

NORTHWESTERN UNIVERSITY

Intersectional Categorization Theory:  
A Compartmentalization Model of Social Stereotyping

SUBMITTED TO THE GRADUATE SCHOOL  
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Psychology

By

Christopher D. Petsko

EVANSTON, ILLINOIS

June 2020

© Copyright by Christopher D. Petsko 2020  
All Rights Reserved

## Abstract

How do perceivers stereotype people who are at the margins of multiple, interlocking social identities? The present dissertation proposes and tests a model of intersectional stereotyping called intersectional categorization theory (ICT). This model advocates the following core ideas: (a) that perceivers use one lens at a time for making sense of other people; (b) that the lenses perceivers use can be singular and simplistic (e.g., viewing an older East-Asian woman as an *old person*), or intersectional and complex (e.g., viewing the same individual as an *older East-Asian woman* specifically); and (c) that different lenses can prescribe categorically distinct sets of stereotypes that perceivers use as frameworks for thinking about targets. The first chapter of this dissertation reviews the phenomenon that ICT seeks to explain (intersectional stereotyping), and it describes ICT's utility for providing order to the psychological literature on this topic. Subsequent chapters describe empirical tests of ICT's core ideas. Experiments 1a, 1b, and 2 provide evidence that when perceivers are viewing targets through the lens of gender, they focus their attention on gender so strongly that they barely attend—at least in these moments—to targets' age groups (Experiments 1a and 2) or racial groups (Experiment 1b). Experiment 3 reveals that as perceivers switch from one lens (e.g., age) to another lens (e.g., gender) for viewing social targets, so too do they come to associate different stereotypic attributes with targets. Experiment 4 provides evidence that perceivers occasionally attend to *intersections* of identities themselves (e.g., race *and* gender) in lieu of singular identities (e.g., gender alone). Finally, Experiments 5a and 5b suggest that the stereotypes perceivers associate with targets can vary dramatically as a function of whether perceivers are using intersectional vs. singular lenses

for thinking about targets. Collectively, these experiments provide strong support for ICT's utility as a framework for examining intersectional stereotyping.

*Keywords:* intersectional stereotyping, person perception, crossed categorization, impression formation, intergroup relations

## Table of Contents

Chapter I: Introduction.....	p. 9
Intersectional Stereotyping and Prejudice.....	p. 11
Origins of Intersectional Stereotyping Research.....	p. 13
Approaches to Studying Intersectional Stereotyping.....	p. 14
The Success of Integration Models in Intersectional Stereotyping Research...	p. 20
Advances in Intersectional Stereotyping Research.....	p. 21
The Limits of Integration Models.....	p. 27
Conclusion.....	p. 29
Intersectional Categorization Theory (ICT).....	p. 30
A Brief History of the Social Identity Tradition.....	p. 32
Overview of Intersectional Categorization Theory.....	p. 34
Assumptions of Lens Selection.....	p. 36
Assumptions of Lens-Inflected Stereotyping.....	p. 41
Assumptions of Lens-Based Discrimination.....	p. 44
The Benefits of ICT as a Model of Intersectional Stereotyping.....	p. 47
ICT as a Tool for Explaining Psychological Findings.....	p. 49
Conclusion.....	p. 50
The Present Experiments.....	p. 52
Chapter II: Initial Evidence of Compartmentalized Lens Usage.....	p. 54
Experiments 1a and 1b.....	p. 55
Experiment 1a.....	p. 55
Method.....	p. 56

Results.....	p. 61
Experiment 1b.....	p. 67
Method.....	p. 68
Results.....	p. 72
Discussion.....	p. 78
Experiment 2.....	p. 80
Method.....	p. 80
Results.....	p. 84
Discussion.....	p. 92
Chapter III: Initial Evidence of Lens-Dependent Stereotyping.....	p. 93
Experiment 3.....	p. 94
Method.....	p. 94
Results.....	p. 96
Discussion.....	p. 99
Chapter IV: Evidence of Intersectional Lens Usage and Stereotyping.....	p. 100
Experiment 4.....	p. 101
Method.....	p. 101
Results.....	p. 103
Discussion.....	p. 106
Experiments 5a and 5b.....	p. 108
Experiment 5a.....	p. 108
Method.....	p. 109

Results.....	p. 111
Experiment 5b.....	p. 113
Method.....	p. 114
Results.....	p. 116
Discussion.....	p. 118
 Chapter V: General Discussion.....	 p. 120
Summary of Findings and Contributions.....	p. 122
Limitations.....	p. 124
Future Directions.....	p. 126
Concluding Remarks.....	p. 128
 References.....	 p. 130
 Appendices.....	 p. 154
Appendix A.....	p. 154
Appendix B.....	p. 155
Appendix C.....	p. 156
Appendix D.....	p. 158
Appendix E.....	p. 159

## Figure and Table Locations

Figures

<i>Figure 1.</i> Pictographic Illustration of the Psychological Processes ICT Explains...	p. 35
<i>Figure 2.</i> Example of Stimulus Matrix (Exp. 1a).....	p. 58
<i>Figure 3.</i> Social Categorization by Experimental Condition (Exp. 1a).....	p. 62
<i>Figure 4.</i> Checklist Attribute Ratings by Experimental Condition (Exp. 1a).....	p. 65
<i>Figure 5.</i> Ratings of Old Women’s Faces by Experimental Condition (Exp. 1a)....	p. 67
<i>Figure 6.</i> Example of Stimulus Matrix (Exp. 1b).....	p. 70
<i>Figure 7.</i> Social Categorization by Experimental Condition (Exp. 1b).....	p. 73
<i>Figure 8.</i> Checklist Attribute Ratings by Experimental Condition (Exp. 1b).....	p. 76
<i>Figure 9.</i> Ratings of Black Women’s Faces By Experimental Condition (Exp. 1b).	p. 78
<i>Figure 10.</i> Social Categorization by Experimental Condition (Exp. 2).....	p. 85
<i>Figure 11.</i> Checklist Attribute Ratings by Experimental Condition (Exp. 2).....	p. 87
<i>Figure 12.</i> Ratings of Old Women’s Faces by Experimental Condition (Exp. 2)....	p. 89
<i>Figure 13.</i> Brief-IAT Association Latencies by Experimental Condition (Exp. 2)..	p. 91
<i>Figure 14.</i> IAT Association Latencies by Experimental Condition (Exp. 3).....	p. 98
<i>Figure 15.</i> Anger Detection Speed (Exp. 4).....	p. 105
<i>Figure 16.</i> Probability of Detecting Anger (Exp. 4).....	p. 106
<i>Figure 17.</i> IAT Association Latencies by Experimental Condition (Exp. 5a).....	p. 112
<i>Figure 18.</i> IAT Association Latencies by Experimental Condition (Exp. 5b).....	p. 117

Tables

<i>Table 1.</i> Most-Nominated Stereotypes for Old Women (Exp. 1a).....	p. 64
<i>Table 2.</i> Most-Nominated Stereotypes for Black Women (Exp. 1b).....	p. 75
<i>Table 3.</i> Most-Nominated Stereotypes for Old Women (Exp. 2).....	p. 86

## Chapter I: Introduction

## Intersectional Categorization Theory:

### A Compartmentalization Model of Social Stereotyping

An older Korean woman stops to admire Cloud Gate, a tourist attraction in Chicago, as she makes her way through Millennium Park. At the attraction, this woman stands out from the crowd. The tourists surrounding her are exclusively White. One of the tourists, a middle-aged White man, walks up to the older Korean woman and asks her to take a photo of him and his family. She denies his request. (It's a bit too cold outside to be handling cell phones, she thinks. And she'd have to take off her gloves.) In response to her denial, the man turns his back to the woman. He scans the crowd for someone else he can ask. "Of course," he mutters to himself, "I found the *one* person here who doesn't speak English." ... Later that day, the woman heads back to her apartment in Lincoln Park. She is carrying with her the shopping bags she acquired from earlier. As she makes her way down a populated street, a group of teenagers approaches her. One of the teenagers hops off his skateboard and extends his phone in her direction. The boy's request is similar to the one from earlier: will she take a photo of him and his friends? Again, the older woman replies with a definitive "no." (She stands by her decision from earlier. And she is beginning to wonder why people find her so approachable.) In response to her denial, the boy scoffs. "Old people," he says to his friends, "they can't use technology to save their lives."

The present dissertation describes a compartmentalization model of intersectional stereotyping. This model, called intersectional categorization theory (ICT), posits that in some contexts, perceivers will think of an older Korean woman as an *old person*. In these moments, they will view her through the lens of her age, which will cause her to seem quite "old" to perceivers, but perhaps not very "Korean" or stereotypic of women. Yet in other contexts,

perceivers will think of her as a *Korean woman*, specifically. In these moments, they will view her through the lens of her intersecting gender and ethnic identities, which will cause her to seem very stereotypic of Korean women, but perhaps not very “old.” A central assumption of ICT is that these lenses *trade off* in the minds of perceivers. As different lenses come into focus, different acts of discrimination become more or less likely. The lens a perceiver uses should predict when it is that they exhibit ageism, but not racism or sexism toward an older Korean woman, and when it is that they exhibit gendered racism—but in this moment, no ageism—toward her instead.<sup>1</sup> Chapter I begins with a review of the phenomenon that ICT seeks to explain (i.e., intersectional stereotyping) as well as the dominant theoretical assumptions that are used to study this phenomenon. The second half of Chapter I describes ICT itself: its tenets, its utility for providing order to the psychological literature on intersectional stereotyping, its predictions, and a discussion of how it compares with other models. It is against this backdrop that the aims of subsequent chapters—which provide empirical tests of ICT’s assumptions—are introduced.

### **Intersectional Stereotyping and Prejudice**

Before delving into the origins and tenets of ICT, it is important to discuss what it is that ICT seeks to explain: how perceivers stereotype, evaluate, and behave toward targets as a function of these targets’ interlocking social identities. In ICT, social identities can be any group membership that is meaningful to perceivers. In this view, social identities can be religious groups (*atheists, Muslims*), sexual orientation groups (*lesbians, bisexuals*), political groups (*Republicans, Democrats*), nationality groups (*Moroccans, Japanese*), and so on. In addition,

---

<sup>1</sup> This example presupposes that the perceiver is prejudiced against older Korean women. Of course, not all perceivers are. Perceivers who themselves are women, older, Korean, or any combination of the three, for example, may exhibit various kinds of in-group favoritism toward the target. This issue will be considered further in the discussion of ICT itself.

ICT allows for the possibility that social identities can be intersections of groups themselves. Thus, social identities can also include *Black women*, *gay school teachers*, and *older Mexican immigrants*, just to name a few. Of note, ICT does not presume that these intersectional social identities are merely subgroups—from the perspective of perceivers—of their “parent” categories. Rather, ICT presumes that intersectional social identities can imply their own sets of stereotypes, which may or may not be overlapping with those of their “parent” categories. ICT begins from the observation that each person belongs to an infinite set of perceivable social identities, and can therefore be viewed, in principle, through an infinite set of social lenses. The aim of ICT is to explain what causes perceivers to view people through one of these lenses (e.g., the lens of *Blackness*) when they could just easily view them through alternatives (e.g., the lenses of *Black womanness*, or *womanness*). In addition, ICT aims to elucidate the consequences of differential lens selection: how these lenses shape not just perceptions of other people, but the behaviors perceivers enact toward other people as well.

The following sections begin by considering the origins of intersectional stereotyping research. The first section covers where the term ‘intersectional’ comes from, as well as what fueled its popularity in social psychology. The second section diverges from the intersectional stereotyping literature to discuss a conceptually overlapping literature on crossed categorization. This second section describes three sets of theoretical assumptions that have been used to describe how perceivers engage in crossed categorization, and it describes what these assumptions imply about the operation of intersectional stereotyping. The third section describes which of these three sets of assumptions became most popular in intersectional stereotyping research—paying attention to the merits of this set of assumptions, but as well, to its limitations.

The final section of this review makes the case for using a different approach to studying intersectional stereotyping: one that characterizes it as a compartmentalized phenomenon.

### **Origins of Intersectional Stereotyping Research**

The psychological literature on intersectional stereotyping borrows its terminology from Black feminist scholars (Crenshaw, 1991; King, 1988; hooks, 1984). Many of these scholars were critical race theorists who wanted to break free from contemporary narratives on racial oppression—which emphasized Black men’s oppression—and from contemporary narratives on gender oppression—which emphasized White women’s oppression. Kimberlé Crenshaw, a legal scholar who is often credited with coining the term ‘intersectionality,’ wrote that “the experiences of women of color are frequently the product of *intersecting* [emphasis added] patterns of racism and sexism” (Crenshaw, 1991, p. 1243). Crenshaw observed that the sexism that Black women experience, for example, is often tinged with a version of racism that makes it distinct from the sexism that White women experience. Crenshaw also observed that all of us experience and perceive social identities *in light of* other social identities. In other words, “each person has a gender, race, sexual orientation, and so on... [and] the meaning of each social group is constructed through the lens of the others” (Ghavami, Katsiaficas, & Rogers, 2016, pp. 34-35). It is important to underscore that intersectionality was initially formulated to characterize the experience of contending with more than one marginalized identity. The concept was not formulated to characterize *the perception* of those contending with more than one marginalized identity. Nevertheless, the principles of intersectionality bore relevance to the study of person perception—and it was in this direction that many psychologists took it.

The implication that stereotyping and prejudice may manifest differently for people at different intersections of social identities caught wind quickly in psychological science (Cole, 2009; Plaut, 2010). In 2009, Elizabeth Cole wrote an influential *American Psychologist* piece that encouraged psychologists to move beyond one-size-fits-all understandings of identity and prejudice. Around this time, social psychologists began publishing many findings that demonstrated the dynamism of social stereotyping, often with ‘intersectionality’ as a keyword. For example, Goff, Thomas, and Jackson (2008) showed that perceivers’ conception of women’s femininity depends on whether these women are Black or White. Remedios and colleagues (2011) found that negative evaluations of Black men relative to White men completely reverse when the men in question are homosexual rather than heterosexual. Social psychologists asked many questions that carried a similar flavor: questions that broadly asked whether the conception of one social identity is influenced in some way by the presence of another (see Kang & Bodenhausen, 2015). The issue, however, is that psychologists did so from a largely singular theoretical approach—an approach that, as explained below, emphasized *integration models* over alternatives.

### **Approaches to Studying Intersectional Stereotyping**

Prior to the expansion of research on intersectional stereotyping, social-identity and social-cognition researchers spent ample time addressing questions related to crossed categorization: that is, how we make sense of people in light of their multiple social category memberships (see Nicolas, de la Fuente, & Fiske, 2017). This literature diverged from the intersectional stereotyping literature in emphasis, but not in substance. Both literatures were concerned with questions about how we conceptualize, feel about, and behave toward people

whose multiple identities are salient to us. However, whereas the intersectional stereotyping literature tended to focus on the cross sections of oppressed demographic groups (e.g., Asian-American women who are homosexual; Ghavami & Peplau, 2018), the literature on crossed categorization was concerned with social identity combinations more broadly (e.g., Harvard-educated carpenters; Kunda, Miller, & Claire, 1990). Psychologists developed many theories and models to make sense of the sprawling literature on crossed categorization. But generally speaking, these theoretical approaches shared one of three sets of assumptions: assumptions of social identity *dominance*, of social identity *compartmentalization*, or of social identity *integration* (Bodenhausen, 2010; Petsko & Bodenhausen, 2020). This taxonomy can of course also be used to characterize viable approaches to studying intersectional stereotyping. Below, I briefly review each approach.

**Dominance models.** Dominance models assume that certain social identities take precedence over others in the minds of perceivers (Urban & Miller, 1998). In the domain of intersectional stereotyping, a dominance model would imply that certain forms of prejudice manifest more often than others, or that intersectional stereotypes are shaped more by one social identity than by others. Many perspectives in social psychology resonate with the assumptions of dominance models. For example, the dual process model of impression formation (Brewer, 1988), along with the continuum model of impression formation (Fiske & Neuberg, 1990), argue that categorization on the basis of age, sex, and race is more or less inevitable—implying by extension that categorization on the basis of other features, like sexual orientation, is not (cf. Rule, Macrae, & Ambady, 2009). Evolutionary psychology also advocates for the perceptual dominance of some social identities over others (Sidanius & Pratto, 2012; Tooby & Cosmides,

2010). In this perspective, the tendency to categorize others by age and sex is regarded as an obligatory psychological tendency that cannot be attenuated. The tendency to categorize others by race, in this perspective, occurs not because race categorization is inevitable, but because perceivers use race as a proxy for inferring who is in coalition with whom. Empirical work in this vein tends to show that when patterns of coalition cross-cut with race, perceivers stop engaging in race categorization. But when patterns of coalition cross-cut with sex, or with age, perceivers nevertheless continue to engage in sex and age categorization (Kurzban, Tooby, & Cosmides, 2001; Pietraszewski, Curry, Petersen, Cosmides, & Tooby, 2015).

A benefit of dominance models is that they often make precise claims about what should not happen (e.g., perceivers should *not* be able to refrain from engaging in sex categorization). This is good for science, as it provides room for falsification. [In Popper's (1959) view, theories are only useful to the extent that they make falsifiable claims.] But an issue with dominance models is that they often contradict one another. For example, just as the evolutionary models above argue for the dominance of sex and age categorization over race categorization, still other models argue for the dominance of *race* categorization over sex and age categorization. The ethnic prominence hypothesis, for example, advocates for this latter possibility (Levin, Sinclair, Veniegas, & Taylor, 2002).

