researchers estimate a series of models, starting with the simplest model under consideration. They extend this model—for example, by adding additional coefficients to a linear model—and compare the resulting models to each other. Based on these comparisons, researchers decide which parameters are necessary, and which parameters are superfluous. To compare models, researchers need a criterion to evaluate each model's performance. One criterion used by researchers is R^2, defined as the variance of predicted values divided by the variance of observed values. This criterion, often used alongside maximum likelihood estimation, illustrates a common challenge in model comparison. Figure 2.7 shows that R^2 increases with each predictor variable added to a linear regression model, even when all observations are random draws from a probability distribution. This shows that R^2, and similar criteria for model comparison, are sensitive to the irregular features of a sample of observations (overfitting). Below, I introduce cross validation as a strategy to balance underfitting and overfitting in Bayesian data analysis.

From the posterior predictive distribution, researchers can determine the joint predictive probability of their observations for a fitted model. Specifically, they can calculate the log pointwise predictive density:

$$\text{lpd} = \sum_{i=1}^{n} \log \Pr(y_i|y)$$

where $\log \Pr(y_i|y)$ is the logarithm of the posterior predictive probability of observed value i, and where $\sum_{i=1}^{n} \log \Pr(y_i|y)$ is the sum of these probabilities across all n observed values.[10] This sum indicates how accurately a fitted model predicts the observed data.

[10] Using the logarithm of the pointwise predictive density is more convenient because calculating the pointwise predictive density as $\prod_{i=1}^{n} \Pr(y_i|y)$ is computationally challenging.

Again, this criterion does not distinguish between regular and irregular features of the observed data. If a feature is regular, researchers should expect to observe the same feature in other samples from the same data-generating process. Cross validation exploits this logic to evaluate statistical models. In the simplest case, a researcher would divide her observations into two subsets, a training set and a test set. She would then fit a model to the training set, and evaluate how accurately the fitted model predicts the observations in the test set. Cross validation thus estimates a model's out-of-sample prediction accuracy. This estimate makes a useful criterion for model evaluation as out-of-sample prediction accuracy decreases both when the fitted model does not learn enough from the observed data (underfitting), and when it learns too much from the observed data (overfitting).

Still, this estimate depends on the specific observations in each subset, and thus on each subset's irregular features. The researcher could improve her estimate by dividing their sample into three (instead of two) equal subsets. She would then train a model on two thirds of her sample, test the model's predictions on the other third, and repeat this procedure for all combinations of subsets. For the best estimate, the researcher could train a model on $n - 1$ of n observations, test the model's prediction accuracy on the left-out observation, and repeat this procedure for all n divisions of the observations (leave-one-out cross validation). For any model, researchers can use leave-one-out cross validation to estimate the model's expected log pointwise predictive density:

$$\text{elpd} = \sum_{i=1}^{n} \log \Pr(y_i|y_{-i})$$

where $\log \Pr(y_i|y_{-i})$ is the logarithm of the posterior predictive probability of observed value y_i given a model fitted to the other $n-1$ observations, and where $\sum_{i=1}^{n} \log \Pr(y_i|y_{-i})$ is the sum of these probabilities across all n divisions of the observations. In practice, exact cross validation, which requires refitting a given model n times, is not feasible.[11] Instead, leave-one-out cross validation can be approximated using the log likelihood of all observations derived from the posterior distribution of parameter values (for details, see Vehtari, Gelman & Gabry, 2017). This procedure also provides standard errors for a model's estimated out-of-sample predictive accuracy, reflecting the uncertainty of cross-validation as an estimate of predictive accuracy. Below, I return to the example to illustrate how to use out-of-sample predictive accuracy for comparing statistical models.

Again, suppose a researcher has collected $n = 150$ observations from four continuous random variables, and wants to test whether three predictor variables (x_1, x_2, x_3) relate to one outcome variable (y). To that end, she estimates a series of linear regression models with the following regression equations:

$$\mu_i = \beta_0 \quad \text{(Model 0)}$$
$$\mu_i = \beta_0 + \beta_1 x_{1,i} \quad \text{(Model 1)}$$
$$\mu_i = \beta_0 + \beta_2 x_{2,i} \quad \text{(Model 2)}$$
$$\mu_i = \beta_0 + \beta_1 x_{1,i} + \beta_2 x_{2,i} \quad \text{(Model 3)}$$
$$\mu_i = \beta_0 + \beta_1 x_{1,i} + \beta_2 x_{2,i} + \beta_3 x_{3,i} \quad \text{(Model 4)}$$

where μ_i is the estimated value of outcome variable y for observation i, β_0 is the fixed intercept across observations, $x_{1,i}$ to $x_{3,i}$ are the observed values of the predictor variables

[11] As an example, consider one of the simpler models presented in Chapter 4. This model takes about 5 minutes to estimate, and would, with $n = 7210$ observations, require about 600 hours to complete exact leave-one-out cross validation.

Table 2.1. Comparison of five linear regression models. R^2 is a Bayesian analogue to R^2 in maximum likelihood estimation (A. Gelman, Goodrich et al., 2017). *ELPD* is the expected log predictive density, with higher numbers indicating that a model is expected to make more accurate out-of-sample predictions (Vehtari et al., 2017). Δ*ELPD* is the difference in *ELPD* between the current and previous model, with positive values indicating that the current model is expected to make more accurate out-of-sample predictions. I preferred a more complex model over a simpler model when $\frac{\Delta ELPD}{SE} \geq 2$. w are stacking weights based on the combined leave-one-out predictive density (Yao et al., 2018).

#	Description	R^2	*ELPD*	*SE*	Δ*ELPD*	*SE*	$\frac{\Delta ELPD}{SE}$	w
0	β_0	.00	-214.3	8.5	-	-	-	.00
1 vs 0	$\beta_0 + \beta_1 x_1$	.25	-193.8	7.8	20.5	6.3	3.3	.12
2 vs 0	$\beta_0 + \beta_2 x_2$	.12	-205.5	8.8	8.8	3.9	2.3	.03
3 vs 1	$\beta_0 + \beta_1 x_1 + \beta_2 x_2$	.29	-190.4	8.1	3.4	3.0	1.1	.85
4 vs 3	$\beta_0 + \beta_1 x_1 + \beta_2 x_2 + \beta_3 x_3$	.30	-191.3	8.1	-1.0	0.5	-0.5	.00

for observation i, and β_1 to β_3 are the corresponding coefficients. Model 0 is an intercept-only model. Models 1 and 2 extend Model 0 by adding, respectively, x_1 and x_2 as predictor variable. Model 3 builds upon Models 1 and 2 by including both x_1 and x_2 as predictor variables. Model 4 extends Model 3 by adding x_3 as another predictor variable.

Table 2.1 compares the expected out-of-sample predictive accuracy of all five models. *ELPD* is the expected log predictive density, with higher numbers indicating that a model is expected to make more accurate out-of-sample predictions (Vehtari et al., 2017). Δ*ELPD* is the difference in *ELPD* between the current and previous model, with positive values indicating that the current model is expected to make more accurate out-of-sample predictions. *SE* is the standard error for each estimate (Vehtari et al., 2017). In this thesis, I prefer a more complex model over a simpler model when the difference in *ELPD* favours the more complex model by at least two standard errors of that difference, $\frac{\Delta ELPD}{SE} \geq 2$.[12] Using this criterion, the researcher concludes that Model 1 and 2 make more accurate out-of-sample predictions than Model 0, showing that both x_1 and x_2 are associated with the outcome variable. Model 1 is estimated to make more accurate out-of-sample predictions than Model 2 ($\Delta ELPD = 11.9$, $SE = 6.1$, $\frac{\Delta ELPD}{SE} = 2.0$), showing that x_1 is a more accurate predictor than x_2. Model 3 does not improve upon Model 1, showing that x_2 does not increase predictive accuracy beyond x_1. Model 4 makes less accurate predictions than simpler models, showing that x_3 does not relate to the outcome variable. Based on this comparison, the researcher concludes that x_1, but not x_2 and x_3, is associated with higher values of y, and reports the posterior distribution for the corresponding coefficient in Model 1, $\beta_1 = 0.50$, $[0.34, 0.64]$.[13] When considering out-of-sample prediction accuracy, the researcher thus reaches a different conclusion than when considering posterior probabilities alone.

In this thesis, I follow similar strategies of model extension and selection based

[12] I have chosen this threshold as a relatively conservative criterion for model comparison—but I could have chosen other thresholds. This threshold resembles the conventional 1.96 × SE criterion used in null hypothesis significance testing.

[13] *ELPD* is sensitive to sample size; with $n = 1000$ observations, but the same intervariable correlations, the researcher would have concluded that Model 3 makes more accurate predictions than Model 1, $\Delta ELPD = 29.2$, $SE = 7.9$, $\frac{\Delta ELPD}{SE} = 3.7$.

on comparing each model's expected out-of-sample prediction accuracy. In addition, I report two additional metrics of model performance. R^2 is a Bayesian analogue to R^2 in maximum likelihood estimation (A. Gelman, Goodrich et al., 2017), and indicates the variance explained by the predictor variables in a model. w are stacking weights based on the combined leave-one-out predictive density of a set of models (Yao et al., 2018). Stacking weights show how to combine predictions from a set of models to make the most accurate out-of-sample predictions (model averaging). I report stacking weights for the interested reader, but focus on model selection in the main text. Thus far, I have discussed relatively simple models, that is, binomial models (without predictors) and linear regression models. In the next section, I introduce various advanced models that I use throughout this thesis.

Extensions

Generalised linear models

Suppose that Miles and Nils, after concluding that their previous observations have been inconclusive, set out to make further observations. Each of them flips the coin ten times. Miles finds that the coin comes up heads $k_1 = 7$ in $n = 10$ times. Nils finds that the coin comes up heads $k_2 = 4$ in $n = 10$ times. Based on these observations, they wonder whether the probability of observing heads depends on who of them is flipping the coin.

Previously, I have introduced linear regression models in which, for each observation, the μ_i parameter in the $y_i \sim \text{Normal}(\mu_i, \sigma)$ likelihood is replaced with a linear combination of predictor variables. Applying this principle, Miles and Nils could estimate a model in which the p_i parameter of the $k_i \sim \text{Binomial}(n, p_i)$ likelihood is replaced with distinct estimates for Miles' and Nils' coin flips:

$$k_i \sim \text{Binomial}(n, p_i) \quad \text{(likelihood)}$$
$$p_i = \beta_i \quad \text{(regression)}$$

where k_i is the observed number of successes for observation i (1 = *Miles*, 2 = *Nils*), p_i is the estimated probability of success for observation i, and β_1 and β_2 are the estimated probabilities of success for Miles and Nils, respectively. In practice, this model would not be a good representation of the underlying data-generating process. As a probability, the parameter p_i is bounded between zero and one. A linear model, however, would consider parameter values above one and below zero, and thus make impossible predictions.

Instead, researchers should include a link function that transforms a model's linear predictions into non-linear parameter values. Thereby, researchers can use linear regression to model likelihoods defined by non-linear parameters (generalised linear models). When parameters are probabilities, the logit transformation makes for a useful link function:

$$\text{logit}(p_i) = \alpha_i = \log \frac{p_i}{1 - p_i}$$

where α_i is the logarithm of the odds of success (that is, the probability of success divided by the probability of failure). The resulting log odds are unbounded (that is, they can range from negative infinity to positive infinity), and can be modelled as a linear combination of predictor variables. Miles and Nils thus restate their regression as a generalised linear

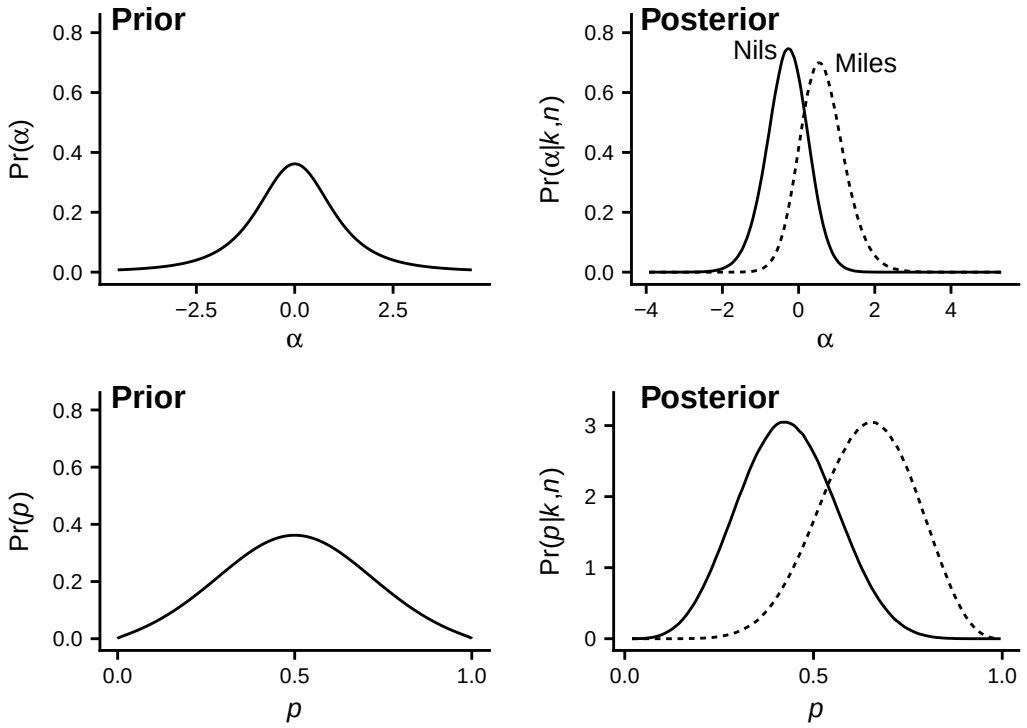

Figure 2.8. Prior and posterior probabilities for log odds α (top row) and probabilities p (bottom row) in the generalised linear model.

model with a logit link function:

$$\begin{aligned} k_i &\sim \text{Binomial}(n, p_i) && \text{(likelihood)} \\ \text{logit}(p_i) &= \alpha_i && \text{(link function)} \\ \alpha_i &= \beta_i && \text{(regression)} \\ \beta_i &\sim \text{Student}(2.5, 0, 1) && \text{(prior)} \end{aligned}$$

where α_i is the estimated log odds for observation i (1 = *Miles*, 2 = *Nils*), and β_1 and β_2 are the estimated log odds for Miles and Nils, respectively.[14]

Figure 2.8 (top row) shows prior and posterior probabilities for Miles' and Nils' log odds of observing heads. Log odds, however, can be difficult to interpret. Whenever feasible, I transform log odds into probabilities using the logistic or inverse-logit function:

$$p_i = \text{logistic}(\alpha_i) = \frac{\exp(\alpha_i)}{1 + \exp(\alpha_i)}$$

where p_i is the estimated probability of success for observation i, and α_i are the estimated log odds for observation i. Figure 2.8 (bottom row) shows prior and posterior distributions for Miles' and Nils' estimated probabilities of observing heads. Miles and Nils note that, though Miles ($p_2 = 0.65$, $[0.37, 0.90]$) has observed heads more often than Nils ($p_1 = 0.43$, $[0.17, 0.70]$), posterior probabilities for both estimates overlap, $\Pr(p_2 > p_1|k, n) = 0.88$. They conclude that their observations have been inconclusive. In Chapter 4, I use logistic regression to model participants' responses in the triple crossed-categorisation task. Below, I introduce multilevel models as another technique used throughout this thesis.

[14] In this formulation, α_i are redundant; I use this formulation to show parallels with linear models.

Multilevel models

Suppose that Miles and Nils enlist another eight people to flip their coin. Each person, including Miles and Nils, flips the coin ten times. In 10 flips, they observe, respectively, 4, 2, 7, 6, 3, 6, 1, 8, 4, and 4 heads. At one extreme, the researchers could use the same strategy as above, and estimate a distinct probability of observing heads for each person (see Figure 2.9, right). This strategy, however, overlooks that coin flips are unlikely to be independent across persons. Everyone flips the same coin, and uses a similar technique when doing so. At the other extreme, the researchers could pool their estimates across persons. By combining information from all observations, this strategy results in more precise estimates. This strategy, however, assumes that all observations come from the same data-generating process. This assumption is not plausible as the resulting model makes inaccurate posterior predictions (see Figure 2.9, left). In this example, neither no pooling nor complete pooling makes for a good model of the observations.

Multilevel models represent a compromise between the two extremes (A. Gelman, Hill & Yajima, 2012). Multilevel models take the hierarchical structure of the observations into account (as does no pooling), but also combine information from all observations to improve individual estimates (as does complete pooling). What sets multilevel models apart is that they include varying intercepts and effects. That is, they assign distinct parameter estimates to each observation, but model these distinct parameter estimates as arising from the same data-generating process.[15] In Bayesian models, this assumption is encoded in the hyperprior for the varying effect. That is, models assign a prior distribution to all varying effects that has itself adjustable parameters with their own prior distributions. In the example, the researchers estimate each person's probability of observing heads with a varying intercept:

$$
\begin{aligned}
k_i &\sim \text{Binomial}(n, p_i) && \text{(likelihood)}\\
\text{logit}(p_i) &= \alpha_i && \text{(link function)}\\
\alpha_i &= \beta_0 + \beta_j && \text{(regression)}\\
\beta_0 &\sim \text{Student}(2.5, 0, 1) && \text{(prior for fixed intercept)}\\
\beta_j &\sim \text{Normal}(0, \sigma_{\text{person}}) && \text{(hyperprior for varying intercept)}\\
\sigma_{\text{person}} &\sim \text{Cauchy}(0, 1) && \text{(prior for } \sigma_{\text{person}}\text{)}
\end{aligned}
$$

where α_i, the estimated log odds for observation i and person j, equals β_0, the fixed intercept across participants, plus β_j, the varying intercept for participant j. In this formulation, β_j is modelled as each person's deviation from the fixed intercept β_0, distributed according to a normal distribution with standard deviation σ_{person}.

Figure 2.9 (middle) shows the resulting posterior probabilities. Because of the shared hyperprior, parameter estimates exhibit partial pooling—that is, they shrink toward each other. Consider Person 7 who has observed only one head in ten coin flips. With no pooling, the researchers might conclude that Person 7 observes heads less often than to be expected by chance, $p_7 = 0.18$, $[0.00, 0.44]$. With partial pooling, the researchers might conclude otherwise, $p_7 = 0.28$, $[0.05, 0.53]$. By considering information from other observations, multilevel models thus prevent premature conclusions from extreme observations (for example, in multiple comparison, see A. Gelman et al., 2012). In this thesis, I rely on partial pooling, alongside regularising priors and posterior predictive checking, to prevent false discoveries.

[15] This is why varying effects are often called "random" effects.

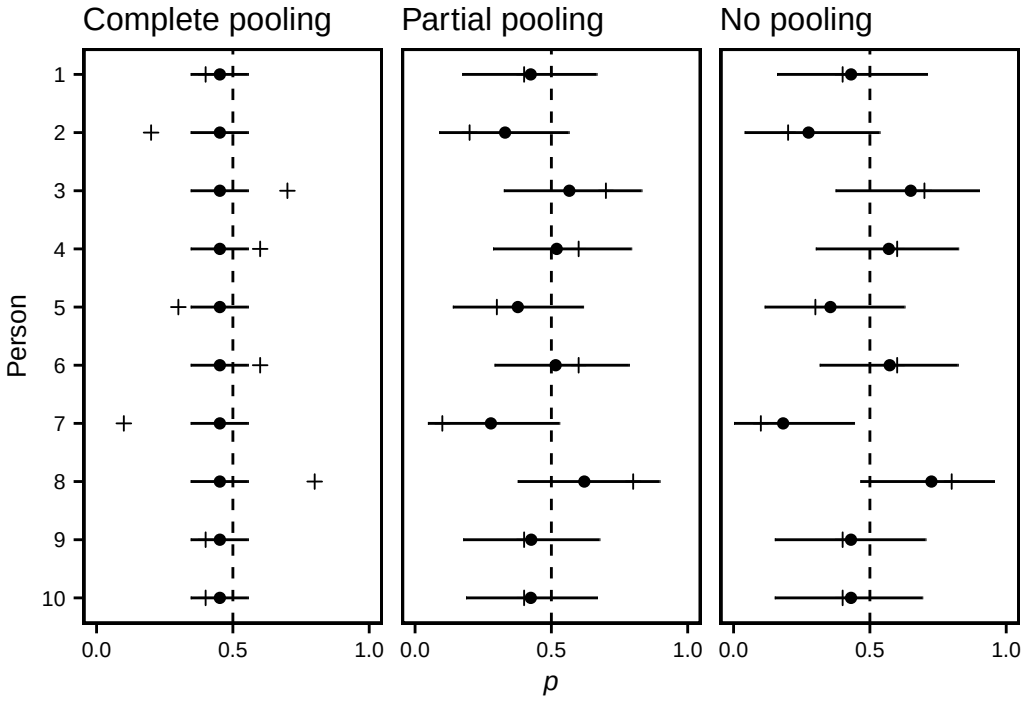

Figure 2.9. Each person's probability of observing heads, estimated with complete pooling (left), partial pooling (middle), and no pooling (right) across persons. Points (•) indicate the most plausible estimate for each person, with intervals around the 97% most plausible estimates. Pluses (+) indicate the observed proportion of heads for each person.

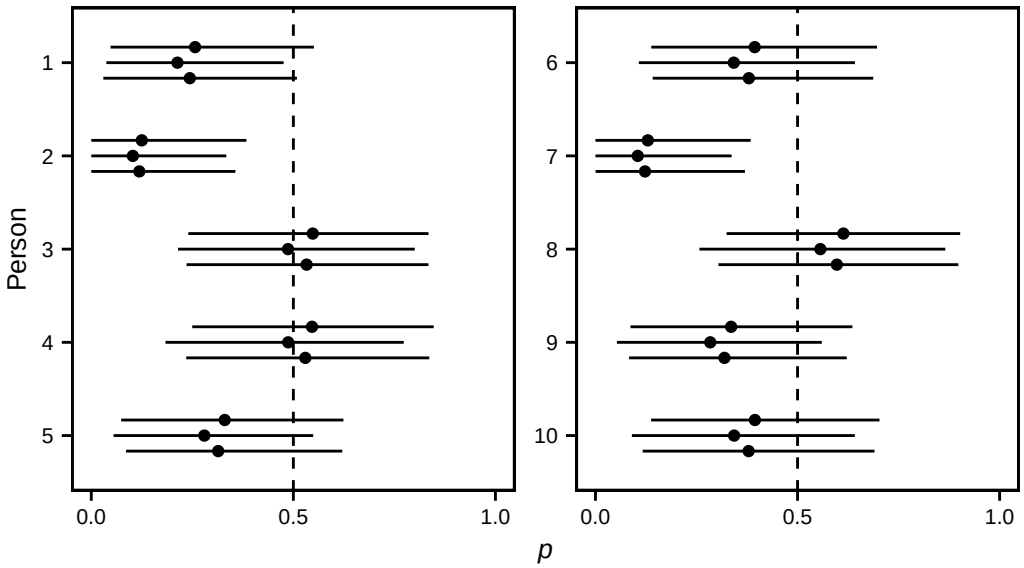

Figure 2.10. Estimated probability of observing heads as a function of who of ten persons is flipping which of three coins. Points (•) indicate the most plausible estimate for each person and coin, with intervals around the 97% most plausible estimates.

Suppose further that Miles and Nils want to test three different coins: a £1 coin, a €1 coin, and a trick coin that is supposed to favour heads. Again, they recruit ten people, each of whom flips each coin three to four times. Each person thus makes ten coin flips. Miles and Nils want to find out whether persons or coins are more important in shaping the outcome of coin flips. To answer this question, they extend the previous model by including an additional varying intercept:

$$\begin{aligned} \alpha_i &= \beta_0 + \beta_j + \beta_l && \text{(regression)} \\ \beta_l &\sim \text{Normal}(0, \sigma_{\text{coin}}) && \text{(hyperprior for varying intercept)} \\ \sigma_{\text{coin}} &\sim \text{Cauchy}(0, 1) && \text{(prior for } \sigma_{\text{coin}}\text{)} \end{aligned}$$

where β_l is the varying intercept for coin j, and $\beta_l \sim \text{Normal}(0, \sigma_{\text{coin}})$ is the corresponding hyperprior. Figure 2.10 shows the resulting posterior probabilities. Person 2 and 7 seem to flip coins in ways that favour tails, suggesting that individuals can influence the outcome of coin flips. Coins, however, do not seem to influence the outcome of coin flips; even the "loaded" coin does not come up heads more often than the other coins. The researchers conclude that a coin cannot be biased, though a coin toss can be.[16] Together, these examples illustrate that Bayesian multilevel modelling is a powerful approach for examining nested observations. In Chapters 4 and 5, I use multilevel models to estimate participants' responses in the triple crossed-categorisation task, and on the associated questionnaire. Below, I introduce models that estimate multiple outcome variables at once.

Multivariate models

Suppose a researcher has collected $n = 150$ observations from continuous random variables, and wants to examine two outcome variables, y_1 and y_2. Sometimes, when faced with two or more outcome variables, researchers might want to estimate all outcome variables as part of the same statistical model. Such multivariate models allow researchers, for example, to examine how outcome variables relate to each other, to compare how predictor variables relate to different outcome variables, or to compare parameter estimates across outcome variables.

In the example, the researcher wants to find out how y_1 relates to y_2, and whether y_2 is, on average, greater than y_1. To that end, she estimates a simple multivariate model:

$$\begin{aligned} \begin{bmatrix} y_1 \\ y_2 \end{bmatrix} &\sim \text{MVNormal}\left(\begin{bmatrix} \mu_1 \\ \mu_2 \end{bmatrix}, \mathbf{S} \right) && \text{(likelihood)} \\ \mathbf{S} &= \begin{pmatrix} \sigma_1 & 0 \\ 0 & \sigma_2 \end{pmatrix} \mathbf{R} \begin{pmatrix} \sigma_1 & 0 \\ 0 & \sigma_2 \end{pmatrix} && \text{(variance-covariance matrix)} \\ \{\mu_1, \mu_2\} &\sim \text{Normal}(0, 1) && \text{(prior for } \mu\text{)} \\ \{\sigma_1, \sigma_2\} &\sim \text{Cauchy}(0, 1) && \text{(prior for } \sigma\text{)} \\ \mathbf{R} &\sim \text{LKJ}(1) && \text{(prior for correlation matrix)} \end{aligned}$$

where MVNormal is the multivariate normal likelihood distribution, μ_1 and μ_2 are the estimated means for variables y_1 and y_2, and $\mathbf{S}$ is the corresponding variance-covariance

[16] Indeed, "you can load a die, but you can't bias a coin" (A. Gelman & Nolan, 2002, p. 308)—as long as it is flipped in the air, instead of being spun or allowed to bounce.

matrix. This model decomposes the variance-covariance matrix into a correlation matrix **R**, and distinct standard deviations for each outcome variable, σ_1 and σ_2. This means that the model estimates distinct means and variances for each outcome, as well as the correlation between the two outcomes. The LKJ(1) distribution is a weakly informative prior for correlation matrices (for details, see Stan Development Team, 2017). Based on this model, the researcher concludes that y_2 ($\mu_2 = 0.99$, $[0.81, 1.17]$) is greater than y_1 ($\mu_1 = 0.00$, $[-0.18, 0.19]$; $\mu_2 - \mu_1 = 0.99$, $[0.78, 1.19]$), and that y_1 and y_2 correlate with each other ($r_{1,2} = 0.34$, $[0.20, 0.50]$).

Above, I have presented a simple multivariate model, without any predictor variables. Instead, the researcher could apply a linear model to μ_1 and μ_2 in order to examine how one or more predictor variables relate to each outcome variable. In this case, σ_1 and σ_2 become residual variances, while **R** becomes the residual correlation matrix. This means that the researcher could, for example, test whether the association between y_1 and y_2 changes after accounting for one or more predictor variables. Another application is missing data imputation. Bayesian models estimate missing values as additional parameters, with their own prior and posterior probabilities. Multivariate models produce accurate estimates of missing values by taking into account intervariable correlations. In Chapters 4 and 5, I use multivariate models to estimate participants' responses to multiple related outcome variables. Below, I introduce structural equation models as a multivariate approach to modelling relationships between latent variables.

Structural equation models

Thus far, I have introduced multivariate models that estimate the variance-covariance matrix for a set of observed variables. Structural equation models take the observed variance-covariance matrix as their starting point. They condense the variance-covariance matrix for a set of observed variables into a smaller set of unobserved variables, and estimate the relationships between these unobserved variables. A structural equation model consists of two parts. The measurement model reduces the covariances between observed variables to a smaller set of unobserved (latent) constructs. By constraining how observed variables relate to latent constructs, researchers compare how well different measurement models correspond to a given set of observations. The structural model estimates the relationships between the latent constructs. By constraining how latent constructs relate to each other, researchers compare how well different structural models correspond to a given set of observations. A detailed explanation of structural equation modeling is beyond the scope of this chapter (for an introduction, see Little, 2013). Below, I outline how I apply structural equation modeling throughout this thesis.

In Chapters 3 and 5, I rely on structural equation models to analyse findings from cross-sectional and longitudinal surveys. Comparing measurement models allows testing how well participants' responses to survey items map onto certain psychological constructs under study. This makes structural equation models useful for scale development. In Chapter 3, I use exploratory factor analysis to determine how many psychological constructs are necessary to represent participants' responses to newly-developed scale items. I use confirmatory factor analysis to test whether the resulting factor structure replicates in other samples. Comparing structural models allows testing whether various hypothesised relationships are necessary to represent associations between multiple psychological constructs. This makes structural equation models useful for testing hypotheses in cross-sectional and longitudinal surveys. In Chapters 3 and 5, I use structural

equation models to examine the hypothesised relationships between intergroup contact and relevant outcome variables. As these models allow estimating relationships between multiple variables, they lend themselves to testing indirect effects from a predictor variable to an outcome variable via one or more intervening variables. In Chapters 3 and 5, I use structural equation models to test various hypothesised mediators of the effects of intergroup contact. Comparing structural equation models allows testing whether scales measure similar constructs in different groups or across multiple time points. In Chapters 3 and 5, I use structural equation models to establish measurement invariance across time points and languages. Below, I describe technical details of the analytic approach used in this thesis.

Above, I have described how specifying measurement and structural models can answer questions on the relationships between latent constructs. After specifying models, researchers need to find parameter estimates that reflect the model specifications and fit the observed responses. Bayesian data analysis—in particular, Markov chain Monte Carlo methods (see below)—would be well suited to this task. Although Bayesian structural equation modeling has been implemented (Merkle & Rosseel, 2016; B. O. Muthén & Asparouhov, 2012), established conventions for scale development and measurement invariance testing are still lacking. Instead, I estimate all structural equation models using robust maximum likelihood estimation. When comparing models, I use multiple model fit indices. I accept model fit if $\chi^2/df < 2$, comparative fit index (CFI) > .90, Tucker-Lewis index (TFI) > .90, and root mean square error approximation (RMSEA) < .08 (Byrne, 2012; Little, 2013). I compare models by comparing χ^2 fit indices (Satorra & Bentler, 2001), but use ΔCFI < .01 as a criterion for establishing measurement invariance (see Little, 2013). For a given model, I report parameter estimates with 95% confidence intervals. When testing for indirect effects, I estimate the relative magnitude of all direct, indirect, and total effects using bias-corrected bootstrap sampling with 5000 resamples (Preacher & Hayes, 2008) and standardised indicators. In the next section, I describe the software used for analysing data in this thesis.

Software

With few exceptions, I rely on the *R* programming language (R Core Team, 2017) for all data analysis tasks. I use *R* packages from the 'tidyverse' family (Wickham, 2017) to import and transform data (Wickham, Francois, Henry & Müller, 2017; Wickham & Henry, 2018) and to visualise observed data and model predictions (Wickham, 2009).

I use specialised *R* packages for implementing the statistical models described above. Simple Bayesian models, such as the simple binomial models introduced above, can be solved analytically. More complex models, however, are analytically intractable because they require solving high-dimensional integrals. Instead, researchers use Markov chain Monte Carlo methods to iteratively approximate the posterior distribution of relevant parameters.[17] In this thesis, I estimate Bayesian models using Hamiltonian Monte Carlo (Betancourt, 2017). I code models in the probabilistic programming language *Stan* (Stan Development Team, 2017), as implemented in the *RStan* package (Stan Development Team, 2018) and the *brms* package (Bürkner, 2017). As explained above, I use robust

[17] This difference is noticeable when comparing, for example, Figure 2.4 with some of the figures in Chapter 4. I have solved the model underlying the former analytically, while the latter is based on draws from Markov chains.

maximum likelihood methods for estimating structural equation models, as implemented in the *lavaan* package (Rosseel, 2012). I use the *Mplus* 7.4 software (L. K. Muthén & Muthén, 1998–2015) for estimating exploratory structural equation models (Asparouhov & Muthén, 2009), which have not yet been implemented in *lavaan*. In the next chapters, I analyse findings from Studies 1–4, using the methods described in this chapter.

3

Bridging binaries: Can intergroup contact foster more fluid and continuous conceptions of sexuality and gender?

In Chapter 1, I have reviewed research by Haslam et al. (2000, 2002) who found that categories varied in the extent to which participants held essentialist beliefs about them. That is, categories varied in the extent to which participants understood them as natural, stable, discrete, immutable, and as defined by necessary characteristics. Research on psychological essentialism, however, often studies how individuals perceive the social categories of others. In this chapter, I examine whether individuals differ in how they understand *their own* category memberships, and whether these differences explain *how* or *why* intergroup contact reduces intergroup bias.

Specifically, I investigate whether individuals differ in how they understand their social identity along a *single* axis of categorisation. Based on the literature on psychological essentialism, I propose that these differences fall along two dimensions. *Social identity continuity* is the extent to which individuals perceive a social identity as their unique position on a continuum, rather than as belonging to one of several discrete categories. *Social identity fluidity* is the degree to which individuals perceive an identity as flexible and changing, rather than as fixed and stable throughout their lifetime. Someone high in both dimensions would, for example, understand their sexual orientation as an individual combination of same-gender and other-gender attraction, which varies across situations, has changed in the past, and may change in the future. As argued in Chapter 1, I hypothesise that intergroup contact fosters more fluid and continuous social identities (*hypothesis 1*) and that, in turn, more fluid and continuous social identities reduce intergroup bias (*hypothesis 2*). If both hypotheses are true, then social identity fluidity and continuity mediate the relationship between intergroup contact and intergroup bias (*bridging-binaries hypothesis*).

Social identity fluidity and social identity continuity are novel constructs. In Studies 1–3, I develop and validate scales for measuring these constructs in the realms of sexuality and gender. Below, I discuss why sexuality and gender make useful case studies for social identity continuity and fluidity.

Sexuality as case study

> *Males do not represent two discrete populations, heterosexual and homosexual. . . . [N]ature rarely deals with discrete categories. Only the human mind invents categories and tries to force facts into separated pigeon-holes. The living world is a continuum in each and every one of its aspects.*
>
> – Kinsey, Pomeroy and Martin (1948, p. 639)

> *Sexual behavior is either normal or abnormal, socially acceptable or unacceptable, heterosexual or homosexual; and many persons do not want to believe that there are gradations in these matters from one to the other extreme.*
>
> – Kinsey, Pomeroy, Martin and Gebhard (1953, p. 469)

How we understand human sexuality has changed a lot in the past century—in particular, how we understand sexuality oriented toward persons of the same gender. In early modern Europe, sodomy laws structured public discourse about same-gender sexuality. Sodomy was a "temporary transgression" (Foucault, 1978, p. 43) for which a person could be punished under the law. In the 19th century, medical discourse redefined same-gender sexuality as a psychiatric condition, and thus invented "the homosexual" as a distinct type of person (Foucault, 1978).[1] This definition underlies essentialist beliefs that frame sexual orientation as discrete, inherent, and immutable (Haslam & Levy, 2006). In the 1950s, Kinsey et al. (1948, 1953) challenged discrete understandings of sexual orientation. They argued that sexual orientation is a continuum between exclusive attraction to men and exclusive attraction to women. In their reports, Kinsey et al. categorised participants' histories of "psychologic responses" and "overt sexual activities" (1953, p. 471) into seven categories, ranging from 0 = *entirely heterosexual* to 6 = *entirely homosexual*. As they categorized only 61–72% of women and 53–78% of men as entirely heterosexual, Kinsey et al. (1948, 1953) concluded that sexual orientation is a continuum rather than a set of two (heterosexual, homosexual) or three (bisexual) distinct types.

Today, most adults agree with Kinsey et al.'s (1948, 1953) view that sexual orientation is continuous rather than discrete. In 2015, YouGov asked a representative sample of British adults about their views on sexuality. Asked which view came closer to their own, 61% agreed that sexuality was a scale, and that it was possible to be near the middle of that scale. Only 27% said that there was no middle ground, and that you were either heterosexual or homosexual. Support for the former position was strongest among younger respondents (see Figure 3.1, right). Most respondents thus believed that sexual orientation is continuous rather than discrete. Respondents not only adopted this belief in principle, but also applied it to describe themselves in practice. In the same survey, YouGov (2015) asked respondents to try to place their sexuality on "a scale of 0 to 6, where 0 is completely heterosexual and 6 is completely homosexual (i.e. gay or lesbian)".[2] On average, 72% of British adults described their sexuality as *completely heterosexual*, while 4% described themselves as *completely homosexual*. The remaining 19% (excluding *don't know* and *no sexuality* responses) placed their sexuality somewhere in between the extremes. Again, younger respondents more often used non-exclusive scale points (1–5) to describe their

[1] In other words, sexuality changed from something you *do* to something you *are*.

[2] This procedure deviates from Kinsey et al.'s (1948, 1953) scale as YouGov (2015) relies on self-report rather than coding respondents' sexual histories.

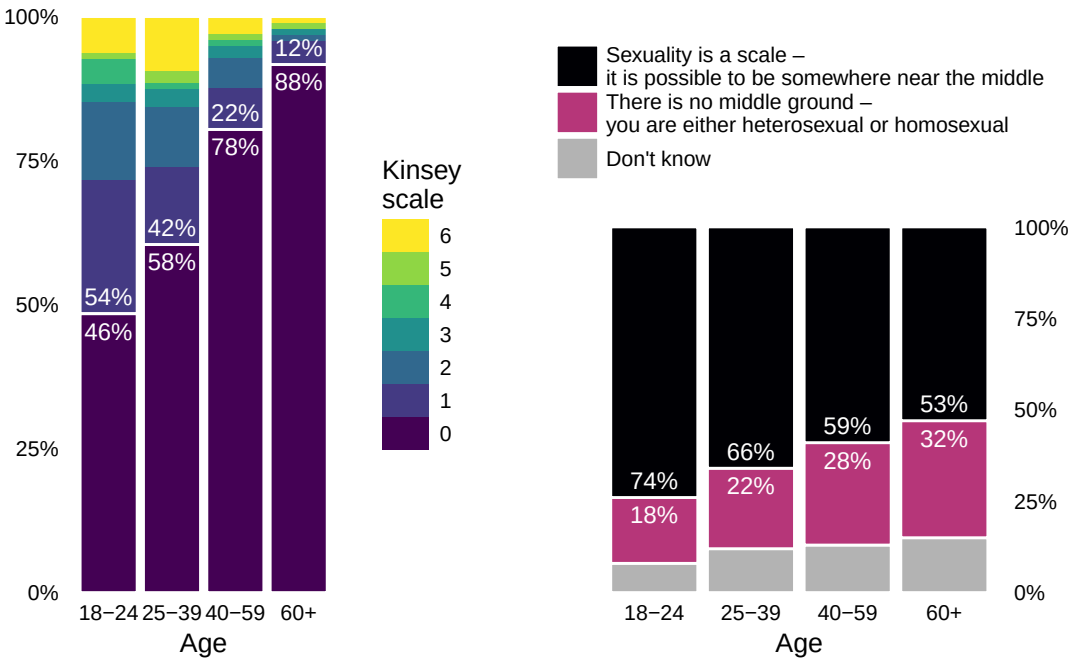

Figure 3.1. Results from a representative survey of British adults (YouGov, 2015). **On the left:** Percentages (in white) of adults who categorised themselves as 0 = *completely heterosexual* (below) or chose any other option on the Kinsey scale (above), excluding *don't know* and *no sexuality* responses. **On the right:** Percentages (in white) of adults who understood sexuality as continuous (above) or discrete (below).

sexuality (e.g., 43% of 18 to 24-year-olds, see Figure 3.1, left).[3]

Kinsey et al.'s (1948, 1953) scale describes sexuality as a *one-dimensional* continuum between exclusive orientation toward men and exclusive orientation toward women. The scale considers both sexual behaviour and psychological responses, but does not differentiate between them. Sexual behaviour and attraction, however, do not always align. A person could, for example be attracted to women and fantasise about having sex with women, without ever having sexual contact with a woman. Another person could have experienced sexual contacts with women, without ever being spontaneously attracted to women. Hence, other researchers argued that sexual orientation should be understood as a *multi-dimensional* continuum. Sell (2007) proposed three dimensions of sexual orientation: *attraction*, *behaviour*, and *identity*. That is, individuals differ in the extent to which they are attracted to men and women, in how often they have sexual contact with men and women, and in how they describe their own sexual orientation. As the examples show, these dimensions often, but not always, align.

In YouGov's (2015) survey, 89% of adults thought that "heterosexual" best described their sexuality (*sexual identity*). Of these, 11% reported having had a sexual experience with someone of the same gender (*sexual behaviour*). Only 56% of self-identified heterosexuals found it inconceivable that they could be attracted to a person of the same gender, if the right person came along at the right time (*sexual attraction*). Similarly, only 60% of self-identified heterosexuals considered it impossible that they could ever have

[3] YouGov's survey is cross-sectional, and thus allows developmental or generational interpretations of these age differences. That said, I would be surprised if differences of this magnitude (46% versus 88%, Figure 3.1, left) were exclusively developmental.

a sexual experience, or a relationship, with a person of the same gender. Again, young people ruled out same-gender attraction, experiences, and relationships less often than older people (e.g., 42%, 45%, and 45% of 18 to 24-year-olds). Together, these findings show that discrete self-descriptions (e.g., heterosexual) are not enough to describe someone's sexual orientation. In addition, these results show that many heterosexual respondents do not rule out future changes to their sexual attraction and behaviour toward people of the same gender.

Sexuality makes a useful case study for social identity continuity and fluidity. First, sexuality *is* continuous and fluid, as shown above. Individuals differ across a multi-dimensional continuum of sexual attraction, behaviour, and identity. Many individuals consider future changes along these dimensions possible. Second, sexuality *is understood* as continuous and fluid by many people. In YouGov's (2015) research, for example, most respondents (61%) viewed sexuality as continuous rather than discrete. Third, sexuality *is not understood* as continuous and fluid by everyone. In YouGov's (2015) research, 27% of respondents viewed sexuality as discrete. Similarly, Haslam and Levy (2006) found that participants differed in how discrete and immutable they believed sexual orientation to be. Together, these observations suggest that there will be sufficient variation across participants in how continuous and fluid they view their own sexuality to be. I study whether contact with sexual minorities explains some of this variation in fluidity and continuity, and whether this variation, in turn, accounts for differences in attitudes toward sexual minorities. If there was no variation in conceptions of sexuality, neither of these hypotheses would make sense.

In this chapter, I develop and validate a scale measuring *social identity continuity* and *social identity fluidity* for sexuality. Above, I have argued that sexual orientation spans multiple dimensions: sexual identity, sexual attraction, and sexual behaviour. In Studies 1 and 2, I formulate items that focus on identity and attraction. I did not include items on behaviour which would relate to observable facts rather than participants' subjective conceptions of sexuality. I generate items for four combinations of the two *construct* dimensions (continuity, fluidity) and the two *sexuality* dimensions (identity, attraction). That is, I measure social identity continuity as the extent to which individuals perceive their sexuality as a continuum between exclusive attraction to women and exclusive attraction to men (Figure 3.2d), and as the extent to which they find the terms 'homosexual' and 'heterosexual' too narrow to describe their sexual identity (Figure 3.2b). I measure social identity fluidity as the extent to which individuals perceive their sexual attraction (Figure 3.2e) and their sexual identity (Figure 3.2c) as fluid and changing, rather than as fixed and stable throughout their lifetime. In Studies 1 and 2, I develop this scale using exploratory factor analysis. In Study 3, I validate this scale using confirmatory factor analysis. In line with the *bridging-binaries hypothesis*, I test whether contact with sexual minorities predicts continuity and fluidity, and whether continuity and fluidity, in turn, predict less sexual prejudice.

Gender as case study

Compared to sexuality, participants rated gender as more natural, discrete, and stable in Haslam et al.'s (2000) study of psychological essentialism. Amongst all category domains, gender was indeed the one most essentialised (see Figure 1.1, Chapter 1). This makes gender an ambitious case study for testing the hypotheses advanced above. Below, I contrast the

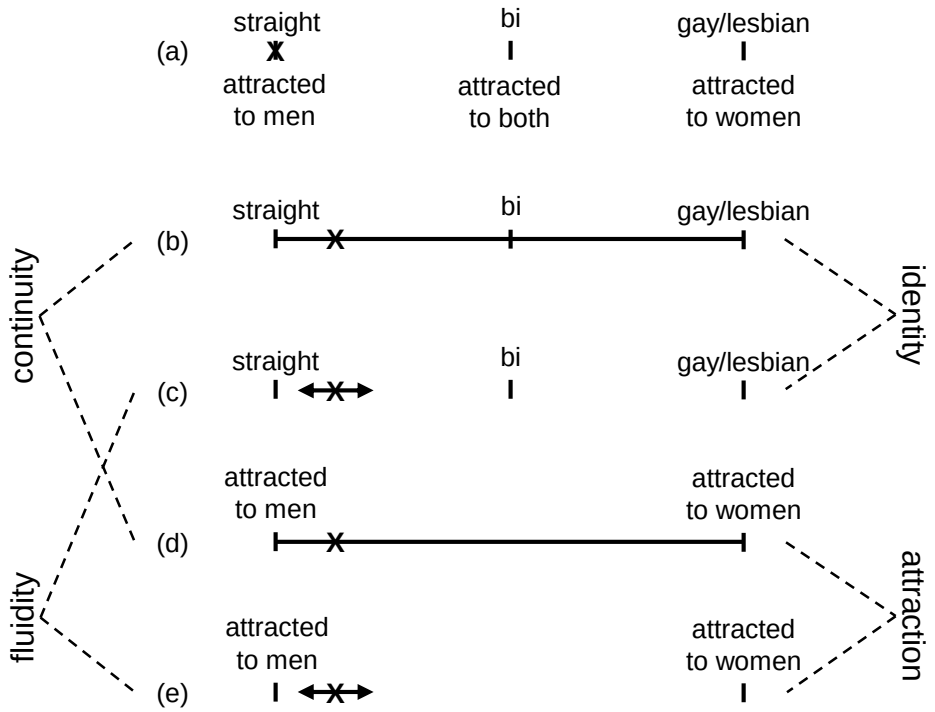

Figure 3.2. Schematic representations of different conceptions of sexuality (for a woman). Row (a) shows a binary, stable understanding of sexuality; (b) and (d) show more continuous conceptions; (c) and (e) show more fluid conceptions; (b) and (c) focus on self-categorisation; (d) and (e) focus on attraction.

essentialist conception of gender to the lived experiences of intersex, transgender, and gender non-conforming people. I outline the harm done by the essentialist conception of gender, before proposing more continuous and fluid conceptions as a means to reduce prejudice against transgender people.

In Western Europe, the received view defines gender as binary, stable, and natural. That is, there are only two genders, man and woman, and every adult is either one or the other. A person who is born as one gender stays that gender for the rest of their life. A person's gender is rooted in their biology, and determined before their birth. In other words, if one is born a girl, then one grows up to be a woman, and stays a woman for the rest of one's life. Importantly, this view holds that biology and *gender identity* align. That is, someone born as one gender also grows up to think of themself as that gender. Further, this view holds that biology and *gender expression* align. That is, someone born as one gender grows up to dress and behave as expected of someone of that gender. Boys and men thus express their gender in masculine ways, while girls and women express their gender in feminine ways. Overall, the received view is an essentialist conception of gender in which biology, identity, and expression align.

The received view, however, is mistaken. Each of its components is contradicted by counterexample. Even assuming that gender was binary, stable, and natural, then gender identity and expression need not align with biology. The necessary correspondence between identity and biology is contradicted by the lived experiences of transgender people. A transgender woman, for example, was assigned the biological sex 'male' at birth, but thinks of herself and lives as a woman. The necessary correspondence between expression and biology is contradicted by gender non-conforming individuals and behaviours. Among Dutch children, for example, between 2.4% and 3.3% of 10-year-olds regularly "behave[d] like opposite sex" or "wishe[d] to be of opposite sex" (van Beijsterveldt, Hudziak & Boomsma, 2006). Among British adults, 35% of women described themselves as "completely feminine", while 28% of men described themselves as "completely masculine" in a representative survey (YouGov, 2016). Like for sexuality, 18 to 24-year-olds were least likely to describe their gender expression as completely aligned with their gender identity (2% of men and 14% of women, see Figure 3.3). Contrary to the essentialist conception, these examples show that biology, identity, and expression need not align.

The received view is also mistaken about gender being natural, binary, and stable. In this view, biological sex—gonads, genitals, and genes—determines whether someone is a woman or a man. Crucially, gender assignment is viewed as unambiguous. Contrary to this view, intersex people—about 1 in 100 people—have chromosomes or anatomies that do not fit the binary conception of biological sex (Ainsworth, 2015). Similarly, gender identity and expression need not be binary or stable. Genderqueer people, for example, identify as neither men nor women, as some combination of genders, or as between or beyond genders (Richards et al., 2016). Though gender seems stable, gender expression, and sometimes gender identity (e.g., for genderfluid people), can change within someone's lifetime. Together, these examples show that gender is continuous and fluid, rather than binary and stable (for a review, see Dunham & Olson, 2016).

The received view of gender is not only mistaken, but also harmful. Tee and Hegarty (2006) found that believing gender to be biological and binary was associated with opposition to transgender persons' civil rights. Roberts, Ho, Rhodes and Gelman (2017) found that essentialist beliefs correlated with support for legislation prohibiting transgender people from using restrooms that correspond to their gender. More generally, the essentialist conception of gender is harmful because it excludes transgender people,

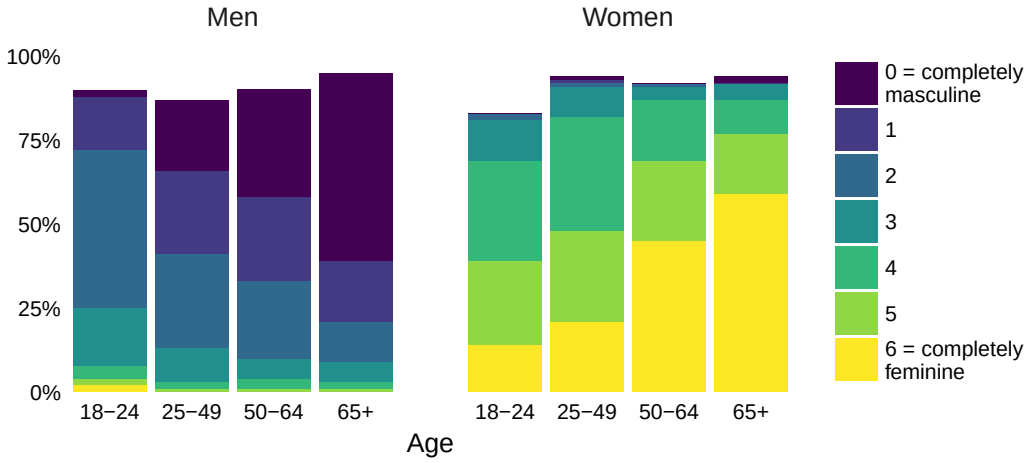

Figure 3.3. Results from a representative survey of British adults (YouGov, 2016). Percentages of men (left) and women (right) who placed themselves along a scale from 0 = *completely mascluine* to 6 = *completely feminine*, excluding *don't know* and *neither masculine nor feminine* responses.

including genderqueer people, from the gender they identify with. Accordingly, Nagoshi et al. (2008) defined gender essentialism as one component of transphobia (that is, anti-transgender prejudice). If binary and stable conceptions of gender harm transgender people, then fostering more continuous and fluid conceptions could be a means to reduce discrimination against transgender people.

In this chapter, I develop and validate a scale measuring *social identity continuity* and *social identity fluidity* for gender. Above, I have argued that gender spans two dimensions: gender identity and gender expression. In Study 2, I formulate items for four combinations of the two *construct* dimensions (continuity, fluidity) and the two *gender* dimensions (identity, expression). That is, I measure social identity continuity as the extent to which individuals perceive their appearance and demeanour as somewhere in between masculine and feminine (Figure 3.4d), and as the extent to which they find the terms 'male' and 'female' too narrow to describe their gender (Figure 3.4b). I measure social identity fluidity as the extent to which individuals perceive their gender identity (Figure 3.4c) and expression (Figure 3.4e) as fluid and changing, rather than as fixed and stable throughout their lifetime. In Study 2, I develop this scale using exploratory factor analysis. In Study 3, I validate this scale using confirmatory factor analysis. In line with the *bridging-binaries hypothesis*, I test whether contact with gender minorities predicts continuity and fluidity, and whether continuity and fluidity, in turn, predict less transphobia.

Present research

The present research has two goals: *developing* a scale for measuring social identity continuity and fluidity, and *testing* the hypothesised relationships of these constructs with intergroup contact and intergroup bias.

In Studies 1–2, I followed DeVellis's (2016) guidelines for scale development. First, I generated a pool of items that reflect the constructs of interest, *social identity continuity* and *social identity fluidity*. I only included items that align with the definitions proposed

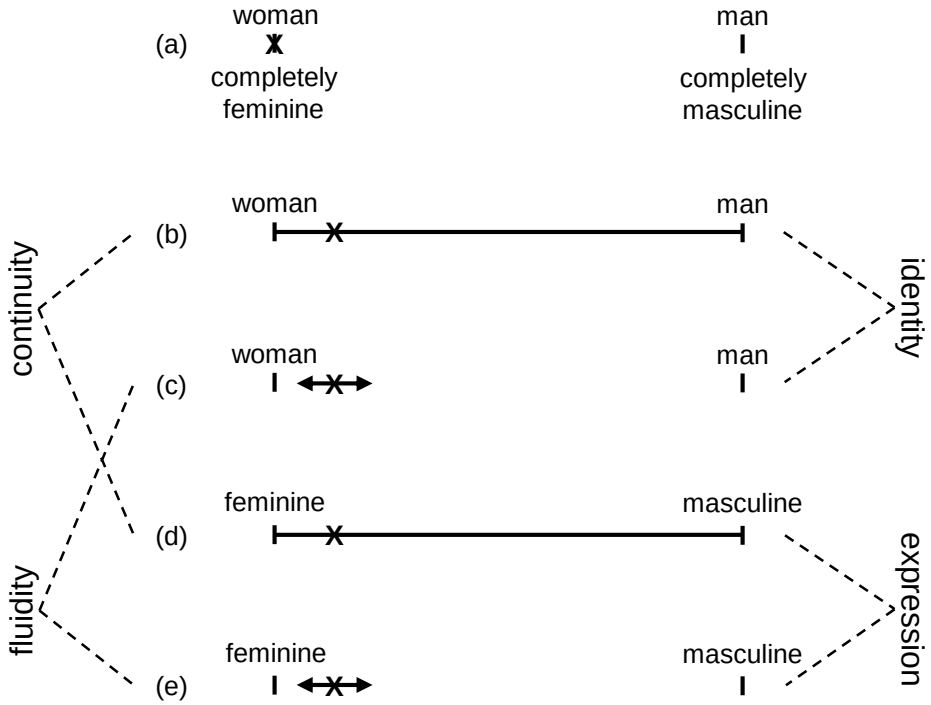

Figure 3.4. Schematic representations of different conceptions of gender (for a woman). Row (a) shows a binary, stable understanding of gender; (b) and (d) show more continuous conceptions; (c) and (e) show more fluid conceptions; (b) and (c) focus on self-categorisation; (d) and (e) focus on expression.

in the preceding sections. I excluded items that focus on participants' perceptions of sexuality and gender *in general*, rather than of *their own* sexuality and gender. I included positively and negatively worded items, testing which better measure the constructs of interests.[4] I followed recommendations for avoiding ambiguous, difficult-to-read, or double-barrelled item wordings (DeVellis, 2016; Lenzner & Menold, 2016). Second, I administered the proposed items to a development sample. Participants responded to the items on a seven-point Likert scale ranging from 1 = *strongly disagree* to 7 = *strongly agree*. I used a seven-point scale to allow sufficient variation in responses. I did not use more scale points to avoid risking participants' ability to discriminate between the meaning of adjacent scale points (Menold & Bogner, 2016). I included 4 = *neither agree nor disagree* as a scale midpoint to accommodate respondents with a 'neutral' or 'moderate' opinion (Menold & Bogner, 2016). I followed DeVellis's (2016) sample size recommendations for scale development, recruiting close to 300 participants per study.

Third, I used responses from the development samples to evaluate the items' performance. Initially, I described and compared items' means, variances, and inter-item correlations. I preferred items with higher inter-item correlations, greater variances, and means closer to the scale midpoint. Next, I used *exploratory factor analysis* to determine the factor structure of the full set of items. I extracted a varying number of factors, as long as at least three items loaded on each factor,[5] retaining factors on the basis of model fit indices, parsimony, and conceptual clarity. I preferred simpler over more complex scale structures to avoid overfitting. Then, I assessed the reliability of each scale factor. I report McDonald's (1999) omega coefficient (ω), and its bootstrap confidence interval (Kelley, 2018), as a measure of *internal consistency* reliability.[6] Finally, I decided which items to include in the next iteration of the scale. I included items on the basis of descriptive statistics, factor loadings, and internal consistency, and removed items that had low inter-item correlations, low or ambiguous factor loadings, or decreased the internal consistency of the scale.

After developing the scale, I estimated how the resulting scale dimensions relate to various other constructs. In Studies 1–2, I assess the scale's *construct validity* by testing whether scale dimensions correlate with variables I would expect the underlying constructs, social identity continuity and fluidity, to correlate with. First, I expect scale dimensions to correlate with related constructs. That is, I expect scale dimensions to correlate *negatively* with essentialist beliefs about sexual orientation, essentialist beliefs about gender, gender norm endorsement, and group identification, and to correlate *positively* with same-gender attraction. Second, I expect scale dimensions to correlate *negatively* with need for closure since an aversion to ambiguity should motivate discrete and stable conceptions of sexuality and gender. Third, if as hypothesised, more continuous and fluid

[4]Some researchers strive to include negatively and positively worded in any scale to counter various response biases. DeVellis (2016), however, argued that opposite-valence items stand to confuse some respondents and often decrease scale performance. I included both types of items, but do not constrain scale dimensions to contain both.

[5]In (confirmatory) structural equation models, three indicators form a just-identified latent variable, while one or two indicators leave a latent variable under-identified. Little, Cunningham, Shahar and Widaman (2002) argued that just-identified latent variables should be referred over under-identified and over-identified ones.

[6]Like Cronbach's alpha, McDonald's omega assumes that a set of items represents a unidimensional factor. Unlike alpha, omega avoids many restrictive and unrealistic assumptions that often bias estimates of internal consistency (Dunn, Baguley & Brunsden, 2014). Like alpha, omega ranges from 0 to 1, with higher values indicating more internal consistency.

conceptions are related to prejudice, then scale dimensions should correlate *negatively* with social dominance orientation, a measure of generalised prejudice. At the same time, I consider whether any of the correlations are strong enough to suggest that scales measure an identical construct, as discussed throughout this chapter. Together, these relationships indicate whether the scale reflects the construct it is supposed to measure. In Studies 1–2, I also test how social identity continuity and fluidity, as measured by the new scales, relate to intergroup contact (hypothesis 1) and to intergroup bias (hypothesis 2). Studies 1–2 provide cross-sectional evidence for or against the hypothesised relationships.

In Study 3, I examine the scales developed in Studies 1–2 in a large-scale, longitudinal survey. First, I use *confirmatory factor analysis* to replicate the factor structure determined in Studies 1–2. Second, I assess the reliability of the resulting scale dimensions by estimating their *internal consistency* (as omega coefficients), by estimating their *temporal stability* (as test-retest correlation coefficients), and by testing for *measurement invariance* (across time and translations). Third, I test whether intergroup contact predicts more social identity continuity and fluidity (hypothesis 1), and whether these constructs, in turn, predict less intergroup bias (hypothesis 2). I measured all variables over three time points to test whether social identity continuity and fluidity mediate the relationship between contact and bias. Study 3 thus seeks to confirm the factor structure of the novel scales, and provides longitudinal evidence for or against the hypotheses proposed above.

Study 1

As part of a wider study, I administered a preliminary set of items to a sample of university students. I recruited sexual-majority and sexual-minority participants, and asked participants to choose whether they wanted to be shown questions and answer options for heterosexual or non-heterosexual people, based on their own sexual orientation. Participants who chose the questionnaire for heterosexual people filled in the measures described below, including ten novel items measuring social identity continuity and fluidity. Participants who chose the questionnaire for non-heterosexual people answered questions on perceived discrimination and collective action (see Chapter 5), and were not included in the present study. In this chapter, I am interested in heterosexual participants' *subjective perceptions* of their sexuality—not in substantive differences in sexual orientation. Accordingly, I included only participants who fulfilled a strict definition of being heterosexual (see *Participants* section).

Method

Participants

Oxford University students of "all sexual orientations" were invited to participate in an online survey for the chance to win up to £100. Of 462 respondents, 305 chose to fill in the questionnaire for heterosexual people. Of these, I excluded 46 respondents who did not identify as heterosexual/straight, had had sexual relations with a same-sex partner, or were attracted to persons of their own gender more than 50% of the time. I excluded another 17 respondents who did not complete any of the continuity/fluidity items. Study 1 thus included 242 heterosexual respondents (156 women, 84 men, 2 "none of these/other"; Mdn_{age} = 22, age range: 18–38 years).

Measures

I measured *sexual orientation* as a multidimensional construct, spanning self-identification, sexual attraction, and sexual behaviour (Sell, 2007). I assessed participants' sexuality with three items (Falomir-Pichastor & Hegarty, 2014): "How do you define your sexual orientation?" (open response), "Have you previously had sexual relationships with a same-sex partner?" (*yes*, *no*), and "How often do you feel sexual attraction to a person of your own sex?" (0 = *never* to 100 = *all the time*). Participants then chose to complete the questionnaire for non-heterosexual or heterosexual students, so that heterosexual respondents completed measures with heterosexual people as the relevant ingroup, and non-heterosexual people as the relevant outgroup.

Intergroup contact was measured as how often, from 1 = *never* to 5 = *very often*, respondents had a variety of positive and negative experiences with non-heterosexual people (Fell et al., 2016; Stephan et al., 2002). *Negative contact* (α = .64) was measured using five items: being verbally abused, intimidated, threatened with harm, ridiculed, and made to feel unwanted. *Positive contact* (α = .85) was measured using five items: being supported, helped, complimented, befriended, and made to feel welcome.

Group identification was measured using four items (adapted from Doosje, Branscombe, Spears & Manstead, 1998): "Being heterosexual is an important part of who I am", "I identify with other heterosexual people", "Overall, being heterosexual has a lot to do with how I feel about myself", and "I see myself as heterosexual" (1 = *strongly disagree*, 7 = *strongly agree*; α = .79).

Essentialist beliefs about sexual orientation were measured as the extent to which participants agreed with ten items from Haslam and Levy's (2006) scale (1 = *strongly disagree*, 7 = *strongly agree*). Five items measured how *discrete* participants' believed sexual orientation to be: "Sexual orientations are categories with clear and sharp boundaries: People are either homosexual or heterosexual", "Homosexual people have a necessary or defining characteristic, without which they would not be homosexual", "Heterosexual and homosexual people are not fundamentally different" (reversed), "Bisexual people are fooling themselves and should make up their minds", and "Knowing that someone is homosexual or heterosexual tells you a lot about them" (α = .72). Five items measured how *immutable* and inherent participants believed sexual orientation to be: "Sexual orientation is caused by biological factors", "Whether a person is homosexual or heterosexual is pretty much set early on in childhood", "People cannot change their sexual orientation", "Homosexuality and heterosexuality are innate, genetically based tendencies", and "Doctors and psychologists can help people change their sexual orientation" (reversed; α = .71).

Outgroup attitudes were measured using Wright, Aron, McLaughlin-Volpe and Ropp's (1997) General Evaluation Scale. Respondents rated how they felt about non-heterosexual people in general on 7-point adjective scales: warm–cold (reversed), negative–positive, friendly–hostile (reversed), suspicious–trusting, respect–contempt (reversed), and admiration–disgust (reversed). Items formed a reliable scale (α = .95) with higher scores indicating more favourable attitudes.

Scale development

I generated ten items for measuring *social identity continuity* and *social identity fluidity*. Five items measured how *continuous* participants understood their sexual identity and sexual attraction to be, e.g., "I am exclusively heterosexual and not at all homosexual" (reversed). Five items measured how *fluid* participants understood their sexual identity

and sexual attraction to be, e.g., "The extent to which I am attracted to men or women may change to some degree between different situations". Participants indicated the extent to which they agreed or disagreed with each item, on a scale from 1 = *strongly disagree* to 7 = *strongly agree*. Table 3.1 shows all items included in Study 1.

Results

I follow the analytic strategy outlined above. First, I describe all scale items in terms of their means, standard deviations, and inter-item correlations. Second, I use exploratory factor analysis to extract a varying number of latent factors. I judge model fit by the indices described in Chapter 2. I retain factors on the basis of model fit, parsimony, and conceptual clarity. Third, I estimate correlations between the scale factors and variables I expected to be associated with social identity continuity and fluidity. Finally, I use exploratory structural equation modelling (Asparouhov & Muthén, 2009) to test the hypothesised relationships between the exploratory scale factors, intergroup contact, and sexual prejudice.

For all items, participants' responses ranged from 1 = *strongly disagree* to 7 = *strongly agree*. Responses to items 4, 6, 7, 8, and 9 averaged around the centre of the scale, $3.53 < Ms < 4.69$. Responses to items 2, 3, and 10 averaged around the upper third of the scale, $5.06 < Ms < 5.76$. Responses to item 1 (reverse coded) averaged in the lower third of the scale, $M = 2.28$. Standard deviations ranged from 1.11 to 1.79 units on the response scale. Figure 3.5 shows coefficients for all inter-item correlations. Correlations ranged from $r = -.11$ to $r = .79$. On average, items 2, 10, and 3 had the lowest correlations with other scale items, $-.02 < Ms_r < .22$, $-.11 < rs < .38$. The other items had mean correlations between $M_r = .30$ and $M_r = .42$ with other scale items, $-.11 < rs < .79$. Together, these findings indicate that the proposed items reflect more than one dimension, that some items do not fit into the overall scale, or both.

In an initial analysis, I estimated one, two, three, and four-factor models for all ten items of the proposed scale. The four-factor model fit better than the three-factor model, $\Delta\chi^2(7) = 42.61$, $p < .001$, $\Delta CFI = .04$, but yielded ambiguous factor loadings. Of the remaining models, only the three-factor model showed acceptable fit, $\chi^2(18) = 49.48$, $\chi^2/df = 2.75$, $CFI = .96$, $TLI = .91$, $RMSEA = .08$. In this model, item 2 did not load on any of the three factors, while items 3, 5, and 10 loaded on a common factor. On closer inspection, all four items concerned sexuality *in general* rather than focusing on participants' *own* sexuality. Items 2 and 10, for example, refer to "an individual" or "some people" rather than to the participants themselves (see Table 3.1). As I was interested in participants' subjective perceptions of their own sexuality, I excluded items 2, 3, 5, and 10 from further analyses.

I estimated models with one or two exploratory factors for the remaining six items.[7] The two-factor solution fit acceptably well, $\chi^2(4) = 11.4$, $\chi^2/df = 2.85$, $CFI = .99$, $TLI = .96$, $RMSEA = .09$, and fit better than the one-factor solution, $\Delta\chi^2(5) = 35.34$, $p < .001$, $\Delta CFI = .04$. As the one-factor solution showed close-to-acceptable model fit, $\chi^2(9) = 46.74$, $\chi^2/df = 5.19$, $CFI = .95$, $TLI = .91$, $RMSEA = .13$, I report results for both solutions. Table 3.1 shows rotated factor loadings for both models. In the two-factor solution, items 1, 4, and 5 loaded on Factor 1, while items 7, 8, and 9 loaded on Factor 2. In line with the hypothesised dimensions, Factor 1 measured how

[7] As discussed under *Present research*, I only consider models with at least three items per factor. This excludes three-factor solutions for only six items.

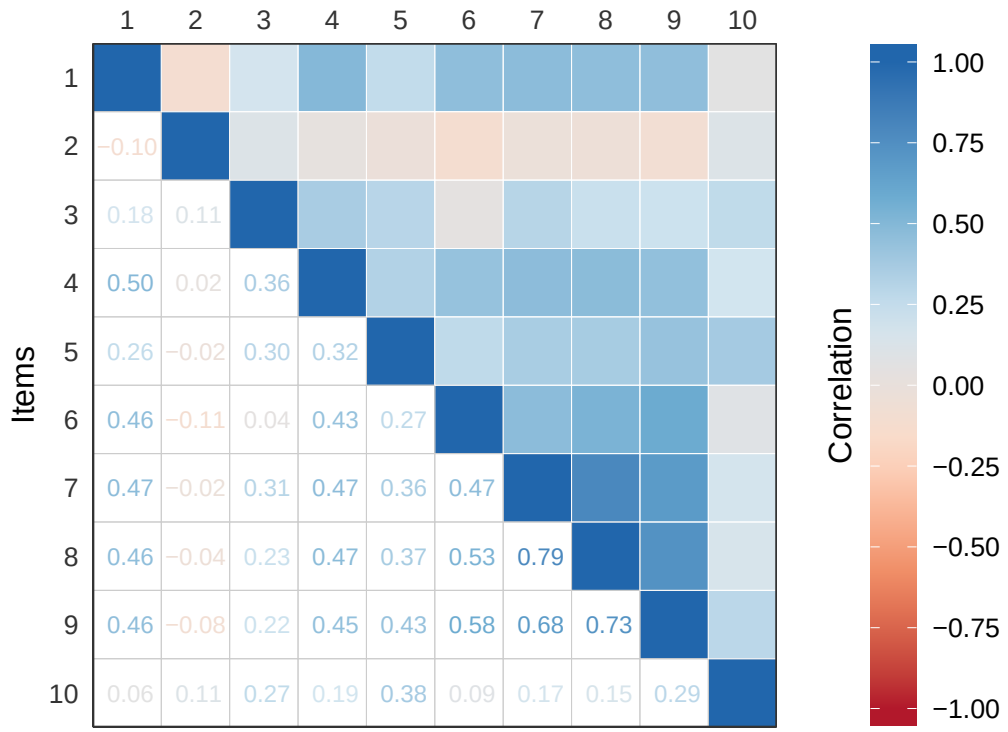

Figure 3.5. Inter-item correlations for social identity continuity and fluidity for sexuality (Study 1).

continuous participants considered their own sexual attraction and categorisation to be. Factor 2 measured how *fluid* participants saw their sexual attraction to men and women to be. In the two-factor solution, Factor 2 (hereafter: *fluidity*) showed very good internal consistency (ω = .90, [.87, .92]). Factor 1 (hereafter: *continuity*) proved less reliable (ω = .67, [.60, .75]), and should be interpreted with caution. The one-factor solution (hereafter: *continuity/fluidity*) showed good reliability (ω = .87, [.84, .90]), assuming that the items form a unidimensional scale.

Next, I estimated latent variable correlations between the one-factor and two-factor constructs and other variables hypothesised to with social identity continuity and fluidity. Figure 3.6 shows (in red) correlation coefficients (ϕ) with 95% confidence intervals. In the two-factor model, the continuity and fluidity factors were strongly correlated, ϕ = .73, [.60, .86]. As expected, factors in both models correlated negatively with discreteness beliefs and group identification. Factors did not correlate with immutability beliefs. As none of these relationships was as strong as the inter-factor correlations, the factors appear to be related to, but distinct from earlier constructs. As expected, all factors correlated positively with same-gender attraction. This correlation, however, was strong enough to call into question whether the proposed scale measured a construct distinct from same-gender attraction. Figure 3.6 also shows (in blue) partial correlation coefficients, controlling for same-gender attraction. Controlling for same-gender attraction did not substantially weaken correlations with other variables, suggesting that the proposed factors measured something beyond sexual attraction.

In two structural equation models, I tested the hypothesised relationships between (exploratory) scale factors, intergroup contact, and outgroup attitudes. A two-factor model, in which I regressed two exploratory factors on positive and negative contact,

Table 3.1. Geomin-rotated factor loadings for the 1-factor and 2-factor solutions in Study 1. Asterisks (*) mark reverse corded items. Only factor loadings for which $p < .05$ are shown.

			2-factor solution	
#	Item	Factor 1	Factor 1	Factor 2
1	I am exclusively heterosexual and not at all homosexual.*	.56	.73	
2	*An individual's sexual orientation is best understood as a combination of the extent he or she is attracted to women and the extent to which he or she is attracted to men.*	-	-	-
3	*I can find a person of the same sex attractive and still be heterosexual.*	-	-	-
4	I cannot imagine being attracted to both men and women.*	.56	.61	
5	*The terms "homosexual" and "heterosexual" are too narrow to describe the full range of sexual attraction and behaviour.*	-	-	-
6	Whether I am attracted to men or women is stable and unchangeable.*	.63	.50	
7	The extent to which I am attracted to men or women may have changed to some degree over the course of my lifetime.	.85		.80
8	The extent to which I am attracted to men or women may change to some degree between different situations.	.90		.95
9	The extent to which I am attracted to men or women may change to some degree in the future.	.82		. 63
10	*Some people may change in the extent to which they are attracted to men or women throughout their lifetime.*	-	-	-

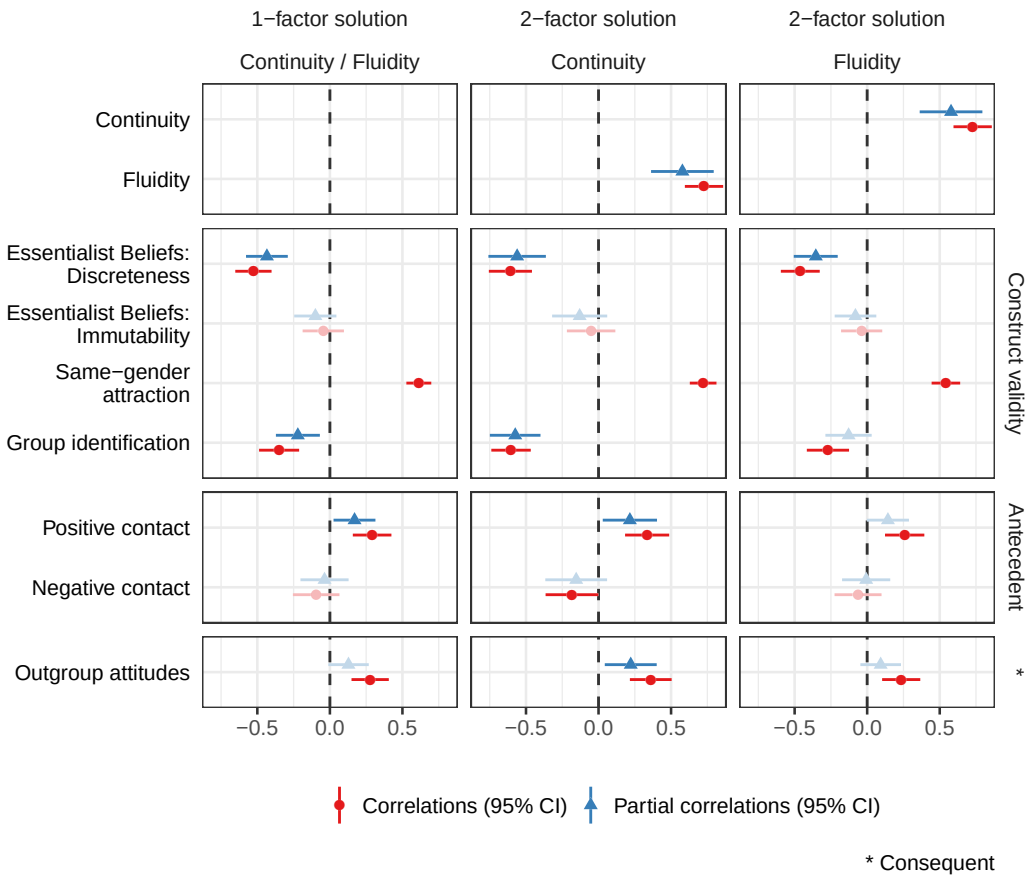

Figure 3.6. Correlations and partial correlations (controlling for same-gender attraction) between scale factors and variables related to construct validity, and with variables hypothesised as antecedent and consequent (Study 1).

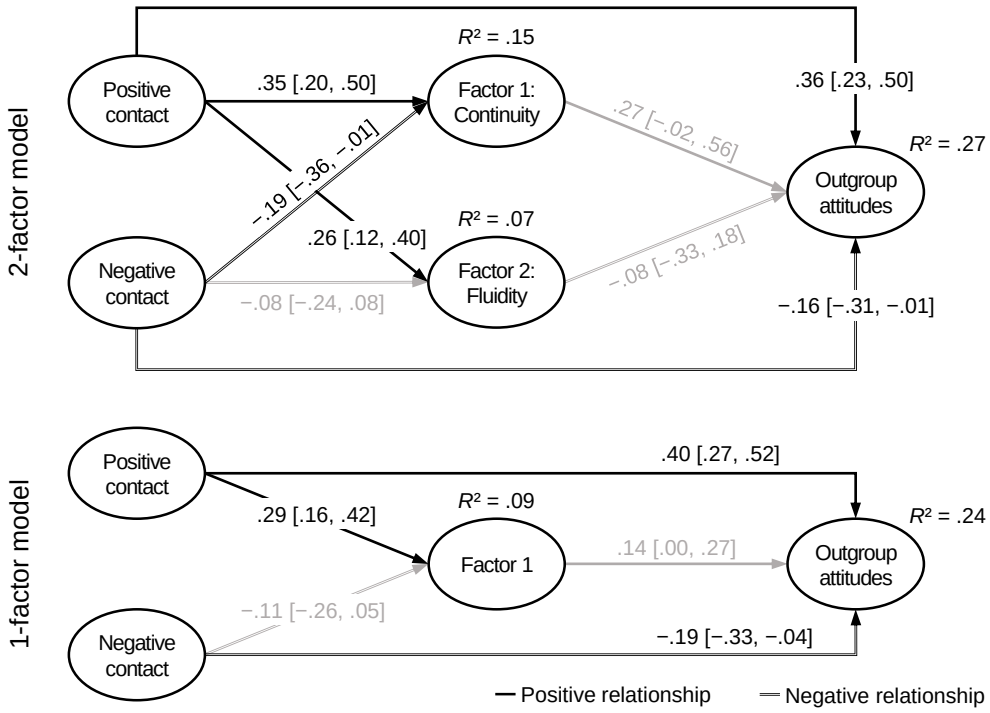

Figure 3.7. Exploratory structural equation models for 1-factor and 2-factor solutions, with standardised coefficients and 95% confidence intervals (Study 1).

provided good fit, $\chi^2(76) = 75.3$, $\chi^2/df = 0.99$, $CFI = 1.00$, $TLI = 1.00$, $RMSEA = 0.00$. Figure 3.7 (top) shows standardised coefficients for all model paths. As predicted, positive contact with non-heterosexual people was associated with higher values of the *continuity* ($\beta = 0.35$, $[0.20,\ 0.50]$) and *fluidity* ($\beta = 0.26$, $[0.12,\ 0.40]$) factors, see hypothesis 1. Negative contact was negatively associated with *continuity* ($\beta = -0.19$, $[-0.36,\ -0.10]$), but not *fluidity* ($\beta = -0.08$, $[-0.24,\ 0.08]$). Other than expected, neither factor was associated with attitudes toward non-heterosexual people, see hypothesis 2. Figure 3.7 (bottom) shows coefficients for a one-factor model. This model also provided acceptable fit, $\chi^2(84) = 120.23$, $\chi^2/df = 1.43$, $CFI = .98$, $TLI = .98$, $RMSEA = .04$. Again, positive contact was associated with higher factor scores ($\beta = 0.29$, $[0.16,\ 0.42]$). However, the model provided little evidence for the hypothesised relationship between continuity/fluidity and outgroup attitudes ($\beta = 0.14$, $[0.00,\ 0.27]$).

Study 2

Study 1 provided preliminary correlational evidence that individuals differ in the extent to which they understand their sexuality as *continuous* and *fluid*, and that individuals with more positive contact with sexual minorities understand their own sexuality as more *continuous* and *fluid* (hypothesis 1). Other than expected, individuals with more continuous and fluid conceptions did not hold more favourable attitudes toward sexual minorities (hypothesis 2).

In Study 1, I developed an initial scale from a small set of items. In Study 2, I generated a larger pool of 28 items, including six items from Study 1, to measure how continuous and fluid participants understood their sexual identity and attraction to be

(see *Sexuality as case study*). I also included a novel set of 27 items to measure how continuous and fluid participants understood their gender identity and expression to be (see *Gender as case study*). I also included a wider set of variables to assess construct validity. I expect essentialist beliefs about gender, gender norm endorsement, social dominance orientation, and need for closure to correlate negatively with the proposed scale. I test the hypothesised relationships of continuity/fluidity with intergroup contact and intergroup bias in two domains: relations with sexual minorities (non-heterosexual people) and gender minorities (transgender people).

As in Study 1, I asked participants to choose whether they wanted to be shown questions and answer options for heterosexual or non-heterosexual people. Participants who chose the questionnaire for non-heterosexual people answered questions on discrimination and collective action (see Chapter 5). As some items referred to their gender, participants also chose whether they wanted to be shown questions and answer options for women or men.

Method

Participants

Oxford University students of "all sexual orientations and genders" were invited to participate in an online survey for the chance to win up to £100. Of 375 respondents, 261 chose to fill in the questionnaire for heterosexual people. Of these, I excluded 36 respondents who did not identify as heterosexual/straight, had had sexual relations with a same-gender partner, or were attracted to persons of their own gender more than 50% of the time. I excluded another 10 respondents who did not complete any of the sexuality-related continuity/fluidity items. Study 2a thus included 215 heterosexual respondents (140 women, 75 men; Mdn_{age} = 22, age range: 18–39 years). Study 2b included 309 cisgender respondents (207 women, 102 men; Mdn_{age} = 21, age range: 18–42 years) who did not report a non-binary gender identity, did not identify as transgender, and completed at least some of the gender-related continuity/fluidity items.

Measures

All participants were asked to describe themselves in terms of age, gender, and sexual orientation. Participants' *gender identity* was recorded as their open response to the question: "What is your gender?". *Sexual orientation* was measured with four items: "How would you define your sexual orientation?" (open response), "Have you previously had sexual relationships with a same-sex partner?" (*yes*, *no*), "How often do you feel attracted to women?" (0–100), and "How often do you feel attracted to men?" (0–100). Participants then chose to complete the questionnaires for heterosexual or non-heterosexual people, so that heterosexual respondents completed measures with heterosexual people as the relevant ingroup and non-heterosexual people as the relevant outgroup. Participants also chose whether they wanted to see questions and answer options for women or men, so that all items conformed to their chosen gender. For both questions, participants could choose that they thought of themselves as "none of these/other", in which case they were not shown any gender-specific or sexuality-specific items.

Measures for heterosexual respondents

In the first part of the study, heterosexual respondents filled in measures concerning sexual orientation. Intergroup contact was measured as contact quantity, positive contact, and negative contact. *Contact quantity* was measured as how many people participants knew who are gay, lesbian, and bisexual (1 = *none*, 5 = *more than ten*; α = .69). *Positive contact* (α = .85), *negative contact* (α = .80), *group identification* (α = .75), *discreteness beliefs* (α = .65), and *immutability beliefs* (α = .74) were measured with the same items as in Study 1.

Outgroup attitudes were assessed with ten feeling thermometer items (Converse, Dotson, Hoag & McGee, 1980). Participants rated how they felt toward criminals, straight/heterosexual men, bisexual people, immigrants, gay men, straight/heterosexual women, women who have sex with women, British people, men who have sex with men, and lesbian women (from 0 = *cold* to 100 = *warm*). I assessed intergroup bias with five indices (α = .91): the difference in participants' ratings of straight men and gay men; straight men and men who have sex with men; straight women and lesbian women; straight women and women who have sex with women; and straight men/women and bisexual people. Higher scores indicated a greater preference for the heterosexual ingroup.

Measures for cisgender respondents

In the second part of the study, participants completed measures concerning gender identity and expression. *Group identification* was measured using four items (adapted from Doosje et al., 1998): "Being a [woman/man] is an important part of who I am", "I identify with other [women/men]", "Overall, being a [woman/man] has a lot to do with how I feel about myself", and "I see myself as a [woman/man]" (1 = *strongly disagree*, 7 = *strongly agree*; α = .84). *Gender norm endorsement* was measured with two items (adapted from Sanchez & Crocker, 2005): "How important is it for you to be similar to the ideal [woman/man]?", and "To what extent is being similar to the ideal [woman/man] an important part of who you are?" (1 = *not at all*, 7 = *a great deal*; r = .82).

Essentialist beliefs about gender were measured with six items from Tee and Hegarty's (pp. 72–73 2006) scale: "There are only two genders, male or female", "All adults identify as either male or female", "If you are either male or female, then you are that gender for all time", "All males have a penis and females all have a vagina", "Gender is determined by biological factors, such as genes and hormones, before birth", "There are only three factors that determine gender: internal gonads (e.g. testes), external genitalia (e.g. penis) and chromosomes (i.e. xx—female and xy—male)" (1 = *strongly disagree*, 7 = *strongly agree*; α = .88).

Participants completed measures about their experiences with and attitudes toward transgender people. Before filling in these measures, participants read definitions of relevant terms. Participants read that "transgender (often abbreviated 'trans') people include anyone who feels that the sex which was assigned to them at birth incompletely describes or fails to describe them". They read that *trans women* were "assigned 'male' at birth but think of themselves and live as women", that *trans men* were "assigned 'female' at birth but think of themselves and live as men", that *genderqueer people* "identify as neither men nor women, as some combination of genders, or as beyond or between genders", and that *gender non-conforming people* "may think of themselves in terms of the gender they were assigned at birth but do not adhere to gender-based norms of society, e.g., cross-dressers". The questionnaire "will refer to all of these people—trans women,

trans men, genderqueer, or gender nonconforming—as transgender people". Participants then chose whether one or more of these labels described them. Only participants who responded that they were not transgender filled in the remaining measures.

Contact with transgender people was measured as contact quantity, positive contact and negative contact. *Contact quantity* was measured as how many trans women, trans men, genderqueer people, and gender non-conforming people participants knew (1 = *none*, 5 = *more than ten*; α = .81). *Positive contact* (α = .95) and *negative contact* (α = .76) were measured as how often respondents had each of five positive and five negative experiences with transgender people (1 = *never*, 5 = *very often*; described in Study 1).

Outgroup attitudes were assessed with ten feeling thermometer items (Converse et al., 1980). Participants rated how they felt toward men, sex workers, trans women, right-wing politicians, women, trans men, cross-dressing people, feminists, genderqueer people, and left-wing politicians (from 0 = *cold* to 100 = *warm*). I assessed intergroup bias with four indices (α = .92): the difference in participants' ratings of men and trans men; women and trans women; men/women and genderqueer people; and men/women and cross-dressing people. Higher scores indicated that participants felt less favourable toward transgender people than toward men and women in general.

Transphobia was measured with six items from Nagoshi et al.'s (2008) scale: "I don't like it when someone is flirting with me, and I can't tell if they are a man or a woman", "I think there is something wrong with a person who says that they are neither a man nor a woman", "I would be upset, if someone I'd known a long time revealed to me that they used to be another gender", "I avoid people on the street whose gender is unclear to me", "When I meet someone, it is important for me to be able to identify them as a man or a woman", and "I am uncomfortable around people who don't conform to traditional gender roles, e.g., aggressive women or emotional men" (1 = *strongly disagree*, 7 = *strongly agree*; α = .86).

In the final part of the study, all participants completed two individual differences measures. *Social dominance orientation* was measured with four items (adapted from Sidanius, Levin, van Laar & Seers, 2008, p. 344): "It's probably a good thing that certain groups are at the top and other groups are at the bottom", "We should do what we can to equalise conditions for different groups" (reversed), "Inferior groups should stay in their place", and "We should increase social equality" (reversed; 1 = *strongly disagree*, 7 = *strongly agree*; α_{2a} = .86, α_{2a} = .84). *Need for closure* was measured with Roets and van Hiel's (2011) 15-item scale (α_{2a} = .84, α_{2a} = .84).

Scale development

I generated 28 items for measuring *social identity continuity* and *social identity fluidity* for participants' perceptions of their own sexuality. Twelve items measured how *continuous* participants understood their sexual identity and attraction to be. Sixteen items measured how *fluid* participants expected their sexual identity and attraction to be across time and situations. Table A.1 shows all items included in Study 2a. I also generated 27 items for studying participants' perceptions of their own gender. Twelve items measured how *continuous* participants understood their gender identity and expression to be. Fifteen items measured how *fluid* participants expected their gender identity and expression to be across time and situations. Table A.2 shows all items included in Study 2b. Participants indicated the extent to which they agreed or disagreed with each item, on a scale from

1 = *strongly disagree* to 7 = *strongly agree*.

Results

I follow the same analytic strategy as in Study 1. First, I describe all scale items in terms of their means, standard deviations, and inter-item correlations. Second, I use exploratory factor analysis to determine the factor structure of the proposed scales. Third, I estimate correlations between the scale factors and variables I expected to correlate with social identity continuity and fluidity. Finally, I use exploratory structural equation modelling to test the hypothesised relationships between the exploratory scale factors, intergroup contact, and intergroup bias. In Study 2a, I analyse heterosexual respondents' perceptions of sexuality. In Study 2b, I analyse cisgender respondents' perceptions of gender.

Study 2a: Sexuality

For most items, participants' responses ranged from 1 = *strongly disagree* to 7 = *strongly agree* (except items 1, 16, 26). Responses to items 1, 7, 11, 12, 15, 16, 17, 18, 20, and 21 (n = 10) averaged in the lower third of the scale, $2.35 < Ms < 2.98$. Responses to items 4 (M = 5.89) and 5 (M = 5.67) averaged in the upper third of the scale. Responses to all other items (n = 16) averaged around the centre of the scale, $3.07 < Ms < 4.78$. Standard deviations ranged from 1.14 to 1.90 units on the response scale. Figure 3.8 shows coefficients for all inter-item correlations. Correlation coefficients ranged from $r = -.04$ to $r = .83$. Items 4, 5, and 8 had the lowest average correlations with other scale items, $.12 < Ms_r < .24$. I excluded these items from further analyses. Without these items, average inter-item correlations ranged from $M_r = .37$ to $M_r = .54$. As in Study 1, inter-item correlations indicate that the items reflect more than one dimension, that some items do not fit into the overall scale, or both.

In an initial analysis, I estimated models with one to eight exploratory factors for the 25 remaining items. I removed 8 items that (substantially) loaded on more than one factor, or that did not have any significant factor loadings in the well-fitting models (see Table A.1). I then estimated models with one to five exploratory factors for the 17 remaining items. Figure 3.9 shows changes in factor structure across the five models. The two-factor solution split items into one factor that measured to what extent participants thought that their sexual identity and attraction might change in the future (F2, *Fluidity: Future*), and one factor that contained the remaining items. The three-factor solutions split off a third factor that measured to what extent participants saw their sexual identity as continuous (F3, *Continuity: Identity*). The four-factor solution split the remaining items into one factor that measured to what extent participants thought of their sexual attraction as continuous (F1, *Continuity: Attraction*), and one factor that measured to what extent participants believed that their sexuality had changed over the course of their lifetime (F4, *Fluidity: Past*). The five-factor solution formed an additional factor from item 10 (F5).

Models with one factor ($\chi^2(119) = 794.06$, $\chi^2/df = 6.67$, $CFI = .74$, $TLI = .71$, $RMSEA = .16$), with two factors ($\chi^2(103) = 398.89$, $\chi^2/df = 3.87$, $CFI = .89$, $TLI = .85$, $RMSEA = .12$), or with three factors ($\chi^2(88) = 236.87$, $\chi^2/df = 2.69$, $CFI = .94$, $TLI = .91$, $RMSEA = .09$) did not represent the observed responses well. The four-factor model ($\chi^2(74) = 110.04$, $\chi^2/df = 1.49$, $CFI = .99$, $TLI = .98$, $RMSEA = .05$) and the five-factor model ($\chi^2(61) = 64.90$, $\chi^2/df = 1.06$, $CFI > .99$, $TLI = 1.00$, $RMSEA = .02$) showed good

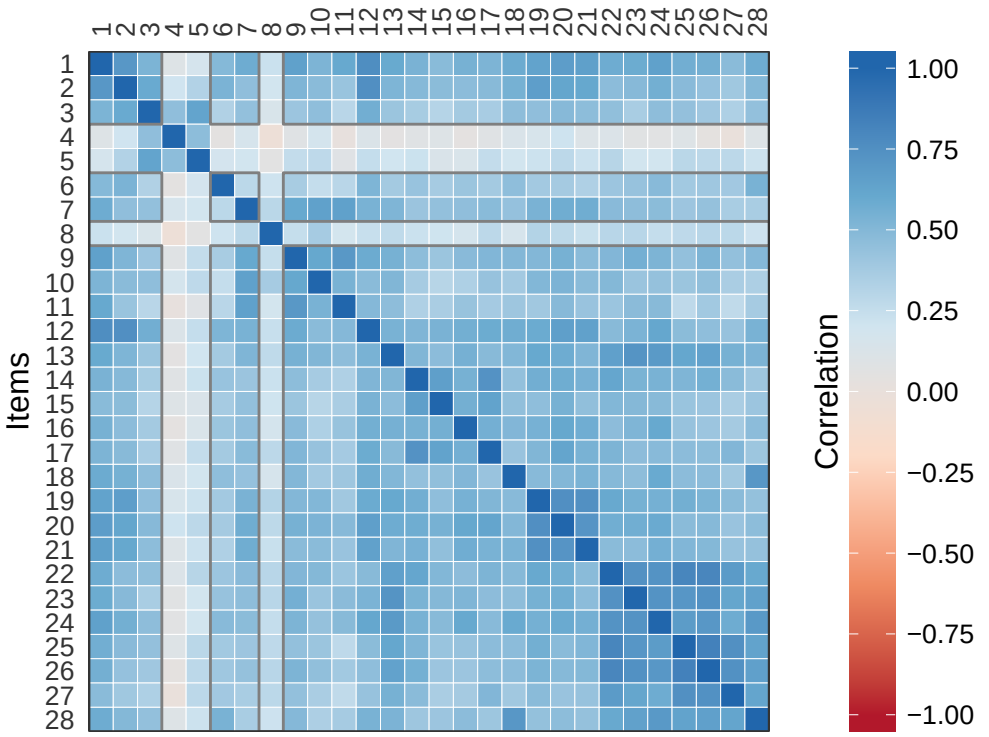

Figure 3.8. Inter-item correlations for social identity continuity and fluidity for sexuality (Study 2a). I excluded items 4, 5, and 8 (in grey) from further analyses.

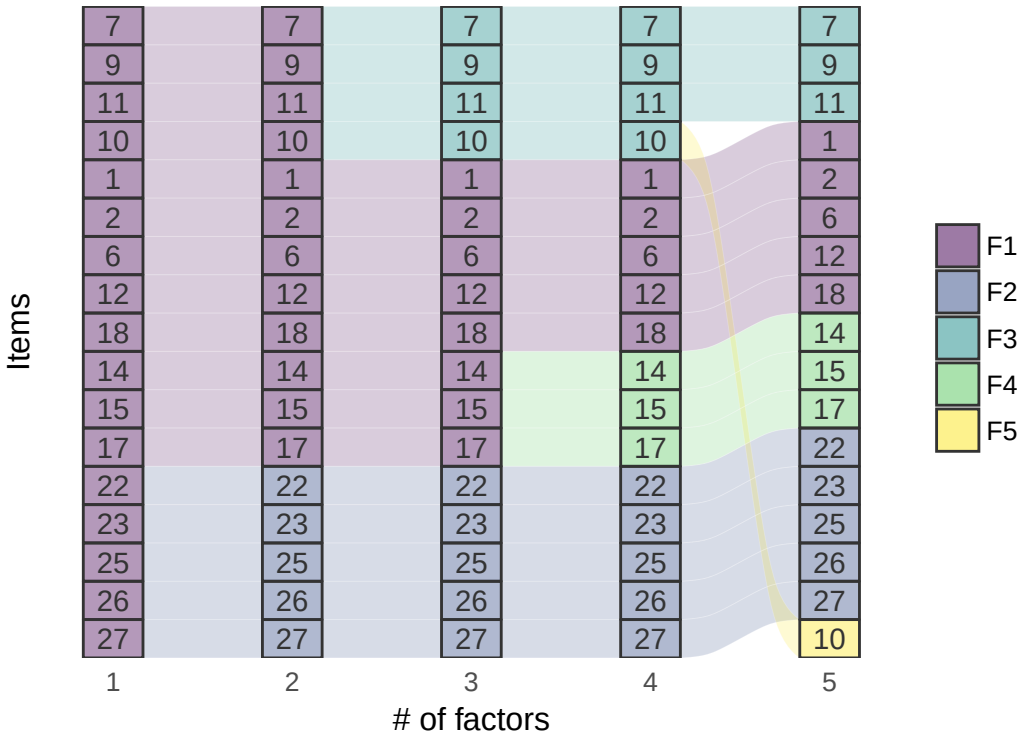

Figure 3.9. Factor structures in exploratory factor analyses (Study 2a). I assigned items to factors according to their highest loading.

model fit. As the five-factor model did not add any substantive factors, I conducted all further analyses for the four-factor structure.[8] Table 3.2 shows all items with Geomin-rotated factor loadings. All factors showed very good internal consistency, *continuity: identity* (ω = .86, [.83, .89]), *continuity: attraction* (ω = .88, [.86, .90]), *fluidity: past* (ω = .86, [.80, .90]), and *fluidity: future* (ω = .94, [.92, .95]). The proposed scale thus proved reliable, assuming that the items formed unidimensional scale dimensions.

Next, I estimated latent variable correlations between the four exploratory factors and variables hypothesised to correlate with social identity continuity and fluidity. Figure 3.10 shows correlation coefficients (in red) with 95% confidence intervals. Scale factors strongly correlated with each other ($.50 < \phi < .65$). As expected, factors correlated negatively with discreteness beliefs, group identification, social dominance orientation, and need for closure ($-.50 < \phi < -.02$). Factors did not correlate, or correlated only weakly, with immutability beliefs. Some factors did not correlate with discreteness beliefs, social dominance orientation, and need for closure. As none of these relationships was as strong as the inter-factor correlations, the factors appear to be related to, but distinct from earlier constructs. As expected, all factors correlated positively with same-gender attraction ($.24 < \phi < .63$). This correlation, however, was strong enough to call into question whether the proposed scales measured constructs distinct from same-gender attraction. Figure 3.10 also shows partial correlation coefficients (in blue), controlling for same-gender attraction. Controlling for same-gender attraction did not substantially weaken correlations with other variables used to assess construct validity, and in some cases strengthened these relationships. This suggests that the proposed factors measured something beyond sexual attraction. As hypothesised, all scale factors correlated with less intergroup bias ($-.17 < \phi < -.25$). Some factors also correlated with contact quantity and/or positive contact.

I tested the hypothesised relationships between scale factors, intergroup contact, and outgroup attitudes in an exploratory structural equation model. I regressed positive and negative contact on contact quantity, the scale factors on all contact variables, and outgroup attitudes on all other variables. The resultant model fit well, $\chi^2(310) = 447.83$, $\chi^2/df = 1.44$, $CFI = .96$, $TLI = .95$, $RMSEA = .04$. Figure 3.11 shows standardised coefficients for all significant model paths. As hypothesised, positive contact was associated with the *continuity: identity* factor, while contact quantity was associated with the *continuity: attraction* and *fluidity: future* factors (hypothesis 1). None of the other paths between the three contact variables and four scale factors was significant. Other than expected, none of the factors was significantly associated with outgroup attitudes (hypothesis 2). Together, all variables explained R^2 = 29%, [15%, 43%] of variance in outgroup attitudes.

Study 2b: Gender

For all items, participants' responses ranged from 1 = *strongly disagree* to 7 = *strongly agree*. Responses to items 9, 12, 14, 17, 23, 25, and 26 (n = 7) averaged in the lower third of the scale, $2.64 < Ms < 2.97$. Responses to items 5 (M = 5.08) and 20 (M = 5.90) averaged in the upper third of the scale. Responses to all other items (n = 18) averaged around the centre of the scale, $3.05 < Ms < 4.79$. Standard deviations ranged from 1.23 to 1.76 units on the response scale. Figure 3.12 shows coefficients for all inter-item correlations.

[8] An additional reason to prefer the simpler model is that the five-factor model has an TLI close to 1.00. TLI values close to or above 1.0 indicate that a model may contain more parameters than necessary to account for the relationships among observed responses (Little, 2013).

Table 3.2. Geomin-rotated factor loadings for the 4-factor solution in Study 2a. Asterisks (*) mark reverse corded items. Only factor loadings for which $p < .05$ are shown. Items are worded for participants who are women.

#	Item	F1	F2	F3	F4
Continuity: Identity					
7	The terms 'homosexual' and 'heterosexual' are too narrow to describe my sexual orientation.			.67	
9	Saying 'I am straight' or 'I am heterosexual' is all there is to say about my sexual orientation.*			.68	
10	The extent to which I am attracted to men or women is more complex than saying 'I am heterosexual' or 'I am homosexual'.			.59	
11	Describing me as 'heterosexual' or 'straight' is enough to describe my sexual orientation.*			.94	
Continuity: Attraction					
1	I am exclusively heterosexual and not at all homosexual.*	.61		.28	
2	I am not at all attracted to women.*	.88			
6	I cannot imagine being attracted to both men and women.*	.55			
12	I am exclusively attracted to men.*	.82			
18	I have never considered going out with a woman.*	.46			
Fluidity: Past					
14	The extent to which I am attracted to men or women may have changed to some degree over the course of my lifetime.				.83
15	Whether I am attracted to men or women has not changed over the course of my lifetime.*				.68
17	My sexual orientation may have changed to some degree since I was younger.				.77
Fluidity: Future					
22	The extent to which I am attracted to men or women may change to some degree in the future.		.74		.16
23	Whether I am attracted to men or women will not change in the future.*		.60	.22	
25	Though I am not attracted to women now, I might be in the future.	.15	.91		
26	Though I think of myself as heterosexual now, I might not do so in the future.		.87		
27	I cannot say that I will always only love men.		.75		

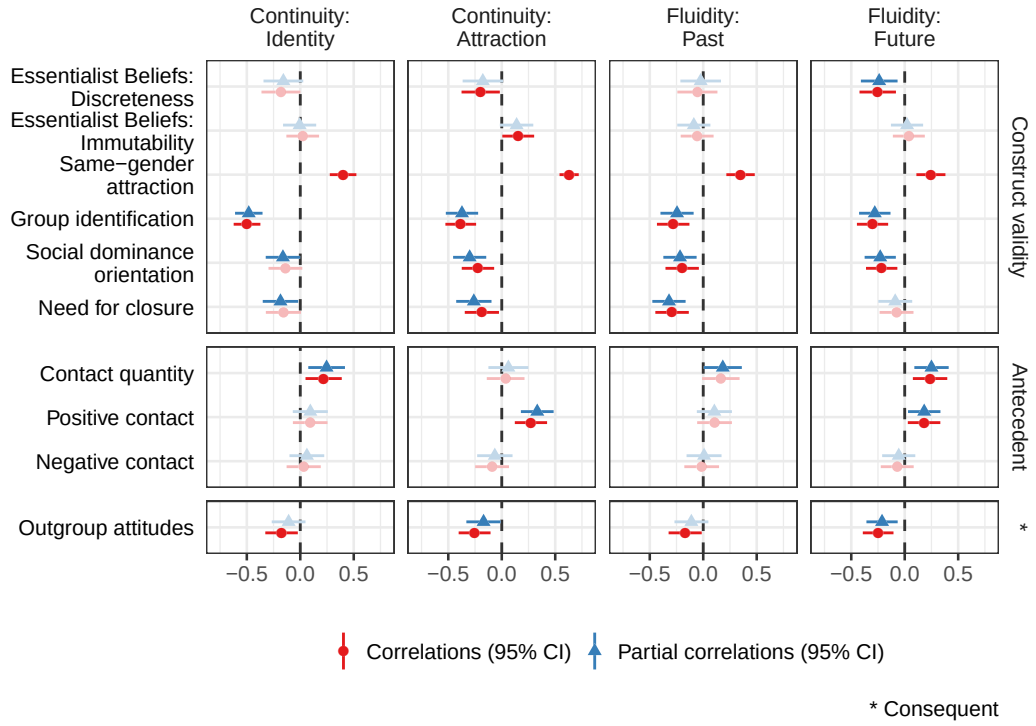

Figure 3.10. Correlations and partial correlations (controlling for same-gender attraction) between scale factors and variables related to construct validity, and with variables hypothesised as antecedent and consequent (Study 2a).

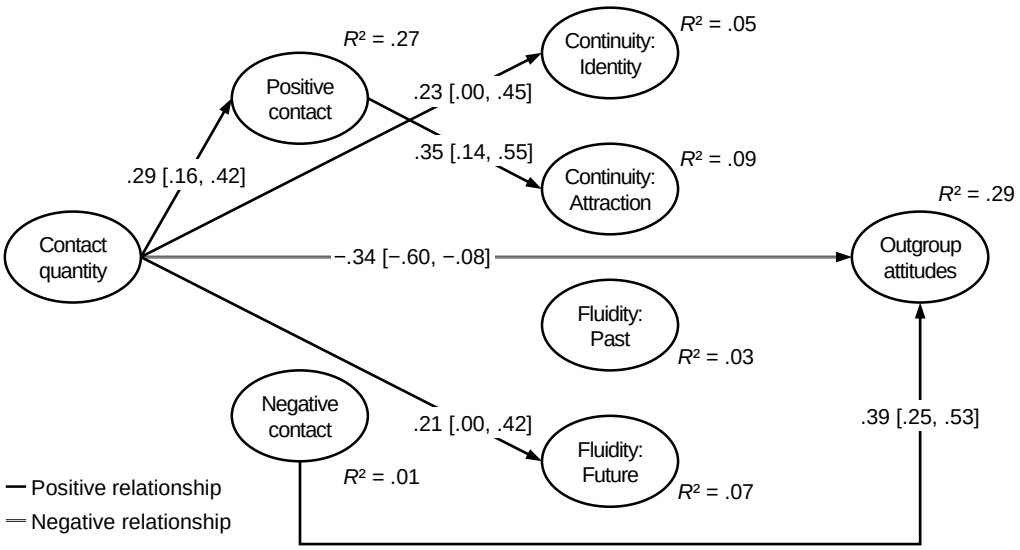

Figure 3.11. Exploratory structural equation models for contact with / attitudes toward sexual minorities, with standardised coefficients and 95% confidence intervals (Study 2a). Only significant paths are shown.

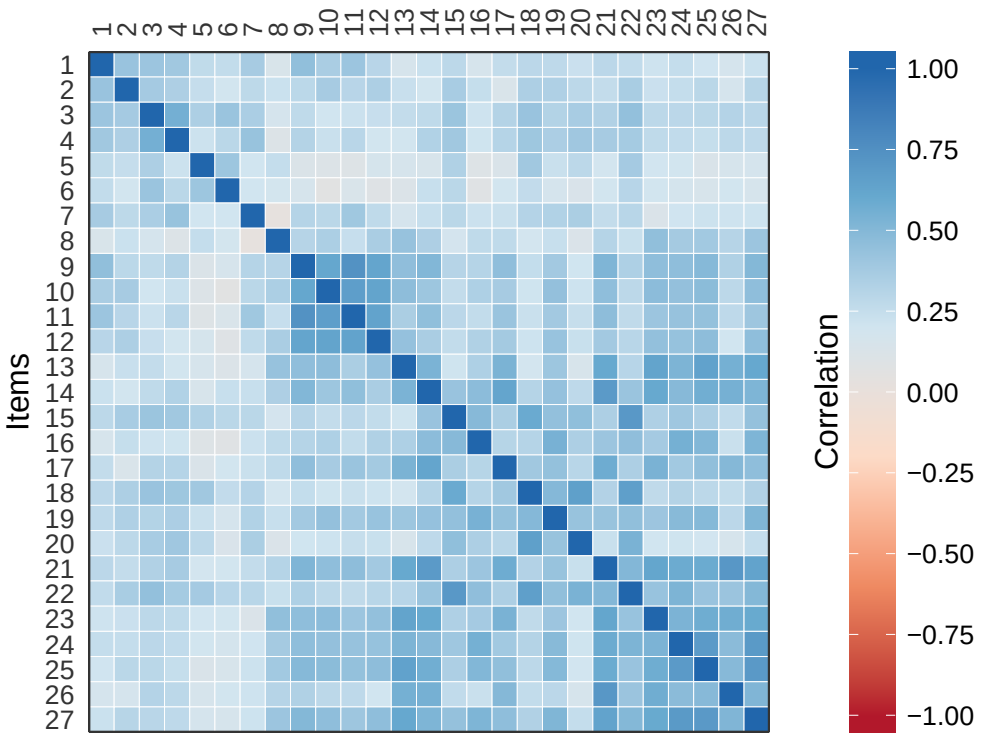

Figure 3.12. Inter-item correlations for social identity continuity and fluidity for gender (Study 2b).

Correlation coefficients ranged from r = .03 to r = .73. Items 1, 2, 5, 6, 7, 8, and 20 had the lowest average correlations with other scale items, .20 < Ms_r < .29. Items 14, 21, 22, 23, 24, 25, and 27 had the highest average correlations with other scale items, .40 < Ms_r < .43. For all other items, average inter-item correlations ranged from M_r = .31 to M_r = .39. As in Study 2a, inter-item correlations indicate that the items reflect more than one dimension, that some items do not fit into the overall scale, or both.

In an initial analysis, I estimated models with one to nine exploratory factors for all 27 scale items. I removed 7 items that did not have any significant factor loadings in the well-fitting models, that (substantially) loaded on more than one factor, or that formed factors that were difficult to interpret (see Table A.2). I then estimated models with one to six exploratory factors for the 20 remaining items. Figure 3.13 shows changes in factor structure across the six models. The two-factor solution split items into one factor about gender identity (F1) and one factor about gender expression (F2). The three-factor solution split the former factor into one factor that measured to what extent participants saw their gender identity as continuous (F1, *Continuity: Identity*) and one factor with the remaining items about gender identity. The four-factor solution split the latter factor into one factor that measured the extent to which participants understood their gender identity as fluid (F2, *Fluidity: Identity*), and one factor that expressed participants' conviction (reverse coded) that their gender identity and expression would remain stable across their lifetime (F4, *Fluidity: Stability*). The five-factor model split items into two additional factors, one factor that measured to what extent participants understood their gender expression as continuous (F5, *Continuity: Expression*), and one factor that measured to what extent participants saw their gender expression as fluid (F2, *Fluidity: Expression*). The six-factor solution formed an additional factor comprising item 14 (F6).

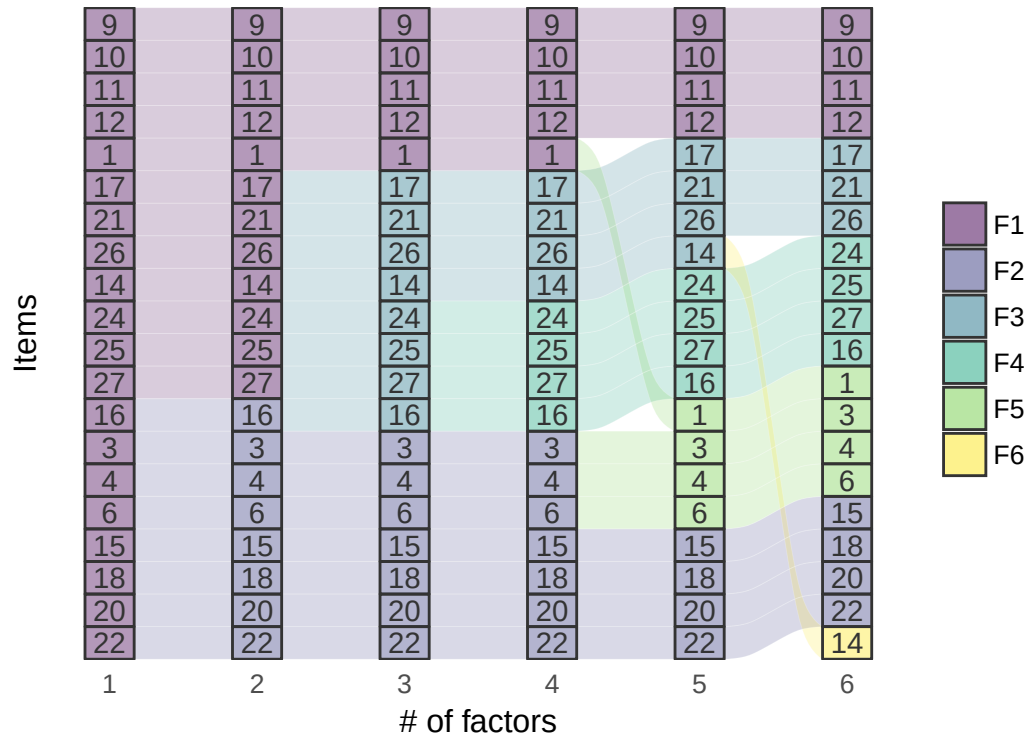

Figure 3.13. Factor structures in exploratory factor analyses (Study 2b). I assigned items to factors according to their highest loading.

Models with one factor ($\chi^2(170) = 1349$, $\chi^2/df = 7.94$, $CFI = .66$, $TLI = .62$, $RMSEA = .15$), with two factors ($\chi^2(151) = 882.86$, $\chi^2/df = 5.85$, $CFI = .79$, $TLI = .73$, $RMSEA = .12$), or with three factors ($\chi^2(133) = 525.25$, $\chi^2/df = 3.95$, $CFI = .89$, $TLI = .84$, $RMSEA = .10$) did not represent the observed responses well. The four-factor model ($\chi^2(116) = 317.92$, $\chi^2/df = 2.74$, $CFI = .94$, $TLI = .90$, $RMSEA = .08$) showed acceptable model fit, but fit less well than the five factor solution ($\chi^2(100) = 221.48$, $\chi^2/df = 2.21$, $CFI = .96$, $TLI = .93$, $RMSEA = .06$), $\Delta\chi^2(96.45) = 16$, $p < .001$, $\Delta CFI = .02$. Though the six-factor solution fit better than the five-factor solution, $\Delta\chi^2(67.89) = 15$, $p < .001$, $\Delta CFI = .02$, it did not add any substantive factors. I conducted all further analyses for the five-factor structure. Table 3.3 shows all items with Geomin-rotated factor loadings. Four factors showed very good internal consistency, *continuity: identity* ($\omega = .88$, [.85, .91]), *fluidity: identity* ($\omega = .86$, [.83, .89]), *fluidity: expression* ($\omega = .86$, [.83, .89]), and *fluidity: stability* ($\omega = .86$, [.82, .89]). One factor showed acceptable internal consistency, *continuity: expression* ($\omega = .72$, [.66, .77]). The proposed scale thus proved reliable, assuming that the items formed unidimensional scale dimensions.

Next, I estimated latent variable correlations between the five exploratory factors and variables I expected to be associated with social identity continuity and fluidity. Figure 3.14 shows correlation coefficients with 95% confidence intervals. Correlations between scale items ranged from $\phi = .18$ to $\phi = .62$. As expected, factors correlated negatively with essentialist beliefs about gender, gender norm endorsement, group identification, social dominance orientation, and need for closure ($-.46 < \phi < .01$). The *fluidity: stability* and *continuity: expression* factors had the most consistent correlations with variables used to assess construct validity. As hypothesised, four of the five scale factors correlated positively with contact quantity, positive contact, and/or negative contact with

Table 3.3. Geomin-rotated factor loadings for the 5-factor solution in Study 2b. Asterisks (*) mark reverse corded items. Only factor loadings for which $p < .05$ are shown. Items are worded for participants who are women.

#	Item	F1	F2	F3	F4	F5
Continuity: Identity						
9	The terms 'male' and 'female' are too narrow to describe how I think about my gender.	.68				
10	Saying 'I am a woman' is all there is to say about how I think about my gender.*	.64			.18	
11	Describing how I think about my gender is more complex than saying 'I am a woman' or 'I am female'.	.80				
12	Describing me as 'female' is enough to describe my gender.*	.66			.26	
Continuity: Expression						
1	The terms 'masculine' and 'feminine' are too narrow to describe how I express my gender in dress and demeanour.	.40				.45
3	Though I am a woman, I do not necessarily dress or behave as people expect women to dress or behave.					.77
4	Though I am a woman, my appearance, expression, and demeanour reflects both 'masculine' and 'feminine' characteristics.					.52
6	How I dress or behave is not defined by my gender.					.46
Fluidity: Identity						
14	The extent to which I think of myself as male or female may have changed to some degree over the course of my lifetime.		.15	.68		
17	The extent to which I think of myself as male or female may change to some degree between different situations.		.18	.62		
21	The extent to which I think of myself as male or female may change to some degree in the future.			.82		
26	Though I think of myself as a woman now, I might not necessarily do so in the future.	-.22		.86		
Fluidity: Expression						
15	The extent to which I speak, dress, and behave in 'masculine' or 'feminine' ways may have changed over the course of my lifetime.		.60		.29	
18	The extent to which I speak, dress, and behave in 'masculine' or 'feminine' ways may change between different situations.		.79			
20	In some situations I express myself in more 'masculine' ways while in others I express myself in more 'feminine' ways.		.76			
22	The extent to which I speak, dress, and behave in 'masculine' or 'feminine' ways may change in the future.	-.10	.61		.34	
Fluidity: Stability						
16	How I express my gender has not changed much since I was younger.*		.30		.60	-.15
24	How I express my gender is unlikely to change much as I grow older.*				.72	
25	How I think about my gender is unlikely to change as I grow older.*			.24	.61	
27	I will always think about my gender in the same way as I do now.*			.25	.62	

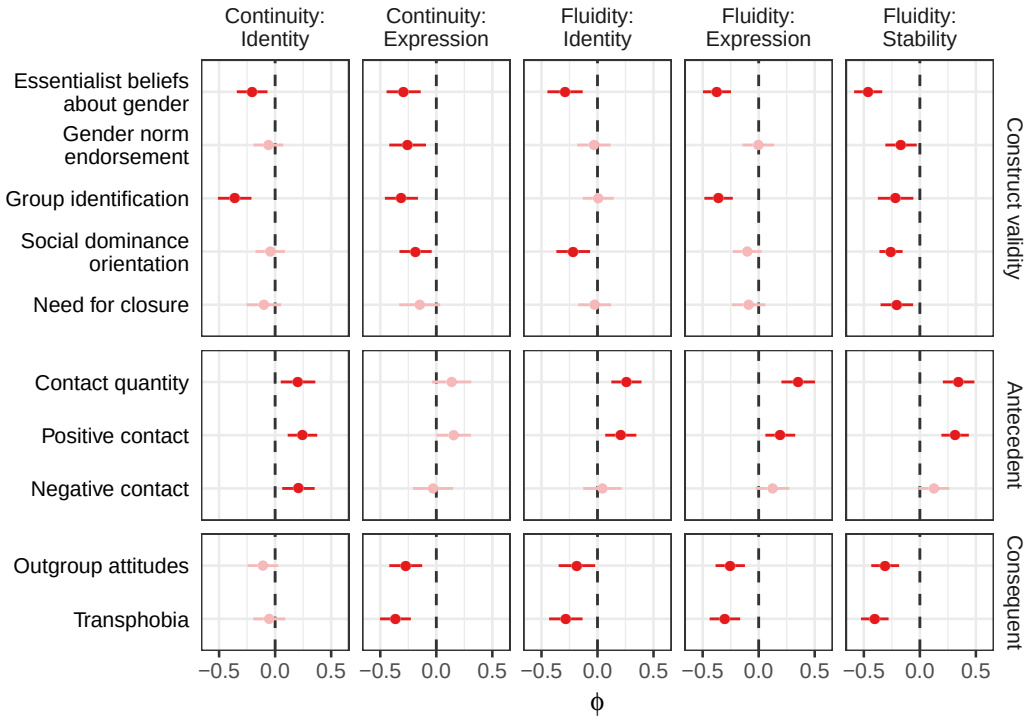

Figure 3.14. Correlations (with 95% confidence intervals) between scale factors and variables related to construct validity, and with variables hypothesised as antecedent and consequent (Study 2b).

transgender people (hypothesis 1). Likewise, four of the five scale factors correlated with more favourable outgroup attitudes and less transphobia (hypothesis 2).

I tested the hypothesised relationships between scale factors, intergroup contact, and transphobia in an exploratory structural equation model.[9] I regressed positive and negative contact on contact quantity, the scale factors on all contact variables, and transphobia on all other variables. The resultant model fit well, $\chi^2(369) = 585.39$, $\chi^2/df = 1.59$, $CFI = .96$, $TLI = .94$, $RMSEA = .04$. Figure 3.15 shows standardised coefficients for all significant model paths. In line with hypothesis 1, contact quantity was associated with the *fluidity: identity* and *fluidity: expression* factors. None of the other paths between the three contact variables and five scale factors was significant. In line with hypothesis 2, the *continuity: expression* and *fluidity: stability* factors were associated with *less* transphobia. Other than expected, the *continuity: identity* factor was associated with *more* transphobia. Paths from the *fluidity: identity* and *fluidity: expression* factors to transphobia were not significant. Together, all variables explained $R^2 = 48\%$, [35%, 60%] of variance in transphobia scores.