**Compartmentalization models.** Compartmentalization models assume *situational* social identity dominance in the minds of perceivers (e.g., Bodenhausen & Macrae, 1998). Applied to the domain of intersectional stereotyping, compartmentalization models imply that certain forms of prejudice can be “switched on” by a social context, and that others can be “switched off.” In self-categorization theory, for example, the perceiver is presumed to search for a social category

that “fits” social reality (Oakes, 1987; Turner et al., 1987). Once a category is chosen, the perceiver then depersonalizes targets in the direction of the category prototype. For example, if race is the category that “fits” social reality, the perceiver will stereotype Asian women as seeming prototypic of Asian people. But if *gender* is the category that “fits” social reality, the perceiver will stereotype Asian woman as seeming prototypic of women (see also Craig & Bodenhausen, 2018). A different model—the stereotype activation-inhibition model—argues as well for the idea that perceivers select lenses for viewing targets in a compartmentalized way (Bodenhausen & Macrae, 1998). But this latter theory adds the assumption that when one lens comes into focus, other lenses are actively inhibited by the mind. From this perspective, thinking of Asian women as “Asian” not only makes traits related to Asianness accessible, but makes traits related to alternative categories (e.g., the category *women*) inaccessible (Macrae, Bodenhausen, & Milne, 1995). Finally, compartmentalist views interweave with research on cultural frame switching (Hong, Chiu, & Kung, 1997). Generally, this tradition argues that cultures themselves can serve as social identities, which, when activated, bring forth wholly distinct sets of norms, behaviors, and attitudes that the perceiver can use to construe the self or others (e.g., Ramírez-Esparza, Gosling, Benet-Martínez, Potter, & Pennebaker, 2006). Of note, compartmentalization models contrast with the evolutionary argument described above (e.g., Tooby & Cosmides, 2010). The stereotype activation-inhibition model, for example, allows for the possibility that age and sex categorization, in addition to race categorization, can be attenuated as other categories come into the perceiver’s focus.

A clear benefit of compartmentalization models is that they can be used to explain discrepant patterns of stereotyping. If in certain moments gender can become the dominant lens

through which perceivers view targets, and if in other moments race can become the dominant lens through which perceivers view targets, then it is no wonder why researchers would argue for dominance models that at times emphasize the perceptual primacy of gender (e.g., Kurzban et al., 2001), and that at other times emphasize the perceptual primacy of race (e.g., Levin et al., 2002). A drawback of compartmentalization models, however, is that they may be more difficult to falsify than dominance models. If researchers presume that all patterns of stereotyping are explainable by situational category salience, then researchers can rationalize any pattern of results they find. This can lead to circular reasoning. If perceivers exhibit racism toward a target, for example, researchers can convince themselves that race *must* have been the lens perceivers were using; but if perceivers instead exhibit sexism toward a target, researchers can convince themselves that sex *must* have been the lens perceivers were using.

**Integration models.** Integration models assume that multiple social identities are perceived *simultaneously*, and that the perceiver somehow integrates them into a coherent mental impression of the target. Thus, these models do not presuppose, as dominance or compartmentalization models do, that the perceiver can respond to Black women solely in terms of these women's race. Instead, they assume that the meaning of race is constructed in light of these women's gender groups, age groups, sexual orientation groups, religious identities, and so on. Integration models vary widely in what it is that they seek to explain, as well as how they assume perceivers integrate information related to multiple social identities. For example, the common in-group identity model (Gaertner & Dovidio, 2000) and the crossed categorization model (Crisp & Hewstone, 1999; 2007) both seek to explain prejudice, and they do so with the assumption that perceivers algebraically combine their feelings of positivity and negativity

toward the multiple identities they perceive. In these models, a perceiver's negative attitude toward gay men, for example, can be made more positive by learning that these men share an in-group identity on some other dimension—say, as affiliates of the same university (Crisp & Beck, 2005). Integration models that explain *stereotype content*, in contrast, tend to have a more complicated view on how identities are integrated. Generally, these latter models emphasize the idea that social identities are perceived as a gestalt (Asch, 1946; Asch & Zukier, 1984), and that impressions of individuals cannot be understood in algebraic terms (but see Anderson, 1971). For example, the parallel-constraint-satisfaction model (Kunda & Thagard, 1996), or the dynamic-interactive theory of person construal (Freeman & Ambady, 2011; Freeman, Stoller, & Brooks, 2020), argue that the mind operates like a connectionist network. In these perspectives, perceivers' stereotypes toward targets are interactively shaped by features of the target, features of the social categories to which the target belongs, as well as by features of the individual perceivers. Applied to intersectional stereotyping, these models suggest that stereotypes toward Black women can *only* be understood as resulting from the way perceivers construe the intersection of Blackness and woman-ness itself.

A benefit of integration models is that they can account for the broadest range of psychological phenomena. They can describe why the stereotype content that emerges for Black women, for example, is distinct from that of White women, Asian women, Hispanic women, and so on (Ghavami & Peplau, 2013; Niemann, Jennings, Rozelle, Baxter, & Sullivan, 1994). But a drawback of these models is that they are even more difficult to falsify than compartmentalization models. Integration models in general, but connectionist integration models in particular, allow for the possibility that infinitely many social identities and concepts

can inform a perceiver's impression of a target. As such, infinitely many impressions of a target are possible outputs of the mind. While this view may closely approximate reality, it has little utility for helping researchers to predict how perceivers will stereotype particular intersections of people (Muslim women, for example). Furthermore, integration models are in some ways antagonistic with the very notion of stereotyping itself. If a target is being perceived in light of all their attributes and group memberships, then they are much closer to being individuated than they are to being stereotyped. Finally, theorizing in integration models tends to center on the features of individual perceivers and targets—making room for, perhaps, but by no means emphasizing the fundamental role that social context plays in shaping impression formation.

### **The Success of Integration Models in Intersectional Stereotyping Research**

Why did integration models become popular during the expansion of intersectional stereotyping research? The answer to this question is that integration models are, at least on their surface, most compatible with the take-home message of intersectional scholarship: that “the meaning of each social group is constructed through the lens of the others” (Ghavami et al., 2016, p. 35). Dominance models are conceptually incompatible with this message because they imply that the meaning of a single group (or handful of groups) dominates how a target is perceived—that Black women inevitably face racism, for example, but not sexism (Levin et al., 2002). Compartmentalization models *seem* incompatible with this message because they intuitively imply that Black women, for example, can only be perceived as *black* or as *women*, but not as *Black women* simultaneously. This ostensible implication is of course false, which I will argue in detail later. For now, suffice it to say that integration models, with their presumption that the social identities are perceived in light of each other, provided a veneer of

compatibility with the take-home message described above. A consequence of this is that researchers used these models' assumptions—whether intentionally or not—to test the promising ideas laid out by Cole (2009): that stereotyping and prejudice against one identity is contingent upon the perception of other identities to which targets belong.

### **Advances in Intersectional Stereotyping Research**

Reasoning on the basis of integrationist assumptions has been exceedingly generative for social psychology. This is true regardless of whether one is embedded in the literature on crossed categorization or the emerging literature on intersectional stereotyping. Indeed, integrationist assumptions helped us to discover that perceivers' impressions of targets' identities can be remarkably dynamic, even when the identities in question are those that we tend to think of in essentialist terms. For example, the speed with which perceivers categorize targets by sex (Johnson, Freeman, & Pauker, 2012), as well as the profile of gender stereotypes perceivers apply to targets (Goff et al., 2008; Hall, Galinsky, & Phillips, 2015), bend flexibly depending on whether the targets are White, Black, or East Asian. Perceptions of sexual orientation are likewise dynamically determined: How “gay” a target seems is affected by targets' race (Asian men seem “gayer” than White or Black men, for example; Johnson & Ghavami, 2011), and so too are the stereotypes that perceivers use to describe these targets (for example, “down-low” is a stereotype that perceivers use to describe *gay Black men*, but not the ‘parent’ categories *gay men* or *Black men*; Calabrese et al., 2018).

Researchers' reliance on integrationist assumptions has resulted in a registry of psychological findings that is both extensive and diverse. Rather than reviewing all findings that fit beneath the umbrella of intersectional stereotyping, I will briefly review those that are related

to racial stereotyping, specifically. This area of research is highly developed and is representative of the kinds of questions psychologists have been asking about intersectional stereotypes. As such, it can be used to illustrate the gains as well as complications that emerge from relying on integrationist assumptions.

**Race-by-age stereotyping.** White perceivers exhibit a tendency to “see” anger on the faces of Black targets more readily than on the faces of White targets (Becker, Neel, & Anderson, 2010; Hugenberg & Bodenhausen, 2003). This bias is typically examined by showing perceivers the faces of White and Black men, whose facial expressions change along a continuum from angry to happy (or vice versa). The perceiver’s task is let the researcher know when the faces stop looking angry, for example, and begin looking happy. Reliably, perceivers report that anger lingers for longer on the faces of Black targets than on the faces of White targets. In an extension of this work, Kang and colleagues (Kang & Chasteen, 2009; Kang, Chasteen, Cadieux, Cary, & Syeda, 2014) manipulated not just the race of social targets, but also whether these targets were young adults *vs.* older adults. Their findings indicate that when the faces are of young adults, perceivers exhibit the conventional race bias. But, when the faces are of older adults, the conventional race bias reverses. In this context, perceivers “see” anger for longer on the faces of White targets than on the faces of Black targets. Thus, these findings indicate that racial bias can be attenuated or exacerbated depending upon targets’ age.

In still other situations, however, racial bias remains immune to any moderation by targets’ age groups. For example, another established racial bias is that the speed with which perceivers distinguish criminal objects (e.g., guns) from innocuous objects (e.g., tools) increases when they are primed with Black relative to White faces (Eberhardt, Goff, Purdie, & Davies,

2004; Payne, 2001). Astonishingly, this bias occurs in equal measure when the target faces are of *children* as when they are of adults (Todd, Thiem, & Neel, 2016; Todd, Simpson, Thiem, & Neel, 2016). This bias also occurs in equal measure when the target faces are of older vs. young adults (Lundberg, Neel, Lassetter, & Todd, 2018). Thus, there is not consistent evidence that the conception of race (and its associated content) is dependent upon the perception of targets' age groups. Instead, there appear to be some contexts in which perceivers *do* integrate targets' race and age groups into their impressions of targets, and other contexts in which they do not.

**Race-by-gender stereotyping.** What about gender? Do perceivers integrate targets' gender into their conceptions of targets' race? It seems so, yes. But there is considerable disagreement over the outcomes of that integration. In particular, there is disagreement over the issue of whether Black women face less racial bias than Black men, or whether they face more. Those who argue that Black women face less racial bias than Black men tend to make this argument on the basis of cultural prototypes (e.g., Purdie-Vaughns & Eibach, 2008). Generally, this argument is that Black men better fit perceivers' prototype of "blackness" than do Black women (Schug, Alt, & Klauer, 2015; Sesko & Biernat, 2010; 2018), and that perceivers consequently direct their racial biases more toward men than toward women (McDonald, Sidanius, & Navarrete, 2011; Sidanius & Pratto, 2012). Support for this possibility comes from several distinct lines of research. For example, perceivers are less likely to categorize Black faces as "Black" when these faces are feminine compared with when they are masculine (Carpinella, Chen, Hamilton, & Johnson, 2015). This suggests that masculine features may be tethered to perceivers' conception of Blackness itself (Galinsky, Hall, & Cuddy, 2013). In addition, the stereotype content that perceivers generate for Black men is more negative than the stereotype

content they generate for Black women (Hosoda, Stone, & Stone-Romero, 2003). This suggests, perhaps, that racial negativity is directed more strongly toward men than toward women (see also, Perszyk, Lei, Bodenhausen, Richeson, & Waxman, 2019). Finally, when Black men exhibit dominant behaviors, perceivers report disliking them to a greater degree than when Black women exhibit dominant behaviors (Livingston, Rosette, & Washington, 2012). These findings all point to the possibility that perceivers' racial biases are indeed attenuated when the targets are women rather than men.

On the other side of the coin, there is evidence that perceivers' racial biases may be exacerbated when the targets are women rather than men (Rosette, Ponce de Leon, Koval, & Harrison, 2018). For example, perceivers harbor a general tendency to underestimate the leadership abilities of Black relative to White Americans, and this bias is compounded against Black women—but not against Black men—when the companies they lead are unprofitable (Rosette & Livingston, 2012). On top of this, Black women in the private sector have lower market rewards (Petrie & Roman, 2004), poorer promotion rates (Yap & Konrad, 2009), and report more frequent racial harassment than Black men do (Buchanan & Fitzgerald, 2008). Thus, these findings imply the exact opposite of those reviewed above. Rather than facing less racial bias than Black men, these latter findings suggest that in some contexts, Black women face more.

**Race-by-orientation stereotyping.** As noted previously, another common racial bias—perceivers' tendency to evaluate Black men more negatively than White men—reverses when the targets are homosexual rather than heterosexual (Remedios et al., 2011; Pedulla, 2014). In a series of experiments, Galen Bodenhausen and I replicated this phenomenon. In these experiments, perceivers nominated the stereotypes that came to their minds most quickly when

thinking of various groups of gay *vs.* heterosexual men, whose races were either White or Black (Petsko & Bodenhausen, 2019b). New participants—who knew nothing of the origins of these stereotypes—then rated these stereotypes on a variety of dimensions, including how positive *vs.* negative they seem, and critically, on how typical they seem of White people *vs.* Black people. We found clear and compelling evidence that, as others have reported (e.g., Wilson, Remedios, & Rule, 2017), perceivers harbored more positive stereotypes toward gay Black men than toward heterosexual Black men. But we also found something else, which is that perceivers characterized gay Black men as “less Black,” and even as “Whiter,” than their heterosexual counterparts.

These data raised an interesting question: if homosexual (*vs.* heterosexual) Black men seem “less Black” to perceivers, might perceivers also evince less racial discrimination against these men? We examined this by testing whether another well-known racial bias—the *race-crime congruency effect* (Jones & Kaplan, 2003)—would replicate when the target men were gay *vs.* heterosexual. The race-crime congruency effect describes a tendency for perceivers to condemn Black men more harshly when accused of “Black” crimes (like drive-by shootings) than when accused of “White” crimes (like insider trading). We reasoned that if gay (*vs.* heterosexual) Black men seemed “less Black” to perceivers, so too should they seem less culpable when accused of stereotypically Black crimes. In these data ( $N = 1314$ ), we replicated race-crime congruency effect: perceivers condemned Black men more harshly when accused of “Black” rather than “White” crimes. However, we found zero evidence that this bias was moderated by defendant sexual orientation (Petsko & Bodenhausen, 2019a, Experiment 1). The upshot of this is that there are clearly situations in which perceivers’ impressions of target men

depend both on these men's sexual orientation groups *and* these men's racial groups. But clearly too, there are situations in which impressions depend *only* on target men's race.

**Race-by-SES stereotyping.** What about race-by-socioeconomic status (SES) stereotypes? Do these stereotypes exhibit the same kinds of variability as those reviewed above? Generally, the answer is no. These stereotypes tend to exhibit a relatively stable pattern. Extant research suggests U.S. respondents associate high SES with White people more than with non-White people (e.g., Lei & Bodenhausen, 2017; Brown-Iannuzzi, Dotsch, Cooley, & Payne, 2017). In an impressive demonstration of this phenomenon, Penner and Saperstein (2008) analyzed data from a 17-year longitudinal study called the National Longitudinal Survey of Youth. These data were recorded by “interviewers” who were themselves tasked with categorizing the survey respondents, during each interview, as either White or non-White. Remarkably, over the course of the 17-year study, 1 in 5 of the more than 12,000 survey respondents switched racial categories at least once (from the perspective of the interviewers). Moreover, these categorization changes were not randomly distributed. Instead, findings indicated that when respondents had been downwardly mobile over time—when they became incarcerated, impoverished, or unemployed over the course of the study—interviewers became substantially more likely to categorize them as non-White, even if interviewers had previously categorized respondents as White. Although these findings are correlational, experimental evidence supports the causal direction they imply. When racially ambiguous targets are depicted wearing business suits (*vs.* janitorial clothing), for example, perceivers are more likely to categorize them as White. But when these same targets are depicted wearing janitorial clothing

(*vs.* business suits), the reverse is true. Under these circumstances, perceivers are more likely to categorize them as Black (Freeman, Penner, Saperstein, Scheutz, & Ambady, 2011).

### **The Limits of Integration Models**

If the section above on ‘advances in intersectional stereotyping research’ seems convoluted and contradictory, that is because the advances themselves have been convoluted and contradictory. Integrationist assumptions allow for the possibility that perceivers’ impressions of targets can be multiplex—at times suggesting one thing, and at times suggesting something else. On top of this, integrationist assumptions do not give much consideration to the role of broader social contexts in shaping impression formation. These two features of integration models limit our ability make sense of discrepant research findings in the literature on intersectional stereotyping.

**Limit 1: Contradictory findings.** An affordance of integration models, as noted previously, is that infinitely many stereotypic judgments are plausible outputs of the mind. This affordance has paved the way for a productive exploration of intersectional stereotyping. However, this affordance has also allowed for a compilation of psychological findings that can seem contradictory and convoluted. Take, for example, the findings on whether racial bias is moderated by whether the targets are old *vs.* young. Clearly, there are situations in which perceivers exhibit less racial bias toward older (*vs.* young) adults (Kang et al., 2014), but clearly too there are other situations in which perceivers’ racial biases are unaffected by whether their targets are older or young (Todd et al., 2016a; 2016b). Take as well the findings on whether racial bias is attenuated when the targets are Black women *vs.* Black men. Occasionally, these

findings suggest that Black women face less racial bias than Black men (Livingston et al., 2012), but at other times these findings suggest the reverse (Rosette & Livingston, 2012).