Study 3

In Studies 1–2, I found that individuals differ in the extent to which they understand their own sexuality (Studies 1, 2a) and gender (Study 2b) as *continuous* and *fluid*, and developed

[9] In the previous model, outgroup attitudes and transphobia were strongly correlated, $\phi = .68$, [.60, .76]. Here, I only report results for transphobia as outcome variable, both for the sake of brevity and because Nagoshi et al.'s (2008) scale is the more established measure.

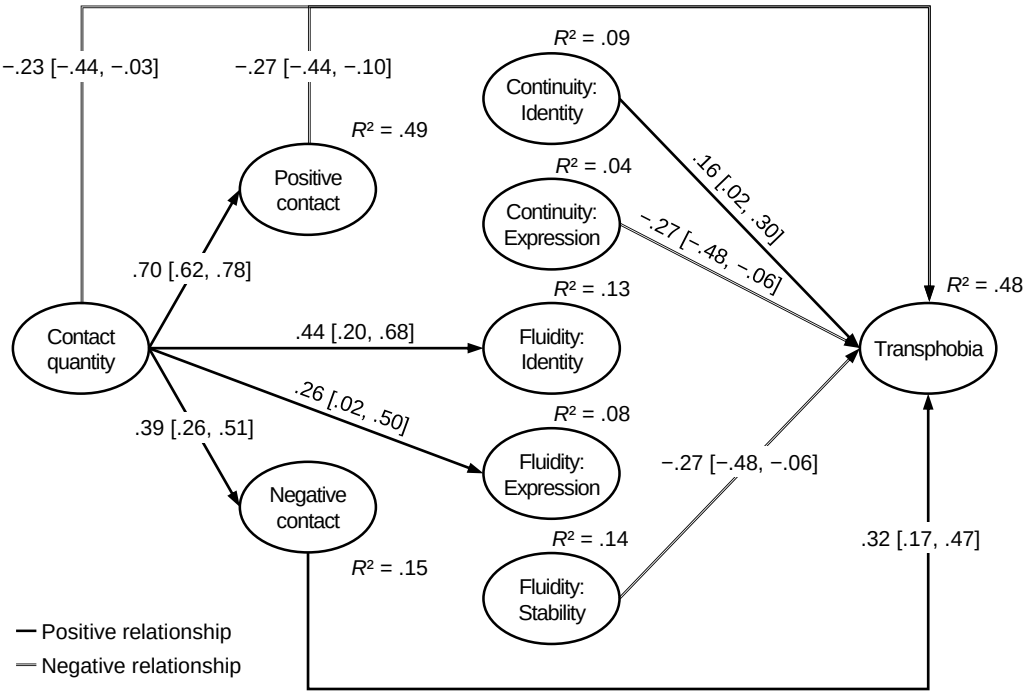

Figure 3.15. Exploratory structural equation models for contact with / attitudes toward transgender people, with standardised coefficients and 95% confidence intervals (Study 2b). Only significant paths are shown.

multidimensional scales that measure these individual differences. Studies 1–2 provided mixed evidence for the hypothesised relationships of social identity continuity and fluidity, as measured by the new scales, with intergroup contact (hypothesis 1) and intergroup bias (hypothesis 2). As expected, intergroup contact was associated with more continuity and fluidity (Studies 1, 2a, 2b), at least for some of the scale dimensions. Likewise, some scale dimensions were associated with less transphobia in Study 2b. Other than expected, none of the scale dimensions was associated with attitudes toward sexual minorities in Studies 1 and 2a, while one scale dimension was associated with *more* transphobia in Study 2b.

Studies 1–2 provided cross-sectional evidence for and against the hypothesised relationships. As such, they did not allow testing the direction of any observed relationships. Study 3 addresses this limitation by testing the hypothesised relationships in a three-wave longitudinal design. I also use Study 3 to confirm the factor structure determined in Studies 1–2, and to test the reliabilities of the resulting scales across three time points (0, 105, 210 days) and two languages (English, German). To reduce the number of questions per survey, I did not include three of the nine subscales in this study (*fluidity: past* for sexuality, and *fluidity: identity* and *fluidity: expression* for gender). I chose the six subscales used because they seemed the least redundant between each other. I reduced the numbers of questions to avoid deterring participants from dropping out of the study.[10]

[10] Tables A.3 and A.4 (in Appendix A) show the translated continuity/fluidity items used in Studies 3a and 3b, respectively.

Method

Participants

Students from six universities, three in the United Kingdom and three in Germany, were invited to participate in a "psychological study on [their] social experiences with people of different social groups" for the chance to win up to £100 (€140). Of 1801 respondents, 1353 (75%) identified as heterosexual. Of these, I excluded 299 respondents who, at any of the time points, did not identify as straight/heterosexual, reported sexual relations with a same-gender partner, or were attracted to people of their own gender more than 50% of the time. Study 3a thus included 1054 heterosexual participants (782 women, 269 men; Mdn_{age} = 22, age range: 18–65 years). Study 3b included 1726 (96%) cisgender participants (1298 women, 426 men; Mdn_{age} = 22, age range: 18–65 years) who, at none of the time points, reported a non-binary gender identity or identified as transgender.[11] Of all respondents, 587 (33%) completed the questionnaire in English, while 1214 (67%) filled in a German translation of the questionnaire.

Procedure

Participants completed the first survey (T1) in the first half of the academic year (September–December). Of the original sample, 73% returned for the second survey (T2) around the middle of the academic year ($Mdn_{T_2\text{-}T_1}$ = 104 days, range: 96–147 days) and 62% for the third survey (T3) toward the end of the academic year ($Mdn_{T_3\text{-}T_1}$ = 205 days, range: 198–250 days).

At each time point, participants answered a range of questions on their sexuality and gender. *Sexual orientation* was measured with four items: "Which of these best describes your sexual orientation?" (*straight/heterosexual*, *gay*, *lesbian*, *bi(sexual)*, *pan(sexual)*, *queer*, *questioning*, *asexual*, *none of these/other*), "Have you previously had sex with a same-sex partner?" (*yes*, *no*), "How often do you feel attracted to women?" and "How often do you feel attracted to men?" (0–100; 0 = *never*, 100 = *all the time*). *Gender identity* was measured with three items: "Which of these best describes your gender?" (*woman*, *man*, *none of these/other*), "Are you transgender?" ("*No, I am not transgender.*", "*Yes, please show me more options.*"), and, if they were transgender, "Which of these best describes your gender identity?" (*trans woman*, *trans man*, *genderqueer/other non-binary*, *gender-nonconforming*, *genderfluid*, *agender*, *none of these/other*). If participants responded "none of these/other" to any of the questions, they were asked to describe their gender, sexual orientation, or gender identity in an open response field.

Participants who described their sexual orientation as "straight/heterosexual" then filled in measures for heterosexual respondents. Participants who were not transgender filled in measures for cisgender respondents. Participants who did not describe themselves as "straight/heterosexual" and/or were transgender instead answered questions on perceived discrimination and collective action (see Chapter 5).

Measures for heterosexual respondents

Intergroup contact was measured with nine items, three for each of *contact quantity* (α_1 = .65, α_2 = .67, α_3 = .71), *positive contact* (α_1 = .58, α_2 = .58, α_3 = .57), and *negative*

[11] Study 3a and 3b included, respectively, three and two respondents who reported different genders at different time points.

contact (α_1 = .63, α_2 = .67, α_3 = .65). Participants reported how many gay men, lesbian/gay women, and bisexual people they knew personally (1 = *none*, 7 = *more than 16*). For each outgroup, if they knew any member of that group, respondents indicated how often, on average, they had "positive/good contact with [them]—e.g., being supported, made to feel welcome, helped, complimented, or befriended by them" and "negative/bad contact with [them]—e.g., being made to feel unwanted, verbally abused, intimidated, threatened, or ridiculed by them" (1 = *never*, 5 = *very often*; adapted from Barlow et al., 2012). If participants did not know any member of a group, positive and negative contact were coded as 1 = *never* for that outgroup.

Participants' perceptions of their sexual identity and attraction were measured with fourteen items from Study 2a. Four items measured the *continuity: identity* dimension, five items measured the *continuity: attraction* dimension, and five items measured the *fluidity: future* dimension (see Table 3.2). Participants indicated the extent to which they agreed or disagreed with each item, on a scale from 1 = *strongly disagree* to 7 = *strongly agree*.

Outgroup attitudes were measured using Wright et al.'s (1997) General Evaluation Scale. Respondents rated how they felt about lesbian, gay, bisexual, and other non-heterosexual people in general on 7-point adjective scales: warm–cold (reversed), negative–positive, friendly–hostile (reversed), suspicious–trusting, respect–contempt (reversed), and admiration–disgust (reversed). Items formed a reliable scale (α_1 = .90, α_2 = .91, α_3 = .91) with higher scores indicating more favourable attitudes.

Measures for cisgender respondents

Participants read definitions of three outgroups: trans women, trans men, and gender-queer people (see Study 2 for wordings). Participants reported how many members of each of these groups they knew personally (1 = *none*, 7 = *more than 16*). If participants knew at least one member of an outgroup, they reported how often they had positive/good and negative/bad contact with members of that outgroup (see above). Intergroup contact was thus measured as *contact quantity* (α_1 = .72, α_2 = .76, α_3 = .80), *positive contact* (α_1 = .79, α_2 = .75, α_3 = .63), and *negative contact* (α_1 = .76, α_2 = .78, α_3 = .67). *Group salience* in interactions with transgender people was measured with three items (see above; α_1 = .71, α_2 = .78, α_3 = .74).

Participants' perceptions of their gender identity and expression were measured with twelve items from Study 2b. Four items measured each of three dimensions, *continuity: identity*, *continuity: expression*, and *fluidity: future* (see Table 3.3). Participants indicated the extent to which they agreed or disagreed with each item, on a scale from 1 = *strongly disagree* to 7 = *strongly agree*. *Transphobia* was measured with the same items as in Study 2b (α_1 = .84, α_2 = .86, α_3 = .86).

Results

I follow the analytic strategy outlined under *Present research*. First, I use confirmatory factor analysis to replicate the factor structure determined in Studies 1–2. Second, I assess the reliability of the resulting scale dimensions by estimating their internal consistency (as omega coefficients), by testing for measurement invariance (across time and languages), and by estimating their temporal stability (as test-retest correlation coefficients). Third, I test the hypothesised relationships in a series of cross-lagged panel models (Little, 2013)

Table 3.4. Model fit indices for configural (0,1), loadings (2,3), intercepts (4,5), means (6,7), and covariances (8,9) invariant models for three time points and two language groups in Study 3a.

#	Invariance	χ^2	*df*	χ^2/df	*CFI*	*TLI*	*RMSEA*	Δ*CFI*
0	Configural (1 factor)	7133	1548	4.61	.78	.76	.09	
1	Configural (3 factors)	2560	1482	1.73	.96	.95	.04	.177
2	Loadings (group)	2621	1524	1.72	.96	.95	.04	-.001
3	Loadings (group, time)	2640	1552	1.70	.96	.96	.04	.000
4	Intercepts (group)	2746	1585	1.73	.96	.95	.04	-.003
5	Intercepts (group, time)	2825	1613	1.75	.96	.95	.04	-.002
6	Means (group)	2884	1616	1.78	.95	.95	.04	-.002
7	Means (group, time)	2879	1618	1.78	.95	.95	.04	.000
8	Covariances (group)	2924	1654	1.77	.95	.95	.04	-.001
9	Covariances (group, time)	2949	1670	1.77	.95	.95	.04	-.000

to establish the direction of any observed effects. I test, for example, whether contact predicts more social identity continuity (forward path), social identity continuity predicts contact (reverse path), or both, while accounting for the effect of each variable on itself over time (autoregressive paths). Finally, I test whether social identity continuity and fluidity mediate any relationships between contact and bias.

Study 3a: Sexuality

I tested the factor structure of the novel scale by comparing a model in which all scale items loaded on one factor (Model 0) to a model in which scale items loaded on the three factors observed in Study 2a (Model 1). Both models allowed loadings, intercepts, and covariances to vary across time points and survey languages. The three-factor model, but not the one-factor solution, fit the data well (see Table 3.4), and improved model fit over the one-factor solution, $\Delta\chi^2(96) = 953.25$, $p < .001$, $\Delta CFI = .18$. Study 3a thus replicated the factor structure observed in Study 2a. Table 3.5 shows omega coefficients for all scale dimensions and time points. All scale dimensions showed very good internal consistency, $.87 \leq \omega \leq .93$.

Next, I tested for measurement invariance across survey languages and time points. Following the recommendations of Little (2013), I compared models that imposed increasingly strict invariance constraints. I used $\Delta CFI < .01$ as criterion to decide whether a more constrained model fit as well as a less constrained model. Table 3.4 shows model fit indices and comparisons. Model 1 (see above) imposed the same factor structure on responses to the English and German version of the questionnaire, but allowed all free parameters to vary across time and language. The model fit well, demonstrating configural invariance across time and language. Models 2 and 3 constrained factor loadings to be equal across language and time. Models 4 and 5 also constrained item intercepts to be equal across language and time. These models fit as well as the configural model, $\Delta CFI < .01$, demonstrating loadings (weak) invariance and intercepts (strong) invariance. Together, these comparisons showed that the scale factors measured the same constructs across survey languages and time points.

Models 6 and 7 constrained latent means to be equal across language and time.

Table 3.5. McDonald's (1999) omega coefficients for internal consistency across scale dimensions and time points in Study 3a.

Survey	Dimension	ω	95% CI
T1	Continuity: Identity	.87	[.85, .89]
T2	Continuity: Identity	.88	[.86, .90]
T3	Continuity: Identity	.89	[.87, .91]
T1	Continuity: Attraction	.89	[.87, .90]
T2	Continuity: Attraction	.88	[.86, .89]
T3	Continuity: Attraction	.89	[.88, .91]
T1	Fluidity: Future	.92	[.91, .93]
T2	Fluidity: Future	.93	[.92, .94]
T3	Fluidity: Future	.93	[.92, .94]

These models fit as well as Model 5, showing that, on average, the extent to which participants agreed with each scale dimension remained stable. Models 8 and 9 also constrained all covariances between the scale dimensions to be equal across language and time. Model fit did not deteriorate, showing that covariances between scale dimensions did not change across survey languages, and that covariances from T1 to T2 were equal to covariances from T2 to T3. All scale dimensions had strong test-retest correlations from T1/T2 to T2/T3 ($\phi_{1,2}$), and from T1 to T3 ($\phi_{1,3}$). *Continuity: identity* correlated with the same construct measured around 103 days $\phi_{1,2}$ = .82, [.79, .85] and 205 days later $\phi_{1,3}$ = .75, [.70, .80]. I found similar test-retest correlations for *continuity: attraction*, $\phi_{1,2}$ = .85, [.82, .89], $\phi_{1,3}$ = .82, [.77, .86], and *fluidity: future*, $\phi_{1,2}$ = .81, [.78, .84], $\phi_{1,3}$ = .73, [.69, .78]. Models 1–9 thus showed that scale dimensions were internally consistent, measurement invariant, and stable across time and translation—confirming the reliability of the newly developed scale.

I then tested the hypothesised longitudinal relationships between scale dimensions, positive and negative contact, and outgroup attitudes. First, I estimated a multi-group measurement model for all six latent variables. I constrained loadings and intercepts (across languages and time points) and covariances (across languages) to be equal. The resultant model fit well: $\chi^2(4457) = 6547.40$, $\chi^2/df = 1.47$, $CFI = .95$, $TLI = .95$, $RMSEA = .03$. I thus combined responses for subsequent analyses. Second, I estimated a measurement model for the combined sample which constrained loadings and intercepts to be equal across time points. This model also fit well: $\chi^2(2138) = 3318.87$, $\chi^2/df = 1.55$, $CFI = .97$, $TLI = .97$, $RMSEA = .02$. Third, I estimated a series of nested cross-lagged models to test the hypothesised relationships. Model 1 estimated autoregressive paths, but constrained all forward and reverse paths to zero. Model 2 added forward paths from positive and negative contact (at T1, T2) to outgroup attitudes (at T2, T3). Model 2 fit better than Model 1, $\Delta\chi^2(4) = 21.51$, $p < .001$, $\Delta CFI = .01$—showing that contact predicted outgroup attitudes.

Model 3 extended Model 2 by estimating forward paths from positive and negative contact to the three scale dimensions. Model 3 did not significantly improve model fit over Model 2, $\Delta\chi^2(12) = 14.22$, $p = .29$, $\Delta CFI = .00$. In contrast to hypothesis 1, intergroup contact did not predict more continuous and fluid understandings of sexuality. Model 4 extended Model 3 by estimating forward paths from the three scale dimensions

to outgroup attitudes. Model 4 did not significantly improve upon Model 3, $\Delta\chi^2(6) = 3.13$, $p = .79$, $\Delta CFI = .00$. Contrary to hypothesis 2, more continuous and fluid understandings of sexuality did not predict more favourable outgroup attitudes. Models 5–7 successively added reverse paths for each of the forward paths described above. Only the difference between Models 5 and 6 was significant, $\Delta\chi^2(12) = 202.90$, $p < .001$, $\Delta CFI = .00$, suggesting that some scale dimensions predicted positive or negative contact.

I estimated a final model in which I included autoregressive paths, forward paths from positive/negative contact to outgroup attitudes, and reverse paths from the three scale dimensions to intergroup contact. The resultant model fit well, $\chi^2(2218) = 4299.55$, $\chi^2/df = 1.94$, $CFI = .95$, $TLI = .94$, $RMSEA = .03$. Figure 3.16 shows standardised coefficients for all significant ($p < .05$) forward and reverse paths in Model 8. As expected, positive contact consistently predicted more favourable outgroup attitudes. Paths from T1 (positive contact) to T2 (outgroup attitudes) and from T2 to T3 did not differ significantly from each other, $\Delta\beta = -.06$, $[-.19, .07]$. Contrary to predictions, *continuity: identity*, *continuity: attraction*, and *fluidity: future* (at T1) predicted, respectively, *less* positive contact, *more* negative contact, and *less* negative contact (at T2). The equivalent paths from T2 to T3, however, were not significant.[12] Together, these findings provided longitudinal evidence *for* the hypothesised relationship between intergroup contact and outgroup attitudes, but *against* the hypothesised relationships between intergroup contact, social identity continuity and fluidity, and outgroup attitudes. Other than expected, social identity continuity and fluidity did not mediate the relationship between intergroup contact and outgroup attitudes.

Study 3b: Gender

I followed the same analytic strategy for participants' perceptions of their own gender identity and expression. I tested the factor structure of the scale by comparing a model in which all scale items loaded on one factor (Model 0) to a model in which scale items loaded on the three factors observed in Study 2b (Model 1). The three-factor model, but not the one-factor solution, fit the data well (see Table 3.6), and improved model fit over the one-factor solution, $\Delta\chi^2(66) = 3966.10$, $p < .001$, $\Delta CFI = .23$. Study 3b thus replicated the factor structure observed in Study 2b. Table 3.7 shows omega coefficients for all scale dimensions and time points. All scale dimensions showed good internal consistency, $.79 \leq \omega \leq .92$.

To test for measurement invariance, I compared models that imposed increasingly strict invariance constraints across survey languages and time points. Table 3.6 shows model fit indices and comparisons. Models 2–5 fit as well as less constrained models, demonstrating loadings (weak) invariance and intercepts (strong) invariance. Constraining latent means (Models 6, 7) and latent covariances (Models 8, 9) to be equal across time and language did not decrease model fit. All scale dimensions had strong test-retest correlations: *continuity: identity* ($\phi_{1,2} = .82$, $[.79, .84]$, $\phi_{1,3} = .77$, $[.73, .80]$), *continuity: expression* ($\phi_{1,2} = .82$, $[.79, .85]$, $\phi_{1,3} = .79$, $[.76, .83]$), and *fluidity: stability* ($\phi_{1,2} = .71$, $[.68, .75]$, $\phi_{1,3} = .65$, $[.60, .70]$). Models 1–9 thus showed that scale dimensions were internally consistent, measurement invariant, and stable across time and translation—confirming the reliability of the newly developed scale.

I then tested the hypothesised longitudinal relationships between scale dimen-

[12] When controlling for multiple comparison, none of the reverse paths was significant, $p > .05/16$. The paths from positive contact to outgroup attitudes remained significant, $p < .05/16$.

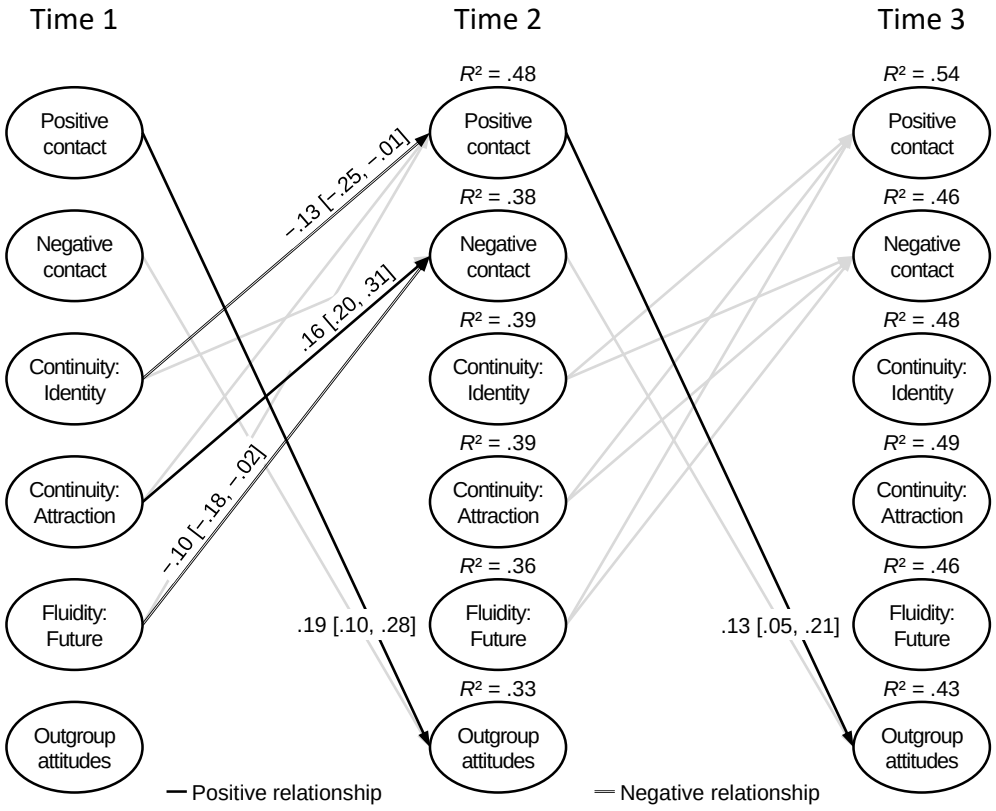

Figure 3.16. Longitudinal structural equation model with standardised coefficients and 95% confidence intervals (Study 3a). Paths in grey were not significant, $p > .05$. Autoregressive paths are not shown.

Table 3.6. Model fit indices for configural (0,1), loadings (2,3), intercepts (4,5), means (6,7), and covariances (8,9) invariant models for three time points and two language groups in Study 3b.

#	Invariance	χ^2	*df*	χ^2/df	*CFI*	*TLI*	*RMSEA*	Δ*CFI*
0	Configural (1 factor)	8664	1110	7.81	.74	.70	.10	
1	Configural (3 factors)	2130	1044	2.04	.96	.96	.04	.225
2	Loadings (group)	2341	1080	2.17	.96	.95	.04	-.006
3	Loadings (group, time)	2401	1104	2.17	.96	.95	.04	-.001
4	Intercepts (group)	2578	1131	2.28	.95	.95	.04	-.005
5	Intercepts (group, time)	2646	1155	2.29	.95	.95	.04	-.001
6	Means (group)	2829	1155	2.45	.94	.94	.04	-.006
7	Means (group, time)	2847	1161	2.45	.94	.94	.04	-.000
8	Covariances (group)	2894	1197	2.42	.94	.94	.04	-.001
9	Covariances (group, time)	2937	1212	2.42	.94	.94	.04	-.001

Table 3.7. McDonald's (1999) omega coefficients for internal consistency across scale dimensions and time points in Study 3b.

Survey	Dimension	ω	95% CI
T1	Continuity: Identity	.88	[.87, .90]
T2	Continuity: Identity	.91	[.90, .93]
T3	Continuity: Identity	.92	[.90, .93]
T1	Continuity: Expression	.82	[.81, .84]
T2	Continuity: Expression	.84	[.82, .85]
T3	Continuity: Expression	.82	[.80, .84]
T1	Fluidity: Stability	.79	[.78, .80]
T2	Fluidity: Stability	.80	[.76, .84]
T3	Fluidity: Stability	.82	[.78, .86]

sions, positive and negative contact, and transphobia. First, I estimated a multi-group measurement model for all six latent variables, which constrained loadings and intercepts (across languages and time points) and covariances (across languages) to be equal. The resultant model fit well: $\chi^2(3675) = 5755.96$, $\chi^2/df = 1.57$, $CFI = .96$, $TLI = .95$, $RMSEA = .03$. I thus combined responses for subsequent analyses. Second, I estimated a measurement model for the combined sample which constrained loadings and intercepts to be equal across time points. This model also fit well: $\chi^2(1749) = 3730.72$, $\chi^2/df = 2.13$, $CFI = .96$, $TLI = .96$, $RMSEA = .03$. Third, I estimated a series of nested cross-lagged models to test the hypothesised relationships. Model 1 estimated autoregressive paths, but constrained all forward and reverse paths to zero. Model 2 added forward paths from positive and negative contact (at T1, T2) to outgroup attitudes (at T2, T3). Model 2 did not fit better than Model 1, $\Delta\chi^2(4) = 1.456$, $p = .83$, $\Delta CFI = .001$. Other than expected, contact with transgender people did not predict more or less transphobia.

Model 3 extended Model 2 by estimating forward paths from positive and negative contact to the three scale dimensions. Model 3 significantly improved model fit over Model 2, $\Delta\chi^2(12) = 59.87$, $p < .001$, $\Delta CFI = .00$. In line with hypothesis 1, intergroup contact predicted how continuous and fluid participants would understand their gender to be around 105 days later. Model 4 extended Model 3 by estimating forward paths from the three scale dimensions to transphobia. Model 4 did not significantly improve upon Model 3, $\Delta\chi^2(6) = 8.34$, $p = .21$, $\Delta CFI = .00$. Contrary to hypothesis 2, more continuous and fluid understandings of sexuality did not predict more favourable outgroup attitudes. Models 5–7 successively added reverse paths for each of the forward paths described above. None of these models improved model fit ($p > .11$, $\Delta CFI < .001$), indicating that reverse paths were not necessary to accurately describe the data.

I estimated a final model in which I included autoregressive paths and forward paths from positive/negative contact to social identity continuity and fluidity. The resultant model showed acceptable fit, $\chi^2(1833) = 5134.88$, $\chi^2/df = 2.8$, $CFI = .93$, $TLI = .92$, $RMSEA = .03$. Figure 3.17 shows standardised coefficients for all significant ($p < .05$) forward paths in the final model. In line with hypothesis 1, positive contact consistently predicted more continuous conceptions of gender identity and expression. Positive contact predicted more fluid understanding of gender from T1 to T2, but not from T2 to T3. Negative contact predicted less fluidity from T1 to T2. None of the other paths from neg-

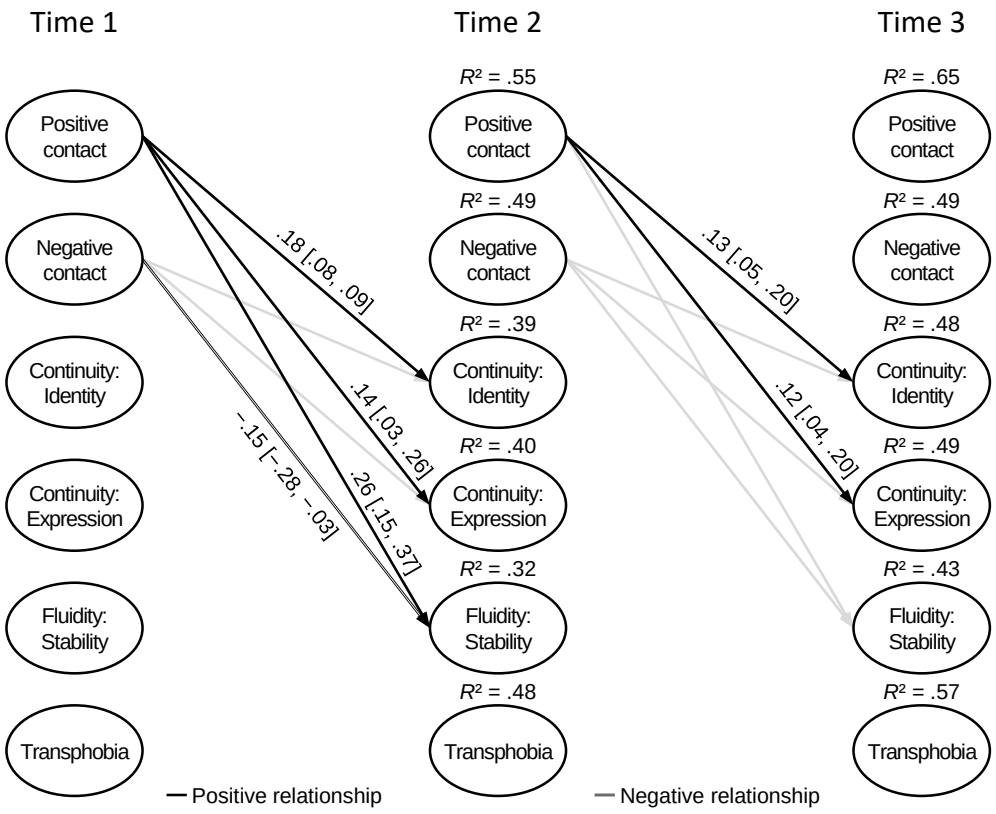

Figure 3.17. Longitudinal structural equation model with standardised coefficients and 95% confidence intervals (Study 3b). Paths in grey were not significant, $p > .05$. Autoregressive paths are not shown.

ative contact to the scale dimensions were significant. Together, these findings provided longitudinal evidence *for* the hypothesised relationships between intergroup contact and social identity continuity and fluidity, but *against* the hypothesised relationships between both social identity continuity and fluidity and outgroup attitudes.

Discussion

In this chapter, I have examined whether individuals differ in how continuous and fluid they understand their sexuality and gender to be, and whether these differences explain how intergroup contact reduces intergroup bias. Studies 1–3 found that intergroup contact was associated with more continuous and fluid social identities (except in Study 3a), but that more continuous and fluid social identities were not, in turn, associated with more favourable outgroup attitudes (except in Study 2b). Below, I summarise findings from the present research, before discussing its strengths, limitations, and implications.

Summary of findings

In this chapter, I have introduced sexuality and gender as categories that are continuous and fluid, but are often understood as binary and stable. Studies 1–3 examined whether individuals differ in the extent to which they understand their own sexuality (Studies 1,

2a, 3a) and gender (Studies 2b, 3b) as continuous and fluid, rather than binary and stable. I developed two novel scales to measure these individual differences. For sexuality, I developed a four-dimensional scale that measured to what extent participants understood their sexual identity and attraction as continuous, and to what extent they believed that their sexuality had changed in the past and would change in the future. For gender, I developed a five-dimensional scale that measured to what extent participants understood their gender identity and expression as continuous, to what extent they thought of their gender identity and expression as fluid, and to what extent they believed that their gender identity and expression would remain stable across their lifetime. Studies 1–3 showed that the resulting scale dimensions were internally consistent, temporally stable, and measurement invariant across time and translations. The present research thus documented reliable differences in social identity continuity and fluidity across individuals.

Studies 1–2 assessed the novel scales' construct validity by testing whether scale dimensions correlated (in the expected direction) with variables I expected the underlying construct to correlate with. As expected, more continuous and fluid identities correlated positively with same-gender attraction, and negatively with gender norm endorsement, group identification, social dominance orientation, and need for closure. Scale dimensions tended to correlate negatively with essentialist beliefs about sexuality and gender, with effect sizes ranging from small to medium. This is evidence that the novel scales and essentialist beliefs scales (Haslam & Levy, 2006; Tee & Hegarty, 2006) measure related, but distinct, constructs. In Chapter 1, I have argued that essentialist beliefs measure how individuals perceive the social categories of others, and that researchers need to investigate how individuals understand their own category memberships. The present research supported this argument, as outlined above. Further, Studies 2–3 found that the novel scales consisted of highly correlated but distinct subscales. The present research thus confirmed that both sexuality (identity, attraction) and gender (identity, expression) span multiple dimensions.

Studies 1–3 tested whether contact with sexual minorities (Studies 1, 2a, 3a) and gender minorities (Studies 2b, 3b) fosters more continuous and fluid conceptions of sexuality and gender (hypothesis 1). Across studies, I found mixed evidence for this hypothesis. In Study 1, positive contact was consistently associated with more continuous and fluid conceptions of sexuality. In Study 2a, contact quantity and positive contact were associated with, respectively, two and one of the four continuity/fluidity subscales. In Study 2b, contact quantity, but not positive contact, was associated with two of the five continuity/fluidity subscales. In Study 3b, but not in Study 3a, positive contact predicted more continuous and fluid identities over time. Across studies, negative contact was not consistently associated with less continuous and fluid identities. Overall, the present research found consistent support for the hypothesised relationships between contact with gender minorities and more continuous and fluid conceptions of gender. However, the present research found inconsistent evidence for the hypothesised relationships between contact with sexual minorities and more continuous and fluid conceptions of sexuality.

Finally, Studies 1–3 tested whether more continuous and fluid conceptions of sexuality (Studies 1, 2a, 3a) and gender (Studies 2b, 3b) reduced sexual prejudice and transphobia (hypothesis 2). As expected, more continuous and fluid identities consistently correlated with less sexual prejudice and transphobia in Studies 1 and 2. Other than expected, more continuous and fluid identities were not associated with sexual prejudice or transphobia in any of the structural equation models estimated for Studies 1–3. This finding suggests that the novel scales did not predict intergroup bias after controlling

for intergroup contact and/or autoregressive paths. Overall, the present research found some evidence *for* the hypothesised relationships between intergroup contact and social identity continuity and fluidity, but *against* the hypothesised relationships between social identity continuity and fluidity and intergroup bias. In contrast to the bridging-binaries hypothesis, more continuous and fluid identities do not explain how intergroup contact reduces intergroup bias. Below, I discuss strengths and limitations of the present research.

Strengths and limitations

A recent survey of social and personality research found that many studies rely on idiosyncratic and poorly validated measurement scales (Flake, Pek & Hehman, 2017). This matters because unreliable measurement, alongside small samples and analytic flexibility, partly explains recent failures to replicate findings from psychological research (Open Science Collaboration, 2015; see Loken & Gelman, 2017). In contrast, the present research followed up-to-date guidelines for scale development and validation (DeVellis, 2016; Lenzner & Menold, 2016; Menold & Bogner, 2016). Thus, the present research resulted in two novel scales that measure reliable and stable differences in individuals' conceptions of sexuality and gender. Prior attempts to measure continuity and fluidity (e.g., YouGov, 2015, 2016) instead used short scales that did not capture all dimensions of sexuality and gender. Another strength of the present research is that it included longitudinal research, which allowed me to establish the direction of any observed relationships. Further, the present research included $N = 2638$ participants in total, which allowed precise estimates of the hypothesised relationships. A final strength of the present research is that it established the validity of the novel scales across two languages and nations.

Despite its strengths, the present research is limited by six methodological limitations. First, Studies 1–3 relied on convenience samples of university students. As such, participants came from the generation that reported the most continuous conceptions of sexual attraction and gender expression in representative surveys (YouGov, 2015, 2016). It is unclear how intergroup contact would affect conceptions of sexuality and gender among older adults and/or less-educated individuals. On the one hand, older adults may be more resistant to continuous and fluid conceptions. On the other hand, older adults may be less accustomed to these ideas, and may stand to gain more from interacting with sexual and gender minorities. Second, Studies 2b and 3b examined contact with and attitudes toward transgender people in samples in which, respectively, 69% and 78% of participants identified as heterosexual. These proportions are slightly lower than the proportions of 18 to 24-year-olds (83%) and 25 to 39-year-olds (81%) who identified as heterosexual in YouGov's (2015) representative survey.[13] Still, Study 3b, but not Study 2b, was close to representative for its age range (18–39 years).

Third, Studies 1–3 were limited by their correlational designs, and could not rule out the confounding effects of unmeasured covariates. Fourth, Studies 2a and 2b used the same sample to establish the factor structure of the proposed scales and to estimate correlations between the resulting scale dimensions and various other constructs. Studies 3a and 3b, however, tested this chapter's substantive hypotheses in a different sample. Fifth, Studies 2a and 2b assessed construct validity with variables that differed from the newly developed scales in the method with which they were measured (for example, by using different response scales). As such, I was not able to use a multitrait–

[13] YouGov (2015) asked participants which of several options best described their sexuality. Participants chose between "heterosexual", "gay or lesbian", "bisexual", "other", and "prefer not to say".

multimethod approach (Campbell & Fiske, 1959) to clearly establish convergent and divergent validity.

Finally, Studies 1–3 aggregated both contact with and attitudes toward binary outgroups (gay/lesbian people, transgender men/women) and non-binary outgroups (bi/pan/queer people, genderqueer people). In line with the arguments presented in Chapter 1, contact with non-binary outgroups should be more effective in fostering continuous and fluid identities, while continuous and fluid identities should matter more for attitudes toward non-binary outgroups. The present research, however, did not consistently differentiate between binary and non-binary outgroups.[14]

Despite mixed findings, the present research resulted in well-validated scales for measuring continuous and fluid conceptions of sexuality and gender. Future research should address limitations of the present research by validating these scales in more representative samples, by manipulating contact and continuity/fluidity experimentally, and by testing the hypothesised relationships for non-binary and binary outgroups. Below, I discuss implications of the present findings for intergroup relations research.

Implications

Findings from the present research offer insights into the relationships between intergroup contact and outgroup attitudes, between intergroup contact and social identification, and between social identification and outgroup attitudes. Below, I discuss implications for each of these relationships in turn.

The present research added to a growing number of studies examining how intergroup contact relates to prejudice against sexual and gender minorities. Past research has consistently found contact with sexual minorities to reduce sexual prejudice (for meta-analyses, see Pettigrew & Tropp, 2006; Smith, Axelton & Saucier, 2009). Studies 1, 2a, and 3a provided additional cross-sectional and longitudinal evidence that positive contact reduces prejudice. In addition, Studies 1 and 2a, but not Study 3a, found that negative contact was associated with more sexual prejudice (as shown, for example, by Barlow et al., 2012). Few studies have examined how contact with gender minorities relates to anti-transgender prejudice. King, Winter and Webster (2009) found, in a study of Hong Kong residents, that people who reported at least some contact with transgender people showed less anti-transgender prejudice than people who did not report any contact (N = 865). Hoffarth and Hodson (2018) found that among Canadians and US Americans, contact quality, but not contact quantity, were associated with less anti-transgender prejudice (N = 451).[15] Studies 2b and 3b included more participants than both prior studies combined (N = 2035), and also included the first longitudinal study of this relationship. Surprisingly, the present research did not replicate earlier studies, as neither positive nor negative contact predicted anti-transgender prejudice over time. Being transgender is concealable and relatively rare, limiting opportunities for intergroup contact. Future research should explore whether these and other factors explain why contact does not seem to reduce anti-transgender prejudice.

As discussed, the present research found that intergroup contact was associated with more continuous and fluid social identities. Alongside other studies (e.g., van

[14]Studies 2 and 3, but not Study 1, measured contact separately for each outgroup. Study 2, but not Studies 1 and 3, measured attitudes separately for each outgroup.

[15]Tee and Hegarty (2006) studied how intergroup contact related to opposition to transgender people's civil rights, but did not measure prejudice. I discuss this study in Chapter 5.

Dommelen et al., 2015; Schmid et al., 2009), the present research heeded Pettigrew's (1997, 1998) call to study whether intergroup contact reduces prejudice by encouraging individuals to reappraise their ingroup (deprovincialisation). Prior research focused either on national pride and other measures of strength of identification (Pettigrew, 1997), or on identification across multiple categories (van Dommelen et al., 2015; Schmid et al., 2009). In contrast, Studies 1–3 tested whether contact can lead individuals to reappraise their social identification along a single category. Across both case studies, I found some support for the hypothesis that contact with diverse others encourages people to rethink their social identities—though the evidence was more consistent for gender than for sexuality. Future research should aim to extend this research to other categories, outside the binaries of sexuality and gender. This research would allow researcher to study whether the observed relationship generalises to other category domains, and could lead toward the development of a generalised social identity continuity and fluidity scale.

In Chapter 1, I have argued that fostering more continuous and fluid identities should blur the boundaries between ingroup and outgroup, and thus undermine the motivational underpinnings of intergroup bias. This hypothesis resounds with older explanations of how contact could reduce prejudice. Brewer and Miller (1984) argued that contact could overcome our tendency to divide others into "us" and "them" (decategorisation). Pettigrew (1997, 1998) argued that contact could encourage us to look beyond our ingroup, and recognise that the outgroup's norms and customs are also valuable (deprovincialisation). In contrast to this tradition, the present research did not find how participants understood their social identities to be associated with intergroup bias. Instead, this finding supports R. Brown and Hewstone's (2005) conclusion. That is, intergroup contact reduces prejudice through affective and cognitive processes (see Pettigrew & Tropp, 2008), and not by blurring boundaries between groups. The present research focused on identification within a single category domain. In the next chapter, I examine whether intergroup contact can change how we think of ourselves and others in terms of multiple category domains.

4

Bridging divides: Can intergroup contact foster more inclusive national and religious identities?

In Chapter 3, I have studied whether intergroup contact can bridge binaries by fostering more continuous and fluid conceptions of sexuality and gender. I found that positive contact predicted more continuous and fluid social identities—which, however, were not associated with more favourable outgroup attitudes. This research focused on social identification within a *single* category domain. However, in many situations, we think of ourselves and others in terms of *multiple* category domains. In this chapter, I investigate how individuals construct their social identities from multiple cross-cutting categories, and how these identities relate to intergroup contact and intergroup bias.

Van Dommelen et al. (2015) developed the triple crossed-categorisation task to study how *individuals* construct their social identities from multiple group memberships. I adapt this paradigm to examine also how *groups* differ in whom their members consider "us" and "not us". Social groups structure their members' lives, especially where groups differ in status and power. Because of these shared experiences, I expect that group differences are as great as, or greater than, individual differences. Shared experiences, however, make only part of an individual's day-to-day experiences. Of these, contact experiences could overcome group differences by encouraging individuals to include the interaction partner's group(s) in their ingroup. As argued in Chapter 1, I hypothesise that intergroup contact could foster more inclusive social identities, and that more inclusive social identities, in turn, reduce intergroup bias. Contact experiences could thus bridge social divides by fostering more inclusive social identities.

I test these hypotheses in the context of caste, religion, and nation in South India. Specifically, I test whether contact with ethnoreligious and caste outgroups fosters more inclusive social identities among South Indian students, and whether more inclusive identities reduce intergroup bias. Below, I describe the social context in which the present research was conducted.

Caste, religion, and nation in South India

For the present research, I collaborated with Professor Shanmukh Kamble at Karnatak University, Dharwad in the South Indian state of Karnataka (see Figure 4.1). South India provides a rich context for studying social identification across multiple social categories. The region, like the rest of the country, is religiously and ethnically diverse. Its inhabitants belong to a multitude of categories that converge, diverge, and intersect in various ways. Of these categories, the present research focuses on caste, religion, and nationality. Below, I introduce these categories, and discuss their current relevance for intergroup relations.[1]

Caste is a system of social relations that divides communities into locally specific hierarchies (Jodhka, 2012). Markers of caste are diverse and include, for example, surname and diet. Historically, the caste system has functioned as a coercive hierarchy that concentrates power and status in the hands of a dominant caste. This hierarchy is maintained by restricting subordinate castes' occupational choices, by barring subordinate castes from owning land, and by enforcing endogamy and segregation. This hierarchy is underpinned by a tradition that ranks castes according to their supposed ritual purity, and condemns contact with less pure castes as polluting.[2] In this tradition, the dominant caste tends to claim the highest ritual status, followed by various intermediate castes. Almost invariably, the caste system places some groups below the 'line of pollution' at the bottom of the ritual hierarchy. A practice of untouchability excludes these castes from public life, and consigns them to landlessness and menial labour, often on the fields of the dominant caste. In India today, *Dalits*—members of castes that have been affected by untouchability—still face economic inequalities that have resulted from this history of exploitation. In addition, Dalits still face interpersonal discrimination that originates from the tradition of ritual status. In the next section, I discuss how legislative and economic changes have shaped current caste relations.

Since Independence, the Indian State has sought to reduce caste-based injustices through various policy measures (Jodhka, 2012). First, the 1950 Indian Constitution prohibited the practice of untouchability (Article 17) as well as discrimination on the grounds of religion, race, caste, sex, or place of birth (Article 15). Since then, governments have introduced further legislative measures to protect Dalits' civil rights and to prevent atrocities against Dalits. Second, the Indian State has instituted one of the most extensive affirmative action programmes in the world. The Constitution recognised communities that had been affected by untouchability (Dalits) as *Scheduled Castes* (SCs), and made room for special provisions to advance these and other disadvantaged groups. Alongside SCs, the Constitution recognised Adivasi—India's indigenous peoples—as *Scheduled Tribes* (STs). Since the 1950s, governments have enacted quotas that reserve seats in state-run educational institutions and state-sector jobs for SCs and STs. Third, in the 1990s, the government extended the reservation system to communities that had been disadvantaged historically, but had been above the 'line of pollution'. The administrative term for these communities is *Other Backward Classes* (OBCs). Today, state-run educational institutions reserve at least 49.5% of available seats—15.0% for Scheduled Castes, 7.5% for Scheduled

[1] I focus on South India since introducing the complexities of caste and religion across the entire country would be impossible in the available space. South Indian states share many linguistic, cultural, and economic similarities, making this region a meaningful unit of analysis.

[2] Orientalist scholars have often reduced caste to ritual status, based in Hindu scripture, that divides Hindus into four varnas—Brahmins, Kshatriyas, Vaishyas, and Shudras—and is, at least to some extent, consensual. Contemporary scholars tend to contest this view, highlighting the complexity of caste, as well as the roles of power, coercion, and resistance in its maintenance (for a discussion, see Jodhka, 2012).

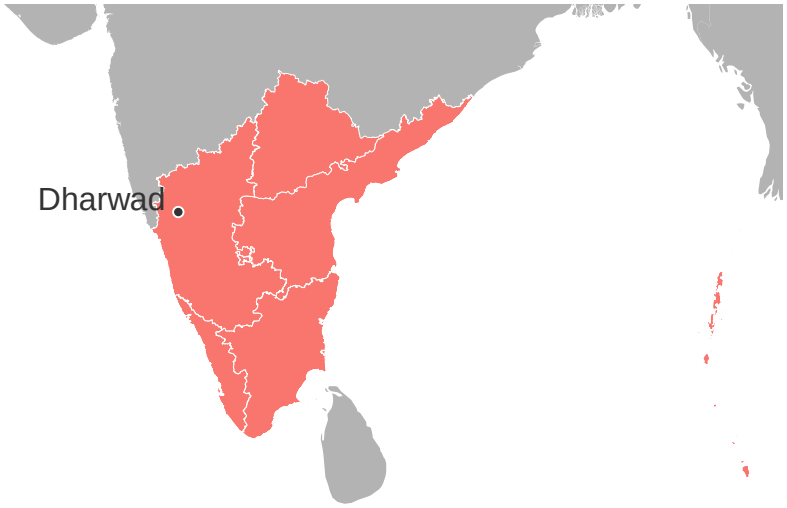

Figure 4.1. Map of South India with state borders.

Tribes, and 27.0% for Other Backward Classes. Together, these changes have widened educational and occupational opportunities for Dalits and other disadvantaged groups. At the same time, these developments have given rise to caste-based identity politics, as discussed below.

Since reservations were introduced, the gap between Dalits' and other caste groups' educational and occupational achievements has narrowed substantially. Still, caste-based inequalities persist. In Karnataka, for example, 17% of the population belong to Scheduled Castes, but make up more than 31% of the state's houseless population and only 9% of its university graduates (Office of the Registrar General & Census Commissioner, 2011). Likewise, interpersonal discrimination persists. The continued practice of untouchability affects Dalits in many rural communities (Shah, Mander, Thorat, Deshpande & Baviskar, 2006), while less overt prejudices hold back educated Dalits in urban job markets (Jodhka, 2012). Dalits, however, have become increasingly organised in asserting their economic and political rights—as evident, for example, in the founding of the Dalit Panthers in 1972 or the Bahujan Samaj Party in 1984. In these movements, Dalits have built a positive identity defined in their resistance against oppression. For Dalits, caste is thus defined by the experience of discrimination and their resistance against it (Jodhka, 2012, p. 169).

For historically dominant castes, Dalit emancipation threatens erstwhile privileges and advantages. Reservations, for example, make it nominally more difficult for those who have to compete for university seats and government jobs in the *General Merit* (GM) category.[3] In rural areas, Dalit movements threaten the economic and social dominance of upper-caste communities. As a result, Dalit assertions are often met with violent backlash (e.g., "15 held for assault on Dalits", 2017). Since 2001, atrocities against Scheduled Castes have increased by 22% (National Crime Records Bureau, 2016). The advantaged thus experience caste as threatened privilege. Other groups have historically occupied an intermediate position in the caste hierarchy. For them, caste is often defined by the question of whether they have access to reserved seats and jobs. Caste groups who are already recognised as Other Backward Classes have an interest in maintaining or

[3]The General Merit category includes anyone who does not claim an Scheduled Caste, Scheduled Tribe, or Other Backward Class background. Other labels used to describe historically advantaged castes are Forward Caste or General Caste.

increasing Other Backward Class reservations. Caste groups who are not yet recognised as OBCs have an interest in gaining access to reservations. As a result, caste identities have remained important in electoral politics and political organising (see Jodhka, 2012).

Caste continues to structure intergroup relations in India, but it is not the only category domain that matters. South India is religiously diverse. In Karnataka, 84% of inhabitants are Hindus. At 13%, Muslims are the largest ethnoreligious minority.[4] Christians (2%), Jains (1%), and various smaller communities make up the remaining proportion (Office of the Registrar General & Census Commissioner, 2011). The present research focuses on Hindu–Muslim relations. Like some caste groups, Muslims also face inequalities of wealth and education. Karnataka, for example, recognises some Muslim communities as OBCs (National Commission for Backward Classes, 2015). India is secular by constitution, but has a long history of communal violence. In the wake of Independence, the Partition of India left 15 million displaced, and 200,000 to 2 million dead (Talbot & Singh, 2009). Since then, communal tensions have led to violence against Muslim minorities in India. In 2002, for example, more than a thousand Muslims were killed in an ethnically-motivated pogrom in Gujarat (Dhattiwala & Biggs, 2012). Recent years have seen an increase in anti-Muslim violence—in particular, mob violence against people suspected of slaughtering cows (Amnesty International, 2017). Below, I discuss Hindu nationalism as an ideological underpinning of these divisions.

Beyond caste and religion, national identity plays a central role in Indian society and politics.[5] Before Independence, the anti-colonial struggle gave rise to two broad currents of nationalism (Menski, 2009). On the one hand, the Indian National Congress party espoused an inclusive nationalism that strove to include Indians of all religions and backgrounds. After Independence, Congress formed all governments until 1977, and instituted many of the secular foundations of the Indian state. On the other hand, some factions of the anti-colonial movement advocated Hindu nationalist ideologies. Broadly speaking, Hindu nationalism (*Hindutva*) is an ideology that equates being Indian with being Hindu, and thus excludes ethnoreligious minorities. Since Independence, Hindu nationalists have opposed the secular nationalism advocated by Congress. Today, the ruling Bharatiya Janata Party (BJP) espouses Hindutva ("Bharatiya Janata Party", 2014). This nationalism is, explicitly or implicitly, defined against the past or present threat from Muslims. A document on the BJP's website, for example, compares the Mughal's historic invasion of India to the Holocaust (Bharatiya Janata Party, n.d.). Hindu nationalism has been linked to recent anti-Muslim violence. Dhattiwala and Biggs (2012), for example, showed how the BJP profited from the 2002 pogrom in Gujarat, which is often linked to the then-Chief Minister and now-Prime Minister Narendra Modi. At the time of this research, the Modi government enjoyed broad support in the Indian population (Pew Research Center, 2017). This means that Indians are faced with two nationalisms that either include or exclude Muslims from the Indian identity. In this section, I have discussed how caste, religion, and nation structure current intergroup relations in South India. In the next section, I introduce the present research.

[4] These proportions resemble India as a whole, where Hindus and Muslims make up 80% and 14% of the population, respectively.

[5] In the present research, for example, 92% of participants *agreed* (11%) or *strongly agreed* (81%) with the statement: "I identify with my nationality."

Study 4

Study 4 investigated how South Indian students construct their social identity from multiple, overlapping group memberships. Specifically, I examined how participants used nationality, religion, and caste to decide whom they considered "us" and "not us". To that end, participants completed a *crossed-categorisation task* (adapted from van Dommelen et al., 2015) in which they reviewed 24 targets, each representing a person with whom participants shared none, one, two, or all of three group memberships. For each target, participants reported whether they considered that person as "us" or "not us". Participants' responses thus showed whom they included in their ingroup—and whom they excluded.

I studied participants' responses in order to answer research questions on individual and group differences in participants' categorisations, and about the consequences of these categorisations for intergroup relations. Before turning to these questions, I explored to what extent participants' responses corresponded to various social identity structures proposed by Roccas and Brewer (2002). In line with current political developments, I was particularly interested in how nationalism would structure participants' categorisations. If participants espoused *inclusive nationalism*, they should include Indians of all castes and religions in their ingroup. If participants instead endorsed *Hindu nationalism*, they should exclude Indian Muslims from their ingroup. I explored to what extent these structures, as well as more and less inclusive patterns, described participants' social identities.

As participants (and targets) held various group memberships, I tested how participants' *objective* group memberships shaped their *subjective* construals of social identity. Broadly, I hypothesised that Hindu participants would not tend to include targets from religious (e.g., Indian Muslims) and national (e.g., Sri Lankan Hindus) outgroups. I expected that participants would be least likely to include targets in their ingroup who differed in terms of both group memberships (e.g., Bangladeshi Muslims). Participants (and targets) came from different caste backgrounds. I hypothesised that participants from higher castes (GM: *General Merit*) would tend to exclude targets from the historically most disadvantaged caste backgrounds (SC/ST: *Scheduled Caste*, *Scheduled Tribe*). Similarly, I expected participants from the most disadvantaged caste backgrounds to exclude targets from upper castes.

I had less clear predictions for participants and targets from so-called *Other Backward Classes* (OBC). If Other Backward Class participants saw themselves as closer to upper-caste people, I would expect them to categorise General Merit but not Scheduled Caste / Scheduled Tribe targets as "us". If Other Backward Class participants saw themselves as closer to Dalits/Adivasis, I would expect them to categorise Scheduled Caste / Scheduled Tribe but not General Merit targets as "us". Alternatively, Other Backward Class participants could occupy the middle ground between the other caste groups, and identify with neither or both caste groups. Conversely, Scheduled Caste / Scheduled Tribe and General Merit participants could categorise Other Backward Class targets in ways that contradict or mirror Other Backward Class participants' own categorisations.

Though I expected to find substantial group differences in participants' ingroup construals, I also expected to find some individual differences. I hypothesised that past experiences with members of religious and caste outgroups would shape whom participants considered "us" and "not us". Specifically, I hypothesised that *positive contact* and *cross-group friendship* would be associated with more inclusive identities, while *negative contact* would be associated with less inclusive identities. I expected that *cross-group*

friendship would show a stronger association with identity inclusiveness than more casual forms of *positive contact*.

In addition, I tested ideological preferences as an alternative explanation of individual differences in ingroup construals. I hypothesised that participants' ideological preference for *social dominance* would be associated with excluding targets of lower-status outgroups from the common ingroup. Sidanius and Pratto (1999; Ho et al., 2015) argued that social dominance orientation measures a general preference for intergroup hierarchy—independent of the status of one's own group—rather than a specific preference for ingroup dominance. Accordingly, I expected that for participants of all caste backgrounds, social dominance orientation would be associated with exclusion of targets of comparably lower social status (i.e. OBC, SC/ST, Muslim, and/or foreign targets for General Merit participants; SC/ST, Muslim, and/or foreign targets for Other Backward Class participants; and Muslim and/or foreign targets for Scheduled Caste / Scheduled Tribe participants) but not targets of comparably higher status. I had no clear prediction about differences across the two subdimensions, dominance and anti-egalitarianism, of social dominance orientation.

Aside from studying why participants categorised some targets as "us" and others as "not us", I examined whether these categorisations were associated with participants' feelings about members of national, religious, and caste outgroups. I hypothesised that categorising a target as "us" would be associated with *more favourable attitudes* and *less desired social distance* toward that target. More broadly, I hypothesised that participants who included Muslim and Scheduled Caste / Scheduled Tribe targets in their ingroup would feel less threatened by members of these groups, in terms of both realistic and symbolic threat.

Aside from research questions about antecedents and consequences, I studied how responses in the crossed-categorisation task related to other measures of social identity. I included these measures to validate the triple crossed-categorisation task, and to relate the findings of the present study to other research in this thesis. I report the relevant measures and analyses in Appendix B. Below, I first describe a pilot study that informed which targets I selected for the crossed-categorisation task, before describing the sample, procedure, measures, and results of the final study.

Pilot study

Informed by local knowledge, I intended to rely on last names to unobtrusively communicate caste membership in the crossed-categorisation task. To find the most prototypical stimuli, I compiled an initial set of 50 surnames, drawing on local knowledge, databases on naming preferences, and publicly available archives of Facebook usernames. I selected 10 surnames for each of five combinations of nationality, religion, and caste: (1) Indian Hindus of upper-caste backgrounds, (2) Indian Hindus of lower-caste backgrounds (Dalits), (3) Indian Muslims, (4) foreign Hindus, and (5) foreign Muslims. From these 50 names, I created 100 stimuli resembling identity cards, 50 with female first names and 50 with male first names. Each card had the target's first and last name, age (21–26 years), and religion (Hindu, Muslim) printed on them, as well as a flag corresponding to the target's nationality (Indian, Nepali, Sri Lankan, Bangladeshi) and a silhouette corresponding to the target's gender (adapted from Ma, Correll & Wittenbrink, 2015). Cards for female and male targets were identical in all attributes except first name and silhouette.

I asked 26 students (19 women, 7 men) of diverse caste groups (*SC/ST*: 8, *GM*: 10,

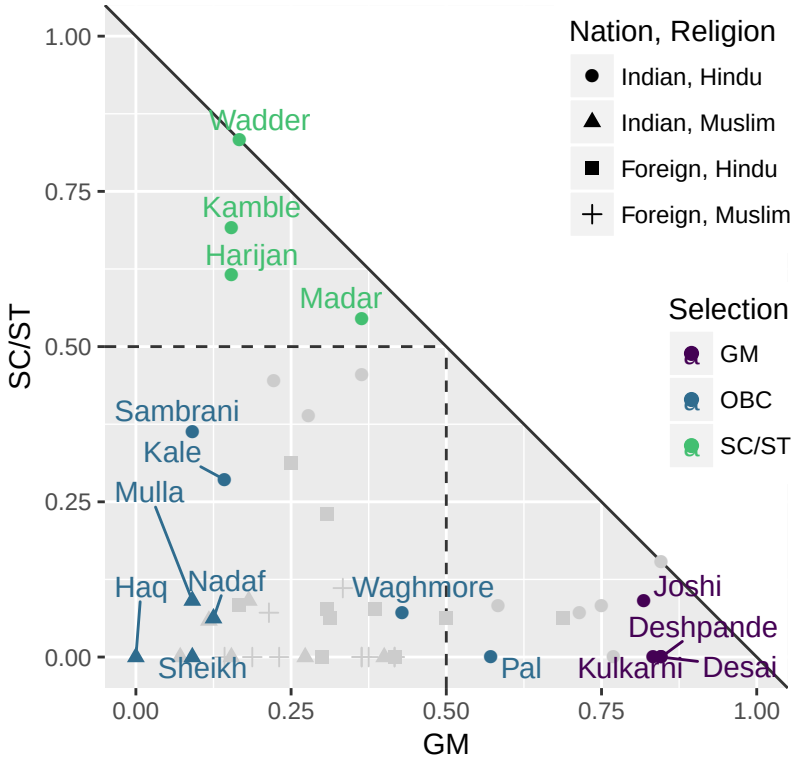

Figure 4.2. Proportion of SC/ST vs GM classifications per target stimulus. Targets in the upper-left quadrant were most often classified as SC/ST, while targets in the lower-right quadrant were most often classified as GM. Targets in the lower-left quadrant were more often classified as GM or OBC, or OBC than as SC/ST or GM, or could not be classified reliably. Colours highlight Indian targets selected as GM, OBC, and SC/ST stimuli, respectively, based on participants' categorisations in the pilot study. GM = General Merit, OBC = Other Backward Class, SC/ST = Scheduled Caste / Scheduled Tribe

OBC: 8) and religious backgrounds (*Hinduism*: 16, *Islam*: 4, *Christianity*: 5, *Jainism*: 1) to review 25 of the 50 stimuli corresponding to their gender. For each stimulus, I asked participants which caste they thought the target most likely belonged to (*SC/ST*, *GM*, *OBC*, *GM or OBC*, *Can't tell/Don't know*), how typical or unusual they considered the target's name for someone of that particular nationality, religion, and (if applicable) caste (1 = *very typical*, 6 = *very unusual*), and whether the target's name reminded them of anyone in particular (*yes*, *no*) and, if so, of whom.

Participants' responses had important implications for designing the crossed-categorisation task. On average, participants regarded all targets' names as at least "somewhat typical" ($1.42 \leq Ms \leq 3.50$). Names, however, were not enough to reliably cue caste membership—only 55% to 83% of participants identified even the most distinctive Dalit targets as Scheduled Caste / Scheduled Tribe (see Figure 4.2). Further, participants consistently categorised Muslim targets as Other Backward Class (see Figure 4.3), introducing a potential confound when comparing Muslim (OBC) and Hindu (GM) targets. Considering these findings, I decided to explicitly state caste membership (as reservation categories: GM, OBC, SC/ST) for Indian targets, and to add *Indian, Hindu, OBC* targets as a sixth category (see Figure 4.4).

For the final set of stimuli, I chose the four most prototypical targets for each

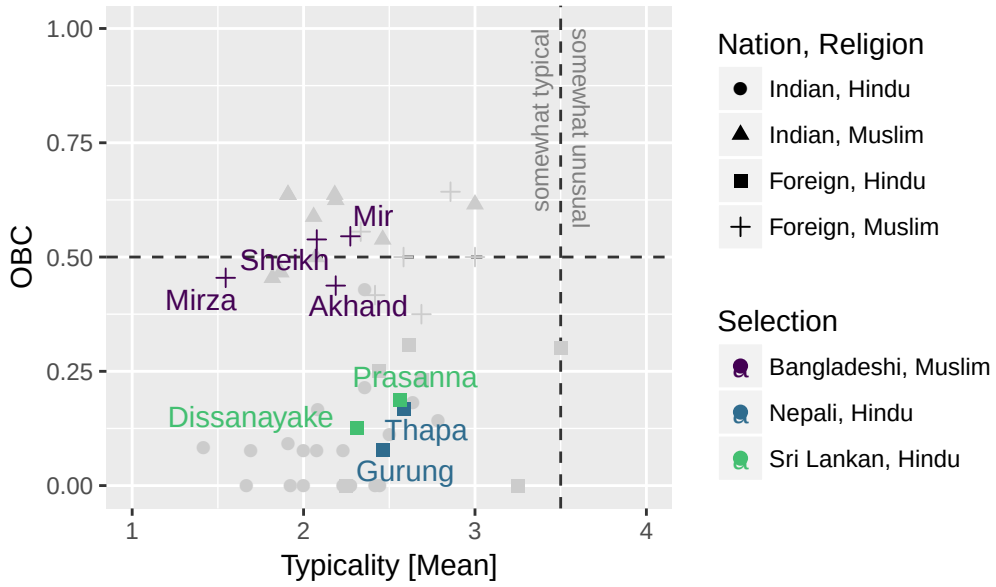

Figure 4.3. Foreign targets selected as Sri Lankan, Hindu (2), Nepali, Hindu (2), and Bangladeshi, Muslim (4) stimuli, respectively, based on participants' average typicality ratings.

of the six categories under study. I ranked Indian targets by the proportion of "correct" categorisations (e.g., Pr(*OBC*) for *Muslim, OBC* targets) weighted by the inversed proportions of "incorrect" categorisations (e.g., 1 – Pr(*SCST*) and 1 – Pr(*GM*) for *Muslim, OBC* targets). I selected the four highest-ranked (Category 1) *Hindu, GM*, (2) *Hindu, OBC*, (3) *Hindu, SC/ST*, and (4) *Muslim, OBC* targets (see Figure 4.2). I ranked foreign targets by their average typicality rating, and excluded targets who were categorised as General Merit or Scheduled Caste / Scheduled Tribe by ≥ 50% of participants (to avoid obvious caste associations). I selected the two most typical (Category 5) *Nepali, Hindu* and *Sri Lankan, Hindu* targets, and the four most typical (6) *Bangladeshi, Muslim* targets (see Figure 4.3). Based on the pilot study, I thus used 6 (categories) × 4 (targets) = 24 stimuli in the final study. Between 0% and 44% of participants (*Mdn* = 10%) recognised each selected name as familiar—in most cases, participants reported that targets reminded them of fellow students, former classmates, or co-workers.

Method

Participants

Three hundred and fifty-one students at Karnatak University (Dharwad, India) participated in exchange for writing materials. Of these, I excluded 49 participants who did not belong to any of the four caste groups (n = 20), failed to indicate their caste group (n = 7), or chose Islam as their own or their family's religion (n = 27).[6] This left 302 participants who chose Hinduism (n = 286), Jainism (n = 8), or Christianity (n = 8) as their or their family's religion, and General Caste (n = 99), Other Backward Class (n = 127), Scheduled

[6] I excluded Muslim participants as this subsample was too small to test how its participants differed from participants of other religions. I did not exclude Jain or Christian participants as their religion was not shared by any of the target stimuli.

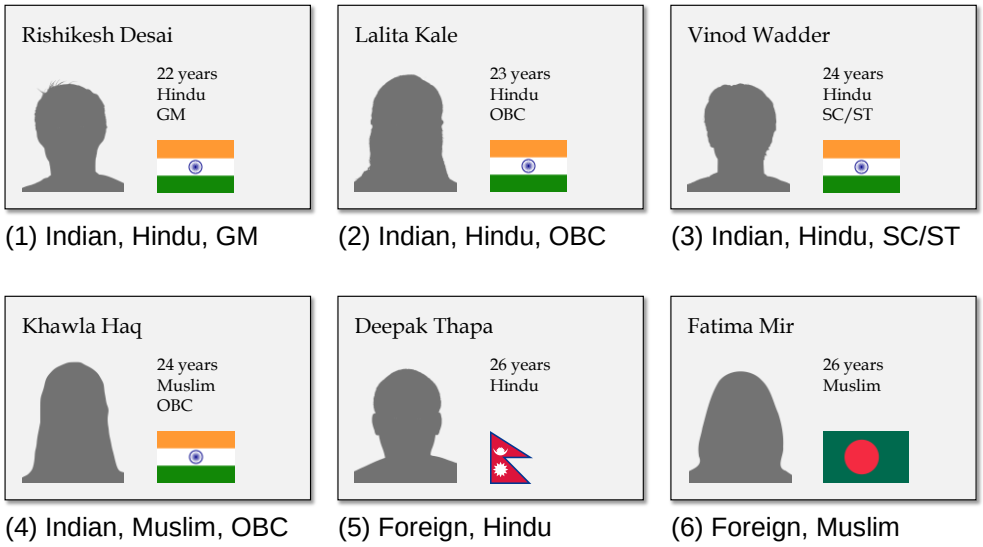

Figure 4.4. Examples of stimuli used for six target categories in Study 4. Ages and silhouettes varied across sessions.

Caste (n = 54), or Scheduled Tribe (n = 22) as their caste group. Table 4.1 summarises all participants by gender, age, nationality, religion, and caste.

Procedure

In the first part of the study, participants completed the *crossed-categorisation task* for the 24 targets selected in the pilot study. Rather than test participants individually, I adapted van Dommelen et al.'s (2015) task to simultaneously test classrooms of 24–71 participants. Instead of asking participants to sort physical cards, the experimenter presented stimuli in a slide-based presentation. Each stimulus represented a (fictitious) person as an identity card showing the person's first and last name, age, religion, nationality, and caste reservation (for Indians), as well as a head-and-shoulders silhouette (for examples, see Figure 4.4). In addition, slides contained a number identifying each target, as well as the response scale(s) corresponding to the question(s) participants were answering at the time (e.g., a feeling thermometer for Question 3). Each slide contained a male target (on the left) and a female target (on the right). The experimenter asked participants to focus on the target that corresponded to their own gender. For a detailed script, see Appendix B.

In each session, participants first familiarised themselves with each target (for 7s) in an automated slideshow. Participants viewed targets (in the same order) for a second time, noting for each target whether they felt that this person was one of their own group (1 = "us"), or not one of their own group (0 = "not us"). Participants viewed targets for a third time, rating how comfortable or uncomfortable they would feel to share a room with this person (social distance; 1 = *very uncomfortable*, 7 = *very comfortable*), and how they felt toward this person (feeling thermometer; 0 = *cold*, *100* = *warm*). Participants then completed a questionnaire containing the measures described below. I report descriptive statistics and inter-variable correlations in Appendix B.

Table 4.1. Participants by gender, age, nationality, religion, and caste. Groups in *italics* were excluded from the final sample. N/A indicates missing responses.

Category		*n*	%
Gender	Woman	215	61
	Man	121	34
	N/A	15	4
Age	18–20	1	0
	21–23	254	72
	24–26	77	22
	27–29	10	3
	30–32	1	0
	36 or older	1	0
	N/A	7	2
Nationality	Indian	339	97
	N/A	12	3
Religion	Buddhism	1	0
	Christianity	11	3
	Hinduism	297	85
	Islam	*27*	*8*
	Jainism	8	2
	Other	2	1
	N/A	5	1
Caste	General Caste	104	30
	Other Backward Class	143	41
	Scheduled Caste	54	15
	Scheduled Tribe	23	7
	Other / Not applicable	*20*	*6*
	N/A	*7*	*2*
Total		351	100

Measures

Intergroup contact was measured as: how often, from 1 = *never* to 5 = *very often*, participants meet outgroup members in their everyday life (*contact quantity*); and how often, on average, they have positive/good contact (*positive contact*) and negative/bad contact (*negative contact*) with outgroup members (adapted from Barlow et al., 2012). To make participants' responses more comparable, I preceded these items with examples of concrete experiences:

> When meeting other people, we can make have positive/good experiences (e.g., being supported, made to feel welcome, helped, complimented, or befriended by them) and negative/bad experiences (e.g., being made to feel unwanted, verbally abused, intimidated, threatened, or ridiculed by them).

Outgroup friendship was measured with two items (adapted from R. N. Turner, Hewstone & Voci, 2007): "How many close friends do you have who are [outgroup members]?" (1 = *none*, 5 = *more than ten*), and "How often do you spend time with [outgroup] friends?" (1 = *never* to 5 = *very often*; $.47 \leq rs \leq .58$). Participants reported their experiences with four groups: Dalits, people from other backward classes, people from general castes, and Muslims.

Social dominance orientation was measured as how much, between 1 = *strongly oppose* and 7 = *strongly favour*, participants endorsed eight statements about social hierarchies ($SDO_{7(S)}$ scale, Ho et al., 2015). Four items measured to what extent participants supported systems of group-based dominance (*SDO-Dominance*), for example: "Some groups of people are simply inferior to other groups". Four items measured to what extent participants opposed egalitarian ideologies (*SDO-Egalitarianism*), for example: "It is unjust to try to make groups equal".[7]

Realistic threat (from Muslims and Dalits) was measured with three items per each outgroup (adapted from Schmid, Hewstone, Küpper, Zick & Tausch, 2014): "The more power [Muslims/Dalits] gain in this country, the more difficult it is for [Hindus/people from my caste group]", "[Muslims/Dalits] have been trying to get ahead economically, at the expense of [Hindus/my caste group]", and "More good jobs for [Muslims/Dalits] means fewer good jobs for [Hindus/people from my caste group]" (1 = *strongly disagree*, 5 = *strongly agree*; $\alpha_{Muslims} = .71$, $\alpha_{Dalits} = .79$). *Symbolic threat* was measured with two items per outgroup: "[Hindus/ people from my caste group] and [Muslims/Dalits] have very different values", and "[Muslims/Dalits] threaten [Hindus'/my caste group's] way of life" (1 = *strongly disagree*, 5 = *strongly agree*; $.44 \leq rs \leq .48$).

Results

I address all research questions in the order I have introduced them in. First, I examined to what extent various *social identity structures* described participants' response patterns. Second, I tested for *group differences* in participants' categorisations of ingroup and outgroup targets. Third, I analysed to what extent past experiences and ideological preferences explained *individual differences* in participants' categorisations. Fourth, I

[7] Ho et al. (2015) proposed a four-factor structure for the SDO_7 scale, with two substantive factors (SDO-D and SDO-E) and two method factors (pro-trait and con-trait) capturing inter-item correlations. Replicating Ho et al.'s findings, I found the proposed four-factor model to fit the data better than the one-factor ($\Delta\chi^2 = 155.14$, $p < .001$) or two-factor ($\Delta\chi^2 = 138.67$, $p < .001$) alternatives. To take this structure into account, I used latent factor scores to represent SDO-D and SDO-E in all analyses.

Table 4.2. Mean *intra*-category (**diagonal**) and *inter*-category correlations between participants' categorisations of targets in the respective categories.

#	Category	1	2	3	4	5	6
1	Indian, Hindu, GM	**.48**	.56	.33	.43	.16	.08
2	Indian, Hindu, OBC	.56	**.41**	.29	.46	.09	.06
3	Indian, Hindu, SCST	.33	.29	**.66**	.35	.03	.16
4	Indian, Muslim, OBC	.43	.46	.35	**.52**	.01	.24
5	Foreign, Hindu	.16	.09	.03	.01	**.50**	.47
6	Foreign, Muslim	.08	.06	.16	.24	.47	**.51**

tested what, if any, *consequences* participants' ingroup construals had for their attitudes and beliefs. Finally, I examined correlates of participants' responses to validate the crossed-categorisation task (see Appendix B).

Social identity structures

Van Dommelen et al. (2015) analysed participants' responses as *Social Identity Inclusiveness*, i.e. the total number of 24 targets a participant had categorised as "us", and as *Social Identity Structures*, i.e. the specific structures proposed by Roccas and Brewer (2002). Neither measure translated well to the present study. As participants had different caste backgrounds, comparing their inclusiveness scores (M = 15.34, SD = 4.88, Mdn = 16) was less meaningful than in van Dommelen et al.'s sample. Coding structures was interesting in principle, but proved difficult in practice. Few (9%, n = 26) participants consistently classified all targets in each category as either "us" or "not us". Relatedly, how participants classified targets within a category was correlated, but only moderately so ($.41 \leq rs \leq .66$, see Table 4.2). The inconsistency in participants' responses made hard and fast coding rules, as used by van Dommelen et al., impractical.

Instead, I quantified the (dis)similarity between a participant's responses and various proposed social identity structures by calculating edit distances. I operationalised *edit distance* as the number of a participant's categorisations that would need to change from "us" to "not us" (and vice versa) to make them identical to the categorisations implied by a certain structure. Figure 4.5 shows the distributions of participants' edit distances from each structure, as well as median edit distances.

Overall, participants' responses most closely resembled the Indian-Hindu merger structure (Mdn = 5.0), in which participants include all targets who are *either* Indian *or* Hindu *or* both in the ingroup. This was closely followed by inclusive nationalism (Mdn = 6.5), i.e. participants include all targets who are Indian, and Hindu nationalism (Mdn = 7.0), i.e. participants include only Indian Hindu targets. In addition, Pan-Hinduism (Mdn = 8.0), i.e. participants include all Hindu targets, and egalitarianism (Mdn = 8.0), i.e. participants include all targets, somewhat described participants' responses. Caste did not seem to dominate participants' responses (Mdn = 12.0). Together, these observations suggested that, on the broadest level of analysis, participants relied on nationality and religion, rather than caste, to define their ingroup.

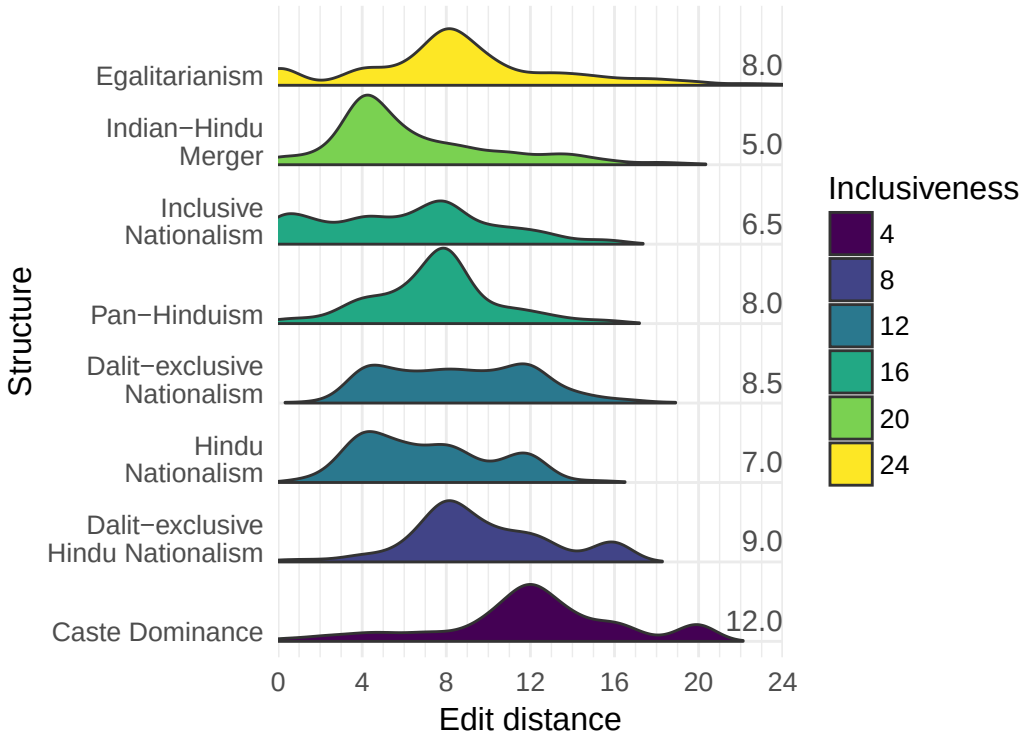

Figure 4.5. Edit distances between each social identity structure and participants' responses. Edit distance is the sum of all categorisations in which a participant's responses differed from the ones implied by the proposed structure. Numbers (on the right) show the median edit distance per structure.

Group differences

As van Dommelen et al. had done, I have so far analysed *which* and *how many* targets participants included in their ingroup as two distinct questions. In the present study, I investigated how participants' (objective) group memberships shaped their (subjective) construals of social identity. I thus compared *how likely* participants were to categorise *which* target as "us" versus "not us"—and how that probability varied across targets' and participants' group memberships.

To that end, I estimated a series of Bayesian multilevel models in *RStan* (Stan Development Team, 2018) with participants' target categorisations (1 = "us", 0 = "not us") as outcome variable. From observed target categorisations, models estimated the log odds α_{ij} that participant j would categorise target i as "us" rather than "not us". Models derived the likelihood of the observed proportions of "us" categorisations from the Bernoulli distribution with a logit link function (see Chapter 2):

$$\begin{aligned} y_{ij} &\sim \text{Bernoulli}(p_{ij}) \\ \text{logit}(p_{ij}) &= \alpha_{ij} \end{aligned}$$

where y is the vector of participants' binary target categorisations, $\text{Bernoulli}(p_{ij})$ is the likelihood distribution, and p_{ij} is the probability of $y_{ij} = 1$ ("us"). Models derived posterior probabilities for all model parameters by drawing samples from the posterior probability distribution:

$$\Pr(\alpha|y) \propto \Pr(y|\alpha)\Pr(\alpha)$$

where $\Pr(\alpha|y)$, the posterior probability of log odds α given the data y, is proportional to $\Pr(y|\alpha)$, the likelihood of the data y given the log odds α, multiplied by $\Pr(\alpha)$, the prior probability of the log odds α (Bayes' theorem). Models assigned weakly informative prior probabilities to all fixed intercepts and coefficients:

$$\beta \sim \text{Student}(2.5, 0, 1)$$

allowing a wide range of plausible effect sizes.[8] As such, priors regularised parameter estimation but had little influence on the resulting posterior probabilities (A. Gelman et al., 2014; McElreath, 2016).

I report point estimates for model parameters based on the median of the posterior probability distribution. For each point estimate, I report an interval spanning the 97% most likely estimates for the corresponding parameter based on the highest posterior density (Plummer et al., 2006). Where useful, I include one-sided hypothesis tests using samples from the posterior probability distribution, where $\Pr(H|M)$ is the posterior probability that hypothesis H (e.g., $\alpha > 0$) is true in model M. As described above, models estimated the log odds α_{ij} that participant j would categorise target i as "us" rather than "not us". As log odds can be difficult to interpret, I transformed log odds, α, into odds, $e^{\alpha} = p/(1-p)$, and probabilities, $e^{\alpha}/(1+e^{\alpha}) = p$, where $p = \Pr(\text{"us"}|M)$ is the estimated probability that participants categorise a target as "us" in model M.

Posterior probabilities tell us which values for a parameter in a model are most likely given said model of the data, but do not tell us whether a parameter is necessary to accurately describe the data. To find out which parameters contribute to a model's

[8] For fixed intercepts, the prior probability distribution, $\alpha \sim \text{Student}(2.5, 0, 1)$, includes log odds $-7.16 < \alpha < 7.16$ among the 99% most plausible estimates. This translates to $.001 < \theta < .999$, where θ is the prior probability of participants categorising a target as "us".

prediction accuracy, I compared a series of increasingly complex models using stratified 10-fold cross-validation (Vehtari et al., 2017). Cross-validation estimates the pointwise out-of-sample prediction accuracy of a fitted Bayesian model, that is, how well a model predicts observations in a new sample. Cross-validation takes both underfitting and overfitting into account, and thus makes a useful criterion for model comparison. When comparing models, I preferred a more complex over a simpler model when the estimated difference in out-of-sample prediction accuracy favoured the more complex model by more than two standard errors of that difference (see Table 4.3).

Model 0 estimated the probability of participants categorising a target as "us" versus "not us" as varying between participants but fixed across targets:

$$\alpha_{ij} = \beta_0 + \beta_j$$

where α_{ij}, the estimated log odds for target i and participant j, equals β_0, the fixed intercept across targets and participants, plus β_j, the varying intercept for participant j.[9] Model 1 estimated distinct odds for all six combinations of targets' nationalities, religions, and castes:

$$\alpha_{ij} = \beta_k + \beta_j$$

where β_k is the varying intercept for all targets i in category k. Each category k ($1 \leq k \leq 6$) spans four targets of the same combination of group memberships. Model 1 made substantially better predictions than Model 0, showing that, as expected, participants used the targets' group memberships to decide who is "us" and "not us".

Models 2 and 3 tested whether participants' own caste memberships shaped how they categorised targets of the same nationality, but of different castes and religions. Model 2 estimated whether, compared to General Merit and Other Backward Class participants, Scheduled Caste / Scheduled Tribe participants were more or less likely to include *Indian, Hindu, GM* ($k = 1$); *Indian, Hindu, OBC* ($k = 2$); *Indian, Hindu, SC/ST* ($k = 3$); and *Indian, Muslim, OBC* ($k = 4$) targets in their ingroup:

$$\alpha_{ij} = \beta_k + \beta_j + \beta_{\text{SCST},k} x_{\text{SCST},j}$$

where $x_{\text{SCST},j}$ codes whether participant j is from a Scheduled Caste / Scheduled Tribe background, and $\beta_{\text{SCST},k}$ estimates the varying odds that Scheduled Caste / Scheduled Tribe participants are more likely than General Merit and Other Backward Class participants to include targets in category k in their ingroup. Model 3 further estimated whether Other Backward Class participants' categorisations differed from General Merit participants':

$$\alpha_{ij} = \beta_k + \beta_j + \beta_{\text{SCST},k} x_{\text{SCST},j} + \beta_{\text{OBC},k} x_{\text{OBC},j}$$

where $x_{\text{OBC},j}$ codes whether participant j is from an Other Backward Class background, and $\beta_{\text{OBC},k}$ estimates the odds that Other Backward Class participants are more likely than General Merit participants to include targets in category k in their ingroup. Model 2, but not Model 3, improved upon the predictions of Model 1. This shows that General Merit and Other Backward Class participants' construals of national identity resembled each other, but differed from Scheduled Caste / Scheduled Tribe participants' construals.

[9] When coding the models, I re-expressed β_j as $\sigma\beta_j$, where σ is the standard deviation of the varying intercept β_j. The prior probabilities of the varying intercept and its standard deviation were: $\beta_j \sim \text{Normal}(0, 1)$ and $\sigma \sim \text{Cauchy}(0, 1)$. This non-centred reparameterization, which I used for all varying intercepts and effects, improves computation (see McElreath, 2016).

Table 4.3. Comparison of models estimating the probability of participants categorising targets as "us" versus "not us". *ELPD* is the expected log predictive density, with higher numbers indicating that a model is expected to make more accurate out-of-sample predictions (Vehtari et al., 2017). Δ*ELPD* is the difference in *ELPD* between the current and previous model, with positive values indicating that the current model is expected to make more accurate out-of-sample predictions. I preferred a more complex model over a simpler model when $\frac{\Delta ELPD}{SE} \geq 2$. *w* are stacking weights based on the combined predictive density (Yao et al., 2018).

#	Description	*ELPD*	*SE*	Δ*ELPD*	*SE*	$\frac{\Delta ELPD}{SE}$	*w*
0	*Varying intercept*	-4703.5	23.7	-	-	-	.00
1 vs 0	*Group differences*	-3853.5	41.6	850.1	37.4	22.7	.00
2 vs 1	*Group differences*	-3791.3	42.0	62.1	11.1	5.6	.00
3 vs 2	*Group differences*	-3792.1	42.2	-0.8	3.6	-0.2	.07
4 vs 2	*Intergroup contact*	-3743.4	42.8	47.9	8.8	5.4	.00
5 vs 4	*Intergroup contact*	-3736.7	42.6	6.7	1.8	3.8	.92
6 vs 5	*Intergroup contact*	-3747.8	42.7	-11.1	3.2	-3.5	.00
7 vs 2	*Social dominance orientation*	-3791.0	42.1	0.4	3.3	0.1	.00

Figure 4.6 shows the estimated probabilities of General Merit, Other Backward Class, and Scheduled Caste / Scheduled Tribe participants categorising targets as "us" according to Model 2. As expected, few participants considered Bangladeshi Muslims as part of their ingroup, Pr("us"|M2) = .17, [.13, .20]. Roughly half of the participants included Sri Lankan and Nepali Hindus in their ingroup, Pr("us"|M2) = .46, [.40, .52], indicating that participants were more likely to consider foreign targets as "us" when they were Hindu rather than Muslim, ΔPr("us"|M2) = .29, [.34, .25]. Still, the (combined) probabilities of General Merit / Other Backward Class and Scheduled Caste / Scheduled Tribe participants categorising Indian, Hindu targets as "us" were, respectively, ΔPr("us"|M2) = .44, [.39, .49] and ΔPr("us"|M2) = .42, [.37, .48] higher than for *foreign, Hindu* targets. In other words, participants were 1.96, [1.75, 2.20] and 1.92, [1.71, 2.16] times more likely to categorise *Indian, Hindu* targets as "us" compared to *foreign, Hindu* targets. In line with earlier observations, participants thus tended to define their ingroup in terms of nationality, though about half of the participants seemed to endorse some form of Pan-Hinduism or altogether more egalitarian social identities.

As Model 2 showed, participants' own caste memberships shaped how they categorised Indians of different castes and religions. General Merit / Other Backward Class participants were equally likely to categorise *Hindu, GM* and *Hindu, OBC* targets as "us" versus "not us", ΔPr("us"|M2) = .01, [−.01, .03]. Almost all General Merit / Other Backward Class participants included targets from both caste groups in their ingroup, Pr("us"|M2) = .94, [.92, .96] and Pr("us"|M2) = .93, [.91, .95]. They were, however, less likely to categorise *Hindu, SC/ST* targets as "us": Pr("us"|M2) = .83, [.79, .87], ΔPr("us"|M2) = −.11, [−.07, −.14] and ΔPr("us"|M2) = −.10, [−.06, −.13]. Of all Indian targets, General Merit / Other Backward Class participants were least likely to include *Muslim, OBC* targets in their ingroup, Pr("us"|M2) = .74, [.69, .79], even less likely than *Hindu, SC/ST* targets, ΔPr("us"|M2) = −.09, [−.14, −.05]. Still, they were 1.61, [1.46, 1.80] and 4.45, [3.64, 5.37] times more likely to include Indian Muslims compared to foreign Hindus and foreign Muslims, respectively. While about 3 of 4 General Merit / Other Backward Class par-

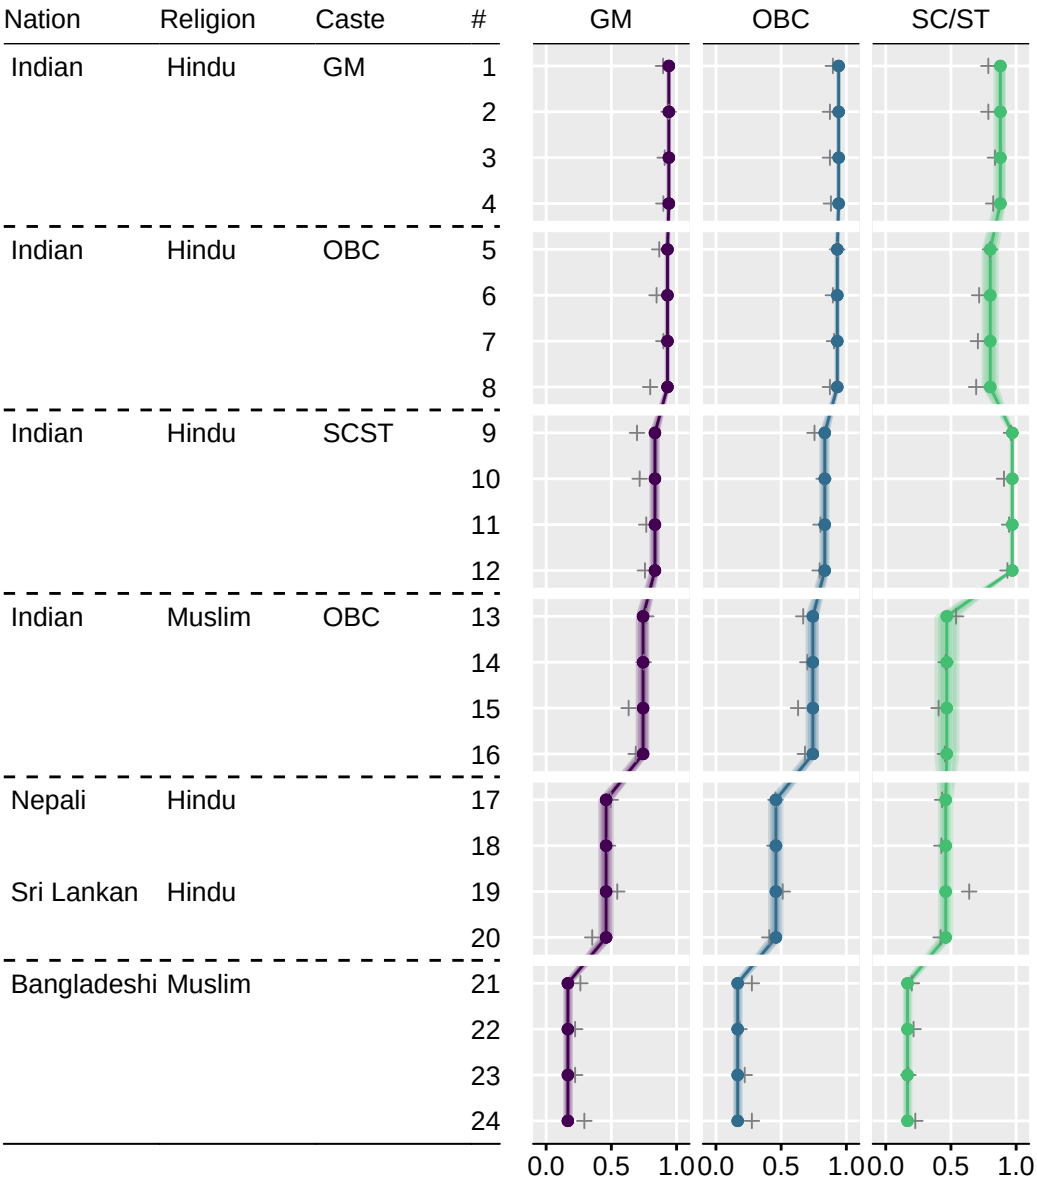

Figure 4.6. Estimated probability of participants categorising a target as "us" vs "not us" by targets' index (#), nationality, religion, and caste (vertical), and participants' caste membership (horizontal). Points (•) indicate the most likely *estimate* for a given target's probability of inclusion (in Model 2, see Table 4.3), while the shaded ribbons encompass the 67% (darkest shade), 89%, and 97% (lightest shade) most likely estimates of that probability. Pluses (+) indicate the *observed* proportion of participants who included a target in their ingroup.

ticipants thus expressed some form of inclusive nationalism or egalitarianism, 1 of 4 excluded targets because of their caste or religion.

Dalit/Adivasi (Scheduled Caste / Scheduled Tribe) participants' responses differed from those of participants from relatively advantaged castes. As expected, Scheduled Caste / Scheduled Tribe participants were more likely to include *Hindu, SC/ST* targets in their ingroup compared to General Merit / Other Backward Class participants, Pr("us"|M2) = .97, [.95, .99], ΔPr("us"|M2) = .14, [.10, .18]. Sticking to their caste group, Scheduled Caste / Scheduled Tribe participants were somewhat less likely to include targets of the other two caste groups relative to General Merit / Other Backward Class participants, ΔPr("us"|M2) = −.06, [−.12, −.02] and ΔPr("us"|M2) = −.13, [−.20, −.07]. Surprisingly, Scheduled Caste / Scheduled Tribe participants were somewhat less likely to include *Hindu, OBC* than *Hindu, GM targets*, ΔPr("us"|M2) = −.08, [−.15, −.02]. Similarly, they were substantially less likely than General Merit / Other Backward Class participants to categorise *Indian, Muslim* targets as "us", Pr("us"|M2) = .47, [.37, .57], ΔPr("us"|M2) = −.27, [−.37, −.18]. Indeed, they were equally (un)likely to consider *Indian, Muslim* and *foreign, Hindu* targets as part of their ingroup, ΔPr("us"|M2) = .01, [−.08, .10]. Participants from historically disadvantaged caste backgrounds were thus more likely to exclude members of religions and (other) caste minorities compared to participants from relatively advantaged backgrounds.

Individual differences

Models 0 to 2 showed that participants' objective group memberships shaped whom they considered "us" and "not us". Though probabilities of "us" and "not us" categorisations differed markedly across targets' and participants' group memberships, none of these associations were without exception. In this section, I examine to what extent individual differences in past experiences (intergroup contact) and ideological preferences (social dominance orientation) explain why some participants excluded targets of national, religious and caste outgroups, and why others did not.

Models 4 to 6 tested whether intergroup contact was associated with how participants categorised Indian targets of other caste or religious backgrounds (see Table 4.3). Model 4 included contact quantity, positive contact, negative contact, and outgroup friendship as predictors of participants' response tendencies:

$$\alpha_{ij} = \beta_k + \beta_j + \beta_{\text{SCST},k} x_{\text{SCST},j} + x_{\text{OG},jk}(\beta_{\text{CQ}} x_{\text{CQ},jk} + \beta_{\text{PC}} x_{\text{PC},jk} + \beta_{\text{NC}} x_{\text{NC},jk} + \beta_{\text{OF}} x_{\text{OF},jk})$$

where $x_{\text{OG},jk}$ codes whether category k is, in terms of caste, religion, or nationality, an outgroup [1] or ingroup [0] member for participant j, $x_{\text{CQ},jk}$ to $x_{\text{OF},jk}$ are the amounts of different forms of contact participant j reported with members of category k ($1 \leq k \leq 4$), and β_{CQ} to β_{OF} are the corresponding coefficients.

Model 4 made more accurate predictions than Model 2, showing that past experiences with members of an (objective) outgroup related to whether participants included targets from that group in their (subjective) ingroup. Looking closer at the model's predictions, negative contact (e^{β} = 0.81, [0.72, 0.91]) and outgroup friendship (e^{β} = 1.50, [1.29, 1.74]) were associated with participants' responses, but neither positive contact (e^{β} = 1.00, [0.87, 1.15]) nor contact quantity (e^{β} = 0.99, [0.86, 1.13]) was. Model 5 included only negative contact and outgroup friendship as predictors of participants' categorisations. Model 5 improved upon the predictions of Model 4, showing that contact quantity and positive contact did not contribute to out-of-sample prediction beyond the

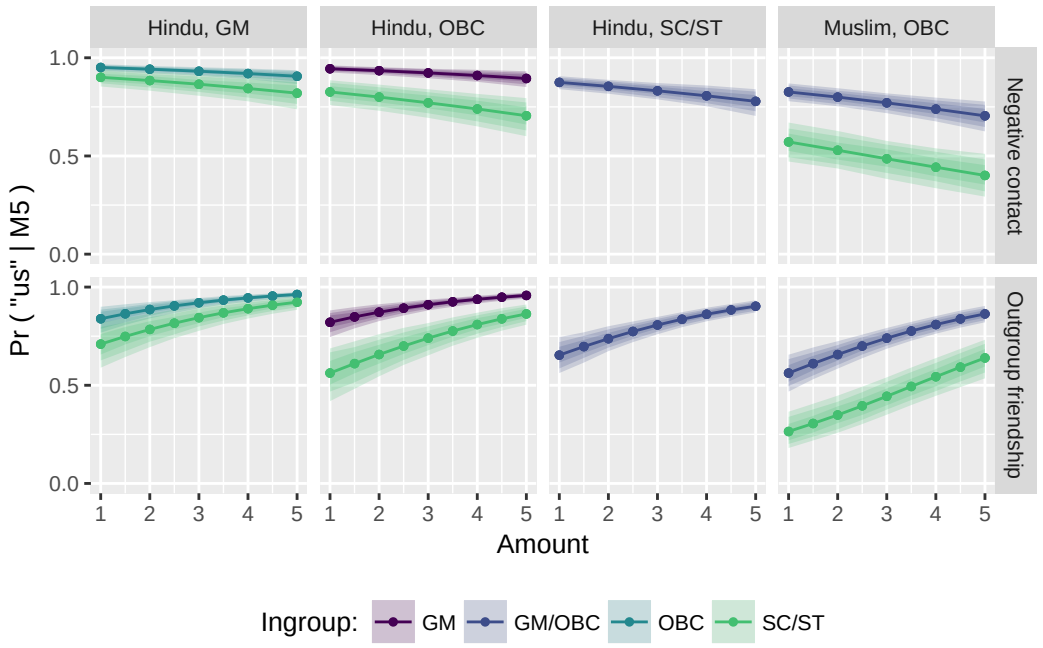

Figure 4.7. Estimated probability of participants categorising a target as "us" vs "not us" as a function of the reported amount of *negative contact* and *outgroup friendship* (in Model 5, Table 4.3).

effects of negative contact and outgroup friendship. Model 6 estimated the relationships between intergroup contact and target categorisations as varying across the four combinations of target caste and religion. As Model 6 made less accurate predictions than Model 5, the association between contact and categorisations did not seem to vary substantially across target categories.

Figure 4.7 shows the estimated probabilities of General Merit, Other Backward Class, and Scheduled Caste / Scheduled Tribe participants categorising targets as "us" versus "not us" as a function of negative contact experiences and outgroup friendship (in Model 5). Across targets and participants, the odds of categorising a target of a religious and/or caste outgroup as "us" were $e^{\beta} = 1.49$, $[1.30, 1.68]$ times higher for each additional standard deviation of outgroup friendship participants reported. Conversely, the odds of categorising an outgroup target as "us" were $e^{\beta} = 0.81$, $[0.72, 0.91]$ times lower for each additional standard deviation of negative contact. This means, for example, that General Merit / Other Backward Class participants who reported "never" having any negative contact with Muslims were more likely to categorise Indian Muslims as "us" compared to participants who reported "very often" having negative contact, $\Delta\Pr(\text{"us"}|M2) = .12$, $[.05, .19]$. General Merit / Other Backward Class participants who reported no friendships with Muslims were a lot less likely to include Indian Muslims in their ingroup than participants who had 2–5 Muslim friends with whom they "sometimes" spent time, $\Delta\Pr(\text{"us"}|M2) = .18$, $[.12, .24]$. Together, these findings suggested that contact experiences were associated with whom participants considered as "us" and "not us".

Models 7 tested whether, beyond intergroup contact, social dominance orientation was associated with how participants categorised targets of other caste or religious backgrounds (see Table 4.3). Specifically, Model 7 tested whether the two subdimensions of social dominance orientation were associated with participants excluding targets of

comparably lower status (as defined under Study 4):

$$\alpha_{ij} = \beta_k + \beta_j + \beta_{\text{SCST},k} x_{\text{SCST},j} + x_{\text{GS},ij}(\beta_{\text{SDO-D}} x_{\text{SDO-D},j} + \beta_{\text{SDO-E}} x_{\text{SDO-D},j})$$

where $x_{\text{GS},ij}$ codes whether target i is a member of a lower-status outgroup [1] or not [0] relative to participant j, $x_{\text{SDO-D},j}$ and $x_{\text{SDO-E},j}$ are the SDO-D and SDO-E scores of participant j, and $\beta_{\text{SDO-D}}$ and $\beta_{\text{SDO-E}}$ are the corresponding coefficients. Looking at the posterior probabilities, the model found little evidence for the relationships between participants' categorisations and their SDO-D scores ($e^{\beta} = 0.88$, $[0.72, 1.05]$) or SDO-E scores ($e^{\beta} = 0.91$, $[0.71, 1.11]$). Accordingly, Model 7 did not make more accurate out-of-sample predictions than Model 2, showing that ideological preferences were not associated with whom participants considered as "us" and "not us".

Together, these findings suggested that both group and individual differences explain whom participants included in their ingroup—and whom they excluded. As expected, personal experiences, positive and negative, with outgroup members explained why some participants included targets of (objective) caste or religions outgroups in their (subjective) ingroup, when others did not. Other than expected, preferences for social dominance and intergroup hierarchy did not seem to motivate participants to exclude lower-status groups.

Consequences

After studying the antecedents of participants' categorisations, I tested what, if any, consequences participants' ingroup construals had for their attitudes and beliefs. Before examining perceptions of *intergroup threat*, I studied how participants' categorisations related to their *desired social distance* (Q2) and *outgroup attitudes* (Q3) ratings for each target in the crossed-categorisation task.

Models 0 to 6 estimated the mean ratings (on the two outcomes) of target i by participant j, and derived the likelihood of the observed ratings from a multivariate normal distribution:

$$\begin{bmatrix} y_{2,ij} \\ y_{3,ij} \end{bmatrix} \sim \text{MVNormal}\left(\begin{bmatrix} \mu_{2,ij} \\ \mu_{3,ij} \end{bmatrix}, \mathbf{S} \right)$$

where y_2 is the vector of participants' responses to question 2, y_3 is the vector of participants' responses to question 3, $\mu_{2,ij}$ and $\mu_{3,ij}$ are the mean ratings of target i by participant j, and $\mathbf{S}$ is the residual variance-covariance matrix for the two outcomes. This means that models estimated distinct means and (residual) variances for each outcome, but also estimated the correlation between the residuals for the two outcomes. Again, models assigned weakly informative priors to all intercepts, coefficients, and (residual) variances:

$$\beta \sim \text{Normal}(0, 1)$$
$$\sigma \sim \text{Cauchy}(0, 1)$$

allowing a wide range of plausible effect sizes. To increase computational efficiency, I used the Cholesky parameterisation of the multivariate normal distribution, with an LKJ(1) prior on the Cholesky-transformed correlation matrix.

Models 0 to 3 tested whether participants' objective group memberships informed how they rated targets with different combinations of group memberships. Model 0 estimated target ratings (on either outcome) as varying between participants but fixed across targets:

$$\mu_{ij} = \beta_0 + \beta_j$$

where μ_{ij}, the estimated mean for target i and participant j, β_0, the fixed intercept across targets and participants, plus β_j, the varying intercept for participant j. Model 1 estimated distinct means for all six combinations of targets' nationalities, religions, and castes:

$$\mu_{ij} = \beta_k + \beta_j$$

where β_k is the varying intercept for all targets i in category k. Model 2 further tested whether, compared to General Merit and Other Backward Class participants, Scheduled Caste / Scheduled Tribe participants responded more or less favourably to targets in the four Indian categories:

$$\mu_{ij} = \beta_k + \beta_j + \beta_{\text{SCST},k} x_{\text{SCST},j}$$

where $x_{\text{SCST},j}$ codes whether participant j is from a Scheduled Caste / Scheduled Tribe backround, and $\beta_{\text{SCST},k}$ estimates the mean difference between Scheduled Caste / Scheduled Tribe and General Merit / Other Backward Class participants' ratings of targets in category k, for $1 \leq k \leq 4$. Model 3 further estimated whether Other Backward Class participants' ratings differed from General Merit participants':

$$\mu_{ij} = \beta_k + \beta_j + \beta_{\text{SCST},k} x_{\text{SCST},j} + \beta_{\text{OBC},k} x_{\text{OBC},j}$$

where $x_{\text{OBC},j}$ codes whether participant j is from an Other Backward Class background, and $\beta_{\text{OBC},k}$ estimates the mean difference between Other Backward Class and General Merit participants' ratings of targets in category k.

Models 4 to 6 tested whether participants who categorised a target as "us" rated that target more favourably than participants who categorised the same target as "not us". Model 4 estimated the mean difference in ratings as constant across targets and participants:

$$\mu_{ij} = \beta_k + \beta_j + \beta_{\text{SCST},k} x_{\text{SCST},j} + \beta_{\text{OBC},k} x_{\text{OBC},j} + \beta_{\text{Q1}} x_{\text{Q1},ji}$$

where $x_{\text{Q1},ji}$ codes whether participant j categorised target i as "us" [1] or "not us" [0], and β_{Q1} estimates the difference in ratings between targets categorised as "us" and targets categorised as "not us". Model 5 tested whether the difference between ratings of targets categorised as "us" and those categorised as "not us" varied across the six target categories:

$$\mu_{ij} = \beta_k + \beta_j + \beta_{\text{SCST},k} x_{\text{SCST},j} + \beta_{\text{OBC},k} x_{\text{OBC},j} + \beta_{\text{Q1},k} x_{\text{Q1},ji}$$

where $\beta_{\text{Q1},k}$ estimates the difference in ratings between targets categorised as "us" and targets categorised as "not us" for category k ($1 \leq k \leq 6$). Model 6 further tested whether the difference in ratings (for Indian targets) differed depending on participants' own caste memberships:

$$\mu_{ij} = \ldots + \beta_{\text{Q1},k} x_{\text{Q1},ji} + \beta_{\text{Q1,SCST},k} x_{\text{SCST},j} x_{\text{Q1},ji} + \beta_{\text{Q1,OBC},k} x_{\text{OBC},j} x_{\text{Q1},ji}$$

where $\beta_{\text{Q1,SCST},k}$ and $\beta_{\text{Q1,OBC},k}$ estimate how much lesser and greater the difference in ratings between targets categorised as "us" and targets categorised as "not us" in category k ($1 \leq k \leq 4$) was for Scheduled Caste / Scheduled Tribe and Other Backward Class participants, respectively.

Table 4.4 compares the estimated out-of-sample prediction accuracy of the multivariate models for *desired social distance* (Q2) and *outgroup attitudes* (Q3). Models 1 to 3 improved upon the predictions of preceding model(s), indicating *i.* that participants'

Table 4.4. Comparison of multivariate models estimating participants' *desired social distance* (Q2) and *outgroup attitudes* (Q3) for each target as a function of group differences and target categorisations. R^2 is a Bayesian analogue to R^2 in maximum likelihood estimation (A. Gelman, Goodrich et al., 2017). As in Table 4.3, I preferred a more complex model over a simpler model when $\frac{\Delta ELPD}{SE} \geq 2$.

#	Description	R^2_{Q2}	R^2_{Q3}	*ELPD*	*SE*	$\Delta ELPD$	*SE*	$\frac{\Delta ELPD}{SE}$	*w*
0	*Group differences*	.23	.26	-16964	71.0	-	-	-	.12
1 vs 0	*Group differences*	.38	.42	-16338	79.6	626.8	40.2	15.6	.02
2 vs 1	*Group differences*	.38	.43	-16317	79.6	20.1	9.8	2.1	.00
3 vs 2	*Group differences*	.39	.43	-16293	80.5	24.2	9.7	2.5	.00
4 vs 3	*Categorisation*	.42	.47	-16028	83.5	264.8	24.5	10.8	.19
5 vs 4	*Categorisation*	.42	.47	-16007	83.6	21.1	8.7	2.4	.67
6 vs 5	*Categorisation*	.42	.47	-16025	84.3	-17.4	5.9	-3.0	.00

ratings of a target differed depending on the group memberships of that target, *ii.* that Scheduled Caste / Scheduled Tribe participants' ratings of Indian targets differed from General Merit / Other Backward Class participants' responses, and *iii.* that, further, Other Backward Class participants' ratings differed from General Merit participants'. Models 4 and 5, but not Model 6, made better out-of-sample predictions than less complex models, showing *iv.* that how favourably participants felt toward a target differed depending on whether they had categorised that target as "us" or "not us", and *v.* that the size of this difference depended on the group memberships of that target—but *vi.* not on the group membership of the participant.

Figure 4.8 shows the estimated mean *desired social distance* ratings for all categories as a function of target categorisations and participants' caste memberships (in Model 5, Table 4.4). For all categories, participants felt more comfortable sharing a room with targets that they categorised as "us" than with targets they categorised as "not us". Across participants, this difference was smallest for *Indian, Hindu, GM* targets ($\beta = 0.46$, $[0.08, 0.81]$; $d = 0.20$, $[0.03, 0.35]$) and *Indian, Hindu, OBC* targets ($\beta = 0.63$, $[0.30, 0.94]$; $d = 0.27$, $[0.13, 0.41]$), followed by *Indian, Hindu, SCST* targets ($\beta = 0.95$, $[0.65, 1.27]$; $d = 0.41$, $[0.28, 0.55]$) and *Indian, Muslim, OBC* targets ($\beta = 1.07$, $[0.82, 1.32]$; $d = 0.46$, $[0.35, 0.57]$). The difference between targets categorised as "us" and targets categorised as "not us" was greatest for *foreign, Muslim* targets ($\beta = 1.45$, $[1.22, 1.70]$; $d = 0.63$, $[0.53, 0.73]$) and *foreign, Hindu* targets ($\beta = 1.24$, $[0.97, 1.52]$; $d = 0.53$, $[0.42, 0.66]$). Effect sizes thus seemed to be greater for categories with lower average ratings (or lower proportions of "us" categorisations, see Figure 4.6).

Figure 4.9 shows the estimated mean *outgroup attitudes* ratings for all categories as a function of target categorisations and participants' caste memberships (in Model 5, Table 4.4). The difference between targets categorised as "us" and targets categorised as "not us" was smallest for *Indian, Hindu, OBC* targets ($\beta = 6.8$, $[2.3, 11.1]$, $d = 0.21$, $[0.07, 0.34]$), followed by *Indian, Hindu, GM* targets ($\beta = 12.4$, $[7.5, 17.3]$, $d = 0.38$, $[0.23, 0.53]$), *Indian, Hindu, SCST* targets ($\beta = 12.3$, $[8.2, 16.4]$, $d = 0.37$, $[0.25, 0.50]$) and *Indian, Muslim, OBC* targets ($\beta = 13.6$, $[10.4, 17.3]$, $d = 0.41$, $[0.32, 0.53]$). Again, this difference was greatest for *foreign, Hindu* targets ($\beta = 21.9$, $[18.9, 25.3]$, $d = 0.67$, $[0.57, 0.77]$) and *foreign, Muslim* targets ($\beta = 16.8$, $[12.8, 20.3]$, $d = 0.51$, $[0.39, 0.62]$). Effects were

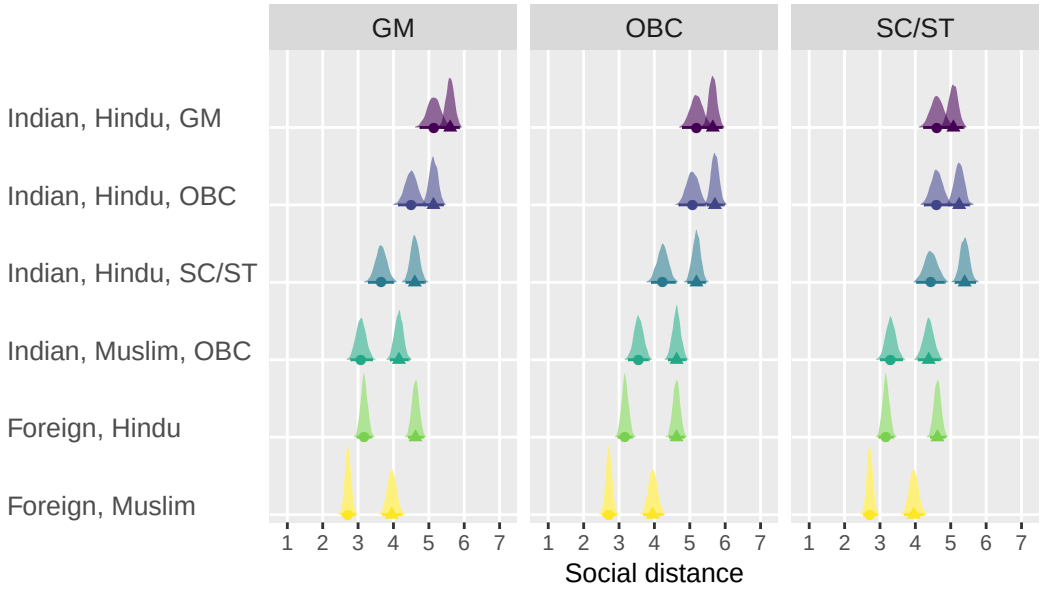

Figure 4.8. Posterior probabilities of *desired social distance* (Q2) ratings as a function of target categorisation (in Model 5, Table 4.4). Points are mean ratings for targets categorised as "us". Triangles are mean ratings for targets categorised as "not us".

thus of similar direction, size, and order as desired social distance ratings. Participants' ratings of desired social distance and outgroup attitudes were highly correlated (r = .58, [.56, .60]), even after taking into account all predictors included in Model 5.

Finally, I tested whether participants' perceptions of *realistic* and *symbolic threat* differed depending on the inclusiveness of their ingroup construals. Specifically, I examined whether participants reported feeling less threatened by Muslims and Dalits if they categorised more targets from these outgroups as "us" versus "not us". I excluded responses from Scheduled Caste / Scheduled Tribe participants to Scheduled Caste / Scheduled Tribe targets as I was interested in perceptions of *inter*group threat. Models 0 to 7 estimated participants' mean responses, between 1 = *strongly disagree* and 5 = *strongly agree*, to each of the 5(items) × 2(outgroups) = 10 items. Models derived the likelihood of the observed responses from the normal likelihood distribution:

$$y_{ij} \sim \text{Normal}(\mu_{ij}, \sigma)$$

where y is the vector of participants' responses to each item, and μ_{ij} is the estimated mean response to item i by participant j, σ is the estimated residual variance (expressed as standard deviation). I used the same priors as for the multivariate models.

Models 0 to 3 tested whether participants' responses differed across the two subscales (realistic vs symbolic threat) and across the two outgroups (Muslims vs Dalits). Model 0 estimated item responses as varying between participants but fixed across all items:

$$\mu_{ij} = \beta_0 + \beta_j$$

where μ_{ij}, the estimated mean response to item i by participant j, equals β_0, the fixed intercept across items and participants, plus β_j, the varying intercept for participant j. Instead of assigning the same mean to all items, Model 1 included distinct intercepts for the two subscales:

$$\mu_{ij} = \beta_k + \beta_j$$

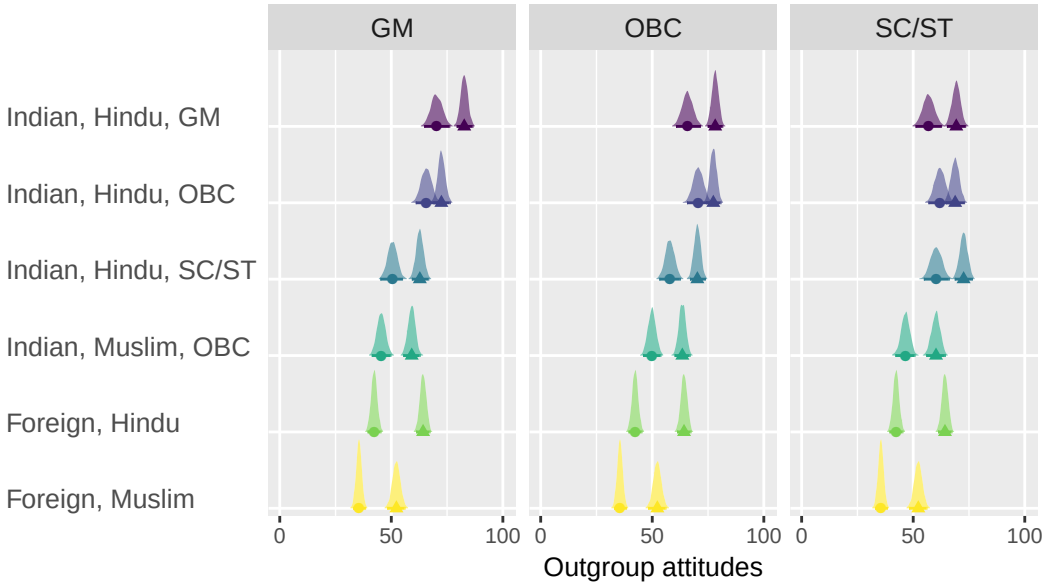

Figure 4.9. Posterior probabilities of *outgroup attitudes* (Q3) ratings as a function of target categorisation (in Model 5, Table 4.4). Points are mean ratings for targets categorised as "us". Triangles are mean ratings for targets categorised as "not us".

Table 4.5. Comparison of models estimating participants' intergroup threat. As in Table 4.3, I preferred a more complex model over a simpler model when $\frac{\Delta ELPD}{SE} \geq 2$.

#	Description	R^2	*ELPD*	*SE*	$\Delta ELPD$	*SE*	$\frac{\Delta ELPD}{SE}$	*w*
0	*Varying intercept*	.31	-3368.1	32.9	-	-	-	.11
1 vs 0	*Subscales*	.31	-3368.9	32.9	-0.8	0.6	-1.3	.00
2 vs 0	*Outgroups*	.32	-3363.7	33.0	4.4	3.3	1.3	.00
3 vs 0	*Subscales* × *Outgroups*	.33	-3335.1	34.0	33.0	8.4	3.9	.19
4 vs 3	*Categorisations*	.33	-3335.5	34.0	-0.3	0.6	-0.5	.00
5 vs 3	*Categorisations*	.34	-3330.5	34.0	4.7	3.6	1.3	.71
6 vs 3	*Categorisations*	.33	-3336.7	34.0	-1.5	0.7	-2.1	.00
7 vs 3	*Categorisations*	.34	-3332.8	34.0	2.3	3.6	0.6	.00

where β_k is the fixed intercept for all items i in subscale k, with $k = 1$ for realistic threat and $k = 2$ for symbolic threat. Model 2 instead estimated distinct intercepts for items concerning the two outgroups:

$$\mu_{ij} = \beta_l + \beta_j$$

where β_l is the fixed intercept for all items i for outgroup l, with $l = 1$ for Dalits and $l = 2$ for Muslims as the relevant outgroup. Model 3 tested the interaction of the two factors, and estimated distinct intercepts for all four combinations of subscales and outgroups:

$$\mu_{ij} = \beta_m + \beta_j$$

where β_m is the fixed intercept for realistic threat from Dalits (when $m = 1$), for symbolic threat from Dalits (when $m = 2$), for realistic threat from Muslims (when $m = 3$), and for symbolic threat from Muslims (when $m = 4$).

Model 3, but not Models 1 and 2, improved upon the predictions of Model 0 (Table 4.5), showing that participants' perceptions of symbolic and realistic threat depended on whether the relevant outgroup was Muslims or Dalits. Specifically, participants reported more realistic (β_1 = 3.62, [3.49, 3.75]) than symbolic (β_2 = 3.21, [3.09, 3.33]) threat from (same-religion) Dalits, $\Pr(\beta_1 > \beta_2|M3) > .99$, but less realistic ($\beta_3$ = 3.23, [3.06, 3.38]) than symbolic (β_4 = 3.47, [3.34, 3.61]) threat from (different-religion) Muslims, $\Pr(\beta_3 < \beta_4|M3) > .99$.