The affordance of contradictory findings does not have to be a limitation of a social psychological perspective. Indeed, situationism, social psychology's guiding assumption, advocates for the idea that social situations ought to evoke flexibility—and even contradiction—in the phenomena we study. But contradictory findings such as these *become* a limitation when explanations for their coexistence lack parsimony. As generative as the assumptions of integration models have been, they do not provide a parsimonious account for why gay (vs. heterosexual) Black men, for example, seem “less Black” to perceivers in some contexts (Petsko & Bodenhausen, 2019b), yet virtually indistinguishable from heterosexual Black men in others (Petsko & Bodenhausen, 2019a). A more useful set of assumptions would help us explain when homosexuality *will* be integrated into perceivers' conception of race, and importantly, when it will not be.

**Limit 2: Failure to consider the broader social context.** As noted previously, integration models center their theorizing, usually, on the features of perceivers and targets (Freeman & Ambady, 2011; Hall, Hall, Galinsky, & Phillips, 2019; Kunda & Thagard, 1996). For example, the integrative framework of social categorization (Kawakami, Amodio, & Hugenberg, 2017) argues that categorization results from what targets look like (e.g., facial cues, body movement) as well as from perceivers' idiosyncratic beliefs and attributes (e.g., prejudices, social status). This illustrates what seems to be a common feature of integration models, which is that they tend not to consider the extra-dyadic factors that can influence social categorization. For example, they tend not to consider that a target's race may feel more relevant to perceivers at

a Black Lives Matter rally than at the post office. Moreover, they tend not to consider that an Asian woman's gender might "pop out" to perceivers more when she is surrounded by Asian men compared with when she is surrounded by White women.

Because integration models do not consider the role of broader social contexts in shaping intersectional stereotypes, attempts to explain discrepant findings have likewise lacked a focus on broader social contexts. For example, attempts to explain discrepant patterns of stereotyping have often involved cataloguing findings by which of several hypotheses they support: the double-jeopardy hypothesis (Barnum, Liden, & DiTomaso, 1995), the subordinate-male target hypothesis (McDonald et al., 2011), the intersectional invisibility hypothesis (Purdie-Vaughns & Eibach, 2008), the ethnic prominence hypothesis (Levin et al., 2002), among others. The issue with cataloguing findings in this way is that doing so implies there are "winners" and "losers." It implies for example that if the majority of findings reveal disadvantages for Black women relative to Black men and White women, that the double-jeopardy hypothesis must be the "winner." The truth, of course, is that there probably are not whole-sale prescriptions psychologists can make about which of these hypotheses "wins" or "loses." The truth is instead likely to be there are some contexts in which one hypothesis will "win," and still other contexts in which that same hypothesis will "lose."

## **Conclusion**

Intersectional stereotyping has been a rapidly developing research area in social psychology. No doubt, the speed with which this area developed has been fueled by the enthusiasm that has accompanied integration models (e.g., Freeman & Ambady, 2011; Hall et al., 2019; Kawakami et al., 2017; Kunda & Thagard, 1996). These models argue that our

perceptions of each other can be highly interactive: that the way we conceptualize one social category can be augmented by the way we conceptualize another. But this approach, as generative as it has been, has some notable limitations. Namely, it is not well-equipped to provide order and parsimony to the emerging, often contradictory literature on intersectional stereotyping. Moreover, it tends to focus psychologists' theorizing on the features of perceivers and targets, casting aside the influence of broader contextual factors on impression formation. Finally, when taken to their extreme, integration models can be incompatible with the very notion of stereotyping itself. Stereotyping requires viewing someone as a member of a social category—not as an integrated product of *every* social category to which they belong.

The following portion of this review describes a compartmentalization model of intersectional stereotyping called intersectional categorization theory (ICT). This model argues that social contexts invite perceivers to select particular lenses for viewing social targets. These lenses can be singular and simplistic (e.g., a race-based lens), or intersectional and complex (e.g., a combined race-by-gender lens). The fundamental assumption of ICT is that perceivers use *one lens at a time* for viewing social targets. When perceivers use the lens of *race*, for example, they will stereotype and discriminate against targets on the basis of race, but not necessarily on the basis of gender. In contrast, when perceivers use a *race-by-gender* lens, they will stereotype and discriminate against targets on the basis of targets' intersecting race and gender categories, but not necessarily on the basis of race by itself or gender by itself.

### **Intersectional Categorization Theory**

Part of the reason psychologists gravitated toward integration models, I said before, was that these models' assumptions were most in line with the take-home message of intersectional

feminism: namely, that “the meaning of each social group is constructed through the lens of the others” (Ghavami et al., 2016, p. 35). But what about compartmentalization models? Are these models’ assumptions *incompatible* with the insights of intersectional feminism? At first blush, perhaps. But on second, perhaps not. To return to the insights of a notable black feminist, Crenshaw (1989) wrote, of discrimination in the legal system:

Black women sometimes experience discrimination in ways [that are] similar to white women’s experiences; sometimes they share very similar experiences with Black men ... And sometimes, they experience discrimination as Black women—not the sum of race and sex discrimination, but as Black women. (p. 149)

This sounds an awful lot like compartmentalization.

As noted previously, an issue with compartmentalization models is that they seem to imply that Black women, for example, can only be perceived as *Black* or as *women*, but not as *Black women* simultaneously. Yet as the above quote illustrates, this is not necessarily the case. Instead, it may be the case that perceivers can represent social targets with varying degrees of intersectional complexity. At times perceivers may construe targets on the basis of simple, singular social identity groups to which they belong, whereas at other times perceivers may construe targets on the basis of complex, intersectional social identity groups to which they belong. Critically, the prejudices that follow from intersectional categories may be wholly distinct from the prejudices that follow from more simplistic categories.

The remainder of this review describes in substance a model of intersectional stereotyping called intersectional categorization theory (ICT). This model is based heavily on the insights and propositions of models that came before it, and advocates the following core ideas:

(a) perceivers use one lens at a time for making sense of other people; (b) the lenses perceivers use can be singular and simplistic (e.g., *elderly*), or intersectional and complex (e.g., *elderly East-Asian women*); (c) which lens perceivers use matters because different lenses bring different prototypes to perceivers' minds. These different prototypes, in turn, can prescribe categorically distinct sets of beliefs, attitudes, and behaviors toward even the *same social targets*. The sections that follow begin with a review of the social identity tradition, highlighting in particular the assumptions this tradition holds about how social groups are represented in the mind. Subsequent sections provide an overview of ICT, delineating 1) the factors that guide which lens a perceiver uses in a given moment in time, 2) the factors that influence the contours of a selected lens (that is, how the lens inflects stereotype content), and 3) the factors that moderate how a lens directs perceivers' behavior. The final sections of this chapter compare ICT with alternative models of person perception, and they provide examples of how ICT can be used to provide order to the convoluted, often contradictory literature on intersectional stereotyping.

### **A Brief History of the Social Identity Tradition**

World War II made an important question about human behavior salient to psychologists: What could compel individuals, when situated within social groups, to commit profound acts of violence against one another? Early attempts to answer this question, which were no doubt influenced by psychodynamic thought (e.g., Freud, 1922; Le Bon, 1908), centered upon the minds of individuals. Perhaps individuals harbored deep levels of pent-up frustration that the group “unleashed” by way of causing them to feel anonymized (Adorno, Fenkel-Brunswik, Levinson, & Sanford, 1950; Diener, Fraser, Beaman, & Kelem, 1976; Dollard, Miller, Doob, Mowrer, & Sears, 1939). Or perhaps instead intergroup aggression was a learned behavior

(Bandura, 1977), or an outcome of individuals' desire to comply with the requests of authority figures (Milgram, 1963). It was not until the late 1960s and 1970s that a different perspective emerged in psychological science: one arguing that the way people behave in groups cannot be understood by analyzing their features as individuals (e.g., Sherif, 1966; 1967).

Henri Tajfel elaborated upon this emerging perspective. Tajfel was finding, as Sherif was, that with relative ease, groups of healthy, normal, non-aggressive people could be experimentally induced to compete with each other (Tajfel, Billig, Bundy, & Flament, 1971). But to his surprise he found that they would do so even when the groups were utterly meaningless (i.e., explicitly random: Billig & Tajfel, 1973), and even when participants were explicitly told that no resources were at stake (Tajfel & Turner, 1979). Moreover, Tajfel noticed that when perceivers thought about social relations in terms of groups—in terms of *us vs. them*, or even in terms of *them vs. them*—they would perceptually maximize differences between groups, and they would perceptually minimize differences within groups (this is often referred to as the accentuation principle: Tajfel, 1959; 1969). These observations, among others, paved the way for what was eventually articulated as the social identity theory of intergroup behavior (or “social identity theory” for short; Tajfel & Turner, 1986). Later, this theory was expanded by Turner and his students into the social identity theory of *intragroup* behavior (or “self-categorization theory” for short; Turner et al., 1987). These theories and their assumptions constitute what is commonly referred to as the social identity tradition.

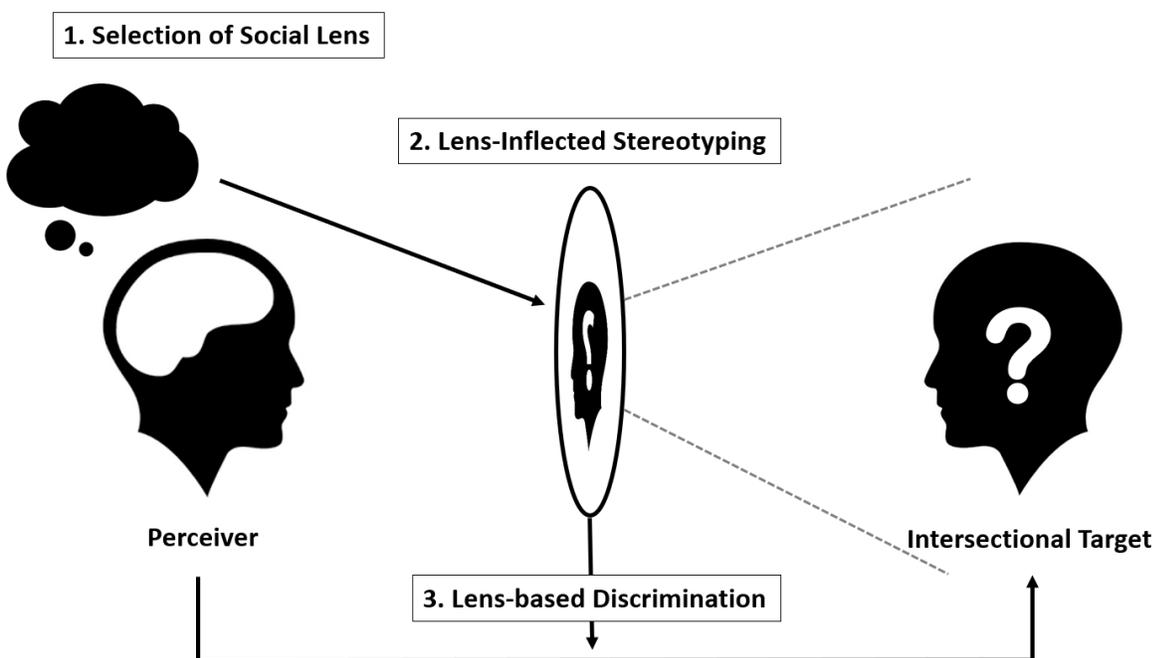
The social identity tradition has evolved over time into a meta-theory: a rich network of assumptions about human psychology that can itself explain other theories (Abrams & Hogg, 2004). Because these assumptions are lengthy and diverse, it would be imprudent to review them

in entirety (but for some notable reviews, see Abrams & Hogg, 1998; 2010; Hogg, Abrams, Otten, & Hinkle, 2004; and Hornsey, 2008). Instead, I will review the assumptions of this tradition that are fundamental to understanding ICT. This begins with the assumption that people cognitively represent social identities as distinct from individual identities (see also: Brewer & Gardner, 1996; Reid & Deaux, 1996). The social identity tradition holds that perceivers cognitively represent social groups in the form of *prototypes*—fuzzy constellations of traits, beliefs, and behaviors that define typical members of a group, and that differentiate that group from other groups. When perceivers categorize a person as a member of a social group, they are thought to depersonalize the categorized person in the direction of the group’s prototype. This is to say that when someone is categorized as a *lawyer*, for example, they stop seeming like the individual that they are and they begin to seem interchangeable with perceivers’ prototype of lawyers (Doise, Deschamps, & Meyer, 1978; Tajfel, Sheikh, & Gardner, 1964). An important assumption of the social identity approach is that “in any given situation, only one identity is psychologically real” to perceivers (Hogg et al., 2004, p. 252). That is, only one social identity, at a given moment in time, is thought to become the salient psychological basis for construing the self or others.

### **Overview of Intersectional Categorization Theory**

The history of the social identity tradition is reviewed above because intersectional categorization theory (ICT) is a derivation from this broader perspective. Like other derivations—for example, the social identity theory of leadership (Hogg, 2001), or uncertainty-identity theory (Hogg, 2007)—ICT maintains the core assumptions of the social identity approach. But also like other derivations, ICT invokes new assumptions, as necessary, for the

explanation of its focal phenomenon: how perceivers stereotype and discriminate against intersectional social targets. Thus, ICT likewise assumes that perceivers cognitively represent groups in the form of prototypes, and that when a target is categorized as a member of a group, they are depersonalized in the direction of the relevant prototype. ICT also presumes that “in any given situation, only one identity is psychologically real” (Hogg et al., 2004, p. 252; Bodenhausen & Macrae, 1998). Finally, ICT explicitly argues that social identities can vary in their complexity. Specifically, it argues that *intersections themselves* can be considered social identities, whose prototypes may be cognitively distinct from the prototypes of non-intersectional, more simplistic “parent” identities. This notion is adapted from social identity complexity theory (Roccas & Brewer, 2002), and it is also articulated in the stereotype activation-inhibition model (Bodenhausen & Macrae, 1998).



*Figure 1.* Illustration of the three psychological processes ICT describes in detail: those relating to lens selection, stereotyping through the inflection of lenses, and discrimination.

The assumptions of ICT will be spelled out across three sections, which are illustrated in Figure 1. The purpose of these sections is to review 1) the factors that guide which lens a perceiver uses at a given moment in time, 2) the factors that influence the contours of a selected lens (that is, how the lens inflects stereotype content), and 3) the factors that moderate whether a lens does or does not direct perceivers' behavior.

### **Assumptions of Lens Selection**

According to Bruner (1957), and later, to Oakes and colleagues (Oakes, 1987; Oakes, Turner, & Haslam, 1991) there are two factors that guide whether or not a category—or in ICT's parlance, a “lens”—comes into focus:

- (1) *Accessibility*, or the ease with which a lens can be retrieved from memory.
- (2) *Fit*, or the extent to which a lens seems to “explain” observed intergroup behavior.

ICT adds a third factor to the mix:

- (3) *Distinctiveness*, or the extent to which a lens-associated identity is rare in a social context.

ICT does not mandate that all three factors need to be simultaneously present in order for a lens to be selected by a perceiver (cf. Oakes et al., 1991); it merely specifies that each of these factors increases the probability that a perceiver will use a lens over alternatives for construing targets. When a lens is selected, targets are thought to be viewed through the inflection of that lens—assimilated toward whatever prototype the lens implies. When a lens is *not* selected, targets are presumed to be viewed as individuals rather than as prototypic members of a group. In these circumstances, a doctor who is both Lesbian and Latina is not viewed through the lenses of *Lesbians*, *doctors*, *Latinas*, or anything else, but is instead construed as an individual with many

attributes (which may or may not include her occupation, her sexual orientation, and her ethnicity). This element of ICT is borrowed from Brewer (1988), who argued that social categories can be represented as attributes of individuals as well as broader groupings toward which individuals can be assimilated. Finally, ICT assumes that lenses can only be selected if they are cognitively available to perceivers (Higgins, 1996; Tulving & Pearlstone, 1966). This is to say that perceivers are not expected to use lenses for which they have no pre-existing beliefs or associations.

**Lens accessibility.** The more accessible a social lens, the more likely a perceiver is to use it for construing social targets. Accessible lenses are those that are readily retrieved from memory (Bruner, 1957; Higgins, 1996). Factors that make a lens accessible to perceivers can be situational or chronic. For example, being situationally primed with a particular identity (e.g., ethnicity) can make that identity more accessible to perceivers and hence more likely to guide perceivers' impressions of targets (Pittinsky, Shih, & Trahan, 2006; Rattan, Steele, & Ambady, 2017; van Rijswijk & Ellemers, 2002; van Twuyver & van Knippenberg, 1995). In addition, accessibility can be increased by situational fluctuations in perceivers' *motivations* (Palma, Garcia-Marques, Marques, Hagá, & Payne, 2019; Volpert-Esmond & Bartholow, 2019; Yamaguchi & Beattie, 2020). For example, when perceivers experience ego-threatening information (like a bad score on an intelligence test), their motivations to restore self-esteem can cause them to view others in a stereotypic light (e.g., Fein & Spencer, 1997), or to cognitively accentuate identity dimensions that distance themselves from threatening others (Mussweiler, Gabriel, & Bodenhausen, 2000). Presumably, this occurs by way of making motivation-serving lenses, like those that allow the perceiver to "see" someone as unintelligent, or as different from

the self, more accessible. As an illustration of this idea, Sinclair and Kunda (1999) informed White perceivers that they had either been evaluated positively or negatively by a Black doctor. When perceivers thought they had been evaluated positively by a Black doctor, stereotypes related to doctors—but not to race—became highly accessible. But when perceivers instead thought they had been evaluated negatively by a Black doctor, stereotypes related to race—but not to doctors—became highly accessible, instead. ICT suggests that these differences in accessibility were driven by motivational desires for perceivers to protect their own egos when faced with unsavory information.