Models 4 to 7 tested whether the proportion of *Indian, Muslim, OBC* and *Indian, Hindu, SC/ST* targets participants had categorised as "us" was associated with less perceived threat from, respectively, Muslims and Dalits. Model 4 estimated this relationship as constant across outgroups and subscales:

$$\mu_{ij} = \beta_m + \beta_j + \beta_{Q_1} x_{Q_1,jl}$$

where $x_{Q_1,jl}$ is the proportion [0–1] of targets belonging to outgroup l that participant j categorised as "us", and β_{Q_1} estimates the difference in threat perceptions between participants who categorised all four targets of outgroup l as "us" and participants who categorised all targets of outgroup l as "not us". Models 5 to 7 instead estimated distinct coefficients $\beta_{Q_1,k}$ for each of the two subscales (M5), $\beta_{Q_1,l}$ for each of the two outgroups (M6), or $\beta_{Q_1,m}$ for each of the four combinations of subscales and outgroups (M7). Other than expected, Models 4 to 7 did not make better out-of-sample predictions than Model 3, indicating that whom participants considered "us" and "not us" was not associated with their perceptions of intergroup threat.

Overall, these findings showed that when participants included a person in their ingroup, they had, on average, more favourable attitudes to and desired less social distance from that specific person. Participants' ingroup construals, however, were not related to more general perceptions of intergroup threat from religious and caste outgroup members.

Discussion

In this chapter, I have examined how South Indian students construct their social identities from multiple cross-cutting categories, and how these identities relate to intergroup contact and intergroup bias. Study 4 found, as hypothesised, that cross-group friendship was associated with more inclusive identities, while more inclusive identities, in turn, were associated with less intergroup bias. These findings show that friendship can bridge social divides by fostering more inclusive social identities. Below, I summarise findings from the present research, before discussing its strengths, limitations, and implications.

Summary of findings

Study 4 adapted van Dommelen et al.'s (2015) crossed-categorisation task to investigate how groups differed in whom their members considered "us" and "not us". As expected, participants from advantaged backgrounds (GM: General Merit) tended to exclude targets from the historically most disadvantaged backgrounds (SC/ST: Scheduled Caste, Scheduled Tribe). Conversely, participants from the most disadvantaged backgrounds tended to exclude targets from more advantaged backgrounds. Participants who occupied an intermediate position in the caste hierarchy (OBC: Other Backward Classes) tended to include and exclude the same targets as General Merit participants. Participants of

all caste groups included Indian Muslims less often than Indian Hindus. Nonetheless, General Merit / Other Backward Class participants, but not Scheduled Caste / Scheduled Tribe participants, included Indian Muslims more often than foreign Muslims and Hindus. This suggests that some participants espoused a form of Hindu nationalism (that excluded Indian Muslims from their ingroup), while others espoused a more inclusive nationalism (that included Indian Muslims in their ingroup). Overall, Study 4 showed that group differences shaped whom participants considered "us" and "not us".

As discussed, none of these observations was without exception—some participants had more inclusive identities than others. Like van Dommelen et al. (2015), Study 4 tested whether past experiences explained why some participants included targets of caste or religious outgroups, when others did not. As hypothesised, participants who reported more friendships with a caste or religious outgroup were more likely to categorise members of that outgroup as "us" rather than "not us". Similarly, van Dommelen et al. (2015) found a significant correlation between intergroup contact and more inclusive identities in Study 2, though not Study 1. In line with recent developments (Barlow et al., 2012; Pettigrew & Hewstone, 2017), Study 4 also tested how negative experiences with caste and religious outgroups affected participants' social identities. As expected, negative contact was associated with less inclusive identities. Beyond intergroup contact, Study 4 tested whether ideological preferences for social dominance motivated participants to exclude lower-status groups. Other than expected, social dominance orientation was not associated with less inclusive identities. These findings show that past experiences, but not ideological preferences, shaped whom participants considered "us" and "not us".

Finally, Study 4 examined what consequences participants' ingroup construals had on their attitudes and beliefs. As hypothesised, participants held, on average, more favourable attitudes to and desired less social distance from targets they had included in their ingroup. At university, Dalit students continue to face interpersonal discrimination, often based on practices and prejudices surrounding untouchability (Pathania & Tierney, 2018). Reducing upper-caste students' desired social distance could help to lessen that discrimination. As far as it fosters more inclusive identities, cross-group friendship thus offers a way to overcome everyday caste discrimination. Contrary to expectation, more inclusive identities were not associated with (lower) perceptions of realistic and symbolic threat. An explanation for this finding may be that perceptions of threat stem from economic anxieties (e.g., fearing affirmative action as a threat to one's career) and ideological beliefs (e.g., Hindutva), rather than from social identity processes. Overall, Study 4 provided evidence that cross-group friendship fosters more inclusive social identities, and that more inclusive social identities, in turn, reduce intergroup bias—though not intergroup threat. Below, I discuss strengths and limitations of the present research.

Strengths and limitations

Alongside van Dommelen et al.'s (2015) research, these findings show the value of the triple crossed-categorisation task as an intuitive and informative method for studying social identification across multiple categories. This method is *intuitive* because respondents do not need to understand "identity" or other abstract constructs in order to complete the task. This makes it a useful tool for cross-cultural or developmental research on social identification. Further, this method is intuitive because it does not limit respondents to certain social identity structures, preconceived by the researcher. Instead, the task

allows respondents to freely construct their ingroup (from the available targets). The crossed-categorisation task is *informative* because it is a rich source of data, as the present research has shown. A single study yielded insights into the effects of target groups, participant groups, and individual differences on social identification, outgroup attitudes, and social distance. Together, these strengths commend the triple crossed-categorisation task as a valuable method for studying social identification.

Beyond van Dommelen et al.'s (2015) initial studies, the present research extended the cross-categorisation task to answer new questions about social identification. First, I recruited respondents from multiple groups, allowing me to compare the relative influence of individual and group differences. Second, I estimated response probabilities as varying across targets and participants using multilevel models. This allowed more fine-grained analyses of ingroup construals than van Dommelen et al.'s qualitative and quantitative summaries. Third, I tested how contact with a *specific* outgroup related to participants' categorisations of and attitudes toward *that* outgroup. Greater specificity may explain why the present research observed more consistent effects of intergroup contact than earlier research.[10] Finally, I adapted the crossed-categorisation task to simultaneously test groups of participants, rather than test participants individually. This change makes testing greater sample sizes more practical. Overall, these changes open the triple crossed-categorisation task to a broader range of research questions on individual and group differences in social identification across multiple categories.

Despite improvements, the present research is qualified by some methodological limitations. Study 4 tested group differences in a quasi-experimental design. That is, participants were recruited from different groups, instead of being randomly assigned to lower-status or higher-status groups. Though ecologically valid, this design does not allow firm conclusions about *why* categorisations differed across groups. Study 4 tested antecedents (e.g., intergroup contact) and consequences (e.g., outgroup attitudes) in a cross-sectional, correlational design. This design does not establish the direction of the observed relationships, for example, whether cross-group friendship leads to more inclusive identities or whether more inclusive identities encourage cross-group friendship. Study 4 presented all participants with the same combination of target groups. This design cannot determine to what extent observed ingroup construals generalise beyond the specific combination of stimuli used. Similarly, Study 4 did not control for factors that correlate with the categories under study, but were not included in the stimuli. Class prejudice, for example, could explain why some participants excluded targets from disadvantaged caste groups. Study 4 measured categorisation, outgroup attitudes, and desired social distance for the same targets. This design cannot rule out that these variables measure the same construct, rather than represent a genuine association across constructs. Intergroup threat, a more distal measure, did not correlate with participants' categorisations.

Future research should address these limitations by varying the set of target categories across participants, by including a wider set of target categories (e.g., class), by assessing intergroup bias with a wider set of distal and proximal measures, and by testing the hypothesised relationships over time. Below, I discuss implications of the present findings for intergroup relations research.

[10] Another explanation may be that the present research included more participants than van Dommelen et al.'s studies: 302 instead of 95 (Study 1) and 123 (Study 2).

Implications

Findings from the present research illuminate how advantaged groups relate to disadvantaged groups, how disadvantaged groups relate to other disadvantaged groups, and how disadvantaged groups relate to advantaged groups.

Among advantaged groups, the present research documented patterns of inclusion and exclusion that map onto the ideologies described in the introduction. For caste, Study 4 found that some upper-caste students discriminated against Dalits when defining their ingroup. Students excluded Dalits, who had historically been affected by untouchability, but not other disadvantaged caste groups. For religion, Study 4 found that some students discriminated against Muslims when defining their ingroup. As discussed, this suggests that some students endorse a form of exclusionary Hindu nationalism (Hindutva). Study 4 did not measure support for this ideology directly. Future studies should assess how ingroup construals relate to both *specific* ideologies (Hindutva, Casteism) and *broader* ideological preferences (see Cotterill, Sidanius, Bhardwajb & Kumarc, 2014). Aside from social dominance, right-wing authoritarianism (Altemeyer, 1981)—especially its traditionalism subdimension (Duckitt & Bizumic, 2013)—could explain why people seek to exclude caste and religious minorities from their ingroup.

Beyond ideology, the hierarchy-attenuating institution (Pratto et al., 2006) of reservation likely shaped advantaged-group members' ingroup construals. On the one hand, reservations (e.g., in higher education) give advantaged-group members opportunities for intergroup contact. As far as these opportunities result in cross-group friendships, reservations could help overcome exclusionary ingroup construals. On the other hand, reservations could give rise to caste-based resentment. As far as advantaged-group members perceive themselves as unfairly disadvantaged (see Norton & Sommers, 2011), they could adopt narrower ingroup construal in response to reservations. Future research should examine how experiences with reservations (e.g., being rejected from a prestigious university) and perceptions of affirmative action (e.g., as 'reverse discrimination') shape ingroup construals in advantaged groups.

Among disadvantaged groups, the present research found a more complex pattern of inclusion and exclusion. Craig and Richeson (2012) raised the question when members of one disadvantaged group *derogate* members of another disadvantaged group, and when they instead strive for a *coalition*. Studying Black–Latino relations in the United States, the researchers found that making discrimination against the ingroup salient decreased derogation of a disadvantaged outgroup. In contrast, Study 4 found that students from Other Backward Classes tended to exclude Scheduled Castes and Scheduled Tribes from their ingroup. Indeed, Other Backward Class students' responses did not differ from General Merit students'. Other Backward Classes occupy an intermediate position in the caste hierarchy—below general castes, but above the 'line of pollution'. Seeing Scheduled Caste / Scheduled Tribe targets in the crossed-categorisation task could remind Other Backward Class participants that their group also faces discrimination. Recognising this disadvantage could pose a threat to Other Backward Class participants' social identity (see Branscombe, Ellemers, Spears & Doosje, 1999; Craig & Richeson, 2012). Excluding the lower-status outgroup could alleviate this threat. However, Study 4 also found that cross-group friendship was associated with Other Backward Class students including Scheduled Caste / Scheduled Tribe students in their ingroup. This shows that intergroup contact could unite disadvantaged groups.

Among the most disadvantaged groups, the present study found that Dalit/Adivasi

(Scheduled Caste / Scheduled Tribe) participants' responses differed from those of participants from relatively advantaged castes, including Other Backward Classes. Prioritising their own caste group, Scheduled Caste / Scheduled Tribe students tended to exclude the two other caste groups. Surprisingly, Scheduled Caste / Scheduled Tribe participants were somewhat less likely to include Other Backward Class than General Merit targets. Further, Scheduled Caste / Scheduled Tribe participants were a lot less likely to include Muslims in their ingroup. At the time of writing, I have no clear hypothesis about this divide. As for the other groups, friendships with members of relatively advantaged groups were associated with Scheduled Caste / Scheduled Tribe students including these groups in their ingroup. As argued in Chapter 1, disadvantaged-group members' support for social change may be contingent on firm group boundaries and distinct identities. In line with the *demobilisation hypothesis*, positive contact could thus undermine disadvantaged-group members support for social change. In the next chapter, I examine this and other possible implications of intergroup contact for social change.

5

Intergroup contact and social change: A darker side of 'we' or a path to solidarity?[1]

In Chapters 1–4, I have argued that intergroup contact could bridge social divides by changing how individuals think and feel about themselves and others. As hypothesised, I have found that positive contact was associated with more continuous, fluid, and inclusive social identities (in Studies 1–4), which, in turn, were associated with more favourable outgroup attitudes (in Studies 2,4). Social divides, however, exist apart from individuals' hearts and minds. Structural inequalities divide individuals into members of advantaged and disadvantaged groups. Changing the attitudes and beliefs of the advantaged is seldom enough to overcome these material differences. Instead, social change often requires the concerted effort of the disadvantaged, and sometimes the advantaged, to challenge the unjust status quo. In this chapter, I study whether intergroup contact helps or hinders this struggle for social justice.

On the one hand, I test whether positive contact with advantaged-group members decreases support for social change among members of disadvantaged groups (*demobilisation hypothesis*). Past research found that contact reduced disadvantaged-group members' support for redistributive policies (Cakal et al., 2011; Dixon et al., 2007; Sengupta & Sibley, 2013), and willingness to engage in collective action (Cakal et al., 2011; Tropp et al., 2012). As reviewed in Chapter 1, positive contact could decrease support for social change by reducing disadvantaged-group members' ingroup identification (*demobilisation by disidentification*), perceptions of discrimination (*demobilisation by perceived equality*), and unfavourable outgroup attitudes (*demobilisation by liking*). I have argued that negative contact could have the opposite effects on these attitudes and beliefs. Accordingly, I also test whether negative contact increases collective action intentions and support for egalitarian policies in disadvantaged groups (*mobilisation hypothesis*).

On the other hand, I study whether intergroup contact can foster support for social change among members of advantaged groups (*solidarity hypothesis*). As argued in Chapter 1, positive contact with disadvantaged-group members could encourage solidarity-based action by reducing ingroup identification (*solidarity by disidentification*), creating awareness of structural discrimination (*solidarity by perceived inequality*), and improving outgroup attitudes (*solidarity by liking*). Fostering more continuous, fluid, and inclusive identities could be another mechanism through which positive contact fosters solidar-

[1]Studies 1 and 3 of this chapter have been published as Reimer et al. (2017).

ity. More continuous and fluid conceptions of identity could reduce advantaged-group members' narrow identification with the ingroup, and thus reduce their resistance to challenges to the ingroup's interests (in line with the *solidarity-by-disidentification hypothesis*). Likewise, more inclusive identities could foster concern for the plight of the disadvantaged group—making injustices faced by "them" into "our" problem (*solidarity by inclusion*). Positive contact could thus help social change by encouraging solidarity-based collective action and support for redistributive policies in advantaged groups. Conversely, negative contact could could have the opposite effects, and discourage solidarity with disadvantaged outgroups.

In this chapter, I pursue three new directions in the study of intergroup contact and social change. Past research has largely overlooked the potential of negative contact to mobilise the disadvantaged, and the potential of positive contact to mobilise the advantaged. Testing the hypotheses advanced above, I study the effects of positive and negative contact on social change in advantaged and disadvantaged groups. Past research has, for the most part, conceptualised social identity as strength of identification (e.g., Cakal et al., 2011). In line with Chapters 3 and 4, I explore the implications of more continuous, fluid, and inclusive identities on social change. I discuss whether more inclusive identities hinder social change by demobilising the disadvantaged (*a darker side of 'we'*; Dovidio et al., 2012) or help social change by mobilising the advantaged (*a path to solidarity*).

Present research

The present research investigated the simultaneous and differential effects of negative and positive contact on advantaged-group and disadvantaged-group members' support for social change. In Studies 1–4, I recruited members of advantaged and disadvantaged groups who filled in measures of intergroup contact, group identification, perceived discrimination, outgroup attitudes, collective action intentions, and support for egalitarian policies. In Chapters 3 and 4, I have already presented findings from Studies 1–4, focusing on intergroup contact, social identification, and intergroup bias in advantaged groups. Here, I present additional results from disadvantaged-group participants to test the demobilisation and mobilisation hypotheses, and from advantaged-group participants to test the solidarity hypothesis.

Studies 1–3 focus on lesbian, gay, bisexual, transgender, and/or otherwise non-heterosexual (LGBT+) students and cisgender/heterosexual students as members of disadvantaged and advantaged groups, respectively. I studied sexual orientation and gender identity because *i.* they had rarely been studied in this context (for an exception, see Becker, Wright, Lubensky & Zhou, 2013, Study 1), *ii.* members of both groups are in frequent contact, and *iii.* the LGBT+ movement has a culture of involving heterosexual people (Russell, 2011), for example, in gay-straight campus alliances. Moreover, intergroup contact is highly effective in reducing sexual prejudice (Pettigrew & Tropp, 2006; Smith et al., 2009) and thus makes a good test case for the demobilisation hypothesis. Studies 1 and 3 measured support for social change among both sexual-minority and heterosexual participants. Study 2, however, measured support for social change only among sexual-minority participants; heterosexual participants instead responded to items related to scale development and validation (see Chapter 3). Among sexual-minority participants, Studies 1 and 2 measured identical constructs. To reach a sufficient sample size, I combined

responses from sexual-minority participants in Studies 1 and 2 (henceforth: Study 1). As a result, the present research included Studies 1 and 3, but not Study 2.

Study 4 examined support for social change among South Indian students of different caste backgrounds. Scheduled Caste (SC) and Scheduled Tribe (ST) students come from historically disadvantaged backgrounds, while General Merit (GM) students come from advantaged backgrounds. Other Backward Class (OBC) students occupy an intermediate position in the caste hierarchy (see Chapter 4). This context allows testing whether intergroup contact can foster solidarity between two disadvantaged groups (SC/ST and OBC students).

As outlined above, I predicted that positive contact would discourage disadvantaged-group members from supporting social change (*demobilisation hypothesis*), that negative contact would encourage disadvantaged-group members to support social change (*mobilisation hypothesis*), and that positive contact would encourage advantaged-group members to support social change (*solidarity hypothesis*). In line with previous research, I studied perceived discrimination, outgroup attitudes, and social identification as potential mediators of the (de)mobilising effects of intergroup contact. In addition, I tested to what extent more fluid, continuous, and inclusive identities mediated the hypothesised relationship between intergroup contact and cross-group solidarity.

Study 1

As part of two wider studies, participants filled in measures of positive and negative contact (predictors), group/movement identification, perceived discrimination, outgroup attitudes (mediators), and collective action intentions (outcome). In both studies, participants chose whether they wanted to be shown questions and answer options for heterosexual or non-heterosexual people. Participants who chose the questionnaire for non-heterosexual people answered questions on contact with the advantaged outgroup, discrimination against and collective action for their disadvantaged ingroup. To reach a sufficient sample size, I combined responses from sexual-minority participants in Studies 1 and 2 (henceforth: Study 1a). In Study 1, self-identified heterosexual participants answered questions on contact with, discrimination against, and solidarity-based collective action for the disadvantaged outgroup (henceforth: Study 1b). In Study 2, heterosexual participants responded to items related to scale development and validation (see Chapter 3), and were not included in the present research.

Study 1 did not include a (well-validated) measure of social identity continuity and fluidity, and thus did not allow testing how more continuous and fluid identities related to support for social change. Study 1b included a measure of advantaged-group members' identification with the LGBT movement. As the LGBT movement is defined by shared goals rather than mere category membership, it leaves room for allies from advantaged outgroups (see Russell, 2011), and thus offers a potential shared identity. In line with the *solidarity-by-inclusion hypothesis*, I test whether positive contact inspires advantaged-group members to identify with the struggle of the disadvantaged, and to thus engage in solidarity-based collective action. Below, I describe samples and measures in Study 1, omitting procedural details that have already been described in Chapter 3.

Method

Participants

Study 1a included 233 sexual-minority respondents (129 women, 93 men, 11 responded "none of these/other"; Mdn_{age} = 21, age range: 18–42 years) who identified as gay or lesbian (n = 88), bisexual or pansexual (n = 112), asexual (n = 12), queer or other (n = 11), or were uncertain about their sexual orientation (n = 10). Study 1b included 249 heterosexual respondents (164 women, 83 men, 2 "none of these/other"; Mdn_{age} = 22, age range: 18–38 years). I excluded 48 respondents who did not identify as heterosexual/straight, had had sexual relations with a same-sex partner, or were attracted to persons of their own gender more than 50% of the time (see Chapter 3).

Measures

Intergroup contact was measured as how often, from 1 = *never* to 5 = *very often*, respondents had a variety of positive and negative experiences with [non-]heterosexual people (Fell et al., 2016; Stephan et al., 2002). *Negative contact* (α_{1a} = .81, α_{1b} = .70) was measured using five items: being verbally abused, intimidated, threatened with harm, ridiculed, and made to feel unwanted. *Positive contact* (α_{1a} = .75, α_{1b} = .85) was measured using five items: being supported, helped, complimented, befriended, and made to feel welcome.

Group identification was measured using four items (adapted from Doosje et al., 1998): "Being [non-]heterosexual is an important part of who I am", "I identify with other [non-]heterosexual people" , "Overall, being [non-]heterosexual has a lot to do with how I feel about myself", and "I see myself as [non-]heterosexual" (1 = *strongly disagree*, 7 = *strongly agree*; α_{1a} = .80, α_{1b} = .78).

Perceived discrimination was measured using items adapted from Mays and Cochran (2001) and Garstka, Schmitt, Branscombe and Hummert (2004). Four items measured *group discrimination* (α_{1a} = .75, α_{1b} = .83): "Nonheterosexual people as a group have been victimized by society", "Historically, non-heterosexual people have been discriminated against more than heterosexual people" (1 = *strongly disagree*, 7 = *strongly agree*), "To what extent does discrimination interfere with non-heterosexual people's pursuit of a full and productive life?", and "To what extent does discrimination make the life of non-heterosexual people harder?" (1 = *not at all*, 4 = *a lot*).

Outgroup attitudes were measured using Wright et al.'s (1997) General Evaluation Scale. Respondents rated how they felt about [non-]heterosexual people in general on 7-point adjective scales: warm–cold (reversed), negative–positive, friendly–hostile (reversed), suspicious–trusting, respect–contempt (reversed), and admiration–disgust (reversed). Items formed a reliable scale (α_{1a} = .92, α_{1b} = .96) with higher scores indicating more favourable attitudes.

Collective action intentions (α_{1a} = .90, α_{1b} = .89) were measured as how often, between 1 = *never* and 6 = *very often*, participants would likely engage in six actions (e.g., participating in demonstrations) to advocate for LGB rights and against LGB discrimination. In addition, I corroborated the scale above with two more tangible behavioural intentions. Participants indicated how many, between 0 and 100, flyers advocating against LGB discrimination they were willing to distribute and what proportion of their potential winnings (0–100%) they agreed to donate to a charity advocating for LGB rights.

In addition, for sexual-minority respondents only, I also measured their sense of personal discrimination (α_{1a} = .77): "I feel like I am personally a victim of society because

Table 5.1. Model fit indices for configural, loadings, intercept, and partially invariant models for minority and majority participants in Study 1. Models showed loadings invariance ($\Delta CFI < .01$), but not intercepts invariance. After freeing one loading and two intercepts, Model 3 showed partial measurement invariance.

#	Model	χ^2	*df*	χ^2/df	*CFI*	*TLI*	*RMSEA*	ΔCFI
0	Configural invariance	395.89	258	1.53	.96	.95	.05	
1 vs 0	Loadings invariance	435.07	271	1.61	.95	.94	.05	-.008
2 vs 1	Intercepts invariance	513.63	284	1.81	.93	.92	.06	-.021
3 vs 0	Partial invariance	425.21	281	1.51	.95	.95	.05	-.002

of my sexual orientation", "I consider myself a person who has been deprived of opportunities that are available to others because of my sexual orientation" (1 = *strongly disagree*, 7 = *strongly agree*), and "To what extent has discrimination made your life harder?" (1 = *not at all*, 4 = *a lot*). For sexual-majority respondents only, I measured how much heterosexual respondents adopted a solidarity-based *movement identity*: "I identify with the values and goals of the LGBT (lesbian, gay, bisexual, and transgender) movement" (1 = *strongly disagree*, 7 = *strongly agree*).

Results

I used the following analytic strategy: First, I conducted latent-mean comparisons between the two groups, and then tested whether the relationships of positive and negative contact with collective action were moderated by group status (*multi-group comparison*). Second, I estimated the predicted direct and indirect (via group identification, personal and group discrimination, and outgroup attitudes) effects of both types of contact on collective action among minority-group members (Study 1a). Third, I estimated the predicted direct and indirect (via group/movement identification, outgroup attitudes, and group discrimination) effects of both types of contact on collective action among majority-group members (Study 1b).

Multi-group comparison

I established partial measurement invariance of the latent factors of positive and negative contact, and collective action (see Table 5.1), indicating that scales measured comparable constructs among minority and majority-group participants. As expected for a numerical minority, non-heterosexual participants reported more positive, $\Delta M_{\text{PC}} = 0.39, [0.24, 0.53]$, and negative, $\Delta M_{\text{NC}} = 0.75, [0.63, 0.87]$, contact than heterosexual participants as well as stronger collective action intentions, $\Delta M_{\text{CA}} = 0.79, [0.44, 1.15]$. Negative and positive contact were negatively correlated among minority respondents, $\phi = -0.31, [-0.49, -0.13]$, but not among majority respondents, $\phi = 0.02, [-0.16, 0.20]$.

Next, I estimated the effects of both types of contact on collective action, and tested whether group status (advantaged vs disadvantaged) moderated these relationships. Positive contact was associated with more collective action for advantaged-group members, $\beta = 0.65, [0.48, 0.82]$, in line with the *solidarity hypothesis*. Contrary to the *demobilisation hypothesis*, however, positive contact had no significant effect for disadvantaged-group members, $\beta = 0.19, [-0.06, 0.45]$, $\Delta\beta_{\text{maj-min}} = 0.46, [0.16, 0.75]$. In line with the *mo-*

bilisation hypothesis, negative contact was associated with more collective action for disadvantaged-group members, $\beta = 0.38, [0.15, 0.60]$. For advantaged-group members, negative contact was associated with less collective action, $\beta = -0.47, [-0.84, -0.09]$, $\Delta\beta_{\text{maj-min}} = -0.84, [-1.28, -0.40]$. For participants of both groups, collective action intentions were associated with the numbers of flyers they agreed to distribute ($\beta_{\text{maj}} = 13.79$, $[10.83, 19.33]$, $R^2 = .31$ and $\beta_{\text{min}} = 11.84, [7.99, 15.69]$, $R^2 = .14$) and the potential winnings they pledged to donate ($\beta_{\text{maj}} = 10.26, [6.88, 13.65]$, $R^2 = .14$ and $\beta_{\text{min}} = 4.70, [0.25, 9.15]$, $R^2 = .02$).

Study 1a: Sexual-minority respondents

Next, I estimated the indirect effects of contact on disadvantaged-group members' collective action intentions via group identification, perceived discrimination, and outgroup attitudes in a structural equation model (see Figure 5.1). The model fit well, $\chi^2(208) = 295.06$, $\chi^2/df = 1.42$, $CFI = .97$, $TLI = .96$, $RMSEA = .04$, explaining 46% of the variance in collective action intentions.

As hypothesised, negative contact was associated with more perceived personal discrimination ($\beta = 0.63$), which, in turn, correlated with perceptions of group discrimination ($\beta = 0.71$). Perceived group discrimination was associated with collective action intentions ($\beta = 0.35$). Similarly, negative contact was associated with higher identification ($\beta = 0.27$), which was associated with more personal discrimination ($\beta = 0.40$) and collective action ($\beta = 0.63$). Contrary to predictions, positive contact was related to neither group identification ($\beta = 0.03, [-0.15, 0.20]$) nor perceived personal ($\beta = -0.09$, $[-0.24, 0.07]$) or group ($\beta = -0.08, [-0.32, 0.16]$) discrimination. Outgroup attitudes were not significantly associated with either perceived group discrimination ($\beta = -0.01$, $[-0.23, 0.21]$) or collective action ($\beta = 0.07, [-0.12, 0.26]$).

For ease of comparison, I summarised all indirect paths that share the same immediate predictor of collective action intentions (Figure 5.2). Indirect paths via group identification, point estimate (PE) $\beta = 0.22, [0.07, 0.51]$, and perceived group discrimination, PE $\beta = 0.23, [0.07, 0.66]$, accounted for 17% and 18% of the absolute association between negative contact and collective action, PE $|\beta| = 1.27, [0.35, 3.27]$.[2] Negative contact was directly associated with collective action, PE $\beta = 0.44, [0.05, 1.15]$, accounting for a further 35% of the relationship.[3] The negative but non-significant indirect path via perceived personal discrimination, PE $\beta = -0.39, [-1.21, 0.01]$, accounted for the remaining 31%. Contrary to the *demobilisation hypothesis*, none of the direct or indirect paths from positive contact to collective action were significant.

These findings contradict earlier research demonstrating the demobilising effects of intergroup contact. Importantly, however, the current study measured contact valence as a bidimensional construct while most previous research relied on unidimensional measures such as contact quality (e.g., Dixon et al., 2010) or number of outgroup friends. When negative and positive contact are negatively correlated, however, measuring contact valence as a unidimensional construct risks conflating the presence of positive contact with the absence of negative contact. To test whether differences in measurement can

[2] As indirect effects in Study 1a differed in sign (Figure 5.2), I computed the absolute effect, that is, the sum of all absolute effect sizes, rather than the total effect, that is, the sum of all positive and negative effect sizes.

[3] Effect sizes are unstandardised coefficients from a model with standardised indicators. Available software does not provide standardised coefficients for bootstrap sampling. As a result, the reported (direct) effect sizes differ from the standardised coefficients shown in Figure 5.1.

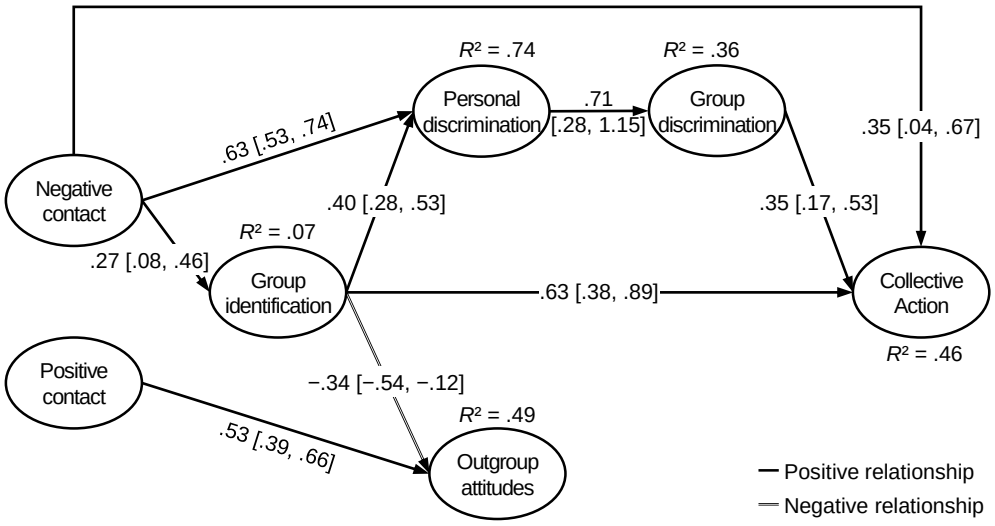

Figure 5.1. Structural equation model showing direct and indirect paths from negative and positive contact to collective action for sexual-minority participants (Study 1a), with standardised coefficients and 95% confidence intervals. Negative and positive contact were correlated ($\phi = -0.32$, $[-0.52, -0.12]$). Only significant paths are shown.

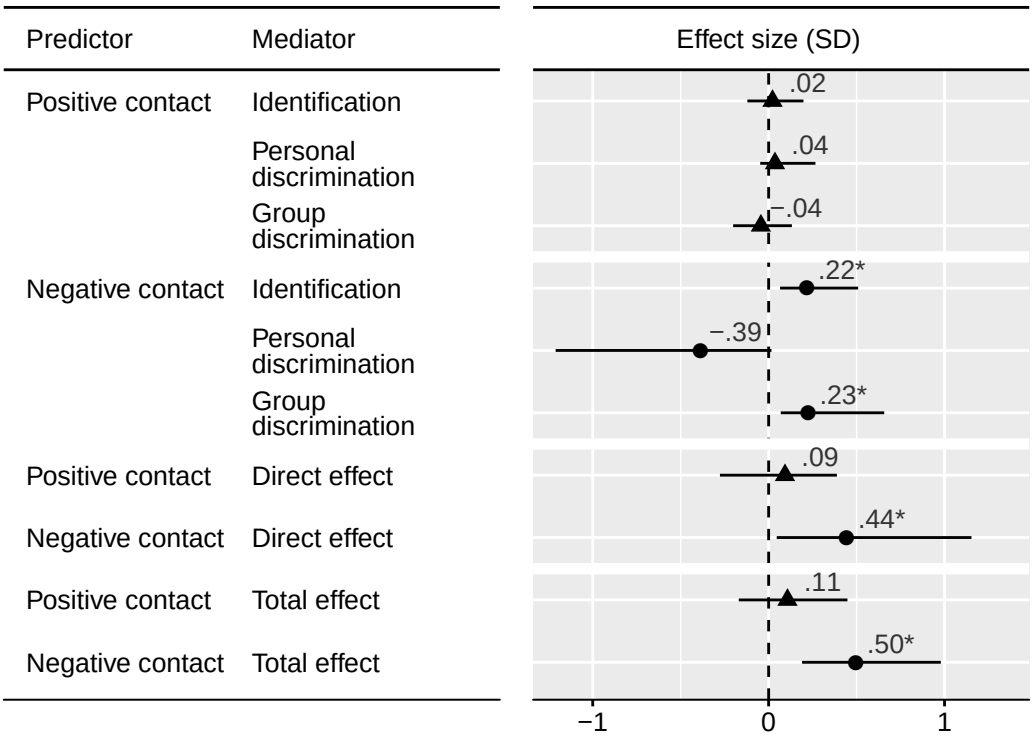

Figure 5.2. Effect sizes with 95% confidence intervals for indirect, direct, and total effects of intergroup contact on collective action among sexual-minority participants (Study 1a).

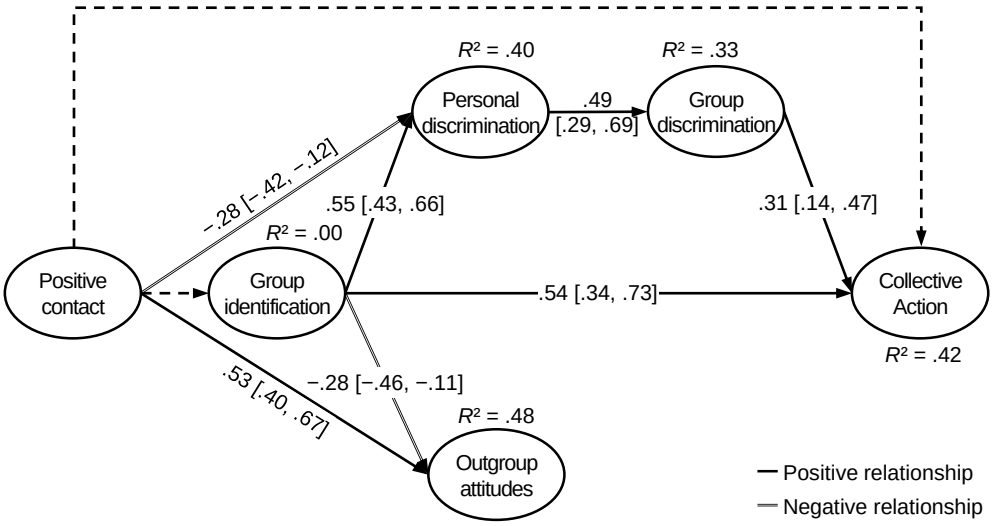

Figure 5.3. Alternative model including direct and indirect effects of positive but not negative contact for sexual-minority students (Study 1a). Previously significant paths are shown as dashed lines. Model fit: $\chi^2(154) = 194.17$, $\chi^2/df = 1.26$, *CFI* = .98, *TLI* = .98, *RMSEA* = .03

explain the difference between previous and present findings, I estimated an alternative model in which I included positive contact as the only predictor variable (Figure 5.3). In line with the *demobilisation hypothesis*, positive contact was associated with lower perceived personal discrimination ($\beta = -0.28$) when negative contact was not accounted for.

Study 1b: Sexual-majority respondents

For advantaged-group members, I estimated a structural equation model with positive and negative contact as predictors, group and movement identification, outgroup attitudes, and perceived group discrimination as mediators, and collective action intentions as outcome. Contrary to predictions, group identification was not associated with collective action intentions ($\beta = -0.03$, $[-0.15, 0.10]$), perceived discrimination ($\beta = 0.04$, $[-0.11, 0.18]$), or outgroup attitudes ($\beta = 0.10$, $[-0.05, 0.24]$), and was omitted from further analyses. The resultant model fit well, $\chi^2(122) = 186.80$, $\chi^2/df = 1.53$, *CFI* = .97, *TLI* = .97, *RMSEA* = .05, explaining 53% of the variance in solidarity-based collective action intentions (Figure 5.4).

In line with the *solidarity-by-inclusion hypothesis*, movement identification ($\beta = 0.43$) was associated with collective action. As expected, outgroup attitudes ($\beta = 0.19$) were associated with collective action though, other than expected, perceived group discrimination was not ($\beta = 0.09$, $[-0.04, 0.22]$). Indirect paths via movement identification, PE $\beta = 0.23$, $[0.15, 0.31]$, and outgroup attitudes, PE $\beta = 0.09$, $[0.03, 0.16]$, accounted for, respectively, 40% and 16% of the total association between positive contact and collective action, PE $\beta = 0.56$, $[0.44, 0.68]$. For negative contact, indirect effects via movement identification, PE $\beta = -0.13$, $[-0.24, -0.06]$, and outgroup attitudes, PE $\beta = -0.05$, $[-0.12, -0.01]$, accounted for, respectively, 61% and 24% of the total effect, PE $\beta = -0.21$, $[-0.41, -0.04]$. In addition, positive, PE $\beta = 0.22$, $[0.09, 0.35]$, but not negative, PE $\beta = 0.02$, $[-0.20, 0.22]$, contact was directly associated with collective action,

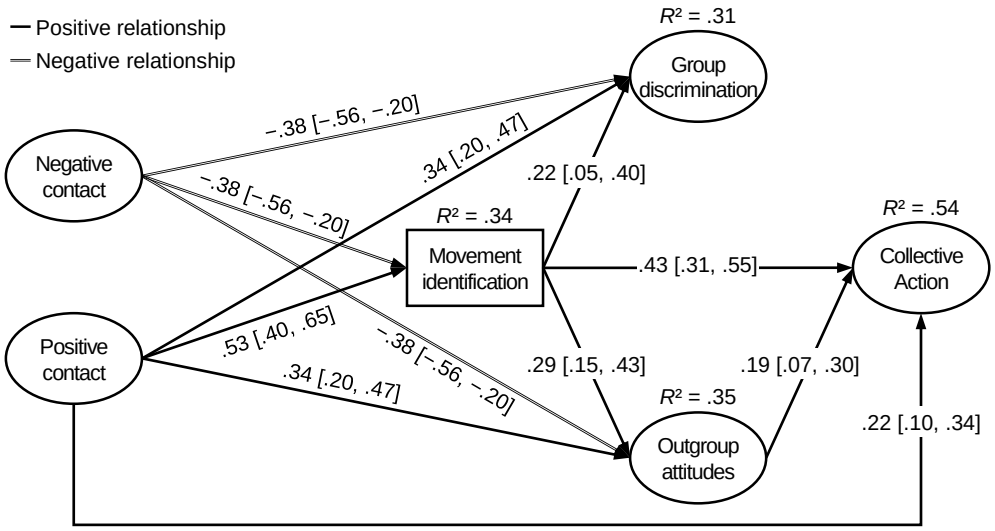

Figure 5.4. Structural equation model showing direct and indirect paths from negative and positive contact to collective action for sexual-majority participants (Study 1b), with standardised coefficients and 95% confidence intervals. Negative and positive contact were not correlated (ϕ = 0.02, [−0.16, −0.20]). Only significant paths are shown.

accounting for 39% of the total relationship. As the combined effect of positive contact surpassed that of negative contact, $\Delta\beta$ = 0.35, [0.13, 0.55], this study highlighted the importance of positive contact in promoting support for social change in disadvantaged groups.

Study 3

Study 1a found that negative contact was associated with more collective action among sexual-minority participants. Contrary to prior research, positive contact was unrelated to collective action or perceived discrimination when negative contact was accounted for. Group identification was positively associated with both perceived discrimination and collective action. Contrary to the darker-side-of-'we' argument, however, positive contact did not weaken group identification. Study 1b found that for sexual-majority participants, positive and negative contact was associated with, respectively, stronger and weaker intentions to defend the rights of sexual minorities. As expected, movement identification mediated these relationships.

Study 3 addressed several shortcomings of Studies 1a and 1b. First, the cross-sectional design of Studies 1a and 1b did not allow me to establish the direction of the observed relationships, for example, whether negative contact with heterosexual people encourages collective action or whether collective action begets negative contact. Study 3 addressed this limitation by using a three-wave longitudinal design. Second, Studies 1a and 1b covered only part of the LGBT+ movement. Study 3 includes transgender people, both as participants and as relevant outgroup members. This allowed me to test whether contact with transgender people encouraged solidarity-based collective action among non-heterosexual, cisgender people (*inter-minority solidarity*) and among heterosexual, cisgender people (*majority-minority solidarity*). Third, Studies 1a and 1b did not include measures of continuity/fluidity. Study 3 included validated scales of social identity

continuity and fluidity (see Chapter 3), and thus tested to what extent the novel scales mediated the relationships between intergroup contact and solidarity-based collective action. Finally, Study 3 tested an additional mechanism through which intergroup contact could (de)mobilise disadvantaged-group members. Tausch, Saguy and Bryson (2015) found that cross-group friendship was associated with less anger about discrimination among Latino Americans, suggesting that positive contact can reduce collective action by quelling group-based anger (*demobilisation by calmness*). Study 3 tested this mechanism, alongside the ones included in the previous studies. This study is comprised of three studies, each focused on a distinct subgroup of participants: sexual-minority (Study 3a), sexual-majority (Study 3b), and gender-majority (Study 3c).

Method

Participants

Study 3a included 433 respondents (288 women, 116 men, 29 responded "none of these/other"; Mdn_{age} = 22, age range: 18–51 years) who identified as gay/lesbian (n = 113), bisexual (n = 158), pansexual (n = 41), queer (n = 32), questioning (n = 67), asexual (n = 11), or otherwise non-heterosexual (n = 11). I excluded 72 participants who had identified as heterosexual at any time point. Forty-three participants were trans women (n = 2), trans men (n = 7), genderqueer/other non-binary (n = 10), gender-nonconforming (n = 5), genderfluid (n = 11), agender (n = 6), or otherwise transgender (n = 2). Study 3b included 1054 heterosexual participants (782 women, 269 men; Mdn_{age} = 22, age range: 18–65 years). I excluded 299 individuals using similar criteria as in Study 1b. Study 3c included 1726 cisgender participants (1298 women, 426 men; Mdn_{age} = 22, age range: 18–65 years) who, at none of the time points, reported a non-binary gender identity or identified as transgender. Study 3c encompassed most of the participants included in Studies 3a and 3b.

Procedure

Participants answered a range of questions on their sexual orientation and gender identity based on which they filled in different sets of items (see Chapter 3). LGBT+ respondents answered questions concerning *heterosexual* (Study 3a) and/or *cisgender people* as the relevant outgroups, and *lesbian, gay, bisexual, or otherwise non-heterosexual* (LGB+; Study 3a) and/or *transgender people* as the relevant ingroups. Heterosexual participants completed measures with *LGB+* as the relevant outgroup (Study 3b). Cisgender participants completed similar measures with *transgender people* as the relevant outgroup (Study 3c). As responses for Studies 3b and 3c were collected as part of a wider study on intergroup contact and social identity (see Chapter 3), I measured fewer variables and used shorter scales than in Studies 1a, 1b, and 3a.

Measures for LGBT+ respondents

Positive and *negative contact* with heterosexual and/or cisgender people, perceived *personal* and *group discrimination* against LGB+ people and/or transgender people were measured with the same scales as in Study 1a. The extent to which participants identified with people of the same sexual orientation (e.g., bisexual people) and/or transgender people was measured using Leach et al.'s (2008) three-item *solidarity* and *centrality* subscales.

Anger was measured as the extent (1 = *not at all*, 5 = *extremely*) to which respondents felt angry, resentful, furious, and displeased about the discrimination against LGB+ and/or transgender people (Tausch et al., 2015). Outgroup attitudes were assessed with a feeling thermometer (Converse et al., 1980). Similar to Study 1a, collective action intentions were measured as how likely, between 1 = *very unlikely* and 7 = *very likely*, participants would be to engage in a variety of actions to advance and defend LGBT+ rights.

Measures for heterosexual/cisgender respondents

Participants reported how many, from 1 = *none* to 7 = *more than 16*, gay men, lesbian/gay women, and bisexual people as well as trans women, trans men, and genderqueer people they knew personally. For each outgroup, if they knew any member of that group, respondents indicated how often, on average, they had "*positive/good contact* with [them]—for example, being supported, made to feel welcome, helped, complimented, or befriended by them" and "*negative/bad contact* with [them]—for example, being made to feel unwanted, verbally abused, intimidated, threatened, or ridiculed by them" (1 = *never*, 5 = *very often*; adapted from Barlow et al., 2012). *Collective action intentions* were measured as how likely, from 1 = *very unlikely* to 7 = *very likely*, respondents would be to "take actions—for example, signing petitions and participating in demonstrations—to advocate for [LGB+/transgender] people's rights and against [LGB+/transgender] discrimination." Social identity *continuity* and *fluidity* were measured as described in Chapter 3 (Study 3). As in Study 1, all participants indicated what proportion of their potential winnings (0–100%) they would donate "to advance the rights of and reduce discrimination against LGBT+ people" at the end of the final survey.[4]

Results

I tested the hypotheses in a series of cross-lagged panel models (see Little, 2013) to establish the direction of the observed effects, for example, whether contact predicts collective action (forward path), collective action predicts contact (reverse path), or both, while accounting for the effect of each variable on itself over time (autoregressive paths). I report reliability indices (all $\alpha > .70$), descriptive statistics, and latent-variable correlations in Appendix C. Below, I report results for sexual-minority (Study 3a), sexual-majority (Study 3b), and gender-majority (Study 3b) participants separately.

For Study 3a, I estimated a two-wave panel model using only responses collected at the beginning (T1) and end (T3) of the academic year. I opted for the simpler (two-wave) model as, in the full (three-wave) model, the relevant constructs remained (mostly) stable across the shorter inter-survey intervals ($\approx$ 100 days) with very few cross-lagged paths in-between.[5] In Studies 3b and 3c, I tested fewer variables with more participants and was able to use all three waves of data. Unlike Studies 1a and 1b, I could not directly test group status as a moderator as I included different measures for the three samples. Moreover, as

[4] As in Studies 1a and 1b, collective action intentions (at time point 3) correlated with the percentage of their potential winnings that LGBT+ (lesbian, gay, bisexual, transgender, and/or otherwise non-heterosexual) participants (β = 7.46, [1.61, 13.32]) and heterosexual/cisgender participants (β = 5.14, [2.80, 7.48]) pledged to donate.

[5] I was concerned about the ratio between the number of participants (433) and the number of estimable parameters when using all three waves (565). Estimating indirect effects over all three waves (with missing data) proved prohibitively demanding computationally. Using the full three-wave dataset, however, did not (qualitatively) change the findings and conclusions (see Figure C.1).

Table 5.2. Model fit indices for configural, loadings intercepts, means, and covariances invariant models for two time points and two language groups in Study 3a.

#	Invariance	χ^2	*df*	χ^2/df	*CFI*	*TLI*	*RMSEA*	Δ*CFI*
0	Configural	3080	1872	1.65	.89	.87	.05	
1	Loadings (group)	3102	1920	1.62	.89	.87	.05	-.002
2	Loadings (group, time)	3120	1944	1.61	.89	.87	.05	-.000
3	Intercepts (group)	3120	1976	1.58	.89	.88	.05	-.003
4	Intercepts (group, time)	3121	2000	1.56	.90	.88	.05	-.002
5	Means (group)	3121	2016	1.55	.90	.89	.05	-.001
6	Covariances (group)	3239	2136	1.52	.90	.89	.05	-.000

only 2% to 3% (n = 43) of respondents were transgender, I did not have a large enough sample for any meaningful (longitudinal) analyses in this subsample.

Study 3a: LGBT+ respondents

In Study 3a, I tested the hypothesised longitudinal relationships between positive and negative contact, group identification, perceived discrimination, anger, outgroup attitudes, and collective action intentions. Covariance missingness ranged from 3% to 46% across items and waves; I used full-information maximum likelihood (FIML) estimation to handle the missing data. I established measurement invariance across languages and time points (Table 5.2). I combined German and English samples, constraining all factor loadings and intercepts to be equal across time points. The resultant cross-lagged model fit well, $\chi^2(1040) = 1604.98$, $\chi^2/df = 1.54$, *CFI* = .94, *TLI* = .93, *RMSEA* = .04, explaining 70% of the variance in collective action intentions (Figure 5.5).

All autoregressive paths were significant ($\beta > 0.48$). Similar to Study 1a, negative contact with heterosexual people at T1 predicted perceived group discrimination ($\beta = 0.13$) and group-based anger ($\beta = 0.13$) $6\frac{1}{2}$ months later (T3). Perceived group discrimination at T1 ($\beta = 0.18$), though not anger ($\beta = 0.01$, $[-0.11, 0.13]$), in turn, predicted collective action intentions at T3. Contrary to the *demobilisation-by-disidentification hypothesis*, neither the centrality ($\beta = -0.06$, $[-0.18, 0.06]$) nor solidarity ($\beta = 0.03$, $[-0.07, 0.14]$) dimensions of participants' sexual identity predicted collective action intentions. Outgroup attitudes did not predict collective action ($\beta = 0.01$, $[-0.02, 0.03]$). Negative contact did not directly predict collective action ($\beta = 0.00$, $[-0.09, 0.08]$). None of the reverse paths were significant (with the exception of the centrality of participants' sexual identity), supporting the hypothesised direction of the relationships. Contrary to the *demobilisation hypothesis*, and in line with Study 1a, none of the direct or indirect paths from positive contact to collective action were significant.

I estimated the mediated effects as the product of the coefficients of the cross-lagged paths from the T1 predictor to the T3 mediator, and from the T1 mediator to the T3 outcome (half-longitudinal mediation). As hypothesized, negative contact had an indirect effect on collective action via group discrimination, PE $\beta = 0.03$, $[> 0.00, 0.07]$. Contrary to Tausch et al.'s (2015) findings, anger did not mediate the paths from positive, PE $\beta = -0.00$, $[-0.03, 0.01]$, and negative, PE $\beta = 0.00$, $[-0.02, 0.03]$, contact to collective action. None of the other direct or indirect paths from T1 contact to T3 collective action were significant after accounting for all other variables (see Figure 5.6). Study 3a thus

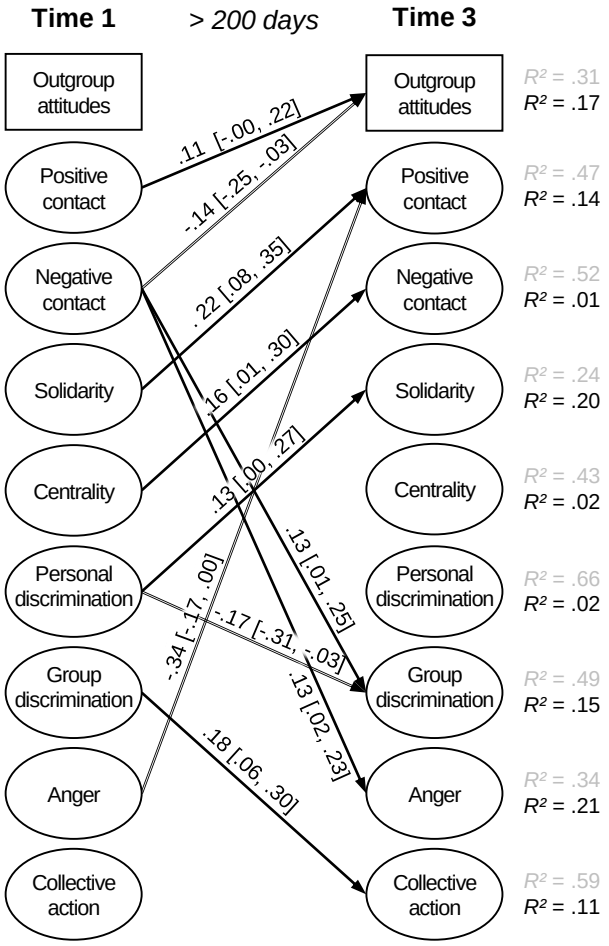

Figure 5.5. Panel model showing cross-lagged paths for sexual-minority participants (Study 3a), with standardised coefficients and 95% confidence intervals. Negative and positive contact were correlated at T1 ($\phi = -0.19$, $[-0.32, -0.05]$). Auto-regressive paths are not shown. R^2 (in grey) indicates the total variance explained; R^2 (in black) indicates the variance explained by the cross-lagged paths alone.

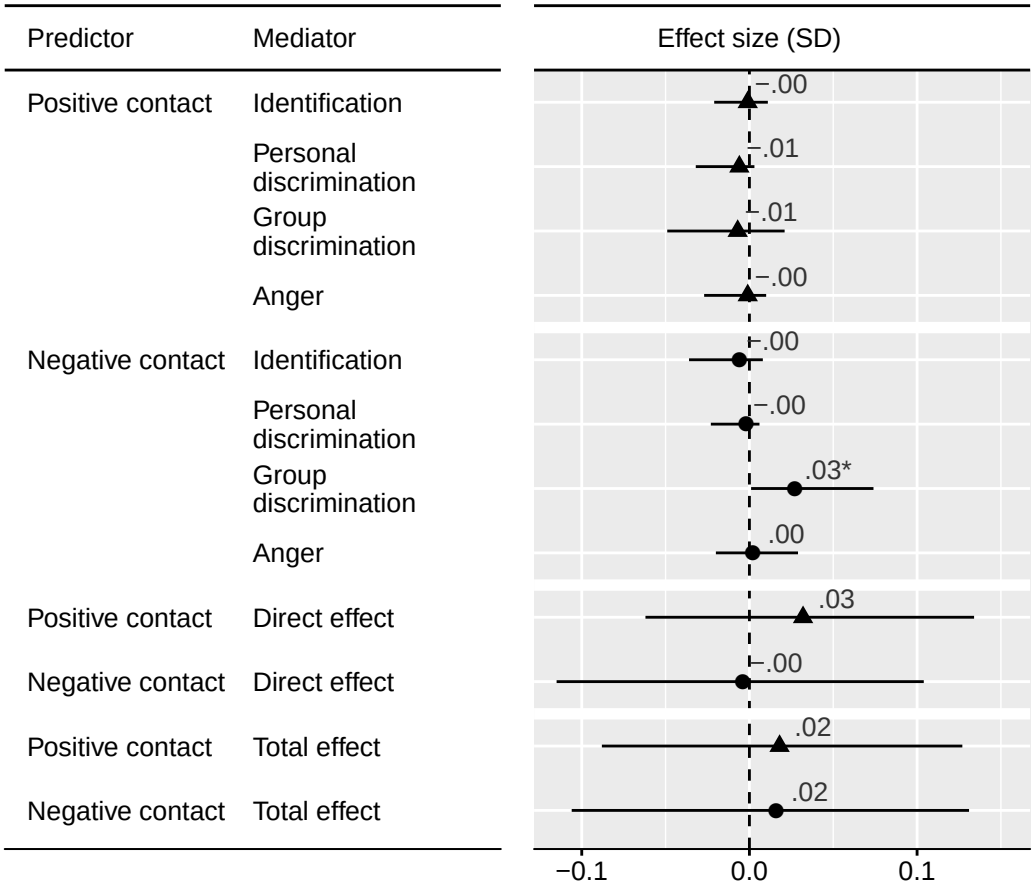

Figure 5.6. Effect sizes with 95% confidence intervals for indirect (half-longitudinal), direct (cross-lagged), and total effects of intergroup contact on collective action among sexual-minority participants (Study 3a).

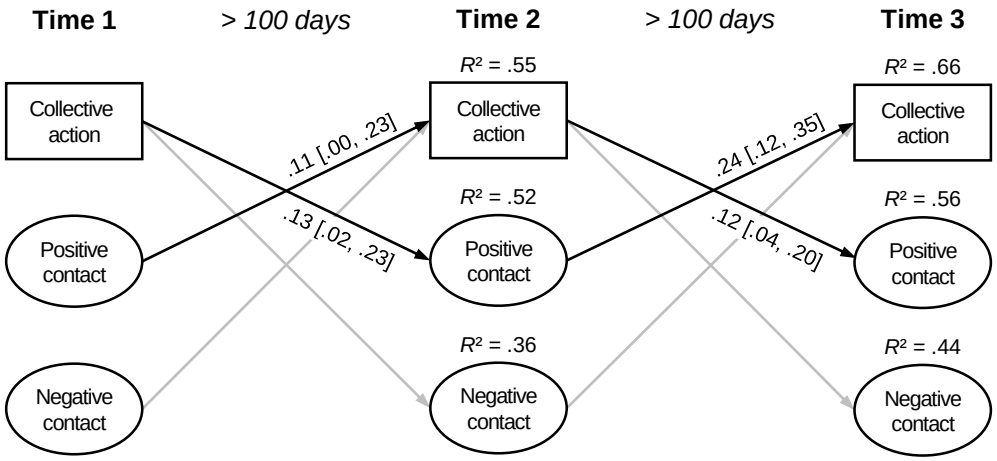

Figure 5.7. Panel model for heterosexual participants (Study 3b), with standardised coefficients and 95% confidence intervals. Negative and positive contact were not correlated at T1 ($\phi = -0.12$, $[-0.29, 0.06]$). Paths in grey were not significant, $p > .05$. Autoregressive paths are not shown.

replicated the indirect association between negative contact and collective action via perceived discrimination.

Study 3b: Heterosexual respondents

In Study 3b, I tested the *solidarity hypothesis* using responses from heterosexual respondents, with lesbian, gay, bisexual, and other non-heterosexual (LGB+) people as the relevant outgroup. First, I estimated a three-wave panel model regressing collective action intentions (for LGB+ rights and against LGB+ discrimination) on positive and negative contact with lesbian, gay, and bisexual people (see Figure 5.7). The resulting model fit well, $\chi^2(166) = 507.33$, $\chi^2/df = 3.06$, $CFI = .94$, $TLI = .93$, $RMSEA = .04$.[6] As hypothesised, positive contact predicted collective action across both $3\frac{1}{4}$ months intervals, $\beta_{T_1T_2} = 0.11$ and $\beta_{T_2T_3} = 0.24$. Likewise, collective action predicted more positive contact across both intervals, $\beta_{T_1T_2} = 0.13$ and $\beta_{T_2T_3} = 0.12$. The relationship between positive contact and (solidarity-based) collective action was thus bi-directional. Negative contact did not predict collective action ($\beta_{T_1T_2} = -0.05$, $[-0.13, -0.03]$; $\beta_{T_2T_3} = -0.02$, $[-0.09, -0.05]$), nor did collective action predict negative contact ($\beta_{T_1T_2} = -0.06$, $[-0.14, 0.02]$; $\beta_{T_2T_3} = 0.05$, $[-0.03, 0.12]$). These findings are consistent with the *solidarity hypothesis* that positive contact fosters support for social change in favour of disadvantaged groups among members of advantaged groups.

Next, I tested to what extent more continuous and fluid understandings of sexuality mediated this relationship. To that end, I replaced outgroup attitudes with collective action intentions in the final longitudinal model of Study 3a in Chapter 3. Model 1 thus included reverse paths from continuity: identity, continuity: attraction, and fluidity: future to positive and negative contact, as well as all autoregressive paths ($\chi^2(1835) = 3851.55$, $\chi^2/df = 2.10$, $CFI = .94$, $TLI = .94$, $RMSEA = .03$). Model 2 added forward paths from

[6] I established measurement invariance for all latent variables (across time points and language groups) as part of Study 3a in Chapter 3. I do not repeat these analyses here.

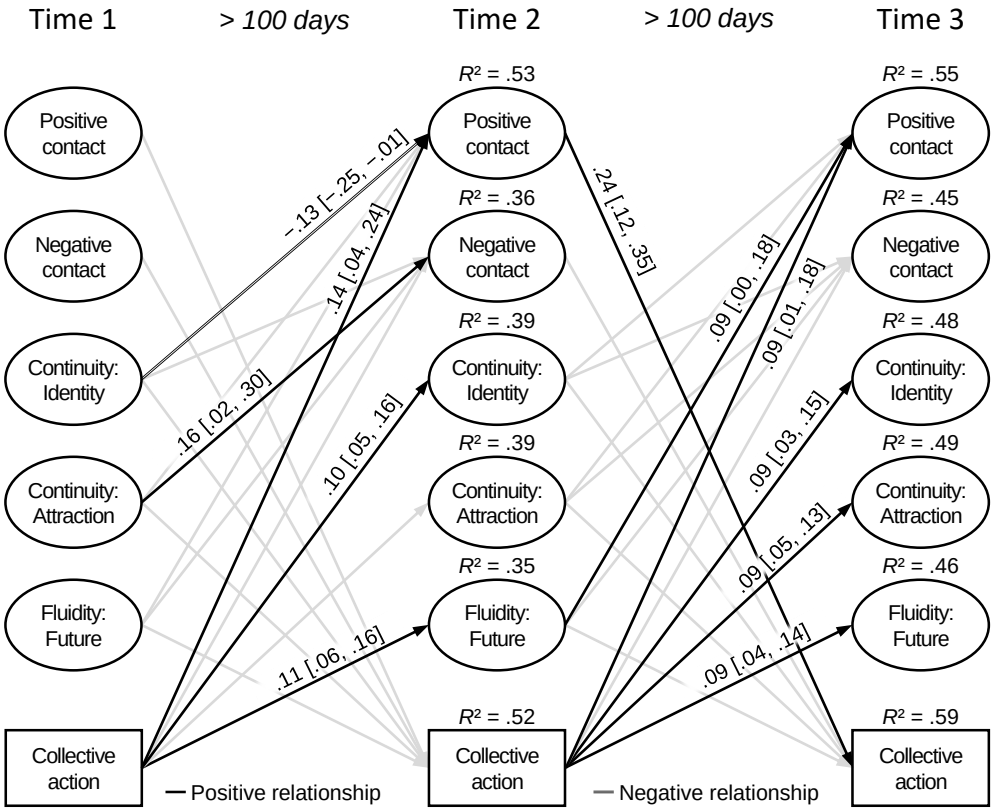

Figure 5.8. Results from Model 5 for heterosexual participants (Study 3b), with standardised coefficients and 95% confidence intervals. Paths in grey were not significant, $p > .05$. Autoregressive paths are not shown.

positive and negative contact (at T1, T2) to collective action (at T2, T3). Model 2 fit better than Model 1, $\Delta\chi^2(4) = 35.35$, $p < .001$, $\Delta CFI = .00$—confirming that contact predicted collective action. Model 3 added reverse paths from collective action to positive and negative contact, $\Delta\chi^2(11.29) = 4$, $p = .023$, $\Delta CFI = .00$. Model 4 added forward paths from the three continuity/fluidity dimensions to collective action. Model 4 improved upon Model 3, $\Delta\chi^2(6) = 18.47$, $p = .005$, $\Delta CFI = .00$—showing that continuity/fluidity factors predicted collective action. Model 5 added reverse paths from collective action to continuity/fluidity, and fit better than Model 4, $\Delta\chi^2(6) = 34.73$, $p < .001$, $\Delta CFI = .00$.

Figure 5.8 shows standardised coefficients for all significant ($p < .05$) forward and reverse paths in Model 5. Contrary to predictions, the forward paths from continuity/fluidity to collective action were not significant, though five of the six reverse paths were significant ($0.8 < \beta < .12$). The observed relationships do not support the hypothesis that more continuous and fluid conceptions of sexuality mediate the relationship between positive contact and cross-group solidarity.

Study 3c: Cisgender respondents

In Study 3c, I tested the *solidarity hypothesis* using responses from cisgender respondents, with transgender people as the relevant outgroup. First, I estimated a multi-group model—regressing collective action intentions on positive and negative contact with transgender people—to compare responses from heterosexual ($n = 1050$) and non-heterosexual

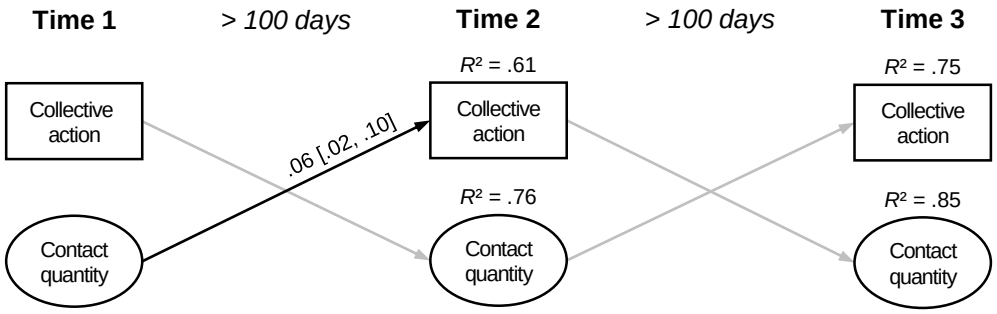

Figure 5.9. Panel model for heterosexual participants (Study 3c), with standardised coefficients and 95% confidence intervals. Paths in grey were not significant, $p > .05$. Autoregressive paths are not shown.

(n = 677) participants. I could not, however, fit the same factor structure to both groups. [7] Without configural invariance, I estimated separate models for heterosexual and non-heterosexual participants.

Only 31% of heterosexual participants reported *any* contact with transgender people. To avoid the resulting collinearity of positive and negative contact, I instead regressed collective action intentions on contact quantity, $\chi^2(46) = 36.15$, $\chi^2/df = 0.79$, $CFI = 1.00$, $TLI = 1.01$, $RMSEA = .00$. As hypothesised, contact quantity (at T1) predicted more collective action (at T2), though the effect was very small, $\beta_{T_1T_2} = 0.06$. None of the other forward or reverse paths was significant (see Figure 5.9). Between 53% and 57% of non-heterosexual respondents reported at least some contact with transgender people. As planned, I regressed collective action intentions on positive and negative contact, $\chi^2(161) = 262.23$, $\chi^2/df = 1.63$, $CFI = .97$, $TLI = .96$, $RMSEA = .03$. As hypothesised, positive contact predicted collective action across both $3\frac{1}{4}$ months intervals, $\beta_{T_1T_2} = 0.19$ and $\beta_{T_2T_3} = 0.26$. Negative contact predicted less collective action from T2 to T3, $\beta_{T_2T_3} = -0.12$, but not from T1 to T2, $\beta_{T_1T_2} = -0.12$, $[-0.15, 0.06]$. None of the reverse paths was significant (see Figure 5.10). On the one hand, these results suggest that positive contact can foster solidarity between disadvantaged minorities. On the other hand, the results offer only mixed evidence on the role of contact in fostering solidarity from an advantaged majority to a disadvantaged minority.

Next, I tested to what extent more continuous and fluid understandings of gender mediated the observed relationships in the sexual-minority sample. To that end, I estimated a series of longitudinal models for positive and negative contact, three dimensions of continuity/fluidity, and collective action intentions. As in Study 3a, I opted for a two-wave panel model, instead of the more complex three-wave model. Models 0–2 established measurement invariance across the two time points (Table C.4). Model 3 estimated autoregressive paths as well as forward paths from positive and negative contact to col-

[7] For heterosexual participants, the estimated covariance matrix was not positive definite, that is, one or more of the variables could be re-expressed as a linear combination of the other variables. In this sample, the observed covariance matrix was most likely affected by collinearity between the latent variables as positive and negative contact had near-perfect correlations at each time point: $\phi_1 = .91$, $\phi_2 = .91$, and $\phi_3 = .92$. The reason for this collinearity was that most participants had not reported *any* contact (see main text), in which case both positive and negative contact were scored as 1 = *never*. For non-heterosexual respondents—who reported more contact overall—positive and negative contact were still correlated positively, but less highly: $\phi_1 = .54$, $\phi_2 = .47$, and $\phi_3 = .40$.

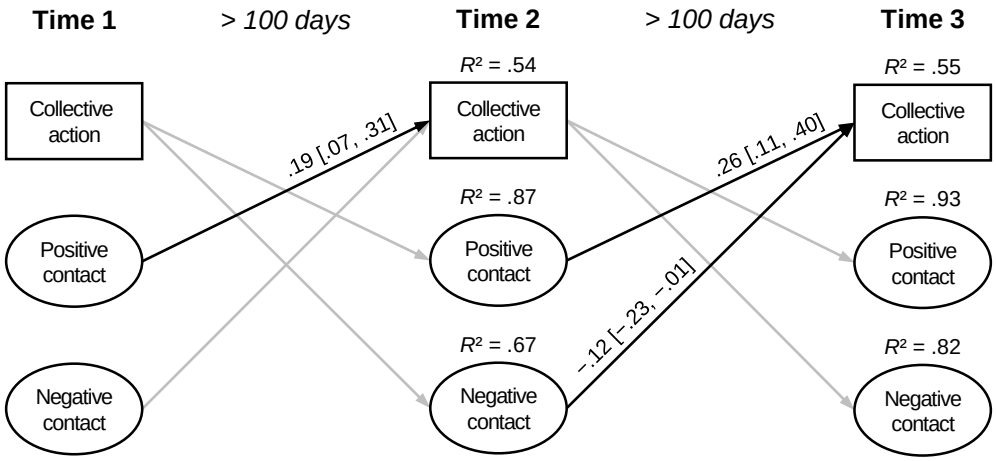

Figure 5.10. Panel model for non-heterosexual participants (Study 3c), with standardised coefficients and 95% confidence intervals. Negative and positive contact were correlated at T1 (ϕ = 0.55, [0.45, 0.66]). Paths in grey were not significant, $p > .05$. Autoregressive paths are not shown.

lective action (see Figure 5.10), but constrained all other paths to zero, $\chi^2(636) = 1056.46$, $\chi^2/df = 1.66$, $CFI = .96$, $TLI = .95$, $RMSEA = .03$. Model 4 added forward paths from positive and negative contact to the continuity/fluidity dimensions, $\Delta\chi^2(6) = 20.76$, $p = .002$, $\Delta CFI = .00$. Model 5 added forward paths from the continuity/fluidity dimensions to collective action, $\Delta\chi^2(3) = 2.92$, $p = .40$, $\Delta CFI = .00$. Model 6 added all reverse paths to Model 5, $\Delta\chi^2(11) = 2.92$, $p = .40$, $\Delta CFI = .00$. Model 4, but not Models 5 and 6, improved upon less complex models. Figure 5.11 shows standardised coefficients for all significant paths in Model 4. These findings do not support the hypothesis that more continuous and fluid conceptions of gender mediate the relationship between positive contact and inter-minority solidarity.

I then repeated these analyses for the sexual-majority sample. I estimated a series of longitudinal models for contact quantity, three dimensions of continuity/fluidity, and collective action intentions. As in Study 3b, I estimated a three-wave panel model. Models 0–2 established measurement invariance across the three time points (Table C.5). Model 3 estimated autoregressive paths as well as forward paths from contact quantity to collective action, but constrained all other paths to zero, $\chi^2(1048) = 2114.39$, $\chi^2/df = 2.02$, $CFI = .95$, $TLI = .95$, $RMSEA = .03$. Model 4 added forward paths from contact quantity to the continuity/fluidity dimensions, $\Delta\chi^2(6) = 8.27$, $p = .22$, $\Delta CFI = .00$. Model 5 added forward paths from the continuity/fluidity dimensions to collective action, $\Delta\chi^2(8.74) = 6$, $p = .19$, $\Delta CFI = .00$. Model 6 added all reverse paths to Model 5, $\Delta\chi^2(52.58) = 14$, $p < .001$, $\Delta CFI = .00$. Model 6, but not Models 4 and 5, fit better than less complex models. Model 7 estimated autoregressive paths, forward paths from from contact quantity to collective action, and all reverse paths from Model 6. Model 6 improved upon Model 3, $\Delta\chi^2(14) = 61.31$, $p < .001$, $\Delta CFI = .00$. Figure 5.12 shows standardised coefficients for all significant paths in Model 7. These findings do not support the hypothesis that more continuous and fluid conceptions of gender mediate the relationship between intergroup contact and majority-minority solidarity.

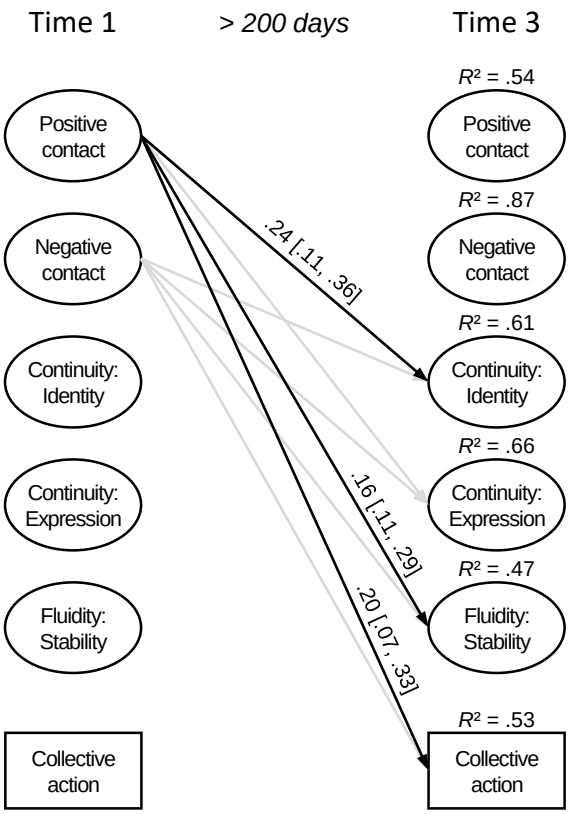

Figure 5.11. Results from Model 4 for non-heterosexual participants (Study 3c), with standardised coefficients and 95% confidence intervals. Paths in grey were not significant, $p > .05$. Autoregressive paths are not shown.

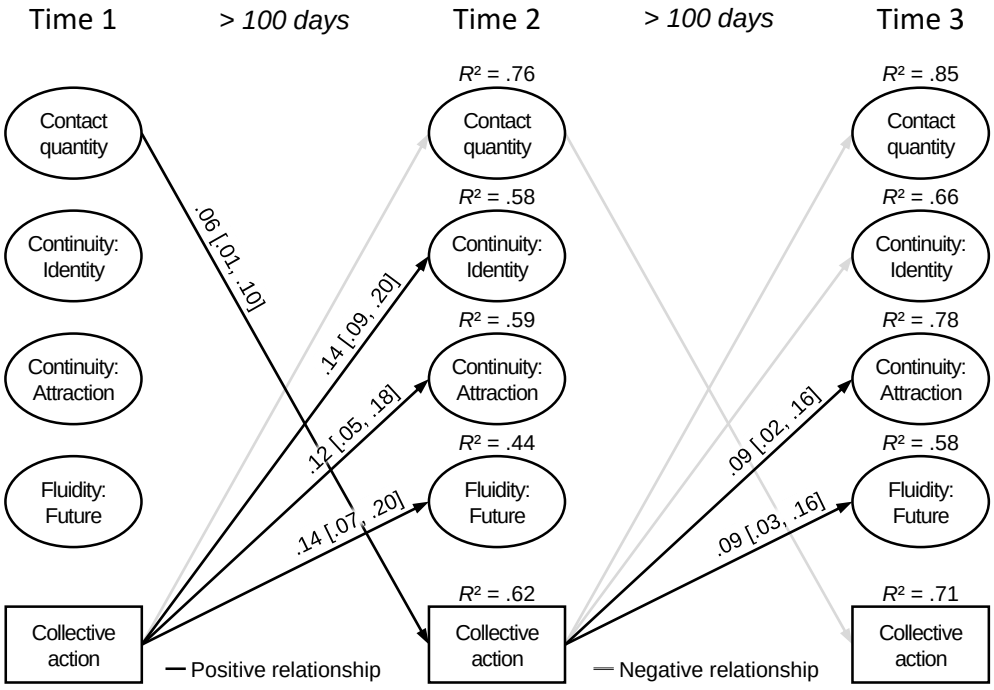

Figure 5.12. Results from Model 7 for heterosexual participants (Study 3c), with standardised coefficients and 95% confidence intervals. Paths in grey were not significant, $p > .05$. Autoregressive paths are not shown.

Study 4

Study 3 replicated findings from Studies 1a and 1b using a longitudinal design. Study 3 found that negative contact predicted more perceived discrimination, which, in turn, predicted more collective action among sexual-minority participants (*demobilisation hypothesis*). Contrary to prior research, positive contact with sexual-majority group members predicted neither perceived discrimination nor collective action (*mobilisation hypothesis*). Among sexual-majority participants, positive contact was associated with more solidarity-based college action for the rights of sexual minorities (*solidarity hypothesis*). Contrary to predictions, more continuous and fluid identities did not mediate any of the observed relationships.

By including transgender people as a relevant outgroup, Study 3 allowed testing whether contact between sexual-minority and gender-minority members fostered *inter-minority solidarity*. As expected, contact with transgender people predicted more solidarity-based collective action among non-heterosexual, cisgender people (but less so among heterosexual, cisgender people). Study 4 studied the hypothesised relationships among South Indian students of different caste backgrounds. Scheduled Caste (SC) and Scheduled Tribe (ST) students come from historically disadvantaged backgrounds, while General Merit (GM) students come from advantaged backgrounds. Other Backward Class (OBC) students occupy an intermediate position in the caste hierarchy. As Study 3, Study 4 thus allowed studying solidarity among two disadvantaged groups. Other than Study 3, Study 4 focused on groups who neither overlap nor share a common political movement, who fall along the same hierarchy, and who compete for resources (see Jodhka,

2012). Sexual and gender minorities do not directly compete for resources,[8] fall in two orthogonal hierarchies (sexuality, gender), share a political movement, and overlap to some extent.

As such, Study 4 aligns with other studies of inter-minority solidarity. Craig and Richeson (2012) studied Black–Latino relations in the United States. As discussed in Chapter 4, the researchers investigated *when* members of one disadvantaged group form a coalition with, rather than derogate, members of another disadvantaged group. They found that making discrimination against the ingroup salient decreased derogation of the disadvantaged outgroup. The researchers explained this finding as disadvantaged-group members forming a common identity as 'disadvantaged racial minority'. Dixon et al. (2015) studied Black–Indian relations in South Africa. The researchers found that, among Indian South Africans, positive contact with Black South Africans was associated with more perceived discrimination against Black South Africans and greater support for social change. In line with these findings, Study 4 tested whether support for reservations for another disadvantaged group was associated with including members of that group in one's ingroup, and with positive contact with members of that group.

In addition, Study 4 tested the *(de)mobilisation* and *solidarity* hypotheses in a novel setting. If intergroup contact fosters support for social change in advantaged groups, then contact with disadvantaged ethnoreligious and caste outgroups should increase support for affirmative action benefiting these groups. If positive contact demobilises the disadvantaged, then positive contact with advantaged caste outgroups should discourage support for reservations among disadvantaged caste groups. If negative contact mobilises the disadvantaged, then negative contact with advantaged caste outgroups should instead encourage support for reservations. In line with the *solidarity-by-inclusion* hypothesis, Study 4 tested whether including a group in one's ingroup increased support for reservations benefiting that group. In line with the *(de)mobilisation* and *solidarity-by-perceived-inequality* hypotheses, Study 4 tested to what extent positive and negative contact affected perceptions of relative life difficulties.

Method

Participants, procedure, and measures have been described in Chapter 4. As in Chapter 4, I report results from 302 participants, of whom 99 (33%) chose General Merit (GM), 127 (42%) chose Other Backward Class (OBC), and 76 (25%) chose Scheduled Caste (SC, n = 54) or Scheduled Tribe (ST, n = 22) as their caste group. Most participants chose Hinduism as their or their family's religion (n = 286, 95%); none of the participants identified as Muslim. Below, I describe additional measures that have not been presented in Chapter 4.

Measures

Perceived life difficulty was measured as how easy or hard participants thought it was, on average, for people from various groups to succeed in India today (1 = *very hard*, 7 = *very easy*). Participants rated seven groups: People from your own background, Scheduled Caste, Scheduled Tribe, Other Backward Class, General Caste, Hindus, and Muslims. Participants were reassured that groups were not necessarily mutually exclusive.

[8] Advancing transgender rights, for example, does not harm sexual minorities' rights, or vice versa. Extending reservation quotas for one caste groups, however, could mean reducing reservation for another.

Policy support was measured with five items. First, participants read that "[c]urrently, 22.5% of seats in central-government funded universities are reserved for Scheduled Caste (SC) and Scheduled Tribe (ST) students". Participants indicated to what extent they opposed or supported reservation for Scheduled Caste / Scheduled Tribe students in higher education (1 = *strongly oppose*, 5 = *strongly support*), and whether they thought that reservation for Scheduled Caste / Scheduled Tribe students in higher education should increase, decrease, or remain unchanged (1 = *decrease a lot*, 5 = *increase a lot*; r = .67). Second, participants read that "[a]n additional 27.5% of seats in central-government funded universities are reserved for students from Other Backward Classes (OBC)". Participants filled in the same two items but for Other Backward Class students (r = .67). Finally, participants read that "[n]o seats in central-government funded universities are reserved for Muslim students nationally, though some states have introduced quotas for Muslim students". Participants indicated to what extent they opposed or supported reservation for Muslim students in higher education.

Results

I used the following analytic strategy: First, I tested for *group differences* in participants' perceptions of life difficulty and in their support for affirmative action. Second, I analysed to what extent contact experiences and ingroup construals explained *individual differences* in participants' responses. I examined participants' responses to measures concerning outgroup members to test the solidarity hypotheses. I examined participants' responses to measures concerning their ingroup to test the (de)mobilisation hypotheses.

Group differences

I investigated how participants' caste memberships shaped their perceptions of life difficulty and their support for reservations. To that end, I estimated two Bayesian multivariate models in *RStan* (Stan Development Team, 2018) with participants' responses to six perceived life difficulty and three policy support measures as outcome variables.[9] Both models estimated how participant j responded to each of the nine outcomes. Models derived the likelihood of the observed responses from a multivariate normal distribution:

$$\mathbf{Y} \sim \text{MVNormal}(\mu_{jk}, \mathbf{S})$$

where $\mathbf{Y}$ is the matrix of observed variable scores (columns) across participants (rows), μ_{jk} is the estimated response to item k by participant j, and $\mathbf{S}$ is the variance-covariance matrix across variables. This means that models estimated distinct means and residual variances for each outcome, as well as residual correlations between all outcomes. Models assigned weakly informative priors to all intercepts, coefficients, and residual variances:

$$\beta \sim \text{Normal}(0, 1)$$
$$\sigma \sim \text{Cauchy}(0, 1)$$

allowing a wide range of plausible effect sizes. Model 0 estimated distinct means for each outcome, but did not consider participants' group memberships: $\mu_{jk} = \beta_k$, where μ_{jk}

[9] For this analysis, I aggregated items so that target groups corresponded across the two outcome measures. For perceived life difficulty, I aggregated items for Scheduled Caste and Scheduled Tribe groups, resulting in six variables. For policy support, I aggregated items for Scheduled Castes / Scheduled Tribes reservations, and for Other Backward Classes reservations, resulting in three variables.

equals β_k, the fixed intercept for outcome k. Model 1 estimated distinct means across participants' ingroups: $\mu_{jk} = \beta_k + \beta_{k,\text{OBC}} + \beta_{k,\text{SCST}}$, where μ_{jk} equals β_k, the fixed intercept for outcome k, plus $\beta_{k,\text{OBC}}$ and $\beta_{k,\text{SCST}}$, the coefficients estimating by how much Other Backward Class and Scheduled Caste / Scheduled Tribe participants' responses differed from General Merit participants'. Model 1 made better out-of-sample predictions than Model 0, $\Delta ELPD = 94.7$, $SE = 14.1$, $\frac{\Delta ELPD}{SE} = 6.7$. This comparison shows that participants' caste memberships shaped their perceived discrimination and policy support ratings.

Figure 5.13 shows estimated mean ratings of perceived life difficulty for all target groups as a function of participants' caste memberships in Model 1. For most target groups, General Merit and Other Backward Class participants' responses closely resembled each other. On average, General Merit / Other Backward Class participants thought that it was somewhat hard (below the midpoint) for "people from [their] own background" to succeed. General Merit / Other Backward Class participants agreed with each other that it was a lot easier for people from Scheduled Caste / Scheduled Tribe backgrounds to succeed, compared to people from their own background. General Merit / Other Backward Class participants also agreed that, of all target groups, life was most difficult for people from General Merit backgrounds. Other Backward Class participants disagreed with General Merit participants when rating life difficulty for people from Other Backward Class backgrounds. General Merit participants rated life difficulty for Other Backward Class people as somewhat easy, while Other Backward Class participants tended toward "neither easy nor hard" and "somewhat hard". Scheduled Caste / Scheduled Tribe participants' responses differed markedly from General Merit and Other Backward Class participants'. For all target groups, Scheduled Caste / Scheduled Tribe participants rated life difficulty as "neither easy nor hard", on average. General Merit, Other Backward Class, and Scheduled Caste / Scheduled Tribe participants all tended to rate life difficulty for Muslims as around the midpoint. Overall, participants' responses seemed to contradict, or even reverse, persisting social inequalities (see Chapter 4, *Caste, religion, and nation in South India*).

Figure 5.14 shows estimated policy support—that is, support for reservations benefiting Scheduled Caste / Scheduled Tribe, Other Backward Class, and Muslim students—as a function of participants' caste memberships in Model 1. General Merit and Other Backward Class participants' responses did not provide evidence for cross-group solidarity. General Merit participants tended to oppose reservations for Scheduled Caste / Scheduled Tribe ($M = 1.87$, $[1.63, 2.12]$), Other Backward Class ($M = 2.70$, $[2.44, 2.96]$), and Muslim students ($M = 2.75$, $[2.46, 3.01]$). Other Backward Class participants tended to support reservations for their own group ($M = 3.72$, $[3.50, 3.95]$), but not for Scheduled Caste / Scheduled Tribe ($M = 2.31$, $[2.10, 2.53]$) or Muslim ($M = 2.79$, $[2.55, 3.04]$) students. Scheduled Caste / Scheduled Tribe participants tended to support reservations for their own group ($M = 4.16$, $[3.88, 4.45]$), for Other Backward Class students ($M = 3.57$, $[3.27, 3.86]$), and, to a lesser extent, for Muslim students ($M = 3.24$, $[2.92, 3.57]$). Overall, participants' responses provided evidence for solidarity from Scheduled Caste / Scheduled Tribe participants to Other Backward Class students, but not for solidarity from General Merit and Other Backward Class participants to Scheduled Caste / Scheduled Tribe students.

Table 5.3 shows the estimated residual correlations before (Model 0) and after (Model 1) accounting for group differences. Perceiving a group to be disadvantaged correlated weakly with more support for reservations benefiting Scheduled Caste / Scheduled Tribe ($r = -.12$, $[-.24, .01]$), Other Backward Class ($r = -.17$, $[-.30, -.05]$), and Muslim ($r = -.15$, $[-.28, -.03]$) students in Model 1. Support for reservations for one

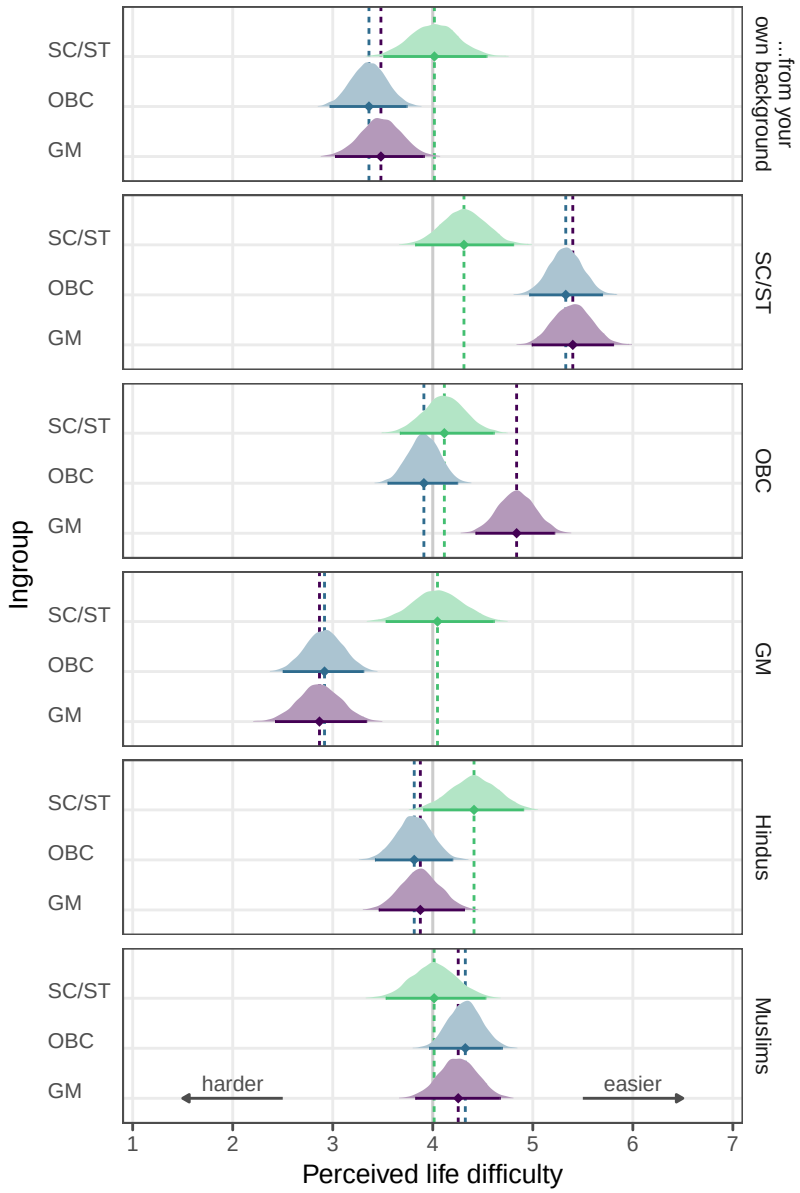

Figure 5.13. Posterior probabilities of mean perceived life difficulty ratings for different target groups (right) by participants' caste ingroup (left) in Model 1. Diamonds mark the most likely estimate; intervals encompass the 97% most likely estimates.

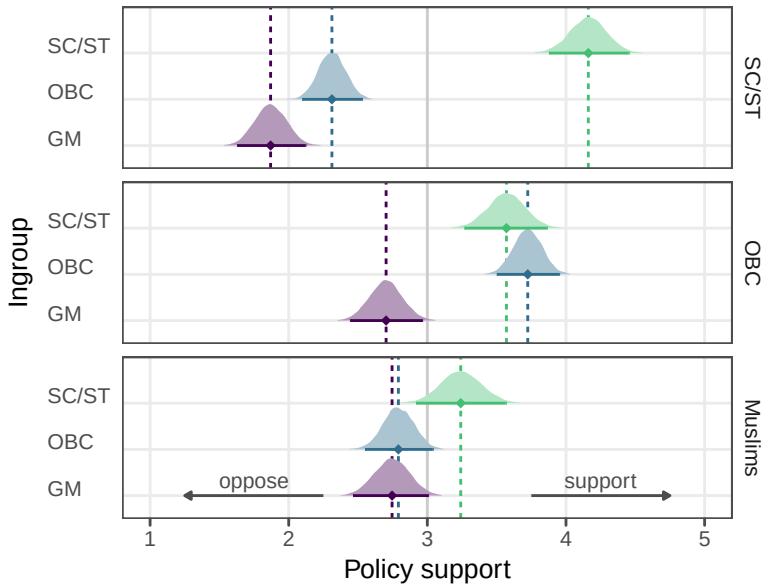

Figure 5.14. Posterior probabilities of mean policy support ratings for different target groups (right) by participants' caste ingroup (left) in Model 1. Diamonds mark the most likely estimate; intervals encompass the 97% most likely estimates.

Table 5.3. Estimated residual correlations between six *perceived life difficulty* (PD) ratings and three *policy support* (PS) ratings in Model 0 (below the diagonal) and Model 1 (above the diagonal) for Study 4.

#	Variable	1	2	3	4	5	6	7	8	9
1	PD, …from own background		.14	.30	.45	.43	.05	.28	.12	.14
2	PD, SC/ST	.10		.33	-.13	-.09	.28	-.12	-.10	-.05
3	PD, OBC	.28	.33		.20	.16	.32	.05	-.17	-.13
4	PD, GM	.47	-.18	.18		.61	-.08	.24	.15	.02
5	PD, Hindu	.44	-.12	.15	.62		-.01	.29	.11	.03
6	PD, Muslim	.04	.28	.31	-.10	-.02		.02	-.17	-.15
7	PS, SC/ST	.29	-.24	-.03	.33	.30	-.03		.24	.34
8	PS, OBC	.11	-.13	-.25	.17	.12	-.16	.28		.47
9	PS, Muslim	.16	-.09	-.13	.06	.05	-.16	.36	.45	

target group correlated moderately with support for reservations for the other target groups, $.24 < r < .47$. These correlations provided evidence for systematic variation in policy support ratings across participants, even after accounting for group differences. In the next section, I examine to what extent contact experiences and ingroup construals explain this variation across individuals.

Individual differences

I examined to what extent contact experiences and ingroup construals (measured as participants' responses in the crossed-categorisation task) accounted for individual differences in perceived life difficulty and policy support. To that end, I estimated a series of Bayesian multilevel models with either perceived life difficulty or policy support as outcome variable. Models derived the likelihood of the observed responses from a normal distribution:

$$y_{ij} \sim \text{Normal}(\mu_{ij}, \sigma)$$

where y_{ij} is the observed response to item i by participant j, and μ_{ij} is the estimated response to item i by participant j, σ is the estimated residual variance (expressed as standard deviation). I used the same weakly informative priors as for the multivariate models.

First, I estimated participants' responses to each of five perceived life difficulty items, with General Caste (GM), Scheduled Caste (SC), Scheduled Tribe (ST), Other Backward Class (OBC), and Muslims as target groups. Model 0 estimated item responses as varying across participants and target groups, and estimated differences across participants' ingroups:

$$\mu_{ij} = \beta_j + \beta_{kl}$$

where μ_{ij}, the estimated mean response to item i by participant j, equals β_j, the varying intercept for participant j, and β_{kl}, the varying intercept for target group k (1 = GM, 2 = OBC, 3 = SC/ST, 4 = Muslims) and ingroup l (1 = GM, 2 = OBC, 2 = SC/ST). Model 0 thus resembled the multivariate model estimated above.

Model 1 extended Model 0 by estimating to what extent outgroup friendship and negative contact with outgroup members were associated with perceptions of life difficulty for outgroups:

$$\mu_{ij} = \beta_j + \beta_{kl} + \beta_{\text{OF}} x_{\text{OG},k} x_{\text{OF,jk}} + \beta_{\text{NC}} x_{\text{OG},jk} x_{\text{OF,jk}}$$

where $x_{\text{OG},jk}$ codes whether target group k is an outgroup [1] or ingroup [0] for participant j, $x_{\text{OF,jk}}$ and $x_{\text{NC,jk}}$ are, respectively, participant j's outgroup friendship and negative contact with target group k, and β_{OF} and β_{NC} estimate the increase in participants' ratings for each additional standard deviation of reported contact. Model 2 estimated the effects of outgroup friendship and negative contact on outgroup ratings as varying across target outgroups:

$$\mu_{ij} = \beta_j + \beta_{kl} + \beta_{\text{OF},k} x_{\text{OG},k} x_{\text{OF,jk}} + \beta_{\text{NC},k} x_{\text{OG},jk} x_{\text{OF,jk}}$$

where $\beta_{\text{OF},k}$ and $\beta_{\text{NC},k}$ estimate effects of contact on ratings for outgroup k. Neither Model 1 nor Model 2 improved upon the out-of-sample predictions of Model 0, showing that intergroup contact neither increased nor decreased participants' perceptions of discrimination against ethnoreligious or caste outgroups (see Table 5.4).

Models 3 and 4 tested whether participants' ingroup construals were associated with their perceptions of life difficulty for various outgroups. Model 3 estimated this

relationship as constant across target groups:

$$\mu_{ij} = \beta_j + \beta_{kl} + \beta_{\text{Q1}} x_{\text{OG},k} x_{\text{Q1,jk}}$$

where $x_{Q1,jk}$ is the proportion [0–1] of targets belonging to outgroup k that participant j categorised as "us", and β_{Q1} estimates the difference in perceived life difficulty between participants who categorised all four targets of an outgroup as "us" and participants who categorised all targets of an outgroup as "not us". Model 4 estimated this relationship as varying across outgroups:

$$\mu_{ij} = \beta_j + \beta_{kl} + \beta_{\text{Q1,k}} x_{\text{OG},k} x_{\text{Q1,jk}}$$

where $\beta_{Q1,k}$ estimates the difference in perceived perceived life difficulty between participants who categorised all four targets of outgroup k as "us" and participants who categorised all targets of outgroup k as "not us". Neither Model 3 nor Model 4 improved upon the out-of-sample predictions of Model 0, showing that more inclusive identities neither increased nor decreased participants' perceptions of life difficulty for ethnoreligious an caste outgroups.

Models 5 to 8 tested to what extent contact with (Models 5, 6) and inclusion of advantaged-group members (Models 7, 8) was associated with perceptions of life difficulty for the ingroup. Model 5 estimated how contact with General Merit people affected Scheduled Caste / Scheduled Tribe participants' perceptions of life difficulty for their ingroup:

$$\mu_{ij} = \beta_j + \beta_{kl} + \beta_{\text{OF,GM}} x_{\text{IG,jk}} x_{\text{SCST,j}} x_{\text{OF,GM,j}} + \beta_{\text{NC,GM}} x_{\text{IG,jk}} x_{\text{SCST,j}} x_{\text{NC,GM,j}}$$

where $x_{\text{IG},jk}$ codes whether target group k is an ingroup [1] or outgroup [0] for participant j, $x_{\text{SCST,j}}$ codes whether participant j is from an Scheduled Caste / Scheduled Tribe background [1] or not [0], $x_{\text{OF,j,GM}}$ and $x_{\text{NC,j,GM}}$ are, respectively, participant j's outgroup friendship and negative contact with General Merit people, and $\beta_{\text{OF,GM}}$ and $\beta_{\text{NC,GM}}$ estimate the increase in participants' ratings for each additional standard deviation of reported contact with General Merit people. Model 6 extended Model 5 by estimating how contact with General Merit people affected Other Backward Class participants' perceptions of life difficulty for their ingroup:

$$\mu_{ij} = \beta_j + \beta_{kl} + \ldots + \beta_{\text{OF,GM}} x_{\text{IG,jk}} x_{\text{OBC,j}} x_{\text{OF,GM,j}} + \beta_{\text{NC,GM}} x_{\text{IG,jk}} x_{\text{OBC,j}} x_{\text{NC,GM,j}}$$

where $x_{\text{OBC,j}}$ codes whether participant j is from an Other Backward Class background [1] or not [0]. Neither Model 5 nor Model 6 improved upon the out-of-sample predictions of Model 0, showing that intergroup contact with advantaged-group members did not increase or decrease disadvantaged-group members' perceptions of life difficulty for their ingroup. Model 7 estimated whether including General Merit targets in their ingroup reduced perceptions of life difficulty among Scheduled Caste / Scheduled Tribe participants:

$$\mu_{ij} = \beta_j + \beta_{kl} + \beta_{\text{Q1,GM}} x_{\text{IG,jk}} x_{\text{SCST,j}} x_{\text{Q1,GM,j}}$$

where $x_{Q1,GM,j}$ is the proportion [0–1] of General Merit targets that participant j categorised as "us", and $\beta_{Q1,GM}$ estimates the difference in perceived life difficulty between participants who categorised all four General Merit targets as "us" and participants who categorised all General Merit targets as "not us". Model 8 extended Model 7 by estimating this relationship among Other Backward Class participants:

$$\mu_{ij} = \beta_j + \beta_{kl} + \beta_{\text{Q1,GM}} x_{\text{IG,jk}} x_{\text{SCST,j}} x_{\text{Q1,GM,j}} + \beta_{\text{Q1,GM}} x_{\text{IG,jk}} x_{\text{OBC,j}} x_{\text{Q1,GM,j}}$$

Table 5.4. Comparison of models estimating participants' group-wise perceived life difficulty ratings. R^2 is a Bayesian analogue to R^2 in maximum likelihood estimation (A. Gelman, Goodrich et al., 2017). *ELPD* is the expected log predictive density, with higher numbers indicating that a model is expected to make more accurate out-of-sample predictions (Vehtari et al., 2017). Δ*ELPD* is the difference in *ELPD* between the current and previous model, with positive values indicating that the current model is expected to make more accurate out-of-sample predictions. I preferred a more complex model over a simpler model when $\frac{\Delta ELPD}{SE} \geq 2$. *w* are stacking weights based on the combined leave-one-out predictive density (Yao et al., 2018).

#	Description	R^2	*ELPD*	*SE*	Δ*ELPD*	*SE*	$\frac{\Delta ELPD}{SE}$	*w*
0	*Group differences*	.35	-1789.5	24.6	-	-	-	.00
1 vs 0	*Solidarity hypothesis*	.35	-1791.2	24.5	-1.8	1.5	-1.2	.00
2 vs 1	*Solidarity hypothesis*	.35	-1788.3	24.4	3.0	2.5	1.2	.33
3 vs 0	*Solidarity hypothesis*	.35	-1787.0	24.7	2.5	2.4	1.0	.32
4 vs 3	*Solidarity hypothesis*	.35	-1787.6	24.7	-0.7	2.4	-0.3	.23
5 vs 0	*(De)mobilisation hypotheses*	.35	-1791.5	24.6	-2.0	1.2	-1.7	.00
6 vs 5	*(De)mobilisation hypotheses*	.35	-1792.6	24.6	-1.1	0.9	-1.2	.00
7 vs 0	*(De)mobilisation hypotheses*	.35	-1788.9	24.5	0.5	1.4	0.4	.00
8 vs 7	*(De)mobilisation hypotheses*	.35	-1788.4	24.5	0.5	1.4	0.4	.12

where $x_{\text{OBC,j}}$ codes whether participant *j* is from an Other Backward Class background [1] or not [0]. Neither Model 7 nor Model 8 improved upon Model 0, showing that including advantaged-group members in the ingroup did not increase or decrease disadvantaged-group members' perceptions of their ingroup's life difficulty (see Table 5.4).

Second, I estimated participants' responses to each of five policy support items, with Scheduled Caste and Scheduled Tribe (SC/ST), Other Backward Class (OBC), and Muslims as target groups. Model 0 estimated item responses as varying across participants and target groups, and estimated differences across participants' ingroups:

$$\mu_{ij} = \beta_j + \beta_{kl}$$

where μ_{ij}, the estimated mean response to item *i* by participant *j*, equals β_j, the varying intercept for participant *j*, plus β_{kl}, the varying intercept for target group *k* (1 = SC/ST, 2 = OBC, 3 = Muslims) and ingroup *l* (1 = GM, 2 = OBC, 2 = SC/ST). Model 0 thus resembled the multivariate model estimated in the previous section. Models 1 to 8 were identical to the models for perceived discrimination (described above). Models 1 to 4 did not improve upon the out-of-sample predictions of Model 0, showing that neither intergroup contact (Models 1, 2) nor inclusive identities (Models 3, 4) foster support for policies benefiting a disadvantaged outgroup. Models 5 to 8 did not improve upon the out-of-sample predictions of Model 0, showing that neither contact with (Models 5, 6) nor inclusion of (Models 7, 8) General Merit people affected Scheduled Caste / Scheduled Tribe and Other Backward Class participants' support for policies benefiting their ingroups (see Table 5.5).

Overall, these findings show that neither intergroup contact nor ingroup construals could explain individual differences in support for social change. Contrary to the *demobilisation hypothesis*, positive contact did not decrease disadvantaged-group

Table 5.5. Comparison of models estimating participants' group-wise policy support ratings. As in Table 5.4, I preferred a more complex model over a simpler model when $\frac{\Delta ELPD}{SE} \geq 2$.

#	Description	R^2	*ELPD*	*SE*	$\Delta ELPD$	*SE*	$\frac{\Delta ELPD}{SE}$	*w*
0	*Group differences*	.49	-1741.0	26.8	-	-	-	.00
1 vs 0	*Solidarity hypothesis*	.50	-1736.1	26.7	4.9	2.7	1.8	.60
2 vs 1	*Solidarity hypothesis*	.50	-1739.9	26.8	-3.8	2.0	-1.9	.00
3 vs 0	*Solidarity hypothesis*	.49	-1741.0	26.9	0.0	0.7	0.0	.00
4 vs 3	*Solidarity hypothesis*	.49	-1741.6	26.9	-0.6	1.1	-0.5	.00
5 vs 0	*(De)mobilisation hypotheses*	.49	-1742.0	26.9	-0.9	1.7	-0.5	.03
6 vs 5	*(De)mobilisation hypotheses*	.49	-1743.3	26.9	-1.3	0.8	-1.6	.00
7 vs 0	*(De)mobilisation hypotheses*	.49	-1741.6	26.8	-0.6	0.8	-0.8	.00
8 vs 7	*(De)mobilisation hypotheses*	.50	-1740.0	26.8	1.6	2.0	0.8	.37

participants' support for social change. Contrary to the *mobilisation hypothesis*, negative contact did not increase disadvantaged-group participants' support for social change. Contrary to predictions, neither positive contact (*solidarity hypothesis*) nor more inclusive identities *(solidarity-by-inclusion hypothesis*) were associated with support for policies benefiting (other) disadvantaged outgroups.

Discussion

In this chapter, I have examined whether positive and negative contact increase or decrease support for social change in advantaged and disadvantaged groups. Studies 1a and 3a found that for disadvantaged-group members, negative contact with the advantaged group was associated with more perceived discrimination and, in turn, more collective action. Contrary to prior research, positive contact was not associated with perceived discrimination or collective action. Studies 1b and 3b found that for advantaged-group members, positive contact was associated with more solidarity-based collective action. Contrary to predictions, more continuous and fluid identities did not mediate any of these relationships. Study 4 found neither intergroup contact nor more inclusive identities to be associated with support for or opposition to social change. Below, I discuss whether, in light of the present findings, intergroup contact helps or hinders social change.

A darker side of 'we'?

Studies 1a and 3a investigated whether positive contact with an advantaged group can hinder social change by discouraging disadvantaged-group members from engaging in collective action (Dixon et al., 2012; Wright & Baray, 2012). Contrary to a growing body of cross-sectional (e.g., Cakal et al., 2011), longitudinal (Tropp et al., 2012), and experimental (Becker et al., 2013) research, positive contact was not associated with (less) collective action among sexual-minority participants. Rather than contradict previous findings, the present research could change their interpretation. Thus far, researchers interpreted their observations as support for the *demobilisation hypothesis*, the idea that positive contact undermines disadvantaged-group members' motivation for collective action. In contrast,

Study 1a found positive contact to be negatively related to perceived discrimination *only when* negative contact was not included in the analysis. As both dimensions of contact were, at least in the minority-group samples, negatively correlated, prior correlational evidence for the demobilising effects of positive contact could have conflated the presence of positive contact with the absence of negative contact.

Dixon et al. (2012) rightly stressed that discrimination manifests itself in subtle ways beyond hostile encounters with dominant outgroups. In the absence of direct hostilities, they argued, (positive) contact can have system-justifying consequences (Jost, Banaji & Nosek, 2004) by lulling the oppressed into a false sense of equality. Our research, however, suggests that direct confrontations with advantaged-group members make discrimination visible, and that it is the absence of such experiences that impedes collective action (*mobilisation hypothesis*). Relatedly, Friedman and Leaper (2010) found that how often sexual-minority participants had experienced discrimination first-hand was an important correlate of collective action. Likewise, Poore et al.'s (2002) research in an Inuit community suggested that isolation not contact breeds ignorance about systemic discrimination. In line with this converging evidence, the present findings call on researchers to re-conceptualise the problem of demobilisation as the absence of (negative) contact experiences with advantaged-group members.

Study 4 examined whether positive and negative contact, respectively, decreases and increases support for policies benefiting one's disadvantaged ingroup. Contrary to the *demobilisation hypothesis*, positive contact was not associated with (less) support for reservation policies. Contrary to the *mobilisation hypothesis*, negative contact was not associated with (more) support for reservation policies. Study 4 thus contradicted both prior research (e.g., Cakal et al., 2011, Study 1) and conclusions drawn from Studies 1a and 3a. I suggest that features of the participants' situation may explain this discrepancy. First, Scheduled Caste / Scheduled Tribe and Other Backward Class students have personally experienced how reservation policies can help them to get admitted into state-funded universities. Unlike LGBT activism, the impact of reservation policies seems immediate and concrete rather than distant and vague. Second, Scheduled Caste / Scheduled Tribe and Other Backward Class students answered questions about perceived discrimination and policy support while being reminded of their reservation status. Unlike sexual orientation, reservation status is an administrative category *defined* by historical inequalities. Together, these features make for a "strong situation" (Mischel, 1977) in which situational features limit the influence of individual differences. This may explain why individual differences in contact experiences did not relate to support for reservation policies in Study 4.

In previous chapters, I have argued that positive contact can foster more continuous, fluid, and inclusive social identities. In this chapter, I have examined how social identification relates to support for social change in disadvantaged groups. Study 1a, but not Study 3a, found strength of identification to be associated with more collective action among sexual-minority participants (see van Zomeren et al., 2008). Contrary to predictions, intergroup contact was not consistently associated with social identification in these studies. Dovidio et al. (2012) argued that more inclusive identities could discourage disadvantaged-group members from supporting social change. Contrary to this prediction, Study 4 found more inclusive identities to be unrelated to perceived discrimination and policy support among students from disadvantaged backgrounds. As reviewed in Chapter 1, Dovidio et al. (2012) argued that common ingroup identities, but not dual identities, would undermine support for social change among the disadvantaged

(see Ufkes et al., 2016). Study 4 measured whether Scheduled Caste / Scheduled Tribe and Other Backward Class students included people from General Merit backgrounds into their ingroup, but not whether more inclusive identities came at the expense of identifying with their ingroup. Thus, the present research does not directly contradict Dovidio et al.'s argument. However, Studies 1a, 3a, and 4 did not provide evidence for the hypothesised "darker side" of more inclusive identities.

A path to solidarity?

Studies 1b and 3b investigated whether positive contact with disadvantaged groups can help achieve social change by encouraging advantaged-group members to engage in solidarity-based collective action. In line with prior research (e.g., Fingerhut, 2011), positive contact was associated with majority-group members' conviction to defend and advance the rights of sexual minorities. Although more favourable attitudes partially explained this relationship, our findings (Study 1b) indicated that solidarity-based collective action is more than simply another expression of outgroup attitudes. Rather, the present research bolsters earlier findings that positive contact helps to close the so-called "principle-implementation gap" between dominant-group members' support for the principle of equality and their opposition to its implementation in policies such as affirmative action (Cakal et al., 2011; Dixon et al., 2007). Negative contact, on the contrary, correlated with less collective action in Study 1b (cross-sectional). Compared with positive contact, however, it was rare and its association with collective action weaker. Furthermore, negative contact did not predict collective action in Study 3b (longitudinal). This pattern of results suggests that the mobilising effects of positive contact outweigh the demobilising effects of negative contact for members of advantaged groups.

Study 4 examined whether cross-group friendship and negative contact, respectively, increases and decreases support for policies benefiting a disadvantaged outgroup. Contrary to the *solidarity hypothesis*, friendship with people from Scheduled Caste / Scheduled Tribe and Other Backward Class backgrounds was not associated with (more) support for reservation policies among General Merit participants. Likewise, negative contact was not associated with (less) support for reservation policies. Study 4 thus contradicted both prior research (e.g., Cakal et al., 2011, Study 2) and findings from Studies 1b and 3b. Again, I suggest that social context may explain this discrepancy. Like Scheduled Caste / Scheduled Tribe and Other Backward Class students, General Merit students have personally experienced the impact of reservation policies. Unlike Scheduled Caste / Scheduled Tribe and Other Backward Class students, General Merit students' experiences have likely been negative. That is, reservation policies made it more difficult for them to get admitted into state-funded universities. Echoing observations by Norton and Sommers (2011), this experience may explain why, in stark contrast to societal realities, General Merit students saw themselves at a disadvantage relative to Scheduled Caste / Scheduled Tribe and Other Backward Class students. Contact experiences did not seem to be sufficient to change this perception—perhaps because Scheduled Caste / Scheduled Tribe students whom General Merit students met at university were often from relatively well-off families.[10] In contrast, LGBT activism centres on civil rights rather than affirmative action and redistribution. Unlike General Merit students, heterosexual/cisgender

[10] Reservations for Scheduled Castes and Scheduled Tribes are not means tested. Reservations for Other Backward Classes exclude applicants from the so-called "creamy layer", that is, families with annual incomes above the equivalent of £8,752 ("OBC creamy layer income limit raised to ₹8 lakh per annum", 2017).

participants had thus little to lose by supporting LGBT advocacy. This may explain why contact experiences did not relate to support for social change in Study 4.

Studies 3c and 4 examined whether intergroup contact fostered solidarity between disadvantaged groups. The present research thus revisited the question whether members of one disadvantaged group derogate members of another disadvantaged group, or whether they instead strive for a coalition (Craig & Richeson, 2012). Study 3c found that, in line with the *solidarity hypothesis*, positive contact was associated with sexual-minority participants' conviction to defend and advance transgender rights. Study 4, however, found that positive contact with people from Scheduled Caste / Scheduled Tribe backgrounds was not associated with Other Backward Class students' support for reservations benefiting Scheduled Caste / Scheduled Tribe students. Indeed, Other Backward Class students broadly opposed reservations for disadvantaged groups other than their own. Scheduled Caste / Scheduled Tribe students tended to support reservations for their own and other caste groups—perhaps because Dalits have a long history of advocating for and benefiting from reservation policies.

Other Backward Class and Scheduled Caste / Scheduled Tribe may be at odds over reservation policies for two reasons. First, reservation quotas are a zero-sum game. That is, increasing one group's quota decreases the seats available to another. Thus far, increases in reservation quotas have always come at the expense of seats in the General Merit category—but that need not be the case in the future. Second, Other Backward Class and Scheduled Caste / Scheduled Tribe students are organised in different social movements (see Chapter 4). Sexual and gender minorities both benefit from challenging restrictive norms around sexuality and gender, and are often organised in a shared social movement under the LGBT umbrella. Together, these differences may explain why the present research found evidence for inter-minority solidarity between sexual and gender minorities (Study 3c), but not between disadvantaged caste groups (Study 4).

Studies 3b, 3c and 4 examined whether positive contact increases support for social change by fostering more continuous, fluid, and inclusive social identities. Other than expected, more continuous and fluid conceptions of sexuality (Study 3b) and gender (Study 3c) did not predict more solidarity-based collective action. Likewise, more inclusive identities were not associated with support for hierarchy-attenuating policies (Study 4). Instead, the present research found that solidarity-based collective action predicted more continuous and fluid identities. This could suggest that engaging with the LGBT movement changed how participants understood their own sexuality and gender. However, Study 1b provided (cross-sectional) evidence that positive contact could encourage advantaged-group members to not only *support* the disadvantaged-groups' movement but to *identify* with it. Movement identification was associated with more collective action. A cross-sectional study, however, cannot differentiate to what extent "identifying with the values and goals of the LGBT movement" measures a different construct than collective action intentions. Future research should examine whether positive contact predicts movement identification over time, and whether movement identification, in turn, predicts more collective action. Overall, Studies 3b, 3c, and 4 did not find that more continuous, fluid, and inclusive identities offered a "path to solidarity" for advantaged groups.

Strengths and limitations

Recent research (e.g., Barlow et al., 2012; Fell, 2015) has shown the importance of considering all types of contact experiences, both positive and negative, for understanding intergroup relations. Alongside earlier studies, the present research highlights the difference between the *reality* of (both positive and negative) contact and its rarefied *ideal*. I studied the former, and found that when considering contact as it occurs in everyday life, negative not positive interactions with majority-group members predicted perceived discrimination and collective action (Study 3a). Dixon et al.'s (2012) argument, however, speaks more to the latter, in particular, to the use of contact in planned interventions. In experimental research, cross-group interactions that were designed to be positive and to avoid the subject of inequality discouraged collective action among disadvantaged-group members (Becker et al., 2013). Manipulations that encouraged disadvantaged-group members to identify with a common ingroup rather than their ethnic/racial ingroup had a similar effect (Ufkes et al., 2016). These studies showed how well-meaning interventions could still have the ironic effect of dampening social change by emphasizing harmony at the expense of openly discussing and acknowledging inequality. Future research should carefully study how intervention programs aimed at reducing prejudice may unwittingly impact perceptions of discrimination and collective action.

Study 3 is one of only two longitudinal studies of the relationships between intergroup contact and collective action. Tropp et al. (2012), who conducted the other longitudinal study, did not measure negative contact, and found mixed results across three ethnic minorities in the United States. Both studies were, however, limited by their correlational design. That is, although longitudinal research can determine the direction of a relationship (e.g., whether contact predicts collective action, or whether collective action predicts contact), it cannot rule out that an unobserved covariate underlies the relationship. Further, the choice of samples limits to what extent the present findings generalise to other groups. Above, I have argued that past research may have conflated the presence of positive contact with the absence of negative contact. For this argument to hold, however, the present findings (focusing on sexual and gender minorities) need to apply to past research (focusing on race relations in South Africa and the United States). Findings from Study 4 highlight that findings from studies on sexual and gender minorities may not generalise to circumstances such as caste relations in South India. Future research should therefore seek to replicate the present findings by studying the effects of positive and negative contact on support for social change in different social contexts. This research would help us understand *when* intergroup contact helps or hinders social change.

Comparing the effects of negative and positive contact is further limited by the lack of a "common unit" of contact valence: being insulted is not necessarily as "bad" as being befriended is "good" (see Fell, 2015). Future research in this emerging field of research should explicate the relative sensitivity and predictive validity of measures of valenced contact. In a similar vein, the present research operationalised collective action as intentions to engage in various activities, for example, joining a protest or activist group. Although additional measures, that is, participants' pledges to donate money, had tangible consequences for their potential winnings, it is unclear to what extent these intentions predicted actual behaviour. Likewise, willingness to undertake all actions assessed required relatively little commitment as the present research focused on normative collective action. By inciting contempt, however, negative contact could

inspire actions beyond normatively accepted expressions of discontent (for a review, see Becker & Tausch, 2015). Future research should consider how negative contact relates to both normative and non-normative (actual) collective action. In the next chapter, I discuss, among other issues, the present research's broader implications for the potential tension between prejudice reduction and collective action approaches to social change.

6

Conclusion

Since Allport (1954/1979) first hypothesised that contact with members of other groups reduces prejudice, a growing literature has transformed this hypothesis into a detailed theory of *when* and *how* intergroup contact reduces prejudice (R. Brown & Hewstone, 2005; Hewstone & Swart, 2011). Extending this literature, recent research has raised questions such as how negative contact affects intergroup relations, and whether intergroup contact helps or hinders social change (for a recent review, see Dovidio, Love, Schellhaas & Hewstone, 2017). In this chapter, I discuss how the present research contributes to these ongoing debates in intergroup relations research.

This thesis has examined whether intergroup contact can foster more continuous, fluid, and inclusive social identities—and whether more continuous, fluid, and inclusive social identities can explain how intergroup contact reduces prejudice. Chapters 1 and 2 have introduced the relevant literature and methods. Chapter 3 has reported research developing a novel measure of social identity continuity and fluidity, and has examined whether contact with sexual and gender minorities fosters more continuous and fluid conceptions of sexuality and gender. Studies 1–3 found that intergroup contact was associated with more continuous and fluid social identities, but that more continuous and fluid social identities were not, in turn, associated with more favourable outgroup attitudes. Chapter 4 has reported research testing whether contact with caste or religious outgroups fosters more inclusive social identities in South India. Study 4 found that, as hypothesised, cross-group friendship was associated with more inclusive identities, while more inclusive identities were associated with more favourable outgroup attitudes and less social distance. As such, the present research provides insights on *how* intergroup contact reduces prejudice.

Throughout, this thesis has studied how both positive and negative contact relate to prejudice in various social settings. Study 1 found that among heterosexual participants, positive and negative contact was associated with, respectively, less and more sexual prejudice. Study 2 found that contact quantity and negative contact, but not positive contact, was associated with, respectively, less and more sexual prejudice. Study 2 also found that among cisgender participants, positive and negative contact with transgender people was associated with, respectively, less and more transphobia. Study 3 found that positive contact, but not negative contact, predicted (less) sexual prejudice over time. Study 3 also found that among cisgender participants, neither positive nor negative contact predicted transphobia over time. Study 4 found that cross-group friendship and negative contact were associated with more and less inclusive identities, and that more inclusive identities were associated with more favourable outgroup attitudes and less social distance.

Chapter 5 has reported additional results from these studies, examining how intergroup contact relates to outgroup attitudes among sexual-minority participants. Study 1 found that positive contact, but not negative contact, was associated with more favourable attitudes toward heterosexual people. Study 3 found that positive and negative contact predicted, respectively, more and less favourable attitudes toward heterosexual people over time. As such, the present research provides insights on *when* intergroup contact reduces prejudice.

Beyond prejudice, this thesis has examined whether intergroup contact discourages or encourages support for social change. Chapter 5 reports research testing how intergroup contact relates to support for social change in advantaged and disadvantaged groups. Studies 1 and 3 found that for disadvantaged-group members, negative contact, but not positive contact, was associated with greater collective action intentions. For advantaged-group members, positive contact predicted more solidarity-based collective action over time. More continuous and fluid identities, however, did not mediate any of these relationships. Study 4 found neither intergroup contact nor more inclusive identities to be associated with support for or opposition to social change. As such, the present research provides insight on whether intergroup contact *helps* or *hinders* social change.