As noted above, the accessibility of social lenses can also be influenced by chronic rather than situational differences that exist across perceivers (Zarate & Smith, 1990). Thus, chronic motivations, such as epistemic needs for cognitive closure (Kruglanski, 2012), desires for system legitimization (Altemeyer, 1988; Jost, Banaji, & Nosek, 2004), and the like, are all expected to influence the probability that perceivers spontaneously use motivation-serving lenses when viewing social targets. If a person is chronically motivated to be racist (Forscher, Cox, Graetz, & Devine, 2015), for example, the lens of *race* is expected to be more accessible in that person's mind than alternative, potentially competing lenses (e.g., *gender*). Support for this idea comes from the fact that chronic racism does seem to correlate with spontaneous use of *race* during impression formation tasks (Stangor, Lynch, Duan, & Glass, 1992).

**Lens fit.** The more a lens appears to “fit” a social context, the more likely a perceiver is to use it for construing social targets. In general, a lens is said to fit a social context if it provides the perceiver with a meaningful explanation for who is doing what (Bruner, 1957). In the social realm, Oakes and colleagues argued that fit comes in one of two forms (e.g., Oakes et al., 1991):

*normative fit* and *comparative fit*. Normative fit describes the extent to which a category or lens is stereotypically associated with a context. For example, in the United States, the criminal justice system is a stereotypically racialized social context (e.g., Eberhardt et al., 2004). As such, perceivers in criminal justice contexts will tend to feel like *race* provides a better normative fit to social reality than other lenses (like *sexual orientation*: Petsko & Bodenhausen, 2019a). ICT adheres to this conceptualization of normative fit, but also expands it to include features of targets. In particular, ICT assumes that the normative fit of a lens can be influenced by how visually stereotypic a target is of a particular social group. For example, it is well-established that targets who look more Afrocentric tend to experience more racial discrimination than targets who look more Eurocentric (e.g., Blair, Judd, & Chapleau, 2004; Eberhardt, Davies, Purdie-Vaughns, & Johnson, 2006). In the view of ICT, this pattern can be explained by the fact that targets' visual Afrocentricity enhances the extent to which the lens of *race* appears to “fit” social reality. In sum, a lens can be normatively “fitting” either because a context is stereotypically linked with an identity, or because features of a target are stereotypically linked with an identity.

Comparative fit describes the extent to which a lens explains patterns of intergroup behavior in a social context. Thus, comparative fit is not influenced by the properties of stereotypes, as normative fit is, but rather by the properties of what groups of people are actually doing. If a perceiver were to walk into a room where a bunch of older adults were arguing with a bunch of young adults, for example, the lens of *age* would provide good comparative fit to the context. This is because in this moment, age would correlate with who's doing what. In contrast a perceiver were to walk into a room where women (regardless of age) were arguing with men (regardless of age), then the lens of *gender* would provide good comparative fit to the

context. In this latter context, perceivers would be expected to stereotype older women as *women*, whereas in the former context they would be expected to stereotype older women as *old*. Although this particular prediction has not been tested, the premise that comparative fit influences which social categories perceivers use has received robust support. For example, when comparative fit emphasizes targets' *political parties* (Biernat & Vescio, 1993), *basketball teams* (Kurzban et al., 2001), or *prisons* (Klauer, Hölzenbein, Calanchini, & Sherman, 2014), perceivers view targets as relatively interchangeable with other members of their political parties, basketball teams, or prisons. However, when targets' behavior is no longer correlated with these group memberships, perceivers' cease to view targets as interchangeable with other members of their social groups.

**Distinctiveness.** The more distinctive a lens-associated identity is in a social context, the more likely a perceiver is to use that lens for construing social targets. A well-known property of the mind is that distinctive information tends to be attention-grabbing (Chapman, 1967; Hamilton & Gifford, 1976). Thus, if a target person has a social identity that is distinctive in a given social environment, it should be more likely than competing identities to grab the perceiver's attention. Suppose for example that the focal target in a given social environment is a South-Asian woman. According to the principle of distinctiveness, this South-Asian woman would be more noticeably "womanly" when surrounded by South-Asian men than when surrounded by White women. In the former context, *gender* would be the lens that optimizes perceptions of this woman's distinctiveness. In the latter context, *ethnicity* would be the lens that optimizes perceptions of this woman's distinctiveness. Although social identity theorists have tended to reject this hypothesis (e.g., Oakes, 1994), there is substantial support for the idea that

distinctive identities activate identity-relevant stereotypes in the minds of perceivers (e.g., Biernat & Vescio, 1993; Nelson & Miller, 1995; Taylor, Fiske, Etcoff, & Ruderman, 1978). Moreover, the role of distinctiveness in shaping perceptions of targets reverberates with the psychological literature on tokenism. This literature reveals that individuals whose identities are distinctive in a social context tend to *feel* that those identities are highly visible to those around them (Kanter, 1977; Yoder, 1991). This is as true of women at West Point military academy as it is of first-generation students at Stanford (Stephens, Markus, & Phillips, 2014; Yoder, Adams, & Prince, 1983).

### **Assumptions of Lens-Inflected Stereotyping**

Once a lens is selected by perceivers—due to accessibility, “fit”, distinctiveness, or some combination of the three—the next question of relevance is how, exactly, this lens influences perceptions of social targets. As previously stated, ICT presumes that lenses invite perceivers to assimilate targets toward whatever prototype the lens implies. In general, the content of prototypes is thought to stem from learned cultural beliefs about social groups (e.g., Dasgupta, 2013; Devine, 1989). However, there are two additional factors that can influence the features of prototypes, and hence, the content of the stereotypes that are applied to targets:

- (1) *Contrast*, or to which group(s) a target’s group is being compared.
- (2) *Favorable self-conception*, or the propensity of a lens to reflect favorably upon the perceiver.

**Contrast.** Prototypes are mentally constructed in ways that perceptually maximize the ratio of between-group to within-group differences. In the social identity tradition, this is referred to as the *meta-contrast principle* (Haslam, Oakes, Reynolds, & Turner, 1999; Hogg et al., 2004).

Applied to ICT, this means that the same targets who are being viewed through the same lens can “look” different, depending on who else is being viewed through that lens (e.g., Haslam & Turner, 1992; Haslam, Turner, Oakes, McGarty, & Hayes, 1992). As a brief example, suppose the perceiver is viewing Latina women through an *ethnicity-by-gender* lens. ICT assumes that in these moments, Latina women will be depersonalized in the direction of what it means (to the perceiver) to be a prototypic Latina woman. But the principle of meta-contrast suggests that the meaning of “Latina womanness” will itself depend on whom these women are being contrasted against (see also van Rijswijk & Ellemers, 2002). Consider a context in which the perceiver is comparing Latina women against Black women *vs.* a context in which the perceiver is comparing Latina women against East-Asian women. When Latina women are being compared against East-Asian women, U.S. perceivers may be inclined to characterize them as relatively low in socioeconomic status (SES; Ghavami & Peplau, 2013). When Latina women are being compared against Black women, however, this may not be the case. This is because in the United States, stereotypes related to SES strongly differentiate between Latina and East-Asian women (Zou & Cheryan, 2017). But these stereotypes do not as strongly differentiate between Latina women and Black women. Thus, the same lens is expected to bring different stereotypes to perceivers’ minds depending on which particular groups are being viewed and contrasted through that lens.

ICT presumes that prototypes are always cognitively represented in ways that accentuate between-group contrasts. Generally, these contrasts are thought to be determined by the intergroup make-up of a social context. To return to the above example, if Latina women are disagreeing with East-Asian women in a social context, then perceivers’ prototype of Latina women will accentuate the features of Latina women that make them differentiable from Black

women. But if Latina women are disagreeing with three groups of women—say, White, Black, and East-Asian women—and if these groups of women all disagree with each other, then perceivers' prototype of Latina women will accentuate the features of Latina women that make them differentiable from *all three* of these other groups. In this way, the prototypes that lenses imply are thought to be highly sensitive to social context—bending and flexing in ways that maximize the perception that no two groups are alike.

**Favorable self-conception.** Prototypes are mentally constructed in ways that allow for favorable self-conception. This idea is prevalent both in social identity theory and in self-categorization theory (Tajfel & Turner, 1979; Turner et al., 1987). ICT argues that favorable self-conception can be achieved by modifying target prototypes in ways that either directly or indirectly benefit the self. Situations in which target prototypes are modified to *directly* benefit the self are those in which perceivers categorize targets as in-group members rather than out-group members. If a lens causes perceivers to “see” targets as in-group members (e.g., women perceiving other women through the lens of *gender*), the features of the prototype that become accentuated in perceivers' minds will be those that are relatively positive. This is because in this context, the perceivers' prototype of the target is directly associated with perceivers' prototype for themselves. In contrast, if a lens causes perceivers to “see” targets as out-group members (e.g., men viewing women through the lens of *gender*), perceivers' conception of the target's prototype is not directly associated with perceivers' prototype for themselves. As such, there is not as much of a psychological need to represent the target's prototype positively. Stated differently, women's prototype of *women* is expected to be more positive, on average, than men's prototype of *women*.

Yet even when a target's prototype is not directly related to perceivers' prototypes for themselves, perceivers may still modify target prototypes in ways that *indirectly* benefit the self. Indirect benefits to the self can be achieved by perceivers' use of *social creativity beliefs*. In social identity theory, social creativity beliefs are those that distort *in-group* prototypes in ways that make the in-group look good (Tajfel & Turner, 1979). ICT places a slightly different emphasis on these beliefs. According to ICT, social creativity can also be used to distort *out-group* prototypes in ways that prevent the in-group from looking bad by comparison. As a brief example, consider two perceivers, one rich and one middle-class, who are viewing a struggling first-generation college student through the lens of *SES* (i.e., socioeconomic status). In this scenario, the rich perceiver's identity—as a rich person—can “look bad” by comparison to the target's identity as a poor person (for example, the rich person could conceivably seem *greedy*, *selfish*, and *entitled* by comparison). Because rich perceivers would be motivated to construct target prototypes in ways that indirectly advantage their prototype for themselves, they should, according to ICT, call to mind a prototype of low-SES students that minimizes these particular comparisons. Thus, rather than viewing the low-SES student as hardworking and deserving of success, for example, the rich perceiver may be more motivated than the middle-class perceiver to cognitively minimize these characteristics. What results might look an awful lot like system justification (e.g., the belief that poor people get what they deserve: Jost et al., 2004).

### **Assumptions of Lens-Based Discrimination**

Once the contours of a selected lens are defined—by the principles of contrast, self-relevance, or both—the perceiver depersonalizes targets in the direction of the lens-associated prototype. When this occurs, ICT assumes that perceivers use this lens, but not alternative lenses,

to guide their behaviors toward targets. This is to say that if prejudiced perceivers are using *gender* as a lens for viewing Black women, they will exhibit gender-based behaviors (e.g., sexism) but not necessarily race-based behaviors (e.g., racism) toward her. That said, two general forces are thought to moderate the expression of lens-associated behaviors:

- (1) *Cultural constrictors*, or cultural constraints on the appropriateness or inappropriateness of lens-associated behaviors.
- (2) *Personal beliefs*, or idiosyncratic notions about the appropriateness or inappropriateness of lens-associated behaviors.

ICT assumes that people have an automatic tendency to rely on lens-associated behaviors when interacting with social targets (whom they have depersonalized toward a prototype). Thus, if a target has been depersonalized toward a prototype that the perceiver evaluates positively, the perceiver will have an automatic tendency to treat that target positively. If a target has been depersonalized toward a prototype that the perceiver evaluates negatively, the perceiver will have an automatic tendency to treat that target negatively. These automated behaviors, however, can be augmented by cultural constrictors, personal beliefs, or both. In general, it is assumed that when cultural constrictors are strong, or when personal beliefs are strong, they will shape perceivers' lens-associated behaviors. But when these forces are weak, or when perceivers do not have sufficient motivation or opportunity to adjust their behaviors in light of them (e.g., Fazio 1990; Fazio & Olson, 2014), then automatic, lens-associated behaviors will ensue.

**Cultural constrictors.** Cultural constrictors are properties of cultures or systems—most notably norms (Cialdini, 2012; Brauer & Chaurand, 2010), values (Schwartz, 1992), and policies (Krieger & Fiske, 2006)—that constrain the variability of acceptable or unacceptable intergroup

behavior. Suppose, for example, that it is normatively inappropriate in a high school to say things that are sexist, but that it is normatively appropriate in that high school to say things that are ageist. ICT presumes that students in this school will be inclined to make prejudiced remarks to an older-woman principal when she is viewed as an *old person*, but not when she is viewed as a *woman*. As another example, suppose a hiring manager works for a company that espouses values relating to diversity, equitable hiring practices, and inclusion. If the values of the organization are strong enough, this hiring manager may be more inclined to hire White men who out themselves as gay than White men who do not—even if, in this example, the hiring manager’s prototype of gay people is more negative than their prototype of heterosexual people.

ICT assumes that cultural constrictors can vary in their strength. This idea is adapted from the notion of cultural tightness and looseness (Gelfand et al., 2011; Gelfand, Nishii, & Raver, 2006). There are likely to be some cultures or systems that strongly prescribe appropriate and inappropriate intergroup behaviors, as well as other cultures or systems that only weakly prescribe appropriate and inappropriate intergroup behaviors. Generally, it is assumed that as the strength of cultural constrictors increases, the variability in lens-associated behaviors decreases. Viewing someone through the lens of sexual orientation in a system with weak constrictors, for example, will allow for a much broader array of lens-associated behaviors than will viewing someone through the lens of sexual orientation in a system with strong constrictors.

**Personal beliefs.** The next major force that can moderate the expression of lens-associated behaviors is the perceiver’s personally held beliefs (e.g., Crandall & Eshleman, 2003; Dunton & Fazio, 1997; Monteith, 1993; Plant & Devine, 1998). As noted previously, lenses will prescribe behaviors toward targets in a fashion that is relatively automatic. Thus, viewing an

older person through the lens of *age* may spontaneously prescribe certain forms of behavior toward the older person (e.g., talking more slowly, making more eye contact). But, provided that the perceiver is motivated and has the opportunity to reflect on his or her beliefs about group-based behaviors, it is expected that *beliefs about intergroup behavior*—rather than one’s prototype of what members of a group are like per se—will be what predict the perceiver’s behavior. To continue with the present example, suppose that a perceiver’s prototype of *old people* includes notions that they are fragile and weak. But suppose as well that this perceiver believes it is condescending and gross to treat older individuals as though they are fragile and weak. ICT proposes that the perceiver will supplant his or her prototype-based behaviors with belief-based behaviors, and that the perceiver will in turn avoid saying anything that might be condescending to older adults—that is, on the condition that they are using the lens of *age*. Of note, ICT predicts that beliefs that correspond to prototypes can come in and out of focus as lenses themselves come in and out of focus. An implication of this is that people who wish to be racially egalitarian, for example, may be nicer to racial outgroup members when viewing them through the lens of *race* than when viewing them through alternative lenses (e.g., *gender*).

### **The Benefits of ICT as a Model of Intersectional Stereotyping**

ICT is designed to make sense of the rapidly expanding literature on intersectional stereotyping. The greatest advantage of ICT over extant models of intersectional stereotyping is that it characterizes social identities as compartmentalized in the perceiver’s mind. Thus, unlike dominance models, which argue that some social identities take precedence over others in perceivers’ minds (e.g., Kurzban et al., 2001; Pietraszewski et al., 2015), ICT argues that intersectional stereotyping is a flexible process. And unlike integration models—which are by far

the most prevalent models in this research area (e.g., Freeman & Ambady, 2011)—ICT argues that this flexible process is constrained by predictable patterns of lens selection. Lenses are not used all at once for conceptualizing social targets, but are instead theorized to come in and out focus as the social environment demands. This is a much-needed perspective in the literature on intersectional stereotyping (Petsko & Bodenhausen, 2020).

ICT is, at its heart, attempting to resurface a series of assumptions that have been overlooked in intersectional stereotyping research. Many of these assumptions are not new (for a review, see Bodenhausen, 2010). The idea that shifting mental categories—or in ICT’s parlance, shifting lenses—exert corresponding shifts in activated stereotypes is featured in the stereotype activation-inhibition model and in self-categorization theory (Bodenhausen & Macrae, 1998; Turner et al., 1987). And the idea that social identities can vary in their complexity is one that pervades not only these earlier models, but many others as well (Hewstone, 1994; Roccas & Brewer, 2002). Still, the combination of these models’ insights into something that can explain intersectional stereotyping is itself a new endeavor. And indeed, ICT diverges from these earlier models in at least a few important ways.

Social identity complexity theory construes the tendency for perceivers to think of others in complex, intersectional ways as a chronic individual difference in perceivers (Roccas & Brewer, 2002). According to this perspective, some perceivers consistently think of the world in terms of complex social identifications, and other perceivers consistently think of the world in terms of more simplistic social identifications. ICT proposes that the tendency to view others through intersectional lenses is variable, coming in and out of focus depending upon the perceiver’s environment, as well as on fluctuations in perceivers’ motivations. The stereotype

activation-inhibition model (Bodenhausen & Macrae, 1998)—which did allow for the possibility that complex lenses can come in and out of focus—paid relatively little attention to fluctuations in lens-inflected stereotyping. In particular, the stereotype activation-inhibition model made few assumptions about how a selected lens, like the lens of *ethnicity*, might imply distinct sets of stereotypes about targets depending on whom a target is being contrasted against, or on whether the target's category reflects poorly (*vs.* favorably) on the self. Finally, although ICT is a derivation of the social identity tradition, there are a few ways in which it diverges from its lineage. For example, earlier identity theorists tended to reject the idea that the distinctiveness of a social identity could play a role in whether perceivers use that identity for construing others (e.g., Oakes, 1994). ICT, however, argues that it does (Biernat & Vescio, 1993; Taylor et al., 1978). As another example, the social identity tradition has not paid much attention to the role of target prototypicality (for example, Afrocentricity) in shaping whether or not a social category is likely to be activated in the minds of perceivers. ICT acknowledges the possibility that features such as target Afrocentricity, for example, can play an important role in guiding whether a lens seems to normatively fit social reality. Finally, social identity theory did not explicitly discuss the role of what moderates the expression of prototype-based behaviors. ICT has attempted to make clear that the strength of cultural constrictions, as well as the strength of personally-held beliefs, can and often do influence the behaviors perceivers enact toward social targets.

### **ICT as a Tool for Explaining Psychological Findings**

ICT can reconcile discrepant research findings in the literature on intersectional stereotyping. Experiments are social contexts that experimenters build. As such, experiments vary in the very dimensions that ICT describes. Take, for example, the principle of normative fit.

As noted previously, a lens normatively “fits” a social context (and should take precedence in the mind over other lenses), if that lens is stereotypically linked with the context itself. Perhaps one of the reasons why perceivers take gay Black men’s race—but not their sexual orientation (Petsko & Bodenhausen, 2019a)—into account when they see these men in a criminal sentencing context is because, as discussed previously, criminal sentencing contexts are normatively aligned with race much more than with sexual orientation (e.g., Eberhardt et al., 2004; 2006). Perhaps this can also explain why, when categorizing weapons as quickly as possible following exposure to Black vs. White faces (Payne, 2001), perceivers exhibit pronounced racial biases, but virtually no bias whatsoever on the bases of age (Todd et al., 2016a; 2016b; Lundberg et al., 2018).

As illustrated in the above examples, an intriguing property of ICT is that it can account for null results. In fact, ICT argues that every positive result should be accompanied by a suite of non-positive (or at least very close to non-positive) results. When a social context is built that normatively fits the lens of *age*, for example, it should bring the lens of age into focus *at the expense* of alternative lenses. As such, in contexts such as these, the stereotypes that perceivers exhibit should vary systematically in response to targets’ age, but not in response to alternative features of targets, like gender, sexual orientation, and so on. The recent uptick in psychologists’ use of Bayes’ factors (Kruschke & Liddell, 2018), and of equivalence testing (Lakens, 2017)—both of which allow the researcher to weigh evidence in favor of null over alternative hypotheses—yields much promise for exploring the implications of ICT.

## **Conclusion**

The social identity tradition contains a rich set of theoretical assumptions that can be used to explain a great many psychological phenomena—from conformity, to leadership, to

stereotyping (Hogg & Abrams, 1988). ICT is couched within this tradition, but is tailored specifically to the explanation of intersectional stereotyping. ICT argues that perceivers use one lens at a time for making sense of social targets, and that these lenses can vary in their complexity. The lens a perceiver uses in a social context is thought to be guided by three factors: how *accessible* a lens is, how much a lens *fits* social reality (normatively or comparatively speaking), and whether a lens-associated identity is *distinctive* in a social environment.

Critically, ICT assumes that as one lens comes into focus, others fall out of focus—and that the selected lens is what shapes perceptions of social targets. Once selected, lenses are thought to inflect perceptions of social targets in ways that maximize the perception of intergroup *contrast*. In addition, lenses are thought to inflect perceptions in ways that allow the perceiver to maintain a *favorable self-conception*. As a lens inflects perceptions of a target, so too does it shape the perceiver's behaviors toward targets. Behaviors are assumed to follow from lens-inflected stereotyping in ways that are relatively automatic. However, lens-associated behaviors can be subject to modification by *cultural constrictors*—shared norms, values, or policies—and by *personal beliefs* about acceptable intergroup behavior (e.g., racial egalitarianism).

Integrationist assumptions have predominated the development of research on intersectional stereotyping and prejudice. These assumptions have been highly generative, but they have allowed for the development of research findings that inconsistent with one another. ICT emphasizes a different set of assumptions—compartmentalist assumptions—and is designed to provide a parsimonious account for why inconsistent findings co-exist in this literature. ICT suggests that because perceivers' minds operate in a compartmentalized way, it is not surprising that the same targets, for example gay Black men, would be perceived as lacking “Blackness” in

some contexts (Petsko & Bodenhausen, 2019b), but lacking “gayness” in others (Petsko & Bodenhausen, 2019a). Of course, these studies are being interpreted from the perspective of ICT post hoc. And to be sure, studies such as these tend to vary from one another in a multitude of ways. As such, a series of experiments is needed to test the basic premises of ICT.

### **The Present Experiments**

The following experiments lay a foundation for ICT by critically examining each of three core ideas: (a) that perceivers use one lens at a time for making sense of other people; (b) that the lenses perceivers use can be singular and simplistic (e.g., viewing an older East-Asian woman as an *old person*), or intersectional and complex (e.g., viewing the same individual as an *older East-Asian woman* specifically); and (c) that different lenses can prescribe categorically distinct sets of stereotypes that perceivers use as frameworks for thinking about targets. The experiments in Chapter II provide initial evidence that lens usage can be compartmentalized. In particular, Experiments 1a, 1b, and 2 explore whether focusing perceivers’ attention on the lens of gender causes them to cease to attend—at least in these moments—to targets’ age groups (Experiments 1a and 2) or racial groups (Experiment 1b), and vice-versa. The next chapter—Chapter III—provides initial evidence that different lenses can bring different sets of stereotypes to perceivers’ minds. Specifically, Experiment 3 investigates whether shifting perceivers’ attention from one lens (e.g., age) to another lens (e.g., gender) correspondingly changes the stereotypic attributes that they come to associate with older women. Chapter IV presents experiments that examine intersectional lens usage: Experiment 4 investigates whether perceivers occasionally attend to intersections of identities themselves (e.g., race *and* gender) in lieu of singular identities (e.g.,

gender alone); Experiments 5a and 5b investigate whether the stereotypes perceivers associate with targets vary as a function of whether perceivers are using intersectional vs. singular lenses.

## Chapter II: Initial Evidence of Compartmentalized Lens Usage

## Experiments 1a and 1b

Experiments 1a and 1b were designed to test the idea that perceivers use one lens at a time for making sense of other people. Experiment 1a tested whether perceivers' impressions of older women become less gendered when they attend to these women's age (and vice versa). Experiment 1b tested whether perceivers' impressions of Black women become less gendered when they are attending to these women's race (and vice versa). A notable contribution of these experiments—above and beyond testing whether intersectional stereotyping can be compartmentalized—is that their predictions compete with a prevalent argument in evolutionary psychology. Specifically, these experiments compete with the argument that perceivers cannot help but to engage in sex and age categorization (Kurzban et al., 2001; Pietraszewski et al., 2015; Sidanius & Pratto, 2012). Whereas as Tooby and Cosmides (2010) argue that sex and age categorization are chronically implemented by perceivers, ICT argues that this is not the case. Instead, ICT argues that perceivers will cease to attend to targets' age groups when paying attention to targets' gender (Experiment 1a), and that perceivers will cease to attend to targets' gender groups when paying attention to targets' race (Experiment 1b).

### Experiment 1a

Participants in Experiment 1a saw 12 interaction partners (half young-adult, half older-adult, half men, and half women) having a disagreement that either comparatively “fit” the interaction partners' age groups (age-fit condition), gender groups (gender-fit condition), or neither their age nor gender groups (control condition). As a reminder, *comparative fit* describes the extent to which a lens correlates with patterns of intergroup behavior (Oakes, 1987). Comparative fit was manipulated by varying patterns of agreement and disagreement among the

interaction partners. For example, in the age-fit condition, the older adults agreed with one another, but disagreed with the young adults. It is worth mentioning that *normative fit*—the extent to which a context is stereotypically associated with particular lenses—was held constant in this experiment. This was done by pre-testing the conversation topics to ensure that they were neutral with respect to age and gender. After the conversation concluded, participants listed stereotypes toward one of the older women who had participated. At the end of the experiment, participants completed a memory task for who said what during the conversation from earlier.

Participants in the age-fit condition were expected to show evidence of viewing targets through the lens of *age*. As such, participants in the age-fit condition were expected to engage in higher levels of age categorization and age stereotyping than participants in the other two conditions. Participants in the gender-fit condition were expected to show evidence of viewing targets through the lens of *gender*. As such, participants in the gender-fit condition were expected to engage in higher levels of gender categorization and gender stereotyping than participants in the other two conditions. In addition, these patterns of categorization and stereotyping were expected to *trade off*: perceivers were expected to exhibit less age categorization and stereotyping when relying on the lens of *gender*, and they were expected to exhibit less gender categorization and stereotyping when relying on the lens of *age*.

**Method.** In Experiment 1a, participants viewed 3 older women (ages 70-80), 3 young women (ages 20-30), 3 older men (ages 70-80), and 3 young men (ages 20-30) discussing one of two conversation topics (an internal replication factor) in a 3 (comparative fit: age, control, gender)  $\times$  2 (conversation topic) between-person design. Experiment 1a had an a priori decision rule to include approximately 100 people per “fit” condition.

**Participants.** A total of 302 U.S. citizens were recruited from MTurk to complete Experiment 1a in exchange for \$3.00. Of these,  $n = 7$  (2.32%) were excluded for not responding “yes” to the question, “Did you take this study seriously?” The remaining participants were mostly male (170 male, 125 female), mostly White (235 White, 23 Black, 14 Asian, 13 Latinx, 4 American Indian, 4 multi-racial, 2 non-respondent), and had ages spanning from 19 to 72 ( $M = 36.47$ ,  $SD = 10.17$ ). In addition, this sample was well-educated (53.58% held at least a bachelor’s degree), and they skewed toward political liberalism ( $M = 3.92$ ,  $SD = 3.04$ , on an 11-point scale from  $0 = extremely liberal$  to  $10 = extremely conservative$ ).

**Procedure.** After providing consent, participants learned that they were going to view a conversation between 12 people who would be discussing their answers to one of two questions: either “Are introverts more intelligent than extraverts,” or the question, “If a person mirrors whomever they’re talking to, are they phony?” These topics were chosen because people can plausibly disagree over them, and critically, because they are normatively neutral with respect to gender and age (according to pre-testing; see Appendix A). Participants were then shown the following instructions:

Please pay attention to who says what. At the end of the conversation, we’ll ask you questions about what you remember from the conversation, and we’ll ask you about your impressions of some of the people who participated.

After this prompt, participants were shown all 12 interaction partners, side by side in an image matrix (see Figure 2). Each matrix featured 12 standardized headshots of White men and women, of whom half were older adults (i.e., in their 80s), and of whom half were young adults (i.e., in their 20s). Images were used with permission from the FACES database (Ebner, Riediger, &

Lindenberger, 2010). To ensure that Experiment 1a sampled images from a broad set of possible stimuli, each participant saw one of nine possible matrices—that is, one of nine possible sets of 12 interaction partners. Image set was included as an internal replication factor and was therefore not expected to meaningfully influence results.



*Figure 2.* One of nine possible image matrices of interaction partners (an internal replication factor) whom participants saw engage in debate in Experiment 1a.

Participants proceeded through a slideshow that depicted whichever conversation they had been randomly assigned. In each conversation, the 12 interaction partners took turns voicing their opinions on the topic at hand, and they readily divided into two camps. It was always the case that half of the conversation partners took one position in the conversation (e.g., that introverts *are* more intelligent than extraverts), and that the other half of conversation partners took the opposite position in the conversation (e.g., that introverts are *not* more intelligent than extraverts). In the age-fit condition, all the older adults (collapsing across gender) took one position in the conversation, and all the young adults (collapsing across gender) took the other.

In the gender-fit condition, all the women (collapsing across age) took one position, and all the men (collapsing across age) took the other. Who took which position in these conditions was always counterbalanced (i.e., some people saw older adults take pro-introvert position, but others saw older adults take an anti-introvert position). In the control condition, men and women of both age groups took one position as well as the other. Each interaction partner spoke three times, resulting in a total of 36 spoken statements.

After the conversation concluded, participants were given the following instructions: “Now that the conversation is over, we’d like to gather some information on how you think *the average American* might perceive (and stereotype) some of the people from the conversation you viewed.” After a page break in the survey, this point was reiterated: “Please note that we are *not* interested in your personal beliefs. Instead, we want to know how the *average American* might perceive (and stereotype) one or more of the people you saw.” This emphasis was included to help circumvent any social desirability concerns that participants might have when reporting on their stereotypes. Previous research has shown that statements such as these—which exonerate participants from personal responsibility—do indeed increase their likelihood of reporting negative stereotypic beliefs about social groups (Devine & Elliot, 1995; Ghavami & Peplau, 2013).

After these instructions, participants were shown one of the older women from the conversation they viewed, and they listed stereotypes about her and provided ratings of her facial features. Participants were expected to characterize the older woman as possessing “older” traits, and as looking more typical of older adults in the age-fit condition than in the other two conditions. By contrast, participants were expected to characterize the older woman as

possessing “more feminine” traits, and as looking more typical of women in the gender-fit condition than in the other two conditions.

***Application of age and gender stereotypes.*** Stereotype application was measured using a checklist procedure (e.g., Devine & Elliot, 1995; Katz & Braly, 1933). Participants were shown 99 traits on a checklist (*loud, witty, gentle, etc.*), and their task was to “select ALL the traits that the average American might use when stereotyping” one of the older women from the conversation they saw. Unbeknownst to participants, the checklist attributes in this task had been rated by two separate groups of MTurkers on how stereotypically feminine they seem ( $N = 80$ ), and on how stereotypic of old people they seem ( $N = 77$ ; see Appendix B for details). Each of the 99 traits was rated by these samples on scales from  $1 = not\ at\ all\ stereotypically\ feminine\ [old]$  to  $7 = very\ stereotypically\ feminine\ [old]$ . As such, participants’ trait nominations could be scaled numerically in terms of how “feminine” and how “old” they seem, on average. Of note, ratings of trait femininity and oldness were moderately correlated,  $r(97) = .48, p < .001$ .

***Face ratings.*** After listing the stereotypes that could be used to characterize one of the older women from the conversation they saw, participants were shown the face of the same older woman again, and they were asked to rate this face on how typical the older woman’s face looked of older adults, and of women, on scales from  $1 = not\ at\ all$  to  $9 = very\ much$ .<sup>2</sup>

***Who-said-what task.*** Finally, participants completed a memory task of the conversation itself. In a randomized order, participants were shown each of the 36 statements that were spoken during the conversation. These statements were intermixed with 36 distractor statements. For each statement, participants’ task was to first answer the question, “Have you seen this statement

---

<sup>2</sup> Participants also rated these faces on a few additional attributes (e.g., on how happy they look), and they provided ratings as well of, respectively, a young woman’s face and of an older man’s face.

before?”, and if they answered yes, to then indicate who said the statement. Each time participants indicated who said a statement, they were shown the 12 photos of their assigned conversation partners (presentation order was randomized for each trial). Of note, participants were told in advance that half of the statements would be distractor statements, and that half would be real.

**Results.** Results are presented while collapsing across internal replication factors, as including these factors in the models does not meaningfully change their interpretation. Stimulus set (e.g., which set of 12 interaction partners participants were assigned) accounted for up to 5.46% of the variance in face ratings, but less than 1.00% of the variance on all other measures. Conversation topic—a discussion of introverts *vs.* phoniness—occasionally does moderate the magnitude of reported effects, but not in ways that change their interpretation (see Appendix C).

***Who-said-what: Age categorization.*** Participants were hypothesized to be categorizing targets by age more in the age-fit condition—when the lens of age is active—than in the other two conditions. To analyze this, memory errors on the who-said-what task were categorized according to whether they were within-age-group errors (for example, confusing older people for older people) or between-age-group errors (for example, confusing older people for young people). Under this approach, *age categorization* is indexed by the extent to which within-age-group errors are more frequent than between-age-group errors.<sup>3</sup> To examine participants’ levels of age-categorization, errors were subjected to a 2 (error type: within-age-group, between-age-group) × 3 (condition: age fit, control gender fit) mixed ANOVA with repeated measures on the

---

<sup>3</sup> An adjustment was made to these error frequencies to account for the fact that some errors are more probable than others (e.g., for the fact that there are six possible ways to make a between-age-group error but only five possible ways to make a within-age-group error). This adjustment was made following the advice of Pietraszewski (2018).