Below, I discuss each of these contributions in turn—reviewing how the present research informs when intergroup contact reduces prejudice, how intergroup contact reduces prejudice, and whether intergroup contact helps or hinders social change. I conclude by discussing the present research in light of the ongoing reproducibility crisis in psychological research.

When does intergroup contact reduce prejudice?

In this section, I discuss what the present research contributes to understanding when intergroup contact decreases or increases prejudice. I review how the effects of intergroup contact differ across various social contexts, and how the effects of intergroup contact depend on its valence—that is, whether it is experienced as positive or negative.

Positive contact *sometimes* decreases prejudice

Since Allport (1954/1979), hundreds (or even thousands) of empirical studies have examined whether intergroup contact reduces prejudice. Several meta-analytic reviews have sought to summarise this vast literature. Pettigrew and Tropp (2006) found that, across 515 studies with 713 samples, intergroup contact was associated with less prejudice. Davies, Tropp, Aron, Pettigrew and Wright (2011) found that, across 135 studies with 208 samples, cross-group friendship was associated with more favourable intergroup attitudes. Lemmer and Wagner (2015) found that, across 79 studies with 129 samples, contact-based interventions reduced ethnic prejudice. Together, these studies provide overwhelming evidence that intergroup contact reduces prejudice.

Not all studies, however, find intergroup contact to reduce prejudice. In Pettigrew and Tropp's (2006, 2011) meta-analysis, for example, 4% of studies reported intergroup contact to be associated with more prejudice, while 2% of studies found no association between intergroup contact and prejudice.[1] Pettigrew and Tropp (2006) found the strength

[1]This is likely an underestimate, given that publication bias tends to conceal non-significant results, and given that unpublished studies tend to be difficult to find.

of the contact-attitude relationship to vary across social contexts: intergroup contact was most strongly associated with sexual prejudice, and least strongly associated with attitudes toward elderly or mentally ill people. Smith et al. (2009) found that, across 41 studies and 83 hypothesis tests, intergroup contact was associated with less sexual prejudice, though 8% of hypothesis tests were not significant. In line with previous studies, the present research consistently found that the more (positive) contact with sexual minorities heterosexual participants reported, the less sexual prejudice they showed. Whereas the relationship between contact with sexual minorities and sexual prejudice is well documented, other social contexts have been examined in far fewer studies. Below, I review how the present research contributes to our understanding of intergroup contact in two such social contexts.

Few studies have examined how contact with transgender people relates to anti-transgender prejudice. Recent research found some cross-sectional evidence that intergroup contact reduces anti-transgender prejudice. King et al. (2009) found, in a study of Hong Kong residents, that people who reported at least some contact with transgender people showed less anti-transgender prejudice than people who did not report any contact (N = 865). Hoffarth and Hodson (2018) found that among Canadians and US Americans, contact quality, but not contact quantity, were associated with less anti-transgender prejudice (N = 451). Adding to this research, this thesis has reported further cross-sectional evidence that positive contact with transgender people is associated with less anti-transgender prejudice (N = 309). In addition, this thesis reported the first longitudinal study of the relationship between intergroup contact and anti-transgender prejudice. Contrary to previous research, the longitudinal study did not find positive contact with transgender people to predict (less) prejudice over time (N = 1726). Below, I discuss possible explanations for this discrepancy in findings.

This discrepancy raises the question *why* the longitudinal study did not replicate earlier cross-sectional evidence that contact with transgender people reduces prejudice. It is possible that contact with transgender people does not reduce prejudice, and that the study's more rigorous longitudinal design and greater sample size better reflects that reality. This explanation might be plausible as contact with transgender people differs in other ways from many social contexts in which contact has been studied. First, being transgender is concealable and relatively rare, limiting opportunities for contact. I am not convinced, however, that this should make contact with transgender people *less* effective when it nonetheless occurs. Second, transgender people had, until recently, fairly limited public exposure. As such, many cisgender people might be unfamiliar with what it means to be transgender, and, as a result, might hold harmful misconceptions. Contact, however, tends to be more effective in reducing the affective dimensions of prejudice, rather than in dispelling its cognitive dimensions (Tropp & Pettigrew, 2005).[2] Third, as contact is rare and exposure is limited, it might be that someone's first contact with a transgender person matters more than continued contact over time. If the benefits of contact occur close in time, rather than unfold over time, researchers should expect to find contact effects in cross-sectional but less so in longitudinal studies.[3] These observations, however, are

[2]Relatedly, Pettigrew and Tropp (2008) showed that contact reduces prejudice primarily through affective processes, such as decreasing anxiety and increasing empathy, rather than cognitive processes, such as increasing knowledge (see also Swart et al., 2011).

[3]This argument applies to the present research in particular, as intersurvey intervals were relatively short and contact was relatively stable. Note that Pettigrew (1998) argued that optimal intergroup contact needs time to unfold.

idle speculation before we know more about anti-transgender prejudice. Future research should seek to further our understanding of what motivates anti-transgender attitudes, beliefs, and behaviours, and how we can overcome these prejudices.

Few studies have examined intergroup contact in the context of intergroup relations in South Asia. Islam and Hewstone (1993) found that among Bangladeshi Muslim and Hindu students, contact quality—and, for Hindus, contact quantity—was associated with more favourable outgroup attitudes (N = 131). Tausch, Hewstone and Roy (2009) found that among Indian Hindu and Muslim students, contact quantity was associated with less social distance, while contact quality was associated with less ingroup bias (N = 87). Both studies provide cross-sectional evidence that intergroup contact improves intercommunal relations in South Asia. Adding to this research, this thesis has reported further cross-sectional evidence that among Indian Hindu students, cross-group friendship is, via more inclusive identities, associated with more favourable attitudes toward and less social distance to Muslims (N = 302). Together, these studies suggest that intergroup contact could help to reduce intercommunal hostilities. This is an important finding as, in India, anti-Muslim violence has been on the rise in recent years (see Chapter 4). Integrated education, as studied in previous and present research, could be a means to undermine the attitudinal and ideological foundations of anti-Muslim violence in India.

In addition, the present research included, to my knowledge, the first investigation of contact between caste groups. I found that, among students of different caste backgrounds, cross-group friendship was (via more inclusive identities) associated with less social distance to and more favourable attitudes toward other caste groups. Remarkably, the effects of cross-group friendship did not differ between participants of advantaged and disadvantaged caste groups, and between caste and ethnoreligious outgroups. Despite progress, caste discrimination continues to affect the lives of disadvantaged caste communities (Jodhka, 2012; Pathania & Tierney, 2018; Shah et al., 2006). As such, intergroup contact offers a way to diminish interpersonal discrimination, especially as it reduces upper-caste students' desired social distance to students from ex-untouchable communities. Beyond interpersonal discrimination, structural discrimination continues to affect disadvantaged caste groups. Later on, I discuss why intergroup contact does not foster support for policies that mitigate these inequalities.

On a different note, this thesis has reported research that studied intergroup relations outside the Western, Educated, Industrialised, Rich, and Democratic (WEIRD; Henrich, Heine & Norenzayan, 2010) societies that dominate psychological research. On the one hand, such research allows social psychologists to test whether their theories generalise to other social contexts.[4] The present research adds to a growing body of research that shows intergroup contact to be effective in reducing social conflict across the globe. On the other hand, such research helps to dispel misconceptions about social contexts outside the WEIRD norm. Reading the psychological literature, I come across more researchers who use the caste system to make a theoretical argument than researchers who have studied the caste system empirically (with notable exceptions, see Cotterill et al., 2014). As such, researchers risk reproducing the Orientalist conception of the caste system which frames caste as a tradition that is, to some extent, consensual (for a discussion, see Jodhka, 2012).

Reflecting on both the literature reviewed in Chapter 4 and the research discussed above, I instead find the caste system to be much closer to more familiar hierarchies.

[4] Among such contexts, India is particularly important as it comprises more than a sixth of the world's population.

Racism in the United States, for example, shares many similarities with the caste system in India: Black Americans and Dalits both share a history of forced servitude and landlessness that manifests itself in enduring structural inequalities. Considering the similarities makes the differences between the two contexts more interesting. Other than the United States, India has an effective system of affirmative action to address these historical inequalities. Future research should seek to further understand support for and opposition to affirmative action in India. This research would further our understanding of the potential impact of future efforts to introduce or extend affirmative action in other social contexts.

Negative contact *sometimes* increases prejudice

As discussed, not all studies find that intergroup contact reduces prejudice (Pettigrew & Tropp, 2006, 2011). Aside from social context, contact valence—that is, whether contact is experienced as positive or negative—might explain when contact decreases or increases prejudice. Recent research has shown the importance of considering both positive and negative experiences when studying the consequences of intergroup contact (for a discussion, see Pettigrew & Hewstone, 2017). Adding to this literature, the present research measured both positive and negative contact in all four studies. Below, I discuss how the present research compares to the findings of other recent studies of negative contact.

Measuring both positive and negative contact, researchers have made four important observations. First, researchers found that, although less frequent than positive contact, the prejudice-increasing effects of negative contact tend to be stronger than the prejudice-reducing effects of positive contact (Barlow et al., 2012; Graf et al., 2014; Hayward, Tropp, Hornsey & Barlow, 2017). However, the present research did not replicate this finding. Of four studies, none found the effect of negative contact on outgroup attitudes to be stronger than the effects of positive contact and contact quantity. Indeed, negative contact was not significantly associated with heterosexual participants' outgroup attitudes in Studies 1 and 3, with non-heterosexual participants' outgroup attitudes in Study 1, and with cisgender participants' attitudes toward transgender people in Study 3. Contrary to earlier research, Study 4 found the effect of negative contact (on participants' ingroup construals) to be *weaker* than the effect of positive contact.

Second, Fell et al. (2016) found that positive contact decreases the effects of (subsequent) negative contact (see also Paolini et al., 2014), and that negative contact increases the effects of (subsequent) positive contact (see also Birtel & Crisp, 2012, albeit for *imagined* contact only). Hayward et al. (2017), however, did not find consistent evidence for the interaction effects observed by Fell et al. (2016). Although the present research measured both forms of contact, testing for interaction effects was beyond the scope of this thesis. Third, Hayward et al. (2017) compared both forms of contact in advantaged and disadvantaged groups. They found that members of disadvantaged minorities experienced both forms of contact more often than members of the advantaged majority, and that negative and positive contact were (negatively) correlated among disadvantaged-group members, but not among advantaged-group members (see Heitmeyer, 2005, cited in Pettigrew & Tropp, 2011).[5] Broadly, the present research, as reported in Chapter 5, replicated these findings when comparing sexual-minority and sexual-majority participants'

[5]Hayward et al. (2017) did not compare advantaged and disadvantaged groups statistically—but they used the same measures in both groups.

contact experiences.

Fourth, Hayward et al. (2017) argued that different mediators account for the consequences of positive and negative contact. They found that anger mediated the relationship between negative contact and outgroup attitudes, that empathy mediated the relationship between positive contact and outgroup attitudes, and that anxiety mediated the relationships between both forms of contact and outgroup attitudes. In Chapter 4, I found that negative contact was associated with less inclusive social identities, suggesting that negative experiences with an outgroup cause people to exclude that outgroup from a superordinate ingroup. Future research should examine less inclusive identities—alongside intergroup anxiety and anger (Hayward et al., 2017)—as a potential mediator of the effects of negative contact.

By measuring both forms of contact, the present research qualified earlier observations about the effects of positive and negative contact. In addition, the present research opened avenues for future research. In Chapter 5, I found that measuring positive contact on a unidimensional scale can mask the effects of negative contact (on collective action intentions). Future research should measure both forms of contact to avoid confounding the presence of one with the absence of the other. Overall, the present research underscores the importance of considering both positive and negative experiences when studying the consequences of intergroup contact.

How does intergroup contact reduce prejudice?

In this section, I review how intergroup contact affects social identification along a single category and across multiple categories. I discuss whether more continuous, fluid, and inclusive identities can explain how intergroup contact reduces prejudice.

Bridging binaries

Chapter 3 has reported research testing whether intergroup contact changes how participants understand their social identity along a single axis of categorisation—and, if so, whether these changes explain how intergroup contact reduces prejudice. Studies 1 and 2 documented reliable differences in the extent to which individuals understand their own sexuality and gender as continuous and fluid, rather than as binary and stable. Study 3 found that intergroup contact predicted more continuous and fluid social identities over time, but that more continuous and fluid social identities did not, in turn, predict less prejudice. Below, I discuss each of these findings in turn. I conclude by outlining avenues for future research using the newly developed measures.

In Chapter 3, I developed and validated two novel scales measuring differences in the extent to which individuals understand their own sexuality and gender as continuous and fluid, rather than as binary and stable. Studies 1–3 showed that the resulting scale dimensions were internally consistent, temporally stable, and measurement invariant across time and translations. Studies 1 and 2 found that, as expected, the newly developed scales correlated positively with same-gender attraction, and negatively with gender norm endorsement, group identification, social dominance orientation, and need for closure. Together, these observations show that the novel scales consistently measured what they were supposed to measure. The present research thus documented reliable individual differences in social identity continuity and fluidity.

Table 6.1. Example items for each dimension of the two social identity continuity and fluidity scales. Asterisks (*) mark reverse corded items. Items are worded for respondents who are (heterosexual) women.

Category	Dimension	Item
Sexuality	*Continuity: Identity*	Describing me as 'heterosexual' or 'straight' is enough to describe my sexual orientation.*
	Continuity: Attraction	I am not at all attracted to women.*
	Fluidity: Future	Though I am not attracted to women now, I might be in the future.
Gender	*Continuity: Identity*	Describing how I think about my gender is more complex than saying 'I am a woman' or 'I am female'.
	Continuity: Expression	Though I am a woman, I do not necessarily dress or behave as people expect women to dress or behave.
	Fluidity: Stability	How I express my gender is unlikely to change much as I grow older.*

In Chapter 1, I have argued that essentialist beliefs measure how individuals understand the social categories *of others*, and that researchers need to investigate how individuals understand *their own* category memberships. The present research supported this argument. Studies 1 and 2 found that the novel scales tended to correlate negatively with essentialist beliefs about sexuality and gender, with effect sizes ranging from small to medium. This suggests that the novel scales and essentialist beliefs scales (Haslam & Levy, 2006; Tee & Hegarty, 2006) measure related, but distinct, constructs. Specifically, the novel scales—which contain many reverse coded or negatively worded items—seem to reflect the respondents' reluctance to accept binary and stable categorisations (see Table 6.1). What makes this reluctance interesting is that all respondents were, by all other measures, heterosexual and/or cisgender.[6] This means, for example, that self-identified straight people can have continuous and fluid conceptions of *their own* sexuality. This goes beyond rejecting essentialist beliefs, which reflects respondents' belief that sexuality might be continuous and fluid *in principle*. The present research thus confirmed that, as suggested by earlier evidence (YouGov, 2015, 2016), individuals differ in how continuous and fluid they understand their social identities to be.

These findings raise the question why individuals differ in how continuous and fluid they understand their social identities to be—in other words, what the antecedents of social identity continuity and fluidity are. On the one hand, individual differences might reflect variation in some other characteristic related to sexuality or gender. More continuous conceptions of sexuality, for example, might reflect natural variation in same-gender attraction. Indeed, Study 2 found the continuity: attraction subscale to be strongly correlated with same-gender attraction. Controlling for same-gender attraction, however, did not change correlations between more continuous identities and other relevant variables. On the other hand, social experiences might affect how continuous and fluid individuals understand their social identities to be. As hypothesised, Study 2 found more continuous and fluid conceptions of sexuality and gender to be associated with respondents' contact experiences. This suggests that the novel scales measure something

[6] Indeed, I used stringent criteria to exclude respondents with non-binary sexualities or gender identities.

beyond natural variation in, for example, same-gender attraction.[7] Study 3 found contact with transgender people, but not with sexual minorities, to predict more continuous and fluid identities over time. The present research thus found some evidence that intergroup contact fosters more continuous and fluid social identities.

Another explanation of individual differences in social identity continuity and fluidity might be ideological differences. Study 2 found more continuous and fluid identities to be associated with less social dominance orientation. Study 3 (in Chapter 5) found that, among heterosexual and cisgender respondents, collective action intentions predicted more continuous and fluid identities over time. These findings suggest that opposing social inequalities is associated with more continuous and fluid identities. Future research should seek to understand how these and other ideological orientations (e.g., Altemeyer, 1981) relate to how individuals understand their group memberships—especially, how they understand the boundaries between their own and other groups (see Ho et al., 2011).

In Chapter 1, I have argued that essentialist beliefs might shape how we perceive outgroup members, but not how we feel about them or how we treat them. Instead, I have argued that social identification motivates intergroup bias, and that fostering more continuous and fluid identities should reduce intergroup bias. Studies 1 and 2 found cross-sectional evidence for the hypothesised relationship between social identity continuity/fluidity and intergroup bias. Study 3, however, did not find more continuous or fluid identities to predict more favourable outgroup attitudes over time. Overall, the present research found cross-sectional and longitudinal evidence that intergroup contact can change how individuals understand their social identity along a single axis of categorisation—but that these change cannot explain how intergroup contact reduces prejudice. This conclusion runs counter to Pettigrew's (1997, 1998) conjecture that contact could reduce prejudice by encouraging individuals to reappraise their ingroups.

Nonetheless, the present research has resulted in well-validated measures of social identity continuity and fluidity. I conclude this section by highlighting two avenues for future research using these measures. First, researchers could use the newly developed measures to study generational differences in how continuous and fluid individuals understand their sexuality and gender to be (see YouGov, 2015, 2016). Specifically, researchers could conduct longitudinal research to disentangle whether age differences are generational or developmental. Second, researchers could extend the present research to other social contexts. Specifically, researchers could attempt to create a generalised measure of social identity continuity and fluidity. In Appendix D, I present four (untested) items that reflect general features of the sexuality and gender scales, and that researchers could use as a starting point for developing a generalised scale. Together, these endeavours would further our understanding of social identification along a single axis of categorisation.

Bridging divides

Chapter 4 has reported research testing whether intergroup contact fosters more inclusive social identities across multiple cross-cutting categories—and, if so whether more inclusive identities explain how intergroup contact reduces prejudice. Study 4 found, as hypothesised, that cross-group friendship was associated with more inclusive social

[7] One could, of course, argue that continuity and fluidity reflect natural variation in same-gender attraction, but that contact with sexual-minorities helps individuals to admit to whatever degree of same-gender attraction they experience.

identities, and that more inclusive identities were, in turn, associated with more favourable outgroup attitudes and less desired social distance. Below, I discuss implications for measuring social identification across multiple categories, for social cohesion in divided societies, and for intergroup contact theory. I conclude by outlining avenues for future research using the triple crossed-categorisation task.

For the present research, I have adapted the triple crossed-categorisation task (van Dommelen, 2014; van Dommelen et al., 2015) to answer new questions about social identification across multiple categories. First, I recruited respondents from multiple groups, allowing me to compare the relative influence of individual and group differences. Second, I estimated response probabilities as varying across targets and participants using multilevel models. This allowed more fine-grained analyses of ingroup construals than quantitative and qualitative summaries. Third, I tested how contact with a *specific* outgroup related to participants' "us"–"not us" categorisations of and attitudes toward *that* outgroup. Together, these changes open the triple crossed-categorisation task to a broader range of research questions, while maintaining its strengths as an intuitive and informative method for studying how individuals construct their social identities from multiple group memberships.

Using this method, I examined how South Indian students construct their social identities from their caste, religious, and national group memberships. Studying both advantaged and disadvantaged groups, I found correlational evidence that friendship can bridge social divides by fostering more inclusive social identities. That is, contact experiences—not ideological preferences (Ho et al., 2015)—explained why some people included caste and religious minorities in their ingroups, and why others sought to exclude minorities from their ingroup. This research suggests that contact experiences can help to counteract harmful ideologies (e.g., Hindu nationalism) that underpin social divisions and conflicts (e.g., violence against Muslims). Future research should use the triple crossed-categorisation task to study other diverse (or diversifying) societies, for example, the United Kingdom or United States. This research could, for example, answer the question what determines when majority-group members include immigrant communities in their ingroup and when they exclude them.

Overall, this thesis has shown that more inclusive identities, but not more continuous and fluid identities, can help to explain how intergroup contact reduces prejudice. The *bridging-binaries hypothesis* postulates that, by fostering more continuous and fluid identities, contact blurs the boundaries between "us" and "them", and thus undermines the motivational underpinnings of intergroup bias. This hypothesis resounds with older explanations of how contact reduces prejudice. Brewer and Miller (1984) argued that contact could overcome our tendency to divide others into "us" and "them". Pettigrew (1997, 1998) argued that contact could encourage us to reappraise our ingroup, and to recognise the outgroup as valuable. In contrast, the *bridging-divides hypothesis* postulates that contact expands whom we consider "us" to include those whom we used to consider "not us". This hypothesis resounds with Gaertner and Dovidio's (2000) conjecture that, by fostering superordinate identities, contact extends ingroup favouritism to (former) outgroup members. As reviewed, the present research supports the latter explanation (that contact reduces bias by changing who is "us" and "not us") but not the former (that contact reduces bias by blurring the boundaries between "us" and "them"). Aside from how contact reduces bias, theories about intergroup contact postulate under which conditions contact reduces bias (see also R. Brown & Hewstone, 2005). Future research should consider under which conditions contact fosters more inclusive identities.

In Chapter 4, I have discussed some methodological limitations of the (improved) triple crossed-categorisation task. I have argued that future research should address these limitations by showing participants different combinations of targets, and by including a wider set of target categories. This research, however, would be impractical with the current version of the task. As all participants in a classroom view the same stimuli, the present task cannot vary target combinations across participants. As additional categories increase the required number of stimuli exponentially, the present task cannot accommodate more than three crossed categories. To accommodate these features, future research should further adapt the triple crossed-categorisation task by implementing "us"–"not us" categorisations as a web-based card-sorting task, and by using a planned missingness design to include a greater number of target stimuli. Building on the multilevel framework developed in this thesis, these changes would make the (triple) crossed-categorisation task an even more versatile method for studying social identification across multiple categories.

Does intergroup contact help or hinder social change?

Beyond prejudice, researchers have questioned whether intergroup contact helps or hinders social change. Researchers have argued that positive contact with an advantaged group could hinder social change by discouraging disadvantaged-group members from engaging in collective action (Dixon et al., 2012; Wright & Baray, 2012). Chapter 5 has reported findings from Studies 1, 3, and 4 that qualify and expand this literature. Studies 1 and 3 examined collective action intentions among sexual-minority and sexual-majority students. Studies 1a and 3a found that for sexual-minority participants, negative contact with the advantaged group was associated with more perceived discrimination and, in turn, more collective action. Contrary to prior research, positive contact was not associated with perceived discrimination or collective action. Studies 1b and 3b found that for sexual-majority participants, positive contact was associated with more solidarity-based collective action. Negative contact was not (consistently) associated with (less) solidarity-based collective action. Below, I compare the present findings to recent research, before outlining how the present findings may reconcile prejudice reduction and collective action approaches to social change. I conclude by discussing findings from Study 4 as evidence for the potential limits of solidarity across social divides.

Since Studies 1 and 3 were published (Reimer et al., 2017), other research has studied how positive and negative contact relate to support for social change in advantaged and disadvantaged groups. Recent cross-sectional research has examined collective action among advantaged and disadvantaged racial groups in the United States. Hayward, Tropp, Hornsey and Barlow (2018) found that among Black and Hispanic Americans, negative contact with White Americans was associated with more collective action. As in Studies 1a and 3a, positive contact was not consistently associated with (more) collective action. Selvanathan, Techakesari, Tropp and Barlow (2017) found that among White Americans, positive contact with Black Americans was associated with more solidarity-based collective action. As in Study 3b, negative contact was not associated with (less) solidarity-based collective action. Broadly, these studies replicate the findings from Studies 1 and 3 in a different context.

Findings from Studies 1 and 3—alongside more recent studies reviewed above—could change how some researchers think about the relationship between intergroup

contact and social change. As discussed, researchers (Dixon et al., 2012; Wright & Baray, 2012) argued that prejudice reduction and collective action approaches to social change are contradictory, and that in historically unequal societies, intergroup contact can maintain injustice by undermining collective action. In contrast, the present research suggests that contact and collective action can be complementary routes to social change and that intergroup contact can unite social groups in the struggle against social injustice. Below, I discuss each of these implications in turn.

As intergroup contact does not necessarily dampen collective action, prejudice reduction and collective action approaches to social change need not contradict each other. Indeed, as oppression results from power plus prejudice (Operario & Fiske, 1998), the two approaches could complement each other. Contact with minority-group members can, as far as it is positive, reduce prejudice among privileged majority-group members. Improving attitudes of those who hold social power (e.g., over employment or law) should lessen discrimination, both by influencing behaviour (e.g., hiring decisions) and by weakening resistance to legislation (e.g., affirmative action). For minority-group members, on the contrary, intergroup contact can be a necessary opportunity to recognise their relative disadvantage (Poore et al., 2002) and, especially if it entails negative experiences, foment collective action. Intergroup contact could hence both diminish prejudice and promote collective action.

Prejudice reduction and collective action seem to place contradictory demands on intergroup contact: the former requires positive interactions while the latter benefits from negative experiences. This, however, is only a seeming contradiction. Intergroup contact, as it occurs in everyday life, tends to entail both kinds of interactions though minority-group members tend to report more negative experiences than majority-group members (Heitmeyer, 2005, cited in Pettigrew & Tropp, 2011). Differences in status could also mean that the same interaction is perceived as positive by the advantaged participant but ambivalent by the disadvantaged partner. The present research thus does not imply that to achieve social change, disadvantaged-group members should be encouraged to seek out aversive encounters with dominant outgroups. Rather, our findings underscore the importance of facilitating desegregation and, in turn, contact (Hewstone, 2009). As discussed in Chapter 5, intervention programmes should be mindful of potential demobilising effects and find ways to address intergroup injustice.

Beyond traditional prejudice reduction and collective action approaches, intergroup contact can bridge social divides and unite people in the struggle for social justice. Positive contact can, as suggested in the present research, convince members of the dominant group to support movements against inequality. Similarly, positive contact between different minority groups could lead them to unite and form new coalitions in the pursuit of social justice (see Dixon et al., 2015). Supporting this idea, Study 3c found that positive contact with transgender people was associated with more solidarity-based collective action among sexual-minority participants. As intergroup contact could, hence, mobilise the disadvantaged, reduce prejudice among the advantaged, and unite all in the struggle for social justice, Studies 1 and 3 offer compelling evidence for the importance of intergroup contact for social change.

Study 4, however, contradicted conclusions from Studies 1 and 3, as well as findings from other recent research (Hayward et al., 2018; Selvanathan et al., 2017). For participants from disadvantaged caste groups, Study 4 examined whether cross-group friendship and negative contact, respectively, decreases and increases support for policies benefiting the disadvantaged ingroup. For participants from advantaged caste groups, Study 4 examined

whether cross-group friendship and negative contact, respectively, increases and decreases support for policies benefiting a disadvantaged outgroup. Contrary to predictions, cross-group friendship and negative contact were not associated with support for reservation policies in either advantaged or disadvantaged groups. In Chapter 5, I have suggested that social context may explain the discrepancy between the findings of Studies 1, 3, and 4. Below, I discuss this and other explanations for the discrepancy in findings.

Study 4 differs from Studies 1 and 3 in terms of four relevant dimensions. First, social change has different implications in the two social contexts. LGBT activism aims at advancing civil rights, while caste reservation aims at effecting redistribution. Unlike upper-caste students, heterosexual/cisgender participants had little to lose by supporting LGBT advocacy. However, this does not explain why Hayward et al. (2018) and Selvanathan et al. (2017) found intergroup contact to be associated with collective action against racial injustice. Second, Study 4 measured policy support as the relevant outcome variable, while Studies 1 and 3 assessed collective action intentions. This difference may explain the discrepancy between Study 4 and Studies 1 and 3, but not between Study 4 and earlier research. Cakal et al. (2011), for example, found positive contact to be associated with support for redistributive policies among Black and White South Africans.

Third, Studies 1 and 3—alongside other previous and recent research—focused on actions and policies with distant and uncertain outcomes. Study 4, however, focused on reservation policies, that is, policies with outcomes that the participants had only recently experienced in their own lives. Fourth, perhaps as a result of the previous observation, South Indian participants' support for or opposition to reservation policies strictly followed caste interests. As group differences outstripped individual differences, Study 4 may have studied attitudes that are too deeply entrenched to be affected by intergroup contact. Future research should seek to replicate the present research by studying the effects of positive and negative contact on support for social change in different situations and contexts. This research would help us understand *when* intergroup contact helps or hinders social change.

Doing research inside the reproducibility crisis

Psychological science finds itself in a reproducibility crisis. That is, psychological research produces many findings that are not reproduced in subsequent replications of the original research. Replicating 100 studies published in leading psychology journals, researchers found only 36% to 47% of direct replications to reproduce the original findings (Open Science Collaboration, 2015). Similarly, recent research failed to reproduce several prominent findings, including behavioural priming effects (Doyen, Klein, Pichon & Cleeremans, 2012; Shanks et al., 2013), the facial feedback hypothesis (Wagenmakers et al., 2016), and the ego-depletion effect (Hagger et al., 2016). In response, researchers proposed several solutions to making psychological science more reproducible (Munafò et al., 2017). In this section, I discuss reasons for and solutions to the reproducibility crisis, and evaluate the present research in light of these reasons and solutions.

In Chapter 2, I have introduced the distinction that every sample has both regular features—that is, features that generalise across samples—and irregular features—that is, features that do not generalise across sample. In essence, repeated failures to reproduce findings mean that the original findings were based on irregular features of the observed data ("noise"), rather than on its regular features ("signal"). Several observations suggest

that psychological research has an unfavourable signal-to-noise ratio. First, human minds are more complex than other systems studied in scientific research, making psychological observations inherently noisy. Second, psychological studies often use poorly validated measurement scales (Flake et al., 2017), which increases statistical noise in the observed data (Loken & Gelman, 2017). Third, psychological research tends to rely on small samples, which means that parameter estimates fluctuate more widely from sample from sample. Small samples thus increase the probability of both false-positive and false-negative findings (Button et al., 2013). Together, these observations provide some of the reasons why findings from psychological research fail to replicate.

This thesis presents research that addresses at least some of these problems. Wherever possible, this research used measures of intergroup contact (and related constructs) that had been validated in previous research (for a review, see Lolliot et al., 2015). In addition, Studies 1–3 followed up-to-date guidelines for scale development and validation, resulting in reliable measures of the constructs under study. Despite these efforts, the present research included several measures that had not been validated systematically. Most pressing, future research should seek to validate and, if necessary, improve existing measures of negative contact (for a starting point, see Hayward et al., 2017). In terms of sample size, the present research exceeded many other studies in psychological research. In total, the four studies in this thesis included 2953 participants. Together, these features should increase the signal-to-noise ratio of the present research.

As discussed, psychological research tends to search for answers to a research question in situations where any useful information (regular features) is buried in a multitude of irrelevant noise (irregular features). If accounted for, this statistical noise would not lead to false inferences. Unfortunately, psychological research has not consistently accounted for the noisiness of its observations. Here, I consider three obstacles to reproducible inferences. First, psychological research tends to rely on null hypothesis significance testing, which bases statistical inference on long-term error rates, and is vulnerable to multiple comparisons. Second, scientific practice rewards significant results, and punishes non-significant results. Third, researchers tend to have a lot of flexibility while conducting, analysing, and reporting psychological research. Together, these obstacles make false conclusions from noisy observations a likely outcome of psychological research: Given enough flexibility, researchers are almost certain to find at least one significant comparison in a dataset (multiple comparisons), especially if they have an incentive to do so (Simmons, Nelson & Simonsohn, 2011). In a noisy dataset, this significant result is particularly likely to reflect an irregular feature of the dataset. Below, I discuss two solutions that researchers can adopt to avoid false conclusions.

One possible solution would be that researchers preregister each study's proposed design and analysis plan, thus limiting flexibility during data collection and analysis (Nosek, Ebersole, DeHaven & Mellor, 2018). By separating prediction and postdiction, researchers would avoid inadvertent multiple comparisons, and meet the stringent assumptions of null hypothesis significance testing. I set out to conduct the present research in 2014, and did not preregister any of the four studies presented in this thesis. Analysing results from these studies, I took care to not make decisions based on whether they resulted in significant findings (and have, in fact, taken decisions that resulted in non-significant findings). Multiple comparison, however, can be a problem even when a researcher computes only a single statistical test—and despite a researcher's best intentions—as long as any detail of the analysis process is contingent on the observed data (A. Gelman & Loken, 2013). Each of my choices might look eminently sensible and in accord with common

research practices. And yet, had the data looked differently, I might well have arrived at a different conclusion. Cognitive biases, such as confirmation bias and hindsight bias, make me a particularly bad judge in this situation (Munafò et al., 2017). As intentions do not serve as a good guide, one can evaluate research based on the number of decisions taken during data analysis. In this comparison, Studies 1–3 fare worse than Study 4, as structural equation models are far more complex than simple multilevel models, and thus demand far more decisions to be taken by the researcher.

Another solution would be for researchers to abandon null hypothesis significance testing, and to adopt a data analysis approach that distinguishes between regular and irregular features in a dataset. In Chapter 2, I have introduced cross-validation, prior distributions, and multilevel models as techniques to regularise statistical inferences. By focusing on out-of-sample predictions, these techniques should result in more reproducible findings than conventional methods.[8] I used these techniques to analyse results from Study 4, but not Studies 1–3 (for the rationale, see Chapter 2). As a result, I am more confident that results from Study 4 reflect regular features. Indeed, if I had to bet on which of the findings in this thesis will replicate, I would place a higher bet on findings from Study 4 than on findings from Studies 1–3. Nonetheless, had I started this project in 2018 instead of 2014, I would have preregistered all, or almost all, of the studies.

Conclusions

Broadly, this thesis has examined whether contact with diverse others can change not only how we see *others*, but also how we see *ourselves*. Pursuing this question, the present research found cross-sectional and longitudinal evidence that intergroup contact fosters more continuous, fluid, and inclusive social identities—and that more inclusive identities, but not more continuous and fluid identities, can explain how intergroup contact reduces prejudice. Beyond prejudice, the present research found that negative contact predicts collective action in disadvantaged groups, while positive contact predicts collective action in advantaged groups. Together, these findings reaffirm that, at least in some contexts, intergroup contact helps, not hinders, social change.

More broadly, this thesis has shown that the social identity perspective (Tajfel, 1981; J. C. Turner et al., 1987) continues to be important for understanding intergroup relations. Unlike earlier research (Brewer & Miller, 1984; Gaertner & Dovidio, 2000; R. Brown & Hewstone, 2005), the present research did not consider social categorisation as one of the conditions that could determine whether contact changes attitudes. Instead, the present research answered Pettigrew's (1997, 1998) call to consider social identification as one of the mechanisms through which contact could change attitudes. Alongside other recent research (van Dommelen et al., 2015; Schmid et al., 2009), the present research has introduced novel ways of measuring and understanding social identification. As such, this thesis has opened avenues for future research to help us understand a social world beyond binary, stable, and exclusive social identities.

On a personal note, I have found this research to bolster my belief in the enduring value of cross-group friendship. Contact with diverse others changes how we see ourselves and others. As such, contact could help us understand ourselves, our own lives, and the

[8] In this thesis, I have applied these techniques within Bayesian data analysis. However, cross-validation, regularisation, and other techniques to improve out-of-sample predictions are not exclusive to Bayesian data analysis (for a review, see Yarkoni & Westfall, 2017).

lives of others—and thus bridge whatever divides us from our fellow human beings. I can think of few pursuits more worthwhile.

Appendices

Appendix A

Materials for Chapter 3

Table A.1. Complete item pool for scale development in Study 2a. Items *in italics* were excluded from the final scale. Items are worded for participants who are women.

#	Item
	Included
1	I am exclusively heterosexual and not at all homosexual.
2	I am not at all attracted to women.
6	I cannot imagine being attracted to both men and women.
7	The terms 'homosexual' and 'heterosexual' are too narrow to describe my sexual orientation.
9	Saying 'I am straight' or 'I am heterosexual' is all there is to say about my sexual orientation.
10	The extent to which I am attracted to men or women is more complex than saying 'I am heterosexual' or 'I am homosexual'.
11	Describing me as 'heterosexual' or 'straight' is enough to describe my sexual orientation.
12	I am exclusively attracted to men.
14	The extent to which I am attracted to men or women may have changed to some degree over the course of my lifetime.
15	Whether I am attracted to men or women has not changed over the course of my lifetime.
17	My sexual orientation may have changed to some degree since I was younger.
18	I have never considered going out with a woman.
22	The extent to which I am attracted to men or women may change to some degree in the future.
23	Whether I am attracted to men or women will not change in the future.
25	Though I am not attracted to women now, I might be in the future.
26	Though I think of myself as heterosexual now, I might not do so in the future.
27	I cannot say that I will always only love men.
	Excluded
3	*Though I am heterosexual, I can find women attractive sometimes.*
4	*Though I am heterosexual, I can judge whether a woman is attractive or not.*
5	*I can find a person of the same gender attractive and still be heterosexual.*
8	*My sexual orientation is best understood as a combination of the extent to which I am attracted to women and the extent to which I am attracted to men.*
13	*Whether I am attracted to men or women is stable and unchangeable.*
16	*I have always been attracted to men/boys not women/girls.*
19	*The extent to which I am attracted to men or women may change to some degree between different situations.*
20	*Whether I am attracted to men or women does not change between different situations.*
21	*In some situations I am more attracted to men while in others I am more attracted to women.*
24	*I will always be attracted to men not women.*
28	*I would never consider going out with a woman.*

Table A.2. Complete item pool for scale development in Study 2b. Items *in italics* were excluded from the final scale. Items are worded for participants who are women.

#	Item
	Included
1	The terms 'masculine' and 'feminine' are too narrow to describe how I express my gender in dress and demeanour.
3	Though I am a woman, I do not necessarily dress or behave as people expect women to dress or behave.
4	Though I am a woman, my appearance, expression, and demeanour reflects both 'masculine' and 'feminine' characteristics.
6	How I dress or behave is not defined by my gender.
9	The terms 'male' and 'female' are too narrow to describe how I think about my gender.
10	Saying 'I am a woman' is all there is to say about how I think about my gender.
11	Describing how I think about my gender is more complex than saying 'I am a woman' or 'I am female'.
12	Describing me as 'female' is enough to describe my gender.
14	The extent to which I think of myself as male or female may have changed to some degree over the course of my lifetime.
15	The extent to which I speak, dress, and behave in 'masculine' or 'feminine' ways may have changed over the course of my lifetime.
16	How I express my gender has not changed much since I was younger.
17	The extent to which I think of myself as male or female may change to some degree between different situations.
18	The extent to which I speak, dress, and behave in 'masculine' or 'feminine' ways may change between different situations.
20	In some situations I express myself in more 'masculine' ways while in others I express myself in more 'feminine' ways.
21	The extent to which I think of myself as male or female may change to some degree in the future.
22	The extent to which I speak, dress, and behave in 'masculine' or 'feminine' ways may change in the future.
24	How I express my gender is unlikely to change much as I grow older.
25	How I think about my gender is unlikely to change as I grow older.
26	Though I think of myself as a woman now, I might not necessarily do so in the future.
27	I will always think about my gender in the same way as I do now.
	Excluded
2	*Describing me as 'feminine' is enough to describe how I express my gender.*
5	*I can still be a woman even if I dress and behave as people expect men to behave and dress.*
7	*How I express my gender is best described as a combination of 'masculine' and 'feminine' characteristics.*
8	*I cannot imagine what it would mean to be both male and female.*
13	*Whether I think of myself as male or female is stable and unchangeable.*
19	*How I express my gender does not change much between different situations.*
23	*I will always think of myself as only female and not at all male.*

Table A.3. German translations of continuity/fluidity items for Study 3a. Item numbers (#) correspond to those in Table 3.2. Asterisks (*) mark reverse coded items. Items are worded for participants who are women.

#	Item
	Continuity: Identity
7	Die Begriffe 'homosexuell' und 'heterosexuell' sind zu eng gefasst, um meine sexuelle Orientierung zu beschreiben.
9	"Ich bin heterosexuell" sagt alles, was es über meine sexuelle Orientierung zu sagen gibt.*
10	Inwieweit ich mich zu Männern oder Frauen hingezogen fühle, ist komplexer als einfach nur zu sagen, "Ich bin heterosexuell" oder "Ich bin homosexuell".
11	Mich als "heterosexuell" oder "hetero" zu beschreiben reicht aus, um meine sexuelle Orientierung zu beschreiben.*
	Continuity: Attraction
1	Ich bin ausschließlich heterosexuell und kein bisschen homosexuell.*
2	Ich fühle mich kein bisschen zu Frauen hingezogen.*
6	Ich kann mir nicht vorstellen, mich sowohl zu Männern als auch Frauen hingezogen zu fühlen.*
12	Ich fühle ausschließlich zu Männern hingezogen.*
18	Ich habe noch nie in Erwägung gezogen, mit einer Frau auszugehen.*
	Fluidity: Future
22	Inwieweit ich mich zu Männern oder Frauen hingezogen fühle, könnte sich in der Zukunft zu einem gewissen Grad ändern.
23	Ob ich mich zu Männern oder Frauen hingezogen fühle, wird sich in der Zukunft nicht ändern.*
25	Obwohl ich mich derzeit nicht zu Frauen hingezogen fühle, könnte sich das in der Zukunft ändern.
26	Obwohl ich mich derzeit als heterosexuell sehe, könnte sich das in der Zukunft ändern.
27	Ich kann nicht sagen, dass ich immer nur Männer lieben werde.

Table A.4. German translations of continuity/fluidity items for Study 3b. Item numbers (#) correspond to those in Table 3.3. Asterisks (*) mark reverse coded items. Items are worded for participants who are women.

#	Item
Continuity: Identity	
9	Die Begriffe 'männlich' und 'weiblich' sind zu eng gefasst, um zu beschreiben, wie ich über mein Geschlecht denke.
10	"Ich bin eine Frau" sagt alles, was es über meine Geschlechtsidentität zu sagen gibt.*
11	Wie ich mein Geschlecht verstehe, geht über ein einfaches "Ich bin eine Frau" oder "Ich bin weiblich" hinaus.
12	Describing me as 'female' is enough to describe my gender.*
Continuity: Expression	
1	Die Begriffe 'maskulin' und 'feminin' sind zu eng gefasst, um zu beschreiben, wie ich mein Geschlecht in meiner Kleidung und meinem Verhalten ausdrücke.
3	Obwohl ich eine Frau bin, kleide oder verhalte ich mich nicht unbedingt so, wie es von Frauen erwartet wird.
4	Obwohl ich eine Frau bin, spiegelt mein Erscheinungsbild, meine Ausdrucksformen und mein Verhalten sowohl 'maskuline' als auch 'feminine' Eigenschaften wieder.
6	Wie ich mich kleide oder verhalte ist nicht durch mein Geschlecht festgelegt.
Fluidity: Stability	
16	Wie ich mein Geschlecht ausdrücke, hat sich nicht geändert seitdem ich klein war.*
24	Wie ich über meine Geschlechtsidentität denke, wird sich wahrscheinlich nicht ändern, wenn ich älter werde.*
25	Ich werde in Zukunft genauso über meine Geschlechtsidentität denken, wie ich es heute tue.*
27	Es ist unwahrscheinlich, dass sich ändert, wie ich mein Geschlecht ausdrücke.*

Appendix B

Materials for Chapter 4

Instructions for crossed-categorisation task

[After participants had the opportunity to read and sign informed consent:] In this study, I am going to introduce you to 24 different persons, each represented by an identity card that looks something like this: **[click for p. 2]**. Each card has a number (#) that corresponds to a column on page 1 of the booklet in front of you. On each card, you can see the name, age, religion, and nationality of a person, as well as a silhouette and a person's caste reservation (for Indian nationals).

[click for p. 3] On each slide, we will show two people. A man on the left-hand side, and a woman on the right-hand side. Please focus only on the person that matches your own gender. Before I am going to ask you to fill anything out, I will show you each person for 7 seconds. Please put your pen down, and focus on the person on the slide. If you have time after studying the information on the card, think about the person on the card until the next card is shown. Please make sure you pay attention to each card, though there is no need to memorise any of the information. I will now go through the cards. **[Click to start presentation; stimuli will change automatically every 7s.]**

Next, I will ask you to again look at each card (on your side of the screen), and answer question 1 (Q1 on your booklet) for each person. So if you feel this person is one of your own group, or 'us', you write a '1' in the 'Q1' column next to the card number. And if you feel this person is not one of your own group, or 'not us', you write a '0' in the 'Q1' column next to the card number.

It doesn't matter how many cards you assign to each of these categories. If you consider a lot of cards to belong to 'us', and you write down a lot of ones, then that is just fine. Or, if you feel that a lot of these people do not belong to 'us', and write down a lot of zeros, then that is just fine too. Please note that this task is not about how *positive* or how *negative* you feel toward the person. For example, you may see someone as not belonging to your own group and still feel positive toward that person.

Please remember, there are no right or no wrong answers. The only important thing is that you put these cards into the categories you *think* or *feel* that they belong to. You do the same for all the cards. Consider only one card at a time and focus on the card corresponding to your gender. When you have made your choice, please put your pen down to let me know that I can move on to the next card. I will now go through the cards. **[Click through the next 24 slides; wait until all participants have made their choice before continuing to the next slide.]**

[After participants have completed Q1, click for next slide.] Next, I will ask

you to again look at each card (on your side of the screen), and answer questions 2 and 3 (Q2 and Q3 in your booklet).

For question 2, imagine a situation in which you consider sharing a room with the person on the card. Think about how *comfortable* or *uncomfortable* you would feel about sharing a room with that person. For each card, write a number between 1 (that is, *very uncomfortable*) and 7 (that is, *very comfortable*) in the Q2 column next to corresponding card number. Take a moment to familiarise yourself with the scale.

[click for next slide] For question 3, you see something like a thermometer. It's not a regular thermometer, but a 'feeling thermometer'. What I will ask you to do next is to look at each card that I will show you, and to write down a number between 0 and 100, representing degrees on that thermometer. If you feel *warm*—or favourably—toward the person on the card, you write down a high number. If you feel *cold* toward the person—if you feel you don't care too much for a person, you can write down a lower number. You can write down any number between 0 and 100. Please try not to think about each of your answers too long. Take a moment to familiarise yourself with the feeling thermometer, and how it works.

[click for next slide] On each slide, I will repeat the answer options for the two questions. Remember to consider only the card corresponding to your gender, to answer both questions for each card, and to put each answer in the right box. When you have made your choice, please put down your pen to let me know that I can move on to the next card. I will now go through the cards. *[Click through the next 24 slides; for every slide, wait until all participants have made their choice before continuing to the next slide.]*

[After participants have completed Q2–Q3, click for next slide.] Next, I will ask you to fill in the remaining questions in the booklet in front of you. Please read and follow all instructions carefully. The questions you'll answer address a variety of topics, and are not necessarily related in any way to the task you have just completed. Clearly mark your answer according to the instructions shown on the first page of the questionnaire. Questions are printed on both sides of the page, please answer all of them. Try not to think about each of your answers too long; your first feeling or impression is generally best. When you are done, please put down your pen and turn over the booklet. Please start filling in the questionnaire now.

[After participants have turned over the booklet:] Please make sure that you have answered questions on both sides of each page. After you leave this classroom, please do not talk about your experience with students who have not yet taken part in the study. Doing so could change their responses when they participants—so it's very important they come into the study without knowing all that you have experienced now. Thank you very much for participating!

Descriptive statistics

Table B.1. Correlations and descriptive statistics in Study 4, excluding responses by participants who belong to the relevant outgroup.

(a) Outgroup: Dalits

#	Variable	1	2	3	4	5	6	*M*	*SD*
1	Contact quantity		0.63	-0.26	0.62	-0.22	-0.12	4.32	0.94
2	Positive contact			-0.29	0.61	-0.19	-0.07	4.14	0.99
3	Negative contact				-0.32	0.24	0.26	2.17	1.26
4	Outgroup friendship					-0.20	-0.13	4.29	0.82
5	SDO: Dominance						0.48	3.09	1.17
6	SDO: Egalitarianism							3.25	1.15

(b) Outgroup: Other Backwards Classes

#	Variable	1	2	3	4	5	6	*M*	*SD*
1	Contact quantity		0.61	-0.08	0.60	-0.18	-0.10	3.87	1.06
2	Positive contact			-0.18	0.50	-0.01	-0.13	3.83	1.06
3	Negative contact				-0.04	0.24	0.18	2.14	1.26
4	Outgroup friendship					-0.02	0.05	3.78	0.96
5	SDO: Dominance						0.39	3.08	1.19
6	SDO: Egalitarianism							3.21	1.06

(c) Outgroup: General Castes

#	Variable	1	2	3	4	5	6	*M*	*SD*
1	Contact quantity		0.44	-0.02	0.36	-0.20	-0.08	3.99	1.02
2	Positive contact			-0.08	0.32	-0.12	-0.06	3.80	1.24
3	Negative contact				-0.05	0.27	0.20	2.01	1.26
4	Outgroup friendship					-0.11	-0.20	3.84	0.97
5	SDO: Dominance						0.42	3.20	1.18
6	SDO: Egalitarianism							3.30	1.17

(d) Outgroup: Muslims

#	Variable	1	2	3	4	5	6	*M*	*SD*
1	Contact quantity		0.64	-0.11	0.61	-0.19	-0.20	3.51	1.21
2	Positive contact			-0.18	0.65	-0.13	-0.16	3.53	1.20
3	Negative contact				-0.12	0.17	0.09	2.30	1.24
4	Outgroup friendship					-0.08	-0.16	3.30	1.01
5	SDO: Dominance						0.44	3.12	1.18
6	SDO: Egalitarianism							3.26	1.13

Correlates of categorisation

I included additional measures to validate the triple crossed-categorization task. I tested to what extent various social identity measures correlated with responses in the triple crossed-categorization task. Below, I describe these additional measures before reporting the relevant analyses.

Measures

Group identification was measured with one item per ingroup (Postmes, Haslam & Jans, 2013): "I identify with my [nationality/religion/caste group]" (1 = *strongly disagree*, 7 = *strongly agree*).

Social identity complexity (Roccas & Brewer, 2002; Schmid et al., 2009) was operationalised as the extent to which participants perceived their different group memberships as conceptually interrelated (*similarity complexity*) and numerically overlapping (*overlap complexity*). Two items measured similarity complexity for religion/nationality and for nationality/caste: "How similar or different are the typical [Hindu/Indian] and the typical [Indian/person from your caste group] to each other?" (1 = *very different*, 5 = *very similar*, reverse coded), and "Do you think that being [Hindu/Indian] means the same as being [Indian/from your caste group]?" (1 = *means something very similar*, 5 = *means something very different*). Two items measured overlap complexity for religion/nationality: "How many [Indians/ Hindus] do you think are [Hindus/Indian]?" (%, r = .63). Three items measured overlap complexity for nationality and caste: "How many Indians do you think are [GC, OBC, SC/ST]?" (%, .33 ≤ rs ≤ .50). Overlap items were reverse coded, so that higher scores indicated less perceived overlap.

Entitativity (Lickel et al., 2000) was measured with six items: "Being Indian is important to each Indian, no matter their caste or religion", "Indians of all castes and religions interact often with each other.", "Indians of all castes and religions are working towards shared goals", "What happens to one Indian also impacts other Indians", "Indians of all castes and religions depend on one another", and "Indians are similar to one another" (1 = *strongly disagree*, 7 = *strongly agree*, α = .80). Participants were encouraged to "take Indian/Indians to mean Indian citizens of all castes and religions" when reading the items of this scale. This scale thus measured an inclusive conception of entitativity.

Results

I explored how participants' responses in the triple crossed-categorization task related to other measures of social identification. To that end, I estimated the correlations between, on the one hand, the proportions of targets that participants had categorised as "us" in each of the six categories, and, on the other hand, the (aggregated) social identification variables described in the *Measures* section. I derived the likelihood of the observed responses from a multivariate normal distribution:

$$\mathbf{Y} \sim \text{MVNormal}(\mu, \mathbf{S})$$

where $\mathbf{Y}$ is the matrix of observed variable scores (columns) across participants (rows), μ is the vector of all variable means, and $\mathbf{S}$ is the variance-covariance matrix across variables.

Figure B.1 shows the estimated correlations between participants' categorisations in the crossed-categorisation task (from left to right) and various social identification

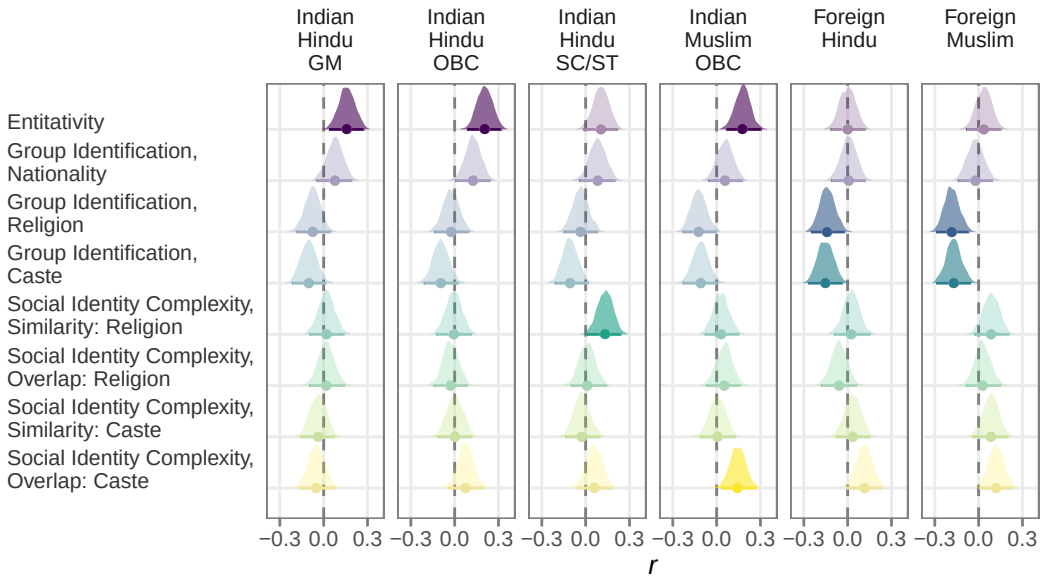

Figure B.1. Correlations between various social identity variables (top to bottom) and the proportion of targets that participants had categorised as "us" in each target category (left to right). Points indicate the most likely estimate for a given correlation, while lines encompass the 97% most likely estimates of that correlation. Curves show the density distributions of the posterior probabilities, based on all samples from the posterior probability distribution. Correlations for which $\Pr(r > 0|M) > .99$ or $\Pr(r < 0|M) > .99$ are shown in a darker shade.

variables (from top to bottom). Entitativity—the extent to which participants saw Indians of *all* castes and religions as similar to one another, as depending on each other, as sharing a common fate, as working towards shared goals, and as valuing their shared identity—was positively correlated with the proportion of *Indian, Hindu, GM* ($r = .16$, $[.04, .27]$, $\Pr(r > 0|M) > .99$), *Indian, Hindu, OBC* ($r = .21$, $[.09, .31]$, $\Pr(r > 0|M) > .99$), *Indian, Hindu, SC/ST* ($r = .10$, $[-.02, .22]$, $\Pr(r > 0|M) = .97$), and *Indian, Muslim, OBC* ($r = .18$, $[.07, .30]$, $\Pr(r > 0|M) > .99$) targets participants categorised as "us".

Group identification was only associated with participants' categorisations of foreign targets, not their categorisations of Indian targets. The extent to which participants identified with their religion was negatively correlated with how many *foreign, Hindu* ($r = -.14$, $[-.25, -.02]$, $\Pr(r > 0|M) > .99$) and *foreign, Muslim* ($r = -.18$, $[-.29, -.07]$, $\Pr(r > 0|M) > .99$) targets participant categorised as "us". I obtained similar correlations for the extent to which participants identified with their caste group, $r = -.15$, $[-.27, -.04]$, $\Pr(r > 0|M) > .99$ and $r = -.17$, $[-.29, -.06]$, $\Pr(r > 0|M) > .99$, respectively. The extent to which participants identified with their nationality was not associated with participants' categorisations—perhaps because there was a little variation in participants responses, with 81% of participants strongly agreeing with the statement. Other variables were not systematically correlated with participants' responses in the crossed-categorisation task (see Figure B.1).

Correlations thus confirmed that participants' categorisations were more than a straight-forward reflection of the strength of their group identification, affirming the value of the crossed-categorisation task for studying identification in contexts with multiple, cross-cutting group memberships. Furthermore, participants' perceptions of entitativity

across caste and religions were associated with more inclusive construals of national identity, providing tentative evidence for its potential for fostering more inclusive identities.

Appendix C

Materials for Chapter 5

Table C.1. Correlations between latent variables for LGBT+ respondents (Study 3a) within T1 (below the diagonal) and within T3 (above the diagonal). Observed mean scores, standard deviations, and reliability indices are shown in the lower half of the table. PD = Personal discrimination; GD = Group discrimination; OA = Outgroup attitudes

#	Variable	1	2	3	4	5	6	7	8	9
1	Positive contact	**.74**	-.23	-.00	-.05	-.13	-.11	-.06	.14	.52
2	Negative contact	-.19	**.71**	.09	.12	.26	.24	.21	.06	-.39
3	Solidarity	.02	.08	**.62**	.63	.41	.35	.34	.47	-.03
4	Centrality	-.08	.08	.54	**.65**	.32	.16	.24	.35	.00
5	PD	-.20	.29	.33	.42	**.82**	.23	.27	.27	-.20
6	GD	-.15	.19	.33	.26	.33	**.77**	.64	.39	-.26
7	Anger	-.07	.13	.36	.35	.33	.59	**.62**	.44	-.14
8	Collective action	.01	.10	.43	.42	.30	.43	.58	**.82**	-.02
9	OA	.46	-.33	-.06	-.17	-.30	-.20	-.15	-.12	**.67**
T1	*M*	3.98	2.09	4.96	4.42	1.95	3.14	3.62	5.23	74.6
T1	*SD*	0.63	0.70	1.26	1.51	0.76	0.64	0.91	1.16	21.2
T1	Cronbach's α	.79	.81	.80	.80	.81	.70	.78	.82	-
T3	*M*	3.92	2.05	4.96	4.47	1.91	3.20	3.62	5.26	73.3
T3	*SD*	0.64	0.72	1.28	1.49	0.72	0.61	0.89	1.13	20.0
T3	Cronbach's α	.81	.86	.84	.84	.81	.73	.80	.85	-

Table C.2. Correlations between latent variables for LGBT+ respondents (Study 3a) between T1 and T3. Above the diagonal, columns represent T1 variables; below the diagonal, rows represent T1 variables. PD = Personal discrimination; GD = Group discrimination; OA = Outgroup attitudes

#	Variable	1	2	3	4	5	6	7	8	9
1	Positive contact	**.74**	-.21	.15	-.07	-.25	-.10	-.11	.08	.40
2	Negative contact	-.18	**.71**	-.03	.12	.17	.17	.13	.05	-.22
3	Solidarity	.02	.06	**.62**	.47	.35	.23	.36	.37	.00
4	Centrality	.03	.09	.34	**.65**	.27	.13	.28	.32	.01
5	PD	-.12	.27	.29	.37	**.82**	.29	.28	.27	-.24
6	GD	-.17	.25	.33	.25	.19	**.77**	.52	.32	-.24
7	Anger	-.13	.23	.33	.24	.28	.48	**.62**	.42	-.16
8	Collective action	.04	.08	.37	.30	.20	.48	.52	**.82**	-.08
9	OA	.40	-.35	-.03	-.10	-.25	-.18	-.14	-.03	**.67**

Table C.3. Correlations between latent variables, observed mean scores and standard deviations, and reliability indices for heterosexual/cisgender respondents (Study 3b) with LGB+ people as the relevant outgroup.

		T1			T2			T3		
#	Variable	1	2	3	1	2	3	1	2	3
1	Positive contact, T1	-	.15	.45	**.83**	.14	.43	**.81**	.12	.44
2	Negative contact, T1		-	-.15	.13	**.62**	-.13	.05	**.65**	-.09
3	Collective action, T1			-	.47	-.15	**.75**	.50	-.09	**.76**
1	Positive contact, T2				-	.03	.49	**.88**	.09	.52
2	Negative contact, T2					-	-.11	-.01	**.75**	-.08
3	Collective action, T2						-	.51	-.01	**.79**
1	Positive contact, T3							-	.03	.56
2	Negative contact, T3								-	-.04
3	Collective action, T3									-
M		4.00	1.56	4.72	4.01	1.64	4.62	4.09	1.63	4.62
SD		0.65	0.53	1.75	0.66	0.55	1.81	0.60	0.55	1.80
Cronbach's α		.72	.72	-	.76	.74	-	.73	.72	-

Table C.4. Model fit indices for configural, loadings, and intercept invariant models for sexual-minority participants in Study 3c.