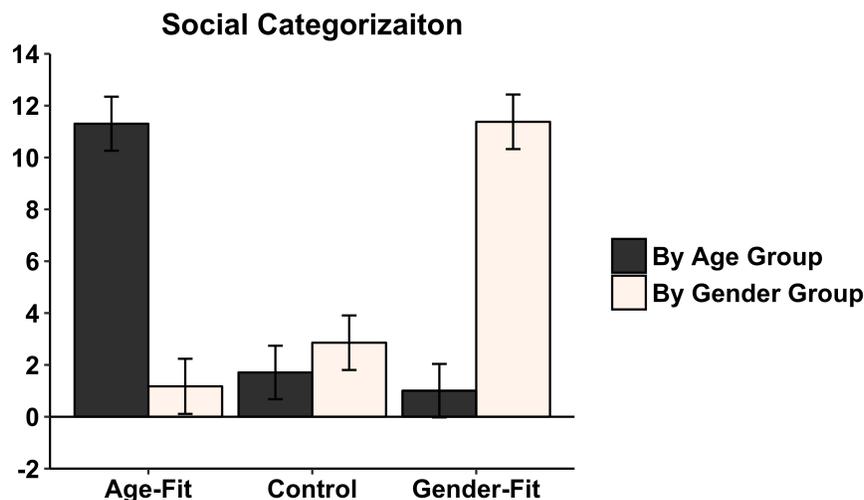


Figure 3. Within- relative to between-group errors (Exp. 1a), broken down by whether participants were in conditions that emphasized the fit of age (left), gender (right), or neither age nor gender (middle). Mean error differences are encompassed by 95% CIs.

first factor.<sup>4</sup> This analysis revealed, first, a main effect of error type, suggesting that across all conditions, participants exhibited a general tendency to categorize targets by age. That is, when participants could not remember who said what, they made significantly more within-age-group errors ( $M = 10.29$ ,  $SE = 0.22$ ) than between-age-group errors ( $M = 5.61$ ,  $SE = 0.22$ ),  $M_{diff} = 4.67$ , 95% CI[4.08, 5.27],  $\beta = 0.93$ ,  $F(1, 295) = 234.98$ ,  $p < .001$ . In addition, and in line with ICT, the degree of age-based categorization was heavily moderated by whether or not participants were in the condition that comparatively emphasized targets' age:  $F(1, 295) = 234.95$ ,  $p < .001$ ,  $\omega_p^2 = 0.44$ . Age categorization was substantially greater in the condition that emphasized age [ $M_{diff} = 11.30$ , 95% CI[10.26, 12.35],  $\beta = 2.25$ ,  $F(1, 295) = 452.16$ ,  $p < .001$ ] than it was in the other two conditions [ $M_{diff} = 1.36$ , 95% CI[0.63, 2.09],  $\beta = 0.27$ ,  $F(1, 295) = 13.32$ ,  $p < .001$ ]. The degree of age categorization was similar whether participants were in the control condition or in the

<sup>4</sup> Mixed-model analyses throughout this dissertation are conducted using the “lme4” and “lmerTest” packages in R (Bates, Mächler, Bolker, & Walker, 2014; Kuznetsova, Brockhoff, & Christensen, 2016). Variability in degrees of freedom is attributable to approximation variability rather than to missing data.

gender-fit condition (see Figure 3),  $F(1, 295) = 0.88, p = .35, \omega_p^2 < .01$ . These data are consistent with ICT: when age was comparatively emphasized (*vs.* not), participants were more likely to use age as a category for organizing their memories. Yet when gender was comparatively emphasized, or when neither age nor gender was emphasized, age categorization plummeted (in raw numbers) by a factor of 10. These findings not only support ICT, but strongly contradict the argument that age categorization is inevitable (e.g., Pietraszewski et al., 2015).

***Who-said-what: Gender categorization.*** Participants were hypothesized to be categorizing targets by gender more in the gender-fit condition—when the lens of *gender* is active—than in the other two conditions. To examine this, participants’ memory errors for who said what were analyzed according to the same  $2 \times 3$  analysis described above. This analysis revealed a main effect of error type, suggesting that on average, participants were using gender as a category for organizing social information. That is, when participants could not remember who-said-what, they were more likely to erroneously attribute a statement to someone of the same gender category ( $M = 10.54, SE = 0.22$ ) than to someone of the opposite gender category ( $M = 5.40, SE = 0.22$ ),  $M_{\text{diff}} = 5.14, 95\% \text{ CI}[4.53, 5.75], \beta = 1.00, F(1, 295) = 272.75, p < .001$ . Furthermore, and as expected, the degree of gender categorization was heavily moderated by whether participants were in the condition that comparatively emphasized gender (see Figure 3):  $F(1, 295) = 202.09, p < .001, \omega_p^2 = 0.40$ . Participants in the gender-fit condition made a substantial number of within- relative to between-gender-group confusions:  $M_{\text{diff}} = 11.38, 95\% \text{ CI}[10.31, 12.44], \beta = 2.22, F(1, 295) = 449.28, p < .001$ . Participants in the other two conditions exhibited the same pattern, but to a much weaker degree:  $M_{\text{diff}} = 2.01, 95\% \text{ CI}[1.27, 2.76], \beta = 0.39, F(1, 295) = 27.90, p < .001$ . Finally, gender categorization was moderated by whether

participants were in the age-fit vs. control conditions,  $F(1, 295) = 4.87, p = .028, \omega_p^2 = .01$ . The nature of this interaction was that gender categorization occurred more in the control condition [ $\beta = 0.55, F(1, 295) = 28.32, p < .001$ ] than in the age-fit condition [ $\beta = 0.23, F(1, 295) = 4.68, p = .031$ ]. This is supportive of the lens-switching dynamics that ICT proposes: of the possibility that when age comes into focus, gender falls out of focus. Moreover, it directly contradicts the argument that gender is in focus at all times and across all contexts (Tooby & Cosmides, 2010).

Table 1.

*Fifteen Most-Selected Traits for Old Women (Exp. 1a), Broken Down by Experimental Condition*

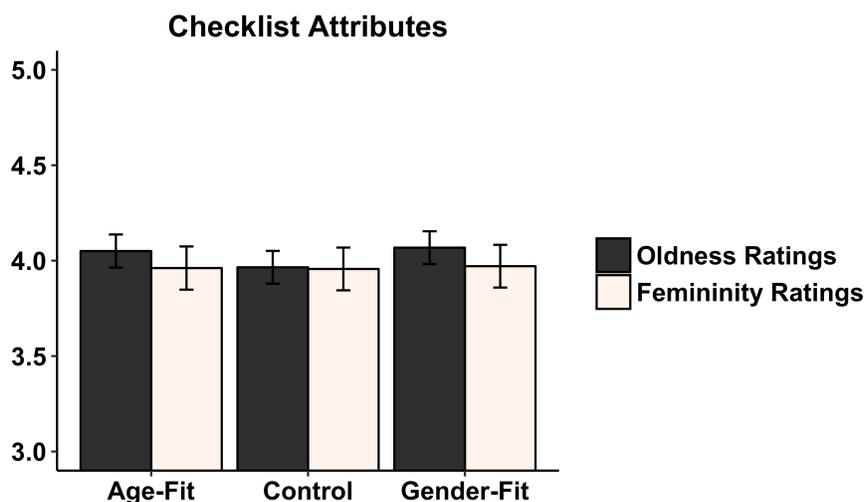
Age ( $n = 97$ )		Control ( $n = 99$ )		Gender ( $n = 99$ )	
Trait	Prop.	Trait	Prop.	Trait	Prop.
Conservative	0.43	Stubborn	0.43	Stubborn	0.46
Stubborn	0.43	Conservative	0.36	Conservative	0.43
Argumentative	0.36	Argumentative	0.33	Tradition-loving	0.35
Conventional	0.35	Quarrelsome	0.33	Very Religious	0.35
Tradition-loving	0.34	Rude	0.32	Argumentative	0.34
Quarrelsome	0.30	Ignorant	0.29	Rude	0.33
Humorless	0.29	Tradition-loving	0.28	Humorless	0.32
<b>Suspicious</b>	0.29	Very Religious	0.28	Conventional	0.28
Rude	0.28	Humorless	0.26	Quarrelsome	0.26
Straightforward	0.27	<b>Low in Intelligence</b>	0.25	<b>Hostile</b>	0.25
Very Religious	0.26	Straightforward	0.25	Ignorant	0.25
Loyal to Family	0.25	<b>Aggressive</b>	0.24	Loyal to Family	0.24
Uneducated	0.21	Conventional	0.24	Straightforward	0.24
Ignorant	0.20	Loyal to Family	0.24	Uneducated	0.24
<b>Quick-tempered</b>	0.20	Uneducated	0.23	Honest	0.22

*Note.* Prop. = the proportion of participants in a condition who chose a given trait to characterize the stereotypes that would be attributed to an old woman from the conversation they viewed. Bold font indicates that a trait appears in the top 15 traits of only one condition.

**Checklist attributes.** Did the stereotypes that perceivers nominate for *individual* older women vary by comparative fit condition? For illustrative purposes, the 15 most-selected traits from each condition (age fit, control, gender fit) are listed in Table 1. Descriptively speaking,

these attributes are highly similar across experimental conditions. And indeed, only five of the 45 traits (marked by bold font) were unique to the 15 most-selected traits of a particular condition.

**Checklist attributes: Stereotypic oldness.** Participants were expected to nominate traits for older women that were rated as “older” in the age-fit condition, when the lens of *age* was activated, than when in the control or gender-fit conditions. To examine this, the “oldness” of participants’ trait nominations was subjected to a one-way ANOVA. Contradicting hypotheses, this analysis yielded null results. Participants’ trait attributions were no “older” in the age-fit condition ( $M = 4.05$ ,  $SE = 0.04$ ) than in the control or gender-fit conditions ( $M = 4.02$ ,  $SE = 0.03$ ),  $M_{diff} = 0.03$ , 95% CI[-0.07, 0.14],  $\beta = 0.07$ ,  $F(1, 292) = 0.39$ ,  $p = .53$ . Moreover, these latter conditions did not differ from each other:  $M_{diff} = -0.10$ , 95% CI[-0.22, 0.02],  $\beta = -0.24$ ,  $F(1, 292) = 2.76$ ,  $p = .10$  (see Figure 4). Thus, participants did not nominate stereotypes for older women that were any “older” in the age-fit condition than in the other conditions.



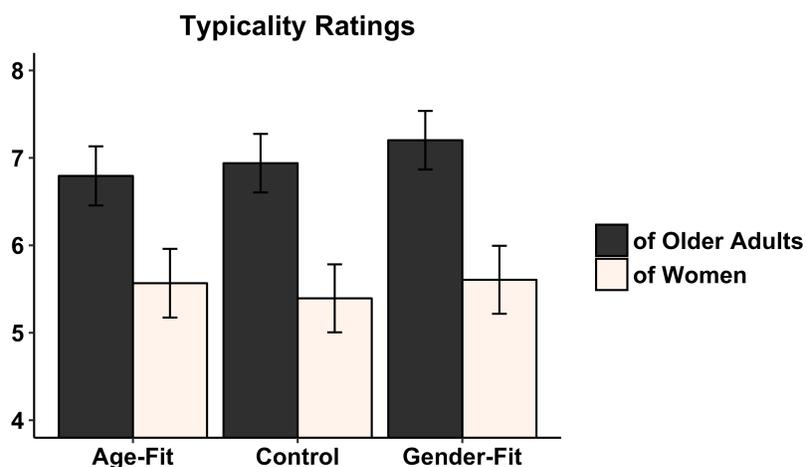
*Figure 4.* How “old” and “feminine” trait nominations for older women were rated to be (Exp. 1a), broken down by whether participants were in conditions that emphasized the fit of age (left), gender (right), or neither age nor gender (middle) categories. Means are encompassed by 95% confidence intervals.

**Checklist attributes: Stereotypic femininity.** Participants were expected to nominate “more feminine” attributes for older women in the gender-fit condition than in the other two conditions. Subjecting the stereotypic femininity of participants’ nominations to a one-way ANOVA, however, revealed that this was not the case. Participants’ trait nominations for older women were no more or less “feminine” in the gender-fit condition ( $M = 3.97, SE = 0.05$ ) than in the other two conditions ( $M = 3.96, SE = 0.04$ ),  $M_{diff} = 0.01$ , 95% CI $[-0.13, 0.15]$ ,  $\beta = 0.02$ ,  $F(1, 292) = 0.03, p = .86$ . Furthermore, these other two conditions did not differ from each other,  $M_{diff} > -0.01$ , 95% CI $[-0.16, 0.15]$ ,  $\beta = -0.01$ ,  $F(1, 292) < 0.01, p = .95$  (see Figure 4).

**Face ratings: Typical of older adults.** Participants were expected to rate the faces of older women as “older looking” when they were in the age-fit condition relative to the other two conditions. A one-way ANOVA indicated, however, that this was not the case: older women were rated as looking no “older” in the age-fit condition ( $M = 6.79, SE = 0.17$ ) than in the control or gender-fit conditions ( $M = 7.07, SE = 0.12$ ),  $M_{diff} = -0.28$ , 95% CI $[-0.69, 0.14]$ ,  $\beta = -0.16$ ,  $F(1, 292) = 1.73, p = .19$ . In addition, ratings in these latter conditions did not differ from each other,  $M_{diff} = -0.26$ , 95% CI $[-0.73, 0.21]$ ,  $\beta = -0.15$ ,  $F(1, 292) = 1.18, p = .28$  (see Figure 5). Thus, in contrast to ICT, there was no evidence that older women were rated as looking “older” in the age-fit condition relative to the other conditions.

**Face ratings: Typical of women.** Finally, participants were expected to regard older women as looking more typical of women in the gender-fit condition than in the other two conditions. Yet a one-way ANOVA on typicality ratings revealed that this was not the case. Participants did not rate older women as looking more typical of women in the gender-fit condition ( $M = 5.61, SE = 0.20$ ) than in the other two conditions ( $M = 5.48, SE = 0.14$ ),  $M_{diff} =$

0.13, 95% CI[-0.35, 0.60],  $\beta = 0.06$ ,  $F(1, 292) = 0.27$ ,  $p = .61$ . Moreover, ratings in these other two conditions did not differ from each other,  $M_{diff} = -0.17$ , 95% CI[-0.73, 0.38],  $\beta = -0.09$ ,  $F(1, 292) = 0.38$ ,  $p = .54$  (see Figure 5). Thus, whereas the predictions of ICT *were* borne out on the who-said-what measure, which indexes social category usage, the predictions of ICT were *not* borne out on measures of stereotype application.



*Figure 5.* Typicality ratings of old women’s faces (Exp. 1a), broken down by whether participants were in conditions that emphasized the fit of age (left), gender (right), or neither age nor gender (middle) categories. Means are encompassed by 95% confidence intervals.

### Experiment 1B

Experiment 1b was launched concurrently with Experiment 1a. Participants in Experiment 1b saw 12 interaction partners (all young adults; half Black, half White, half men, and half women) having a disagreement that either comparatively “fit” the interaction partners’ racial groups (race-fit condition), gender groups (gender-fit condition), or neither their race nor gender groups (control condition). As in Experiment 1a, comparative fit was manipulated by changing the patterns of agreement and disagreement among the conversation partners.

Conversation topics were pre-tested to ensure that they did not normatively fit the lenses of race or gender. After the conversation concluded, participants reported on stereotypes toward one of

the Black women from the conversation they saw. At the end of the experiment, participants completed a memory task for who said what during the conversation from earlier.

Participants in the race-fit condition were expected to show evidence of viewing targets through the lens of *race*. As such, participants in the race-fit condition were expected to engage in higher levels of race categorization and racial stereotyping than participants in the other two conditions. Participants in the gender-fit condition were expected to show evidence of viewing targets through the lens of *gender*. As such, participants in the gender-fit condition were expected to engage in higher levels of gender categorization and gender stereotyping than participants in the other two conditions. In addition, these patterns of categorization and stereotyping were expected to *trade off*: perceivers were expected to exhibit less racial categorization and stereotyping when relying on the lens of *gender*, and they were expected to exhibit less gender categorization and stereotyping when relying on the lens of *race*.

**Method.** In Experiment 1b, participants viewed 3 Black women, 3 White women, 3 Black men, and 3 White men discussing one of two conversation topics in a 3 (fit: race, control, gender)  $\times$  2 (conversation topic) between-person design. Again, conversation topic was designed to serve as an internal replication factor. Experiment 1b, like Experiment 1a, had an a priori decision rule to include approximately 100 people per “fit” condition.

**Participants.** A total of 298 U.S. citizens were recruited to participate in Experiment 1b. Of these participants, the majority ( $n = 168$ ) were undergraduates at Northwestern University who participated in exchange for class credit; the remaining participants ( $n = 130$ ) were recruited from the MTurk website in exchange for \$3.00. Of these participants,  $n = 14$  (4.70%) were excluded for not responding “yes” to the question, “Did you take this study seriously?”

Remaining participants were mostly male (143 male, 134 female, 2 non-binary), mostly White (187 White, 16 Black, 53 Asian, 15 Latinx, 1 American Indian, 10 multi-racial, 2 non-represented), and had ages spanning from 18 to 70 ( $M = 25.81$ ,  $SD = 10.47$ ). Because this sample drew heavily on undergraduate students, Experiment 1b's sample was less educated than Experiment 1a's (only 21.13% held at least a bachelor's degree), and they skewed slightly more toward political liberalism ( $M = 3.72$ ,  $SD = 2.54$ , on the same 11-point scale as before).

**Procedure.** The procedure of Experiment 1b was identical to Experiment 1a's, however the stimulus materials were different. In Experiment 1b, participants learned that they would be watching a conversation either on the topic of whether assisted suicide should be illegal, or on the topic of whether celebrities earn too much money. Again, these topics were chosen because people could plausibly disagree over them, and critically, because they are rated as normatively neutral with respect to the gender and racial groups used in this experiment (according to pre-testing; see Appendix D).

As in Experiment 1a, participants were first shown a matrix of all the conversation partners whom they would be viewing (see Figure 6). This matrix featured 3 standardized headshots each of Black women, Black men, White women, and White men. Images were used with permission from the Chicago Face Database (Ma, Correll, & Wittenbrink, 2015). Each participant saw one of ten possible matrices—one of ten possible sets of 12 interaction partners. Thus, a total of 120 distinct conversation partners were used as stimuli in this study, a design feature that was implemented, again, to ensure that any observed effects would generalize across social targets.



*Figure 6.* One of ten possible image matrices of interaction partners (an internal replication factor) whom participants saw engage in debate in Experiment 1b.

The dynamics of how the conversation unfolded were structurally identical to those described in Experiment 1a. In the race-fit condition, all the Black conversation partners (collapsing across gender) took one position, and all the White conversation partners (collapsing across gender) took the other. In the gender-fit condition, all the women (collapsing across race) took one position, and all the men (collapsing across race) took the other. Stance on position (e.g., whether women *vs.* men thought celebrities earned too much money) was counterbalanced across participants. In the control condition, men and women of both racial groups took one position as well as the other.

After the conversation concluded, participants were told—as in Experiment 1a—that they were to report on how “the average American” would perceive some of the social targets from the conversation they viewed (Devine & Elliot, 1995; Ghavami & Peplau, 2013). Participants

were then shown one of the Black women from the conversation they viewed, and they listed stereotypes about her and provided ratings of her facial features. Participants were expected to characterize the Black women as possessing “Blacker” traits, and as looking more typical of Black Americans, in the race-fit condition than in the other two conditions. By contrast, participants were expected to characterize the Black woman as possessing “more feminine” traits, and as looking more typical of women, in the gender-fit condition than in the other two conditions.

***Application of race and gender stereotypes.*** Participants viewed the same checklist of traits described in Experiment 1a, and they were asked to select all the attributes the average American might use when stereotyping a randomly-selected Black woman from their conversation. Unbeknownst to participants, checklist attributes had been rated on how stereotypically feminine they seem (see Experiment 1a), as well as on how “Black” they seem ( $n = 78$ ; as reported in Petsko & Bodenhausen, 2019b). Thus, participants’ stereotype nominations could be scaled numerically on how “feminine” and how “Black” they seem, respectively. As before, these ratings were made on a 7-point scale (from  $1 = not\ at\ all$  to  $7 = very\ feminine$  [Black]). And of note, the extent to which attributes were rated as “Black” was moderately correlated with the extent to which they were rated as feminine:  $r(97) = -.32, p = .001$  (see also Galinsky et al., 2013).

***Face ratings.*** As in Experiment 1a, participants provided face ratings of their randomly assigned social target. Specifically, they provided ratings of how “typical of Black Americans”

and how “typical of women” their target appeared. Ratings were made on a scale from *1 = not at all* to *9 = very much*.<sup>5</sup>

**Who-said-what task.** Finally, participants completed a memory task for the conversation itself. In a randomized order, participants were shown each of the 36 statements that were spoken during the conversation, intermixed with 36 distractor statements. Thus, participants were shown a total of 72 statements, one by one, and their task was to first answer the question, “Have you seen this statement before?”, and if they answered yes, to then indicate who said it.

**Results.** As in Experiment 1a, the particular conversation topic to which participants were assigned, as well as the particular set of conversation partners participants had been assigned (that is, which set of 12), did not meaningfully influence the interpretation of results. Thus, for simplicity, I again present analyses while collapsing across conversation topic (but see Appendix E) and while collapsing across the random effects of stimuli. Notably, random effects of stimuli accounted for very little variance across outcome measures (8.84% on face rating measures, less than 0.10% on all other measures).

**Who-said-what: Race categorization.** Participants were hypothesized to be categorizing targets by race more in the race-fit condition—when the lens of *race* is active—than in the other two conditions. To analyze this, participants’ memory errors were categorized as either within- or between-racial-group errors, and they were subjected to a 2(error type: within-category, between-category) × 3 (condition: race fit, control gender fit) mixed ANOVA with repeated measures on the first factor. This analysis revealed a main effect of error type, suggesting that

---

<sup>5</sup> Participants also made face ratings of an individual Black man, and of an individual White woman from the conversation they viewed. In addition, faces were rated on how happy, angry, dominant, and trustworthy they appeared.

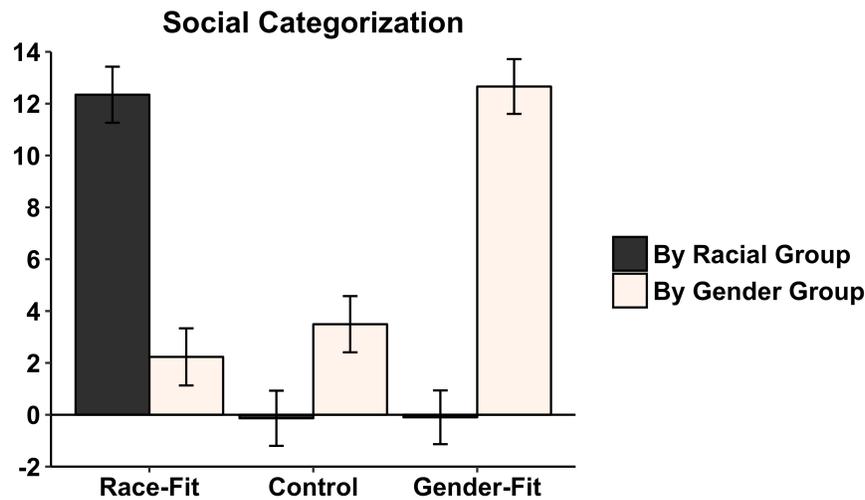


Figure 7. Within- relative to between-group errors (Exp. 1b), broken down by whether participants were in conditions that emphasized the fit of race (left), gender (right), or neither race nor gender (middle). Mean error differences are encompassed by 95% confidence intervals.

participants were indeed using race as a category for organizing social information. That is, when participants could not remember who said what, they made significantly more within-racial-group errors ( $M = 10.62$ ,  $SE = 0.22$ ) than between-racial group errors ( $M = 6.58$ ,  $SE = 0.22$ ),  $M_{diff} = 4.04$ , 95% CI[3.43, 4.65],  $\beta = 0.79$ ,  $F(1, 568) = 167.03$ ,  $p < .001$ . In line with ICT, the degree of race categorization was heavily moderated by whether or not participants were in the condition that comparatively fit targets' race:  $F(1, 568) = 346.78$ ,  $p < .001$ ,  $\omega_p^2 = 0.38$ . The nature of this interaction was that race categorization was substantially greater in the condition that emphasized race [ $M_{diff} = 12.35$ , 95% CI[11.26, 13.43],  $\beta = 2.41$ ,  $F(1, 568) = 500.91$ ,  $p < .001$ ] than it was in the other two conditions [ $M_{diff} = -0.12$ , 95% CI[-0.86, 0.63],  $\beta = -0.02$ ,  $F(1, 568) = 0.09$ ,  $p = .76$ ]. In addition, the degree of race categorization was not moderated by whether participants were in the control condition or in the condition that comparatively emphasized targets' gender (see Figure 7),  $F(1, 568) < 0.01$ ,  $p = .96$ ,  $\omega_p^2 < .01$ . This is highly consistent with ICT: race categorization was not being used at all when participants were in the

control condition or in the gender-fit condition. Notably, this is also consistent the perspective that race, specifically, is not an obligatory feature of targets to which perceivers attend (e.g., Kurzban et al., 2001). Indeed, in both the control condition and gender-fit condition (in which who race is non-diagnostic of who is in allegiance with whom), categorization by race was practically zero.

***Who-said-what: Gender stereotyping.*** Participants were expected to be engaging in gender categorization more in the gender-fit condition—when the lens of *gender* is active—than in the other two conditions. Subjecting participants’ gender-related errors to the same  $2 \times 3$  analysis described above revealed a main effect of error type: when participants could not remember who said what, they were more likely to erroneously attribute a statement to someone of the same gender category ( $M = 11.76$ ,  $SE = 0.23$ ) than to someone of the opposite gender category ( $M = 5.63$ ,  $SE = 0.23$ ),  $M_{\text{diff}} = 6.13$ , 95% CI[5.51, 6.75],  $\beta = 1.12$ ,  $F(1, 568) = 370.83$ ,  $p < .001$ . Furthermore, and as expected, the degree of gender categorization was heavily moderated by whether participants were in the condition that comparatively emphasized targets’ gender (see Figure 7):  $F(1, 568) = 215.47$ ,  $p < .001$ ,  $\omega_p^2 = 0.27$ . When participants were in the gender-fit condition, they became substantially more likely to make within- relative to between-gender errors:  $M_{\text{diff}} = 12.66$ , 95% CI[11.61, 13.72],  $\beta = 2.32$ ,  $F(1, 568) = 552.34$ ,  $p < .001$ . Participants in the other two conditions exhibited the same pattern, but to a much weaker degree:  $M_{\text{diff}} = 2.86$ , 95% CI[2.09, 3.63],  $\beta = 0.52$ ,  $F(1, 568) = 52.75$ ,  $p < .001$ . The degree of gender categorization was not moderated by whether participants were in the race-fit vs. control conditions,  $F(1, 568) = 2.56$ ,  $p = .11$ ,  $\omega_p^2 < .01$ . These findings are consistent with the idea that categorizing targets by gender is more obligatory than categorizing people by race (Tooby & Cosmides, 2010).

However, these findings are also demonstrative that categorization by gender can be substantially attenuated when perceivers are attending to another social identity that targets harbor. In this way, the data are consistent with the idea that the perception of gender, despite its primacy, can be highly compartmentalized—just as the perception of other identities can be.

**Checklist attributes.** Did the stereotypes that perceivers nominate for individual Black women vary by comparative fit? For illustrative purposes, the 15 most-selected traits from each condition (race fit, control, gender fit) are listed in Table 2. As in Experiment 1a, these trait nominations are, descriptively speaking, highly similar across conditions.

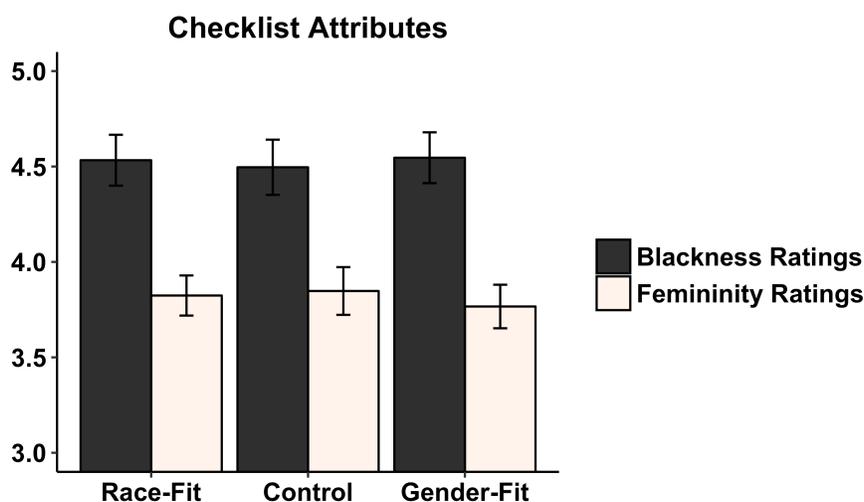
Table 2.

*Fifteen Most-Selected Traits for Black Women (Exp. 1b), Broken Down by Condition*

<b>Race (<i>n</i> = 91)</b>		<b>Control (<i>n</i> = 94)</b>		<b>Gender (<i>n</i> = 99)</b>	
<b>Trait</b>	<b>Prop.</b>	<b>Trait</b>	<b>Prop.</b>	<b>Trait</b>	<b>Prop.</b>
Loud	0.51	Argumentative	0.56	Loud	0.53
Poor	0.49	Loud	0.52	Aggressive	0.49
Argumentative	0.48	Poor	0.52	Argumentative	0.48
Uneducated	0.46	Uneducated	0.50	Stubborn	0.47
Low in Intelligence	0.45	Stubborn	0.44	Quick-tempered	0.45
Stubborn	0.43	Aggressive	0.40	Rude	0.45
Aggressive	0.42	Low in Intelligence	0.38	Low in Intelligence	0.40
Quick-tempered	0.42	Quick-tempered	0.38	Uneducated	0.40
Talkative	0.41	Rude	0.38	Talkative	0.38
Rude	0.37	Talkative	0.37	Quarrelsome	0.35
Quarrelsome	0.35	Lazy	0.36	Hostile	0.34
Ignorant	0.32	Quarrelsome	0.35	Ignorant	0.34
Lazy	0.32	Ignorant	0.34	Poor	0.33
Hostile	0.30	<b>Impulsive</b>	0.29	<b>Persistent</b>	0.29
<b>Unreliable</b>	0.30	<b>Straightforward</b>	0.29	<b>Arrogant</b>	0.29

*Note.* Prop. = the proportion of participants in a condition who chose a given trait to characterize the stereotypes that would be attributed to a Black woman from the conversation they viewed. Bold font indicates that a trait appears in the top 15 traits of only one condition

**Checklist attributes: Stereotypic blackness.** ICT predicts that the traits participants use to characterize Black women should be more stereotypically Black in the race-fit conditions than in the other conditions. However, subjecting the average “Blackness” ratings of participants’ trait nominations to a one-way ANOVA yielded null results. Contrary to hypotheses, participants characterized Black women as no more stereotypically Black in the race-fit condition ( $M = 4.53$ ,  $SE = 0.07$ ) than in the control or gender-fit conditions ( $M = 4.52$ ,  $SE = 0.04$ ),  $M_{diff} = 0.01$ , 95% CI[-0.16, 0.18],  $\beta = 0.02$ ,  $F(1, 280) = 0.02$ ,  $p = .89$ . In addition, Black women were characterized by traits that were equally “Black” across these latter conditions,  $M_{diff} = -0.05$ , 95% CI[-0.24, 0.18],  $\beta = -0.07$ ,  $F(1, 280) = 0.27$ ,  $p = .61$  (see Figure 8).



*Figure 8.* How “Black” and “feminine” trait nominations for Black women were rated to be (Exp. 1b), as a function of which “fit” condition participants had been in: race (left), gender (right), or neither race nor gender (middle). Means are encompassed by 95% confidence intervals.

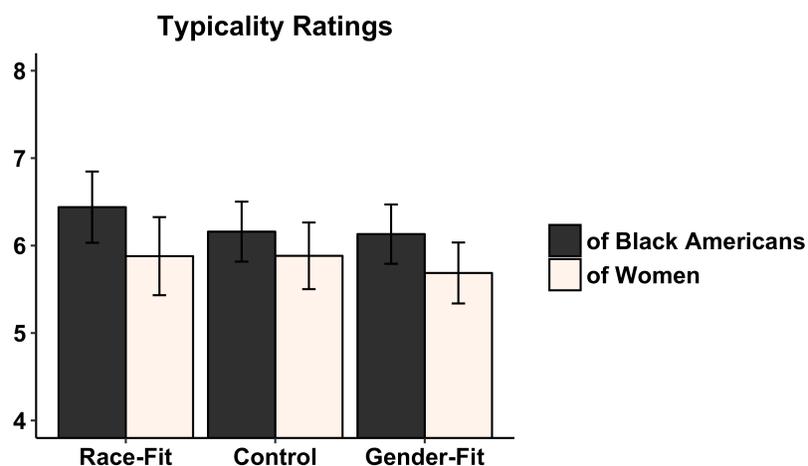
**Checklist attributes: Stereotypic femininity.** ICT predicts that the traits participants use to characterize Black women should be more stereotypically feminine in the gender-fit conditions than in the other conditions. However, subjecting the average “femininity” ratings of participants’ trait nominations to a one-way ANOVA also yielded null results. Contrary to

hypotheses, participants characterized Black women as no more stereotypically feminine in the gender-fit condition ( $M = 3.77$ ,  $SE = 0.06$ ) than in the control or race-fit conditions ( $M = 3.83$ ,  $SE = 0.04$ ),  $M_{diff} = -0.07$ , 95% CI[-0.21, 0.07],  $\beta = -0.12$ ,  $F(1, 280) = 0.97$ ,  $p = .33$ . In addition, Black women were characterized by traits that were equally “feminine” regardless of whether they were from the control condition or the race-fit condition,  $M_{diff} = 0.02$ , 95% CI[-0.14, 0.19],  $\beta = 0.04$ ,  $F(1, 280) = 0.08$ ,  $p = .77$  (see Figure 8).

**Face ratings: Typical of Black Americans.** I next investigated whether participants characterized their assigned Black women as looking “Blacker”—that is, more typical of Black Americans—in the race-fit condition than in the other two experimental conditions. A one-way ANOVA indicated that this was not the case: Black women were rated as looking no “Blacker” in the race-fit condition ( $M = 6.43$ ,  $SE = 0.19$ ) than in the control or gender-fit conditions ( $M = 6.15$ ,  $SE = 0.13$ ),  $M_{diff} = 0.29$ , 95% CI[-0.15, 0.74],  $\beta = 0.17$ ,  $F(1, 281) = 1.69$ ,  $p = .19$ . In addition, face ratings of “Blackness” did not vary across the control vs. age-fit conditions,  $M_{diff} = 0.03$ , 95% CI[-0.48, 0.53],  $\beta = 0.02$ ,  $F(1, 292) = 0.01$ ,  $p = .91$  (see Figure 9). Thus, and again in contrast to ICT, there was no evidence that Black women were rated as looking “more” or “less Black” depending on experimental condition.

**Face ratings: Typical of women.** Finally, a one-way ANOVA of how “typical of women” participants rated the Black women also yielded null results. Participants did not rate Black women as looking more typical of women in the gender-fit condition ( $M = 5.68$ ,  $SE = 0.19$ ) than in the other conditions ( $M = 5.88$ ,  $SE = 0.14$ ), [ $M_{diff} = -0.19$ , 95% CI[-0.66, 0.28],  $\beta = -0.10$ ,  $F(1, 281) = 0.66$ ,  $p = .42$ ]. Likewise, face ratings of Black women did not vary across the control vs. race-fit conditions [ $M_{diff} < 0.01$ , 95% CI[-0.55, 0.56],  $\beta < 0.01$ ,  $F(1, 281) < 0.01$ ,  $p =$

.99] (see Figure 9). Thus, as in Experiment 1a, all measures of stereotype application to individual targets failed to support ICT. Only indices of social categorization supported ICT.