#	Model	χ^2	*df*	χ^2/df	*CFI*	*TLI*	*RMSEA*	Δ*CFI*
0	Configural invariance	966.30	525	1.84	.96	.95	.04	
1 vs 0	Loadings invariance	992.41	538	1.84	.95	.95	.04	-.001
2 vs 1	Intercepts invariance	995.04	556	1.79	.96	.95	.03	.002

Table C.5. Model fit indices for configural, loadings, and intercept invariant models for sexual-majority participants in Study 3c.

#	Model	χ^2	*df*	χ^2/df	*CFI*	*TLI*	*RMSEA*	Δ*CFI*
0	Configural invariance	1842.75	834	2.21	.96	.95	.03	
1 vs 0	Loadings invariance	1871.29	856	2.19	.96	.95	.03	.000
2 vs 1	Intercepts invariance	1886.21	886	2.13	.96	.95	.03	.001

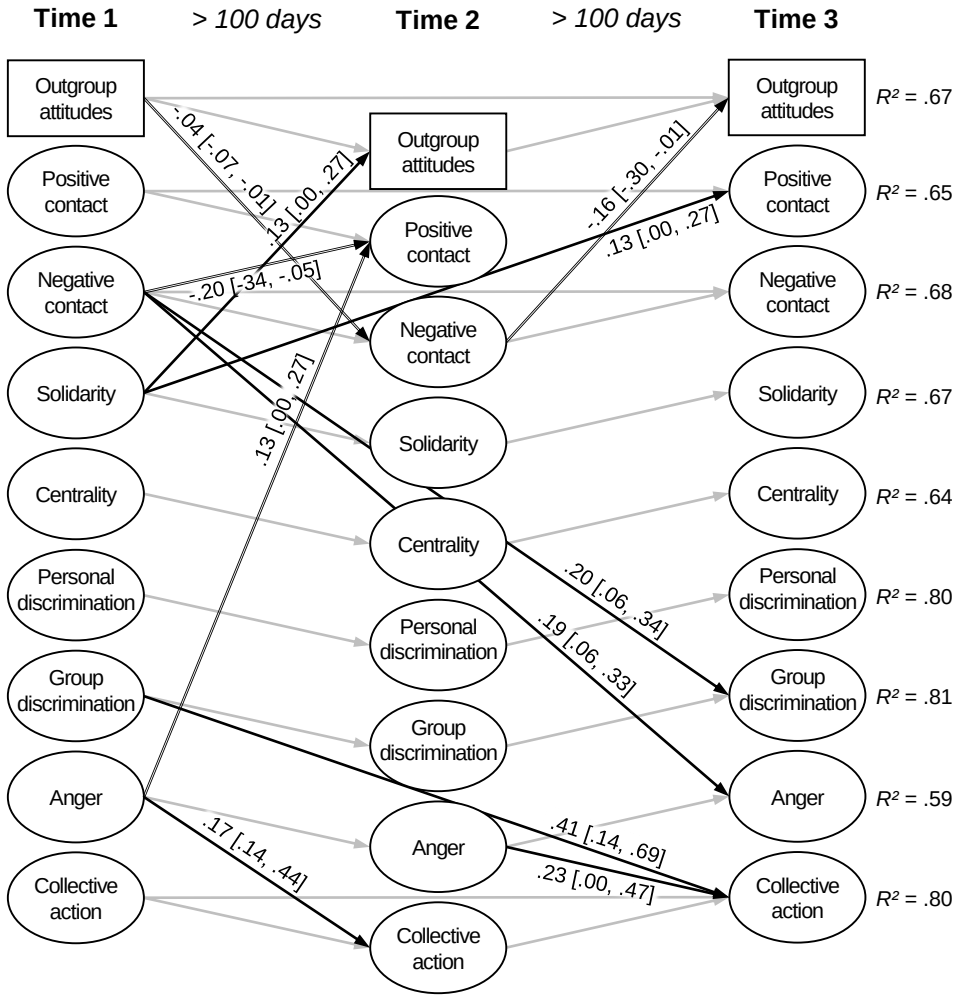

Figure C.1. Three-wave panel model showing autoregressive (in grey) and cross-lagged (in black) paths for sexual-minority participants (Study 3a), with standardised coefficients and 95% confidence intervals. Model fit: $\chi^2(2360) = 3449.20$, $\chi^2/df = 1.46$, $CFI = .94$, $TLI = .92$, $RMSEA = .03$

Appendix D

Materials for Chapter 6

Table D.1. Four (untested) items for a generalised social identity continuity and fluidity scale. For a concrete example, replace [ingroup] with 'Hindu' and [category] with 'religion'. Scale anchors are highlighted in *italics*.

#	Item
1	How *simple* ("saying 'I am [ingroup]' is all there is to say about how I think about my [category]") or *complex* ("how I think about my [category] is more complex than saying 'I am [ingroup]'") is the way you think about your [category]?
2	Consider the way you think about your [category]. To what extent do you think this will *change* or *stay the same*?
3	How *similar* or *different* is the way you think and behave from the way people expect someone of your [category] to think and behave?
4	Consider the way your religion is reflected in your thoughts and behaviour. To what extent do you think this will *change* or *stay the same*?

References

15 held for assault on Dalits. (2017, 7th October). *The Hindu*. Retrieved from http://www.thehindu.com

Ainsworth, C. (2015). Sex redefined. *Nature*, *518*, 288–291. doi: 10.1038/518288a

Al Ramiah, A. & Hewstone, M. (2013). Intergroup contact as a tool for reducing, resolving, and preventing intergroup conflict: Evidence, limitations, and potential. *American Psychologist*, *68*, 527–542. doi: 10.1037/a0032603

Alexander, M. (2012). *The new Jim Crow: Mass incarceration in the age of colorblindness* (Revised ed.). New York, NY: New Press.

Allport, G. W. (1979). *The nature of prejudice*. New York, NY: Addison-Wesley. (Original work published 1954)

Altemeyer, B. (1981). *Right-wing authoritarianism*. Winnipeg, Canada: University of Manitoba Press.

Amnesty International. (2017, 28th June). *India: Hate crimes against Muslims and rising Islamophobia must be condemned.* Retrieved from http://www.amnesty.org/

Ashmore, R. D., Deaux, K. & McLaughlin-Volpe, T. (2004). An organizing framework for collective identity: Articulation and significance of multidimensionality. *Psychological Bulletin*, *130*, 80–114. doi: 10.1037/0033-2909.130.1.80

Asparouhov, T. & Muthén, B. (2009). Exploratory structural equation modeling. *Structural Equation Modeling: A Multidisciplinary Journal*, *16*, 397–438. doi: 10.1080/10705510903008204

Barlow, F. K., Paolini, S., Pedersen, A., Hornsey, M. J., Radke, H. R., Harwood, J., … Sibley, C. G. (2012). The contact caveat: Negative contact predicts increased prejudice more than positive contact predicts reduced prejudice. *Personality and Social Psychology Bulletin*, *38*, 1629–1643. doi: 10.1177/0146167212457953

Bastian, B. & Haslam, N. (2006). Psychological essentialism and stereotype endorsement. *Journal of Experimental Social Psychology*, *42*, 228–235. doi: 10.1016/j.jesp.2005.03.003

Baumeister, R. F. & Leary, M. R. (1995). The need to belong: Desire for interpersonal attachments as a fundamental human motivation. *Psychological Bulletin*, *117*, 497–529. doi: 10.1037/0033-2909.117.3.497

Becker, J. C. & Tausch, N. (2015). A dynamic model of engagement in normative and non-normative collective action: Psychological antecedents, consequences, and barriers. *European Review of Social Psychology*, *26*, 43–92. doi: 10.1080/10463283.2015.1094265

Becker, J. C., Wright, S. C., Lubensky, M. E. & Zhou, S. (2013). Friend or ally: Whether cross-group contact undermines collective action depends on what advantaged group members say (or don't say). *Personality and Social Psychology Bulletin*, *39*, 442–455. doi: 10.1177/0146167213477155

Berry, J. W. (1997). Immigration, acculturation, and adaptation. *Applied Psychology*, *46*, 5–34. doi: 10.1111/j.1464-0597.1997.tb01087.x

Betancourt, M. (2017, January). *A conceptual introduction to Hamiltonian Monte Carlo*. Retrieved from https://arxiv.org/abs/1701.02434

Bharatiya Janata Party. (n.d.). *Hindutva: The great nationalist ideology*. Retrieved April 16, 2018, from http://www.bjp.org/

Bharatiya Janata Party. (2014). In *Encyclopaedia Britannica*. Retrieved from https://www.britannica.com/topic/Bharatiya-Janata-Party

Bird, A. & Tobin, E. (2018). Natural kinds. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. Retrieved from https://plato.stanford.edu/archives/spr2018/entries/natural-kinds/

Birtel, M. D. & Crisp, R. J. (2012). "treating" prejudice: An exposure-therapy approach to reducing negative reactions toward stigmatized groups. *Psychological Science*, *23*, 1379–1386. doi: 10.1177/0956797612443838

Blumer, G. (1958). Race prejudice as a sense of group position. *The Pacific Sociological Review*, *1*, 3–7. doi: 10.2307/1388607

Branscombe, N. R., Ellemers, N., Spears, R. & Doosje, B. (1999). The context and content of social identity threat. In N. Ellemers, R. Spears & B. Doosje (Eds.), *Social identity: Context, commitment, content* (pp. 35–58). Oxford, England: Blackwell.

Brescoll, V. & LaFrance, M. (2004). The correlates and consequences of newspaper reports of research on sex differences. *Psychological Science*, *15*, 515–520. doi: 10.1111/j.0956-7976.2004.00712.x

Brewer, M. B. (1999). The psychology of prejudice: Ingroup love or outgroup hate? *Journal of Social Issues*, *55*, 429–444. doi: 10.1111/0022-4537.00126

Brewer, M. B. (2007). The social psychology of intergroup relations: Social categorization, ingroup bias, and outgroup prejudice. In A. W. Kruglanski & E. T. Higgins (Eds.), *Social psychology: Handbook of basic principles* (2nd ed., pp. 695–715). New York, NY: Guilford Press.

Brewer, M. B. & Miller, N. (1984). Beyond the contact hypothesis: Theoretical perspectives on desegregation. In N. Miller & M. B. Brewer (Eds.), *Groups in contact: The psychology of desegregation* (pp. 281–302). London, UK: Academic Press.

Brewer, M. B. & Pierce, K. P. (2005). Social identity complexity and outgroup tolerance. *Personality and Social Psychology Bulletin*, *31*, 428–437. doi: 10.1177/0146167204271710

Bürkner, P.-C. (2017). brms: An R package for Bayesian multilevel models using Stan. *Journal of Statistical Software*, *80*, 1–28. doi: 10.18637/jss.v080.i01

Brown, M. K., Carnoy, M., Currie, E., Duster, T., Oppenheimer, D. B., Shultz, M. M. & Wellman, D. (2003). *Whitewashing race: The myth of a color-blind society*. Berkeley, CA: University of California Press.

Brown, R. & Hewstone, M. (2005). An integrative theory of intergroup contact. *Advances in Experimental Social Psychology*, *37*, 255–343. doi: 10.1016/S0065-2601(05)37005-5

Butler, J. (1999). *Gender trouble: Feminism and the subversion of identity*. New York, NY: Routledge. (Original work published 1990)

Button, K. S., Ioannidis, J. P., Mokrysz, C., Nosek, B. A., Flint, J., Robinson, E. S. & Munafò, M. R. (2013). Power failure: Why small sample size undermines the reliability of neuroscience. *Nature Reviews Neuroscience*, *14*, 365–376. doi: 10.1038/nrn3475

Byrne, B. M. (2012). *Structural equation modeling with Mplus: Basic concepts, applications, and programming*. Hove, UK: Routledge.

Cakal, H., Hewstone, M., Schwär, G. & Heath, A. (2011). An investigation of the social identity model of collective action and the 'sedative' effect of intergroup contact among Black and White students in South Africa. *British Journal of Social Psychology*, *50*, 606–627. doi: 10.1111/j.2044-8309.2011.02075.x

Campbell, D. T. (1958). Common fate, similarity, and other indices of the status of aggregates of persons as social entities. *Systems Research and Behavioral Science*, *3*, 14–25. doi: 10.1002/bs.3830030103

Campbell, D. T. & Fiske, D. W. (1959). Convergent and discriminant validation by the multitrait–multimethod matrix. *Psychological Bulletin*, *56*, 81–105. doi: 10.1037/h0046016

Converse, P. E., Dotson, J. D., Hoag, W. J. & McGee, W. H. (1980). *American social attitudes data sourcebook, 1947–1978*. Cambridge, MA: Harvard University Press.

Cotterill, S., Sidanius, J., Bhardwajb, A. & Kumarc, V. (2014). Ideological support for the Indian caste system: Social dominance orientation, right-wing authoritarianism and karma. *Journal of Social and Political Psychology*, *1*, 98–116. doi: 10.5964/jspp.v2i1.171

Craig, M. A. & Richeson, J. A. (2012). Coalition or derogation? How perceived discrimination influences intraminority intergroup relations. *Journal of Personality and Social Psychology*, *102*, 759–777. doi: 10.1037/a0026481759

Crisp, R. J. (2010). Prejudice and perceiving multiple identities. In J. F. Dovidio, M. Hewstone, P. Glick & V. M. Esses (Eds.), *The SAGE handbook of prejudice, stereotyping and discrimination* (p. 508—525). London, UK: SAGE Publications.

Crisp, R. J. & Hewstone, M. (1999). Differential evaluation of crossed category groups: Patterns, processes, and reducing intergroup bias. *Group Processes and Intergroup Relations*, *2*, 303–333. doi: 10.1177/1368430299024001

Crisp, R. J. & Hewstone, M. (2007). Multiple social categorization. *Advances in Experimental Social Psychology*, *39*, 163–254. doi: 10.1016/S0065-2601(06)39004-1

Cuddy, A. J. C., Rock, M. S. & Norton, M. I. (2007). Aid in the aftermath of hurricane Katrina: Inferences of secondary emotions and intergroup helping. *Group Processes & Intergroup Relations*, *10*, 107–118. doi: 10.1177/1368430207071344

Cumming, G. (2014). The new statistics: Why and how. *Psychological Science*, *25*, 7–29. doi: 10.1177/0956797613504966

Davies, K., Tropp, L. R., Aron, A., Pettigrew, T. F. & Wright, S. C. (2011). Cross-group friendships and intergroup attitudes: A meta-analytic review. *Personality and Social Psychology Review*, *15*, 332–351. doi: 10.1177/1088868311411103

Deaux, K. (1996). Social identity. In A. W. Kruglanski & E. T. Higgins (Eds.), *Social psychology: Handbook of basic principles* (pp. 777–798). New York, NY: Guilford Press.

Deschamps, J. C. & Doise, W. (1978). Crossed category memberships in intergroup relations. In H. Tajfel (Ed.), *Differentiation between social groups: Studies in the social psychology of intergroup relations* (p. 141—158). London, UK: Academic Press.

DeVellis, R. F. (2016). *Scale development: Theory and applications* (4th ed.). London, UK: SAGE Publications.

Dhattiwala, R. & Biggs, M. (2012). The political logic of ethnic violence: The anti-Muslim pogrom in Gujarat, 2002. *Politics & Society*, *40*, 483–516. doi: 10.1177/0032329212461125

Dhont, K. & Van Hiel, A. (2009). We must not be enemies: Interracial contact and the reduction of prejudice among authoritarians. *Personality and Individual Differences*,

46, 172–177. doi: 10.1016/j.paid.2008.09.022

Dixon, J., Durrheim, K., Thomae, M., Tredoux, C., Kerr, P. & Quayle, M. (2015). Divide and rule, unite and resist: Contact, collective action and policy attitudes among historically disadvantaged groups. *Journal of Social Issues*, *71*, 576–596. doi: 10.1111/josi.12129

Dixon, J., Durrheim, K. & Tredoux, C. (2007). Intergroup contact and attitudes toward the principle and practice of racial equality. *Psychological Science*, *18*, 867–872. doi: 10.1111/j.1467-9280.2007.01993.x

Dixon, J., Durrheim, K., Tredoux, C., Tropp, L., Clack, B. & Eaton, L. (2010). A paradox of integration? Interracial contact, prejudice reduction, and perceptions of racial discrimination. *Journal of Social Issues*, *66*, 401–416. doi: 10.1111/j.1540-4560.2010.01652.x

Dixon, J., Levine, M., Reicher, S. & Durrheim, K. (2012). Beyond prejudice: Are negative evaluations the problem and is getting us to like one another more the solution? *Behavioral and Brain Sciences*, *35*, 411–425. doi: 10.1017/S0140525X11002214

Doosje, B., Branscombe, N. R., Spears, R. & Manstead, A. S. R. (1998). Guilty by association: When one's group has a negative history. *Journal of Personality and Social Psychology*, *75*, 872–886. doi: 10.1037/0022-3514.75.4.872

Dovidio, J. F., Gaertner, S. L. & Saguy, T. (2009). Commonality and the complexity of "we": Social attitudes and social change. *Personality and Social Psychology Review*, *13*, 3–20. doi: 10.1177/1088868308326751

Dovidio, J. F., Love, A., Schellhaas, F. M. & Hewstone, M. (2017). Reducing intergroup bias through intergroup contact: Twenty years of progress and future directions. *Group Processes & Intergroup Relations*, *20*, 606–620. doi: 10.1177/1368430217712052

Dovidio, J. F., Saguy, T., Gaertner, S. L. & Thomas, E. L. (2012). From attitudes to (in)action: The darker side of "we". In J. Dixon & M. Levine (Eds.), *Beyond prejudice: Extending the social psychology of conflict, inequality and social change* (pp. 248–268). Cambridge, UK: Cambridge University Press.

Doyen, S., Klein, O., Pichon, C.-L. & Cleeremans, A. (2012). Behavioral priming: It's all in the mind, but whose mind? *PloS ONE*, *7*, e29081. doi: 10.1371/journal.pone.0029081

Duckitt, J. & Bizumic, B. (2013). Multidimensionality of right-wing authoritarian attitudes: Authoritarianism-conservatism-traditionalism. *Political Psychology*, *34*, 841–862. doi: 10.1111/pops.12022

Dunham, Y. & Olson, K. R. (2016). Beyond discrete categories: Studying multiracial, intersex, and transgender children will strengthen basic developmental science. *Journal of Cognition and Development*, *17*, 642–665. doi: 10.1080/15248372.2016.1195388

Dunn, T. J., Baguley, T. & Brunsden, V. (2014). From alpha to omega: A practical solution to the pervasive problem of internal consistency estimation. *British Journal of Psychology*, *105*, 399–412. doi: 10.1111/bjop.12046

Falomir-Pichastor, J. M. & Hegarty, P. (2014). Maintaining distinctions under threat: Heterosexual men endorse the biological theory of sexuality when equality is the norm. *British Journal of Social Psychology*, *53*, 731–751. doi: 10.1111/bjso.12051

Falomir-Pichastor, J. M. & Mugny, G. (2009). "I'm not gay.... I'm a real man!": Heterosexual men's gender self-esteem and sexual prejudice. *Personality and Social Psychology Bulletin*, *35*, 1233–1243. doi: 10.1177/0146167209338072

Fell, B. (2015). *The interaction of positive and negative intergroup contact* (Doctoral dissertation, University of Oxford). Retrieved from `https://ora.ox.ac.uk/objects/uuid:f346d455-8e16-44b2-9fd1-3b8d332f0983`

Fell, B., Lolliot, S., Schmid, K., Biro, M., Christ, O. & Hewstone, M. (2016). *Interaction of positive and negative contact.* Unpublished manuscript.

Fingerhut, A. W. (2011). Straight allies: What predicts heterosexuals' alliance with the LGBT community? *Journal of Applied Social Psychology*, *41*, 2230–2248. doi: doi:10.1111/j.1559-1816.2011.00807.x

Fiske, S. T. (2015). Intergroup biases: A focus on stereotype content. *Current Opinion in Behavioral Sciences*, *3*, 45–50. doi: 10.1016/j.cobeha.2015.01.010

Fiske, S. T., Cuddy, A. J. C., 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*, 878–902. doi: 10.1037/0022-3514.82.6.878

Flake, J. K., Pek, J. & Hehman, E. (2017). Construct validation in social and personality research: Current practice and recommendations. *Social Psychological and Personality Science*, *8*, 370–378. doi: 10.1177/1948550617693063

Foucault, M. (1978). *The history of sexuality* (Vol. 1: An introduction). New York, NY: Pantheon Books.

Friedman, C. & Leaper, C. (2010). Sexual-minority college women's experiences with discrimination: Relations with identity and collective action. *Psychology of Women Quarterly*, *34*, 152–164. doi: 10.1111/j.1471-6402.2010.01558.x

Gaertner, S. L. & Dovidio, J. F. (2000). *Reducing intergroup bias: The common ingroup identity model.* Hove, UK: Psychology Press.

Gaertner, S. L., Dovidio, J. F., Anastasio, P. A., Bachman, B. A. & Rust, M. C. (1993). The common ingroup identity model: Recategorization and the reduction of intergroup bias. *European Review of Social Psychology*, *4*, 1–26. doi: 10.1080/14792779343000004

Gaertner, S. L., Dovidio, J. F., Guerra, R., Hehman, E. & Saguy, T. (2016). A common ingroup identity: Categorization, identity, and intergroup relations. In T. N. Nelson (Ed.), *Handbook of prejudice, stereotyping, and discrimination* (2nd ed., pp. 433–454). Hove, UK: Psychology Press.

Garstka, T. A., Schmitt, M. T., Branscombe, N. R. & Hummert, M. L. (2004). How young and older adults differ in their responses to perceived age discrimination. *Psychology and Aging*, *19*, 326–335. doi: 10.1037/0882-7974.19.2.326

Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A. & Rubin, D. B. (2014). *Bayesian data analysis* (3rd ed.). Boca Raton, FL: CRC Press.

Gelman, A., Goodrich, B., Gabry, J. & Ali, I. (2017, November). *R-squared for Bayesian regression models.* Retrieved from `http://www.stat.columbia.edu/~gelman/research/unpublished/bayes_R2.pdf`

Gelman, A., Hill, J. & Yajima, M. (2012). Why we (usually) don't have to worry about multiple comparisons. *Journal of Research on Educational Effectiveness*, *5*, 189-211. doi: 10.1080/19345747.2011.618213

Gelman, A. & Loken, E. (2013, November). *The garden of forking paths: Why multiple comparisons can be a problem, even when there is no "fishing expedition" or "p-hacking" and the research hypothesis was posited ahead of time.* Retrieved from `http://www.stat.columbia.edu/~gelman/research/unpublished/p_hacking.pdf`

Gelman, A. & Nolan, D. (2002). You can load a die, but you can't bias a coin. *The American Statistician*, *56*, 308–311. doi: 10.1198/000313002605

Gelman, A. & Shalizi, C. R. (2013). Philosophy and the practice of bayesian statistics. *British Journal of Mathematical and Statistical Psychology*, *66*, 8–38. doi: 10.1111/j.2044-8317.2011.02037.x

Gelman, A., Simpson, D. & Betancourt, M. (2017). The prior can often only be understood in the context of the likelihood. *Entropy*, *19*, 555. doi: 10.3390/e19100555

Gelman, S. A. (2003). *The essential child: Origins of essentialism in everyday life*. Oxford, UK: Oxford University Press.

Gigerenzer, G. (2004). Mindless statistics. *The Journal of Socio-Economics*, *33*, 587–606. doi: 10.1016/j.socec.2004.09.033

Glick, P. & Fiske, S. T. (2001). An ambivalent alliance: Hostile and benevolent sexism as complementary justifications for gender inequality. *American Psychologist*, *56*, 109–118. doi: 10.1037/0003-066X.56.2.109

Goh, J. X., Hall, J. A. & Rosenthal, R. (2016). Mini meta-analysis of your own studies: Some arguments on why and a primer on how. *Social and Personality Psychology Compass*, *10*, 535–549. doi: 10.1111/spc3.12267

Graf, S., Paolini, S. & Rubin, M. (2014). Negative intergroup contact is more influential, but positive intergroup contact is more common: Assessing contact prominence and contact prevalence in five Central European countries. *European Journal of Social Psychology*, *44*, 536–547. doi: 10.1002/ejsp.2052

Graham, J., Waytz, A., Meindl, P., Iyer, R. & Young, L. (2017). Centripetal and centrifugal forces in the moral circle: Competing constraints on moral learning. *Cognition*, *167*, 58–65. doi: 10.1016/j.cognition.2016.12.001

Greenwald, A. G. & Pettigrew, T. F. (2014). With malice toward none and charity for some: Ingroup favoritism enables discrimination. *American Psychologist*, *69*, 669–684. doi: 10.1037/a0036056

Hagger, M. S., Chatzisarantis, N. L., Alberts, H., Anggono, C. O., Batailler, C., Birt, A. R., … Zwienenberg, M. (2016). A multilab preregistered replication of the ego-depletion effect. *Perspectives on Psychological Science*, *11*, 546–573. doi: 10.1177/1745691616652873

Haslam, N. & Levy, S. R. (2006). Essentialist beliefs about homosexuality: Structure and implications for prejudice. *Personality and Social Psychology Bulletin*, *32*, 471–485. doi: 10.1177/0146167205276516

Haslam, N., Rothschild, L. & Ernst, D. (2000). Essentialist beliefs about social categories. *British Journal of Social Psychology*, *39*, 113–127. doi: 10.1348/014466600164363

Haslam, N., Rothschild, L. & Ernst, D. (2002). Are essentialist beliefs associated with prejudice? *British Journal of Social Psychology*, *41*, 87–100. doi: 10.1348/014466602165072

Hayward, L. E., Tropp, L. R., Hornsey, M. J. & Barlow, F. K. (2017). Toward a comprehensive understanding of intergroup contact: Descriptions and mediators of positive and negative contact among majority and minority groups. *Personality and Social Psychology Bulletin*, *43*, 347–364. doi: 10.1177/0146167216685291

Hayward, L. E., Tropp, L. R., Hornsey, M. J. & Barlow, F. K. (2018). How negative contact and positive contact with Whites predict collective action among racial and ethnic minorities. *British Journal of Social Psychology*, *57*, 1–20. doi: 10.1111/bjso.12220

Hegarty, P. (2002). 'It's not a choice, it's the way we're built': Symbolic beliefs about sexual orientation in the US and Britain. *Journal of Community & Applied Social Psychology*, *12*, 153–166. doi: 10.1002/casp.669

Henrich, J., Heine, S. J. & Norenzayan, A. (2010). The weirdest people in the world? *Behavioral and Brain Sciences*, *33*, 61–83. doi: 10.1017/S0140525X0999152X

Herek, G. M. & McLemore, K. A. (2013). Sexual prejudice. *Annual Review of Psychology*, *64*, 309–333. doi: 10.1146/annurev-psych-113011-143826

Hewstone, M. (1990). The 'ultimate attribution error'? A review of the literature on intergroup causal attribution. *European Journal of Social Psychology*, *20*, 311–335. doi: 10.1002/ejsp.2420200404

Hewstone, M. (2009). Living apart, living together? the role of intergroup contact in social integration. *Proceedings of the British Academy*, *162*, 243–300.

Hewstone, M. & Brown, R. (1986). Contact is not enough: An intergroup perspective on the 'contact hypothesis'. In M. Hewstone & R. Brown (Eds.), *Contact and conflict in intergroup encounters* (pp. 1–44). Oxford, UK: Blackwell.

Hewstone, M., Rubin, M. & Willis, H. (2002). Intergroup bias. *Annual Review of Psychology*, *53*, 575–604. doi: 10.1146/annurev.psych.53.100901.135109

Hewstone, M. & Swart, H. (2011). Fifty-odd years of inter-group contact: From hypothesis to integrated theory. *British Journal of Social Psychology*, *50*, 374–386. doi: 10.1111/j.2044-8309.2011.02047.x

Ho, A. K., Sidanius, J., Kteily, N., Sheehy-Skeffington, J., Pratto, F., Henkel, K. E., … Stewart, A. L. (2015). The nature of social dominance orientation: Theorizing and measuring preferences for intergroup inequality using the new SDO_7 scale. *Journal of Personality and Social Psychology*, *109*, 1003–1028. doi: 10.1037/pspi0000033

Ho, A. K., Sidanius, J., Levin, D. T. & Banaji, M. R. (2011). Evidence for hypodescent and racial hierarchy in the categorization and perception of biracial individuals. *Journal of Personality and Social Psychology*, *100*, 492–506. doi: 10.1037/a0021562

Hoffarth, M. R. & Hodson, G. (2018). When intergroup contact is uncommon and bias is strong: The case of anti-transgender bias. *Psychology & Sexuality*. (Advance online publication) doi: 10.1080/19419899.2018.1470107

Insko, C. A. & Schopler, J. (1998). Differential distrust of groups and individuals. In C. Sedikides, J. Schopler & C. A. Insko (Eds.), *Intergroup cognition and intergroup behavior* (pp. 75–107). Mahwah, NJ: Lawrence Erlbaum Associates.

Ioannidis, J. P. A. (2005). Why most published research findings are false. *PLoS Medicine*, *2*, e124. doi: 10.1371/journal.pmed.0020124

Islam, M. R. & Hewstone, M. (1993). Dimensions of contact as predictors of intergroup anxiety, perceived out-group variability, and out-group attitude: An integrative model. *Personality and Social Psychology Bulletin*, *19*, 700–710. doi: 10.1177/0146167293196005

Jodhka, S. S. (2012). *Caste*. New Delhi, India: Oxford University Press.

Jost, J. T., Banaji, M. R. & Nosek, B. A. (2004). A decade of system justification theory: Accumulated evidence of conscious and unconscious bolstering of the status quo. *Political Psychology*, *25*, 881–919. doi: 10.1111/j.1467-9221.2004.00402.x

Jost, J. T. & Kay, A. C. (2005). Exposure to benevolent sexism and complementary gender stereotypes: Consequences for specific and diffuse forms of system justification. *Journal of Personality and Social Psychology*, *88*, 498–509. doi: 10.1037/0022-3514.88.3.498

Kelley, K. (2018). MBESS: The MBESS R package [Computer software manual]. Retrieved from `https://CRAN.R-project.org/package=MBESS` (R package version 4.4.3)

King, M. E., Winter, S. & Webster, B. (2009). Contact reduces transprejudice: A study on attitudes towards transgenderism and transgender civil rights in Hong Kong. *International Journal of Sexual Health*, *21*, 17–34. doi: 10.1080/19317610802434609

Kinsey, A. C., Pomeroy, W. B. & Martin, C. E. (1948). *Sexual behavior in the human male*. Philadelphia, PA: W. B. Saunders.

Kinsey, A. C., Pomeroy, W. B., Martin, C. E. & Gebhard, P. H. (1953). *Sexual behavior in the human female*. Philadelphia, PA: W. B. Saunders.

Knowles, E. D., Lowery, B. S., Chow, R. M. & Unzueta, M. M. (2014). Deny, distance, or dismantle? How white Americans manage a privileged identity. *Perspectives on Psychological Science*, *9*, 594–609. doi: 10.1177/1745691614554658

Kraus, M. W., Rucker, J. M. & Richeson, J. A. (2017). Americans misperceive racial economic equality. *Proceedings of the National Academy of Sciences*, *114*, 10324–10331. doi: 10.1073/pnas.1707719114

Kruschke, J. K. & Liddell, T. M. (2018). The Bayesian new statistics: Hypothesis testing, estimation, meta-analysis, and power analysis from a Bayesian perspective. *Psychonomic Bulletin & Review*, *25*, 178–206. doi: 10.3758/s13423-016-1221-4

Lambert, B. (2018). *A student's guide to Bayesian statistics*. London, UK: SAGE Publications.

Leach, C. W., Snider, N. & Iyer, A. (2007). "Poisoning the conscience of the fortunate": The experience of relative advantage and support for social equality. In I. Walker & H. J. Smith (Eds.), *Relative deprivation: Specification, development, and integration* (pp. 136–163). Cambridge, UK: Cambridge University Press.

Leach, C. W., van Zomeren, M., Zebel, S., Vliek, M. L., Pennekamp, S. F., Doosje, B., … Spears, R. (2008). Group-level self-definition and self-investment: A hierarchical (multicomponent) model of in-group identification. *Journal of Personality and Social Psychology*, *95*, 144–165. doi: 10.1037/0022-3514.95.1.144

Lemmer, G. & Wagner, U. (2015). Can we really reduce ethnic prejudice outside the lab? a meta-analysis of direct and indirect contact interventions. *European Journal of Social Psychology*, *45*, 152–168. doi: 10.1002/ejsp.2079

Lenzner, T. & Menold, N. (2016). *GESIS survey guidelines: Question wording.* Mannheim, Germany: GESIS – Leibniz Institute for the Social Sciences. doi: 10.15465/gesis-sg_en_017

Levine, M., Prosser, A., Evans, D. & Reicher, S. (2005). Identity and emergency intervention: How social group membership and inclusiveness of group boundaries shape helping behavior. *Personality and Social Psychology Bulletin*, *31*, 443–453. doi: 10.1177/0146167204271651

Lickel, B., Hamilton, D. L., Wieczorkowska, G., Lewis, A., Sherman, S. J. & Uhles, A. N. (2000). Varieties of groups and the perception of group entitativity. *Journal of Personality and Social Psychology*, *78*, 223–246. doi: 10.1037/0022-3514.78.2.223

Little, T. D. (2013). *Longitudinal structural equation modeling*. New York, NY: Guilford Press.

Little, T. D., Cunningham, W. A., Shahar, G. & Widaman, K. F. (2002). To parcel or not to parcel: Exploring the question, weighing the merits. *Structural Equation Modeling: A Multidisciplinary Journal*, *9*, 151–173. doi: 10.1207/S15328007SEM0902_1

Loken, E. & Gelman, A. (2017). Measurement error and the replication crisis. *Science*, *355*, 584–585. doi: 10.1126/science.aal3618

Lolliot, S., Fell, B., Schmid, K., Wölfer, R., Swart, H., Voci, A., … Hewstone, M. (2015). Measures of intergroup contact. In G. J. Boyle, D. H. Saklofske & G. Matthews (Eds.), *Measures of personality and social psychological constructs* (pp. 652–683). San Diego, CA: Academic Press.

Ma, D. S., Correll, J. & Wittenbrink, B. (2015). The Chicago face database: A free stimulus set of faces and norming data. *Behavior Research Methods*, *47*, 1122–1135. doi: 10.3758/s13428-014-0532-5

Macrae, C. N. & Bodenhausen, G. V. (2000). Social cognition: Thinking categorically about others. *Annual Review of Psychology*, *51*, 93–120. doi: 10.1146/annurev.psych.51.1.93

Major, B., Kaiser, C. R. & McCoy, S. K. (2003). It's not my fault: When and why attributions to prejudice protect self-esteem. *Personality and Social Psychology Bulletin*, *29*, 772–781. doi: 10.1177/0146167203029006009

Martin, D. & Macrae, C. N. (2007). A face with a cue: Exploring the inevitability of person categorization. *European Journal of Social Psychology*, *37*, 806–816. doi: 10.1002/ejsp.445

Mays, V. M. & Cochran, S. D. (2001). Mental health correlates of perceived discrimination among lesbian, gay, and bisexual adults in the United States. *American Journal of Public Health*, *91*, 1869-1876. doi: 10.2105/AJPH.91.11.1869

McDonald, R. P. (1999). *Test theory: A unified treatment*. Mahwah, NJ: Lawrence Erlbaum.

McElreath, R. (2016). *Statistical rethinking: A Bayesian course with examples in R and Stan*. Boca Raton, FL: CRC Press.

McShane, B. B., Gal, D., A. Gelman, C. R. & Tackett, J. L. (2018, April). *Abandon statistical significance*. Retrieved from `https://arxiv.org/abs/1709.07588v2`

Medin, D. & Ortony, A. (1989). Psychological essentialism. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 179–195). Cambridge, UK: Cambridge University Press.

Menold, N. & Bogner, K. (2016). *GESIS survey guidelines: Design of rating scales in questionnaires*. Mannheim, Germany: GESIS – Leibniz Institute for the Social Sciences. doi: 10.15465/gesis-sg_en_015

Menski, W. (2009). Hindu nationalism. In S. N. Katz (Ed.), *The Oxford international encyclopedia of legal history*. Retrieved from `http://www.oxfordreference.com/view/10.1093/acref/9780195134056.001.0001/acref-9780195134056-e-366`

Merkle, E. C. & Rosseel, Y. (2016, November). *blavaan: Bayesian structural equation models via parameter expansion*. Retrieved from `https://arxiv.org/abs/1511.05604`

Mezulis, A. H., Abramson, L. Y., Hyde, J. S. & Hankin, B. L. (2004). Is there a universal positivity bias in attributions? A meta-analytic review of individual, developmental, and cultural differences in the self-serving attributional bias. *Psychological Bulletin*, *130*, 711–747. doi: 10.1037/0033-2909.130.5.711

Migdal, M. J., Hewstone, M. & Mullen, B. (1998). The effects of crossed categorization on intergroup evaluations: A meta-analysis. *British Journal of Social Psychology*, *37*, 303–324. doi: 10.1111/j.2044-8309.1998.tb01174.x

Mischel, W. (1977). The interaction of person and situation. In D. Magnusson & N. S. Endler (Eds.), *Personality at the crossroads: Current issues in interactional psychology* (pp. 333–352). Hillsdale, NJ: Lawrence Erlbaum.

Morey, R. D., Hoekstra, R., Rouder, J. N., Lee, M. D. & Wagenmakers, E. J. (2016). The fallacy of placing confidence in confidence intervals. *Psychonomic Bulletin & Review*, *23*, 103–123. doi: 10.3758/s13423-015-0947-8

Mummendey, A., Simon, B., Dietze, C., Grünert, M., Haeger, G., Kessler, S., … Schäferhoff, S. (1992). Categorization is not enough: Intergroup discrimination in negative

outcome allocation. *Journal of Experimental social psychology*, *28*, 125–144. doi: 10.1016/0022-1031(92)90035-I

Munafò, M. R., Nosek, B. A., Bishop, D. V. M., Button, K. S., Chambers, C. D., du Sert, N. P., … Ioannidis, J. P. A. (2017). A manifesto for reproducible science. *Nature Human Behaviour*, *1*, 0021. doi: 10.1038/s41562-016-0021

Muthén, B. O. & Asparouhov, T. (2012). Bayesian structural equation modeling: A more flexible representation of substantive theory. *Psychological Methods*, *17*, 313–335. doi: 10.1037/a0026802

Muthén, L. K. & Muthén, B. O. (1998–2015). *Mplus user's guide* (7th ed.). Los Angeles, CA: Muthén & Muthén.

Nagoshi, J. L., Adams, K. A., Terrell, H. K., Hill, E. D., Brzuzy, S. & Nagoshi, C. T. (2008). Gender differences in correlates of homophobia and transphobia. *Sex Roles*, *59*, 521–531. doi: 10.1007/s11199-008-9458-7

National Commission for Backward Classes. (2015). *Central list of OBCs of the state of Karnataka.* Government of India. Retrieved from `http://www.ncbc.nic.in`

National Crime Records Bureau. (2016). *Crime in India* [Data set]. Ministry of Home Affairs, Government of India. Retrieved from `http://ncrb.gov.in/`

Norton, M. I. & Sommers, S. R. (2011). Whites see racism as a zero-sum game that they are now losing. *Perspectives on Psychological Science*, *6*, 215–218. doi: 10.1177/1745691611406922

Nosek, B. A., Ebersole, C. R., DeHaven, A. C. & Mellor, D. T. (2018). The preregistration revolution. *Proceedings of the National Academy of Sciences*, *115*, 2600–2606. doi: 10.1073/pnas.1708274114

Oakes, P. J. (1987). The salience of social categories. In J. C. Turner, M. A. Hogg, P. J. Oakes, S. D. Reicher & M. S. Wetherell (Eds.), *Rediscovering the social group: A self-categorization theory* (pp. 117–141). Oxford, UK: Blackwell.

OBC creamy layer income limit raised to ₹8 lakh per annum. (2017, 13th September). *The Hindu*. Retrieved from `http://www.thehindu.com`

Office of the Registrar General & Census Commissioner. (2011). *2011 census of India* [Data set]. Ministry of Home Affairs, Government of India. Retrieved from `http://censusindia.gov.in/`

Open Science Collaboration. (2015). Estimating the reproducibility of psychological science. *Science*, *349*, aac4716. doi: 10.1126/science.aac4716

Operario, D. & Fiske, S. T. (1998). Racism equals power plus prejudice: A social psychological equation for racial oppression. In J. L. Eberhardt & S. T. Fiske (Eds.), *Confronting racism: The problem and the response* (pp. 33–53)). London, UK: SAGE Publications.

Otten, S. & Moskowitz, G. B. (2000). Evidence for implicit evaluative in-group bias: Affect-biased spontaneous trait inference in a minimal group paradigm. *Journal of Experimental Social Psychology*, *36*, 77–89. doi: 10.1006/jesp.1999.1399

Page-Gould, E., Mendoza-Denton, R., Alegre, J. M. & Siy, J. O. (2010). Understanding the impact of cross-group friendship on interactions with novel outgroup members. *Journal of Personality and Social Psychology*, *98*, 775–793. doi: 10.1037/a0017880

Paolini, S., Harwood, J. & Rubin, M. (2010). Negative intergroup contact makes group memberships salient: Explaining why intergroup conflict endures. *Personality and Social Psychology Bulletin*, *36*, 1723–1738. doi: 10.1177/0146167210388667

Paolini, S., Harwood, J., Rubin, M., Husnu, S., Joyce, N. & Hewstone, M. (2014). Positive and extensive intergroup contact in the past buffers against the disproportionate

impact of negative contact in the present. *European Journal of Social Psychology*, *44*, 548–562. doi: 10.1002/ejsp.2029

Pathania, G. J. & Tierney, W. G. (2018). An ethnography of caste and class at an Indian university: Creating capital. *Tertiary Education and Management*. (Advance online publication) doi: 10.1080/13583883.2018.1439998

Pettigrew, T. F. (1997). Generalized intergroup contact effects on prejudice. *Personality and Social Psychology Bulletin*, *23*, 173–185. doi: 10.1177/0146167297232006

Pettigrew, T. F. (1998). Intergroup contact theory. *Annual Review of Psychology*, *49*, 65—85. doi: 10.1146/annurev.psych.49.1.65

Pettigrew, T. F. & Hewstone, M. (2017). The single factor fallacy: Implications of missing critical variables from an analysis of intergroup contact theory. *Social Issues and Policy Review*, *11*, 8–37. doi: 10.1111/sipr.12026

Pettigrew, T. F. & Tropp, L. R. (2006). A meta-analytic test of intergroup contact theory. *Journal of Personality and Social Psychology*, *90*, 751–783. doi: 10.1037/0022-3514.90.5.751

Pettigrew, T. F. & Tropp, L. R. (2008). How does intergroup contact reduce prejudice? Meta-analytic tests of three mediators. *European Journal of Social Psychology*, *38*, 922–934. doi: 10.1002/ejsp.504

Pettigrew, T. F. & Tropp, L. R. (2011). *When groups meet: The dynamics of intergroup contact*. Hove, UK: Psychology Press.

Pew Research Center. (2017, November). *Three years in, Modi remains very popular.* Retrieved from `http://www.pewresearch.org/`

Pittinsky, T. L. & Montoya, R. M. (2009). Is valuing equality enough? equality values, allophilia, and social policy support for multiracial individuals. *Journal of Social Issues*, *65*, 151–163. doi: 10.1111/j.1540-4560.2008.01592.x

Plummer, M., Best, N., Cowles, K. & Vines, K. (2006). CODA: Convergence diagnosis and output analysis for MCMC. *R News*, *6*, 7–11. Retrieved from `https://journal.r-project.org/archive/`

Poore, A. G., Gagne, F., Barlow, K. M., Lydon, J. E., Taylor, D. M. & Wright, S. C. (2002). Contact and the personal/group discrimination discrepancy in an inuit community. *Journal of Psychology*, *136*, 371–382. doi: 10.1080/00223980209604164

Postmes, T., Haslam, S. A. & Jans, L. (2013). A single-item measure of social identification: Reliability, validity, and utility. *British Journal of Social Psychology*, *52*, 597–617. doi: 10.1111/bjso.12006

Pratto, F., Sidanius, J. & Levin, S. (2006). Social dominance theory and the dynamics of intergroup relations: Taking stock and looking forward. *European Review of Social Psychology*, *17*, 271–320. doi: 10.1080/10463280601055772

Preacher, K. J. & Hayes, A. F. (2008). Asymptotic and resampling strategies for assessing and comparing indirect effects in multiple mediator models. *Behavior Research Methods*, *40*, 879–891. doi: 10.3758/BRM.40.3.879

Prentice, D. A. & Miller, D. T. (2007). Psychological essentialism of human categories. *Current Directions in Psychological Science*, *16*, 202–206. doi: 10.1111/j.1467-8721.2007.00504.x

R Core Team. (2017). *R: A language and environment for statistical computing.* Vienna, Austria. Retrieved from `https://www.R-project.org/`

Reimer, N. K., Becker, J. C., Benz, A., Christ, O., Dhont, K., Klocke, U., … Hewstone, M. (2017). Intergroup contact and social change: Implications of negative and positive contact for collective action in advantaged and disadvantaged groups. *Personality*

and Social Psychology Bulletin, *43*, 121-136. doi: 10.1177/0146167216676478

Reimer, N. K., Hospers, H. & Kok, G. J. (2014). *Prejudice and the 'gay gene': The implications of genetic essentialism for anti-gay prejudice and gay stereotype endorsement.* Unpublished manuscript.

Richards, C., Bouman, W. P., Seal, L., Barker, M. J., Nieder, T. O. & T'Sjoen, G. (2016). Non-binary or genderqueer genders. *International Review of Psychiatry*, *28*, 95–102. doi: 10.3109/09540261.2015.1106446

Richeson, J. A. & Sommers, S. R. (2016). Toward a social psychology of race and race relations for the twenty-first century. *Annual Review of Psychology*, *67*, 439–463. doi: 10.1146/annurev-psych-010213-115115

Richeson, J. A. & Trawalter, S. (2005). On the categorization of admired and disliked exemplars of admired and disliked racial groups. *Journal of Personality and Social Psychology*, *89*, 517–530. doi: 10.1037/0022-3514.89.4.517

Roberts, S. O., Ho, A. K., Rhodes, M. & Gelman, S. A. (2017). Making boundaries great again: Essentialism and support for boundary-enhancing initiatives. *Personality and Social Psychology Bulletin*, *43*, 1643–1658. doi: 10.1177/0146167217724801

Roccas, S. & Brewer, M. B. (2002). Social identity complexity. *Personality and Social Psychology Review*, *6*, 88–106. doi: 10.1207/S15327957PSPR0602_01

Roets, A. & van Hiel, A. (2011). Item selection and validation of a brief, 15-item version of the need for closure scale. *Personality and Individual Differences*, *50*, 90–94. doi: 10.1016/j.paid.2010.09.004

Rosseel, Y. (2012). lavaan: An R package for structural equation modeling. *Journal of Statistical Software*, *48*, 1–36. doi: 10.18637/jss.v048.i02

Rothbart, M. & Taylor, M. (1992). Category labels and social reality: Do we view social categories as natural kinds? In G. R. Semin & K. Fiedler (Eds.), *Language, interaction and social cognition* (pp. 11–36). London, UK: SAGE Publications.

Russell, G. M. (2011). Motives of heterosexual allies in collective action for equality. *Journal of Social Issues*, *67*, 376–393. doi: 10.1111/j.1540-4560.2011.01703.x

Saguy, T., Tausch, N., Dovidio, J. F. & Pratto, F. (2009). The irony of harmony: Intergroup contact can produce false expectations for equality. *Psychological Science*, *20*, 114–121. doi: 10.1111/j.1467-9280.2008.02261.x

Sanchez, D. T. & Crocker, J. (2005). How investment in gender ideals affects well-being: The role of external contingencies of self-worth. *Psychology of Women Quarterly*, *29*, 63–77. doi: 10.1111/j.1471-6402.2005.00169.x

Satorra, A. & Bentler, P. M. (2001). A scaled difference chi-square test statistic for moment structure analysis. *Psychometrika*, *66*, 507–514. doi: 10.1007/BF02296192

Schmid, K. & Hewstone, M. (2011). Social identity complexity: Theoretical implications for the social psychology of intergroup relations. In R. M. Kramer, G. J. Leonardelli & R. W. Livingston (Eds.), *Social cognition, social identity, and intergroup relations: A festschrift in honor of Marilynn B. Brewer* (p. 77—102). Hove, UK: Psychology Press.

Schmid, K., Hewstone, M. & Al Ramiah, A. (2013). Neighborhood diversity and social identity complexity: Implications for intergroup relations. *Social Psychological and Personality Science*, *4*, 135–142. doi: 10.1177/1948550612446972

Schmid, K., Hewstone, M., Küpper, B., Zick, A. & Tausch, N. (2014). Reducing aggressive intergroup action tendencies: Effects of intergroup contact via perceived intergroup threat. *Aggressive Behavior*, *40*, 250–262. doi: 10.1002/ab.21516

Schmid, K., Hewstone, M., Tausch, N., Cairns, E. & Hughes, J. (2009). Antecedents and consequences of social identity complexity: Intergroup contact, distinctiveness threat, and outgroup attitudes. *Personality and Social Psychology Bulletin*, *35*, 1085–1098. doi: 10.1177/0146167209337037

Sell, R. L. (2007). Defining and measuring sexual orientation for research. In I. H. Meyer & M. E. Northridge (Eds.), *The health of sexual minorities: Public health perspectives on lesbian, gay, bisexual and transgender populations* (pp. 355–374). New York, NY: Springer.

Selvanathan, H. P., Techakesari, P., Tropp, L. R. & Barlow, F. K. (2017). Whites for racial justice: How contact with Black Americans predicts support for collective action among White Americans. *Group Processes & Intergroup Relations*. (Advance online publication) doi: 10.1177/1368430217690908

Sengupta, N. K. & Sibley, C. G. (2013). Perpetuating one's own disadvantage: Intergroup contact enables the ideological legitimation of inequality. *Personality and Social Psychology Bulletin*, *39*, 1391–1403. doi: 10.1177/0146167213497593

Shah, G., Mander, H., Thorat, S., Deshpande, S. & Baviskar, A. (2006). *Untouchability in rural India*. New Delhi, India: SAGE Publications.

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*, e56515. doi: 10.1371/journal.pone.0056515

Sherif, M., Harvey, O. J., White, B. J., Hood, W. R. & Sherif, C. W. (1961). *Intergroup conflict and cooperation: The Robbers Cave experiment*. Retrieved from `http://psychclassics.yorku.ca/Sherif/` (Original work published 1954)

Sidanius, J., Levin, S., van Laar, C. & Seers, D. O. (2008). *The diversity challenge: Social identity and intergroup relations on the college campus*. New York, NY: Russell Sage Foundation.

Sidanius, J. & Pratto, F. (1999). *Social dominance : An intergroup theory of social hierarchy and oppression*. Cambridge, UK: Cambridge University Press.

Simmons, J. P., Nelson, L. D. & Simonsohn, U. (2011). False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological Science*, *22*, 1359–1366. doi: 10.1177/0956797611417632

Singer, P. (2011). *The expanding circle: Ethics, evolution, and moral progress*. Princeton, NJ: Princeton University Press. (Original work published 1981)

Smaldino, P. E. & McElreath, R. (2016). The natural selection of bad science. *Royal Society Open Science*, *3*, 160384. doi: 10.1098/rsos.160384

Smith, S. J., Axelton, A. M. & Saucier, D. A. (2009). The effects of contact on sexual prejudice: A meta-analysis. *Sex Roles*, *61*, 178–191. doi: 10.1007/s11199-009-9627-3

Spears, R. & Otten, S. (2012). Discrimination: Revisiting Tajfel's minimal group studies. In J. R. Smith & S. A. Haslam (Eds.), *Social psychology: Revisiting the classic studies* (pp. 160–177). London, UK: SAGE Publications.

Stan Development Team. (2017). *Stan modeling language: User's guide and reference manual.* Retrieved from `http://mc-stan.org/` (Version 2.17.3)

Stan Development Team. (2018). *RStan: The R interface to Stan.* Retrieved from `http://mc-stan.org/` (R package version 2.17.3)

Stephan, W. G., Boniecki, K. A., Ybarra, O., Bettencourt, A., Ervin, K. S., Jackson, L. A., … Renfro, C. L. (2002). The role of threats in the racial attitudes of Blacks and Whites. *Personality and Social Psychology Bulletin*, *28*, 1242–1254. doi: 10.1177/01461672022812009

Stephan, W. G. & Stephan, C. W. (2000). An integrated threat theory of prejudice. In S. Oskamp (Ed.), *Reducing prejudice and discrimination* (pp. 23–45). Mahwah, NJ: Lawrence Erlbaum Associates.

Swart, H., Hewstone, M., Christ, O. & Voci, A. (2011). Affective mediators of intergroup contact: A three-wave longitudinal study in South Africa. *Journal of Personality and Social Psychology*, *101*, 1221–1238. doi: 10.1037/a0024450

Tajfel, H. (1969). Cognitive aspects of prejudice. *Journal of Biosocial Science*, *1*, 173–191. doi: 10.1017/S0021932000023336

Tajfel, H. (1981). *Human groups and social categories: Studies in social psychology*. Cambridge, UK: Cambridge University Press.

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

Talaska, C. A., Fiske, S. T. & Chaiken, S. (2008). Legitimating racial discrimination: Emotions, not beliefs, best predict discrimination in a meta-analysis. *Social Justice Research*, *21*, 263–296. doi: 10.1007/s11211-008-0071-2

Talbot, I. & Singh, G. (2009). *The Partition of India*. Cambridge, UK: Cambrdige University Press.

Tausch, N., Hewstone, M. & Roy, R. (2009). The relationships between contact, status and prejudice: An integrated threat theory analysis of Hindu–Muslim relations in India. *Journal of Community & Applied Social Psychology*, *19*, 83–94. doi: 10.1002/casp.984

Tausch, N., Saguy, T. & Bryson, J. (2015). How does intergroup contact affect social change? [its] impact on collective action and individual mobility intentions among members of a disadvantaged group. *Journal of Social Issues*, *71*, 536–553. doi: 10.1111/josi.12127

Tee, N. & Hegarty, P. (2006). Predicting opposition to the civil rights of trans persons in the United Kingdom. *Journal of Community & Applied Social Psychology*, *16*, 70–80. doi: 10.1002/casp.851

Tropp, L. R., Hawi, D. R., Van Laar, C. & Levin, S. (2012). Cross-ethnic friendships, perceived discrimination, and their effects on ethnic activism over time: A longitudinal investigation of three ethnic minority groups. *British Journal of Social Psychology*, *51*, 257–272. doi: 10.1111/j.2044-8309.2011.02050.x

Tropp, L. R. & Pettigrew, T. F. (2005). Differential relationships between intergroup contact and affective and cognitive dimensions of prejudice. *Personality and Social Psychology Bulletin*, *31*, 1145–1158. doi: 10.1177/0146167205274854

Turner, J. C., Hogg, M. A., Oakes, P. J., Reicher, S. D. & Wetherell, M. S. (1987). *Rediscovering the social group: A self-categorization theory*. Oxford, UK: Blackwell.

Turner, R. N., Hewstone, M. & Voci, A. (2007). Reducing explicit and implicit outgroup prejudice via direct and extended contact: The mediating role of self-disclosure and intergroup anxiety. *Journal of Personality and Social Psychology*, *93*, 369–388. doi: 10.1037/0022-3514.93.3.369

Ufkes, E. G., Calcagno, J., Glasford, D. E. & Dovidio, J. F. (2016). Understanding how common ingroup identity undermines collective action among disadvantaged-group members. *Journal of Experimental Social Psychology*, *63*, 26–35. doi: 10.1016/j.jesp.2015.11.006

van Beijsterveldt, C. E. M., Hudziak, J. J. & Boomsma, D. I. (2006). Genetic and environmental influences on cross-gender behavior and relation to behavior problems: A study of Dutch twins at ages 7 and 10 years. *Archives of Sexual Behavior*, *35*, 647–658.

doi: 10.1007/s10508-006-9072-0

van Dommelen, A. (2014). *The role of social identity inclusiveness and structure in intergroup relations: Individual and contextual differences in ethnic and religious minority group members* (Doctoral dissertation, University of Sydney). Retrieved from `http://hdl.handle.net/2123/12007`

van Dommelen, A., Schmid, K., Hewstone, M., Gonsalkorale, K. & Brewer, M. (2015). Construing multiple ingroups: Assessing social identity inclusiveness and structure in ethnic and religious minority group members. *European Journal of Social Psychology*, *45*, 386–399. doi: 10.1002/ejsp.2095

van Zomeren, M. (2013). Four core social-psychological motivations to undertake collective action. *Social and Personality Psychology Compass*, *7*, 378–388. doi: 10.1111/spc3.12031

van Zomeren, M., Postmes, T. & Spears, R. (2008). Toward an integrative social identity model of collective action: A quantitative research synthesis of three socio-psychological perspectives. *Psychological Bulletin*, *134*, 504–535. doi: 10.1037/00332909.134.4.504

van Zomeren, M., Spears, R., Fischer, A. H. & Leach, C. W. (2004). Put your money where your mouth is! Explaining collective action tendencies through group-based anger and group efficacy. *Journal of personality and social psychology*, 649–664. doi: 10.1037/0022-3514.87.5.649

Vehtari, A., Gelman, A. & Gabry, J. (2017). Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Statistics and Computing*, *27*, 1413–1432. doi: 10.1007/s11222-016-9696-4

Wagenmakers, E.-J., Beek, T., Dijkhoff, L., Gronau, Q. F., Acosta, A., R. B. Adams, J., … Zwaan, R. A. (2016). Registered replication report: Strack, martin, & stepper (1988). *Perspectives on Psychological Science*, *11*, 917-928. doi: 10.1177/1745691616674458

Wasserstein, R. L. & Lazar, N. A. (2016). The ASA's statement on p-values: Context, process, and purpose. *The American Statistician*, *70*, 129–133. doi: 10.1080/00031305.2016.1154108

Weiner, B., Perry, R. P. & Magnusson, J. (1988). An attributional analysis of reactions to stigmas. *Journal of Personality and Social Psychology*, *55*, 738. doi: 10.1037/0022-3514.55.5.738

Wickham, H. (2009). *ggplot2: Elegant graphics for data analysis*. New York, NY: Springer.

Wickham, H. (2017). *tidyverse: Easily install and load the 'tidyverse'*. Retrieved from `https://CRAN.R-project.org/package=tidyverse` (R package version 1.2.1)

Wickham, H., Francois, R., Henry, L. & Müller, K. (2017). *dplyr: A grammar of data manipulation*. Retrieved from `https://CRAN.R-project.org/package=dplyr` (R package version 0.7.4)

Wickham, H. & Henry, L. (2018). *tidyr: Easily tidy data with 'spread()' and 'gather()' functions*. Retrieved from `https://CRAN.R-project.org/package=tidyr` (R package version 0.8.0)

Wohl, M. J. A., Branscombe, N. R. & Klar, Y. (2006). Collective guilt: Emotional reactions when one's group has done wrong or been wronged. *European Review of Social Psychology*, *17*, 1—37. doi: 10.1080/10463280600574815

Wright, S. C. (2010). Collective action and social change. In J. F. Dovidio, M. Hewstone, P. Glick & V. M. Esses (Eds.), *The SAGE handbook of prejudice, stereotyping and discrimination* (pp. 345–360). London, UK: SAGE Publications.

Wright, S. C., Aron, A., McLaughlin-Volpe, T. & Ropp, S. A. (1997). The extended contact effect: Knowledge of cross-group friendships and prejudice. *Journal of Personality and Social Psychology*, *73*, 73–90. doi: 10.1037/0022-3514.73.1.73

Wright, S. C. & Baray, G. (2012). Models of social change in social psychology: Collective action or prejudice reduction? conflict or harmony? In J. Dixon & M. Levine (Eds.), *Beyond prejudice: Extending the social psychology of conflict, inequality and social change* (pp. 225–247). Cambridge, UK: Cambridge University Press.

Yao, Y., Vehtari, A., Simpson, D. & Gelman, A. (2018). Using stacking to average Bayesian predictive distributions. *Bayesian Analysis*. (Advance online publication) doi: 10.1214/17-BA1091

Yarkoni, T. & Westfall, J. (2017). Choosing prediction over explanation in psychology: Lessons from machine learning. *Perspectives on Psychological Science*, *12*, 1100–1122. doi: 10.1177/1745691617693393

YouGov. (2015). *1 in 2 young people say they are not 100% heterosexual* [Data set]. Retrieved from `https://yougov.co.uk/news/2015/08/16/half-young-not-heterosexual/`

YouGov. (2016). *Only 2% of young men feel completely masculine (compared to 56% of over 65s)* [Data set]. Retrieved from `https://yougov.co.uk/news/2016/05/13/low-young-masculinity-britain/`

Yzerbyt, V., Corneille, O. & Estrada, C. (2001). The interplay of subjective essentialism and entitativity in the formation of stereotypes. *Personality and Social Psychology Review*, *5*, 141–155. doi: 10.1207/S15327957PSPR0502_5