*Figure 9.* Typicality ratings of Black women's faces (Exp. 1b), broken down by whether participants were in conditions that emphasized the fit of race (left), gender (right), or neither race nor gender (middle) categories. Means are encompassed by 95% confidence intervals.

## Discussion

Experiments 1a and 1b provide initial support for the idea that perceivers' minds select lenses for viewing others in a compartmentalized way. When participants were in a social context that comparatively fit the lens of *gender*, for example, they came to remember women as interchangeable with other women, and men as interchangeable with other men, but they ceased to view these targets as being interchangeable with other members of their age groups or racial groups (Figures 3 and 7, respectively). This implies that when participants were using *gender* as a lens for making sense of targets, they were no longer using alternative lenses for making sense of targets. On top of this, these categorization findings contradict a notable psychological claim (e.g., Kurzban et al., 2001; Pietraszewski et al., 2015): namely, the claim that perceivers cannot help but to engage in high levels of gender- and age-based categorization, *even* when social contexts comparatively fit alternative social identities. Of note, the paradigm used in

Experiments 1a and 1b closely parallels the paradigm used in the literature that advances this view (e.g., see Kurzban et al., 2001).

A big issue for Experiments 1a and 1b is that they do not provide evidence that participants *apply* stereotypes to individual targets in a compartmentalized way. These null findings may be attributable to the fact that participants were explicitly asked to refrain from reporting on their personal views, and to instead call to mind what “the average American” thinks about intersectional targets. These instructions were included so that participants could candidly report on negative stereotypes toward targets while eschewing personal responsibility for them (Devine & Elliot, 1995; Ghavami & Peplau, 2013). However, these instructions may have also caused participants to call a person other than themselves to mind—potentially, a person who had not viewed these targets in the same social context that they themselves had just viewed them in. As a separate issue, participants may not have felt as though they were in a position to actually judge individual targets. For example, if participants could not remember the contributions of individual older women or Black women from the conversation they saw, they could have easily convinced themselves that they were unfit for the task of characterizing these individuals’ attributes. Social judgeability can play a pronounced role in moderating whether or not activated stereotypes are expressed toward individual targets (Leyens, Yzerbyt, & Shadron, 1992; Schadron & Yzerbyt, 1991). Finally, it is possible that the stereotype application measures showed little variability by condition because these measures were all self-report measures. Responses to self-report measures tend to be highly controllable. If controllability is to account for the non-application of stereotypes, then adding speeded response measures—which are

relatively more difficult for participants to control—may unveil the patterns of lens-based stereotyping that ICT proposes.

## Experiment 2

Experiment 2 was designed to be a replication study of Experiment 1a. As such, it also indexed whether perceivers use one lens at a time for making sense of social targets (specifically, older women). However, Experiment 2 contained several design features that made it an improvement over Experiments 1a and 1b. For one, Experiment 2 included a speeded response measure that was designed to index participants' tendency to associate *age* and *gender*-related stereotypes with individual older women (Nosek, Bar-Anan, Sriram, Axt, & Greenwald, 2014). Second, Experiment 2 contained modified instructions that were a) designed to enhance participants' perception that they were in a position to judge individual older women, and that b) asked them to report on what they personally believed about these women (rather than what “the average American” believes).

### Method

In Experiment 2, participants viewed 3 older women (ages 70-80), 3 young women (ages 20-30), 3 older men (ages 70-80), and 3 young men (ages 20-30) discussing the topic of whether introverts are more intelligent than extraverts. In the service of economy, Experiment 2 included only an age-fit and gender-fit condition (it did not include a control condition as in Experiment 1a). Thus, Experiment 2 had a between-person experimental design with one factor: comparative fit (age, gender).

**Participants.** A total of  $N = 220$  undergraduates from Northwestern University participated in Experiment 2 in exchange for course credit. Of those,  $n = 12$  (5.45%) were

excluded prior to analyses for failing to reply “yes” to the question, “Did you take this study seriously?” The remaining participants were mostly female (139 female, 69 male), mostly White (95 White, 51 Asian, 26 Black, 21 Latinx, 15 multiracial), and had ages spanning from 18 to 22 ( $M = 18.84$ ,  $SD = 0.88$ ). As was the case with previous samples, this sample skewed toward political liberalism ( $M = 3.07$ ,  $SD = 1.71$ , on an 11-point scale as before from  $0 = \textit{extremely liberal}$  to  $10 = \textit{extremely conservative}$ ).

**Procedure.** Participants in Experiment 2 completed a similar procedure to that described in Experiment 1a. As before, participants learned that they would watch a conversation that would unfold between 12 people. To amplify the salience of age or gender, respectively, participants in the age-fit condition were asked, “Which of the following age groups do you belong to,” while participants in the gender-fit condition were instead asked, “Which of the following gender groups do you belong to?” Responses options on these questions were categorical (e.g., *young adults*, *women*), so as to amplify the accessibility of the condition-relevant lens. As before, participants were assigned to see one of 9 possible sets of 12 interaction partners (an internal replication factor). After viewing the conversation—which provided comparative fit either to the lens of *gender* or *age*—participants saw the following instructions: “Now that the conversation is over, we’d like to gather some information on how you perceive some of the individuals who participated.” After these instructions, participants completed the following dependent measures in order:

***Stereotype application.*** Stereotype application in Experiment 2 was measured via a checklist measure (the same one used in Experiments 1a and 1b). Before completing this measure, participants were given instructions that were designed to increase their perceptions

that targets were social judgeable (Leyens et al., 1992; Schadron & Yzerbyt, 1991). In particular, participants in Experiment 2 were told:

Psychologists have recently started exploring a phenomenon known as ‘thin slicing.’ Thin slicing refers to the striking ability of perceivers (like you) to make accurate inferences about other people’s psychological attributes, *even* in situations where they have had only brief exposure to the person they are judging.

Participants were then given some additional examples to drive home this point:

For example, perceivers are able to accurately judge others’ personalities (Ambady, 2010), religious affiliations (Rule et al., 2010), and even their earning potentials after just a few seconds of exposure to their faces or behaviors (Livingston & Pearce, 2009).

Finally, participants were shown the face of a randomly selected older woman from the conversation they viewed. Participants were then asked the question, “What personality traits would you use to describe this person, if you had to infer what this person is like?” Below this question was the list of 99 checklist traits (taken from Petsko & Bodenhausen, 2019b), which participants could use for reporting on their impression of the older woman.

**Face ratings.** After reporting on their impressions, participants were then shown the same older woman’s face again, and they were asked how “psychologically typical of women” they think she might be (based on how her face looks), and how “psychologically typical of old people” they think she might be (based on how her face looks). These ratings were made on a scale from *1 = not at all* to *9 = very much*, and they were embedded among other psychological attributes (e.g., how trustworthy she is, how angry she is). In addition, participants rated the faces

of a randomly selected older man (from the conversation they were assigned), and of a randomly selected young woman on these same attributes.

**Brief IAT.** Participants then completed a brief-IAT that used the older women from the conversation that they had been assigned as its stimuli. This measure compared the speed with which participants associated these older women with age-related stereotypes (e.g., *elderly*, *frail*, *cranky*) vs. gender-related stereotypes (e.g., *maternal*, *feminine*, and *emotional*). Prior to beginning the brief-IAT, participants were shown the faces of all three older women from the conversation they saw, and they were told that these three women were “TARGET PEOPLE.” Participants will then complete four blocks of 20 speeded response trials. In two of the four blocks, participants had to hit the “I” key if they saw either a TARGET PERSON picture or an “OLD PEOPLE word.” In the other two blocks, participants were instructed to hit the “I” key if they saw either a TARGET PERSON picture or a “WOMEN word.”

**Who-said-what task.** Participants then completed the who-said-what task exactly as it was described in Experiment 1a. Each of the 36 statements from the conversation participants viewed, intermixed with 36 distractor statements, were presented to participants in a randomized order. Participants’ task was to first indicate whether they had seen a statement before, and if so, to indicate which of the 12 conversation partners said it.

**Social desirability.** Lastly, participants completed a measure of socially desirable responding: the 10-item short form of the Marlowe-Crowne social desirability scale (Reynolds, 1982: e.g., “I never resent being asked to return a favor.”). This measure was completed on a scale from 0 = *not at all true of me* to 4 = *very true of me*. These items were averaged into a composite index ( $\alpha = .66$ ).

## Results

**Who-said-what: Age categorization.** As in Experiment 1a, participants were expected to exhibit substantially more age categorization in the age-fit condition than in the gender-fit condition. To examine this possibility, memory errors from the who-said-what task were subjected to a 2(fit: age, gender)  $\times$  2 (error type: within-age-group, between-age-group) mixed ANOVA with repeated measures on the second factor. This analyses revealed, first, a main effect of error type, such that participants tended to make significantly more within-age-group errors ( $M = 12.78$ ,  $SE = 0.22$ ) than between-age-group errors ( $M = 5.16$ ,  $SE = 0.22$ ),  $M_{\text{diff}} = 7.62$ , 95% CI[7.02, 8.22],  $\beta = 1.29$ ,  $F(1, 416) = 620.01$ ,  $p < .001$ . This suggests that across conditions, participants were indeed more likely to remember the age groups of targets than not. However, there was also evidence that this tendency was heavily moderated by fit condition:  $F(1, 416) = 450.59$ ,  $p < .001$ ,  $\omega_p^2 = 0.52$ . The nature of this interaction was that within- (relative to between-) age-group errors were significantly more prevalent in the age-fit condition ( $M_{\text{diff}} = 14.12$ , 95% CI[13.26, 14.98],  $\beta = 2.41$ ,  $F(1, 416) = 1043.74$ ,  $p < .001$ ) than they were in the gender-fit condition ( $M_{\text{diff}} = 1.12$ , 95% CI[0.28, 1.96],  $\beta = 0.19$ ,  $F(1, 416) = 6.87$ ,  $p = .009$ ). This is to say participants in the age-fit condition exhibited excellent memory for the age groups of the people they viewed. In contrast, participants in the gender-fit condition exhibited substantially reduced memory for the age groups of the people they viewed (see Figure 10).

**Who-said-what: Gender categorization.** Participants were also expected to exhibit substantially greater gender categorization in the gender-fit condition than in the age-fit condition. To test this idea, data from the who-said-what task were subjected to a 2(fit: age, gender)  $\times$  2 (error type: within-gender-group, between-gender-group) mixed ANOVA with

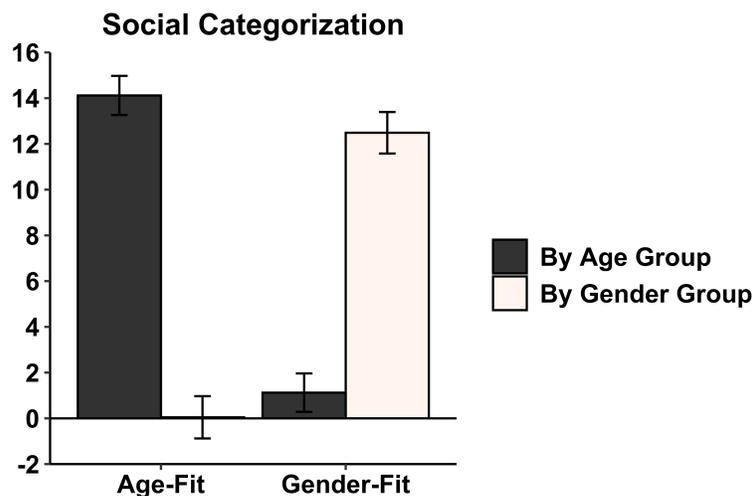


Figure 10. Within- relative to between-group errors (Exp. 2), broken down by whether participants were in conditions that emphasized the fit of age (left) or gender (right). Mean error differences are encompassed by 95% confidence intervals.

repeated measures on the second factor. This analysis revealed that when participants could not remember who said what, they tended to remember the gender of the person who spoke more often than not. That is, the frequency of within-gender-group errors ( $M = 12.05$ ,  $SE = 0.23$ ) significantly exceeded the number of between-gender group errors ( $M = 5.78$ ,  $SE = 0.23$ ):  $M_{diff} = 6.26$ , 95% CI[5.62, 6.91],  $\beta = 1.12$ ,  $F(1, 416) = 359.71$ ,  $p < .001$ . Moreover, and consistent with the predictions of ICT, participants' attention to gender (i.e., their memory for it) was significantly moderated by whether participants had been in the gender-fit vs. the age-fit condition:  $F(1, 416) = 354.49$ ,  $p < .001$ ,  $\omega_p^2 = .46$ . The nature of this interaction was that memory for targets' gender groups was substantially stronger in the gender-fit condition than in the age-fit condition. Whereas within- (vs. between-) gender-group errors were common in the gender-fit condition ( $M_{diff} = 12.49$ , 95% CI[11.58, 13.39],  $\beta = 2.22$ ,  $F(1, 416) = 728.19$ ,  $p < .001$ ), they were not common in the age-fit condition. In fact, participants in the age-fit condition were not more likely to make within- relative to between-gender group errors at all:  $M_{diff} = 0.04$ ,

95% CI[-0.88, 0.97],  $\beta = 0.01$ ,  $F(1, 416) = 0.01$ ,  $p = .92$ . This is to say that in the age-fit condition, participants *were* attending to targets' gender groups—but in these moments, there was no evidence that they were attending to targets' age groups (see Figure 10). This pattern of results is highly consistent with the theoretical perspective laid out by ICT.

**Checklist stereotypes.** The most-selected stereotypes that participants in Experiment 2 nominated in the age- and gender-lens conditions are listed in Table 3. Descriptively speaking, participants exhibited largely similar stereotype nominations regardless of which comparative fit condition they had been assigned. Indeed, only 4 attributes were unique to each condition.

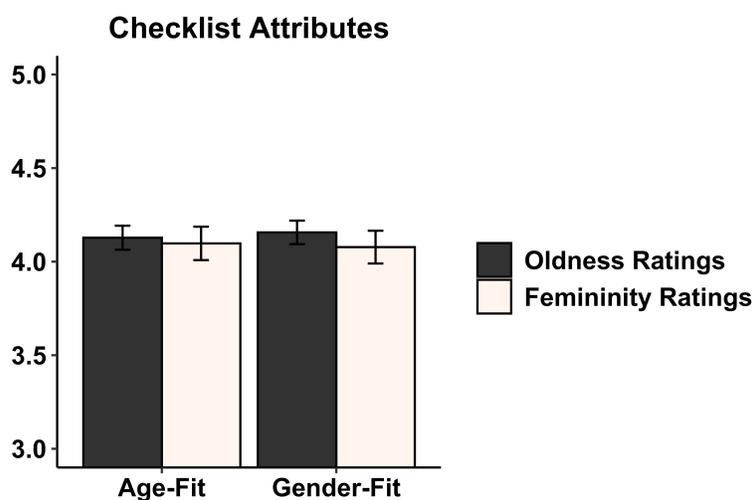
Table 3.

*Fifteen Most-Selected Traits for Old Women (Exp. 2), Broken Down by Experimental Condition*

Age ( $n = 108$ )		Gender ( $n = 112$ )	
Trait	Prop.	Trait	Prop.
Straightforward	0.53	Straightforward	0.57
Stubborn	0.52	Stubborn	0.50
Argumentative	0.51	Argumentative	0.48
Conservative	0.48	Conservative	0.40
Persistent	0.43	Honest	0.37
Intelligent	0.38	Persistent	0.36
Tradition-loving	0.38	Practical	0.36
Practical	0.31	Tradition-loving	0.35
Honest	0.31	<b>Conventional</b>	0.33
<b>Individualistic</b>	0.31	Intelligent	0.33
<b>Arrogant</b>	0.30	<b>Loyal to family ties</b>	0.31
Quick-tempered	0.27	Rude	0.28
<b>Quarrelsome</b>	0.26	<b>Humorless</b>	0.27
Rude	0.24	Quick-tempered	0.27
<b>Ignorant</b>	0.23	<b>Aggressive</b>	0.26

*Note.* Prop. = the proportion of participants in a condition who chose a given trait to characterize the stereotypes that would be attributed to an old woman from the conversation they viewed. Bold font indicates that a trait appears in the top 15 traits of only one condition.

**Checklist stereotypes: “Oldness”.** According to ICT, the words participants select from the checklist measure to characterize older women should be rated as stereotypically older in the age-fit condition than in the gender-fit condition. To investigate whether this was indeed the case, ratings of stereotypic oldness of participants’ checklist nominations were subjected to a one-way ANCOVA (fit: age, gender) that controlled for participants’ social desirability concerns. In contrast to the predictions of ICT, participants did not list attributes for older women that were rated as any “older” in the age-fit condition ( $M = 4.13$ ,  $SD = 0.32$ ) than in the gender-fit condition ( $M = 4.16$ ,  $SD = 0.34$ ),  $M_{diff} = -0.03$ , 95% CI[-0.12, 0.06],  $\beta = -0.09$ ,  $F(1, 205) = 0.38$ ,  $p = .54$ . Thus, as in Experiment 1a, there was no evidence that participants applied “older” stereotypes to old women when placed in a condition that comparatively emphasized her age (vs. her gender; see Figure 11).



*Figure 11.* How “old” and “feminine” trait nominations for older women were rated to be (Exp. 2), broken down by whether participants were in conditions that emphasized the fit of age (left), or gender (right). Means are encompassed by 95% confidence intervals.

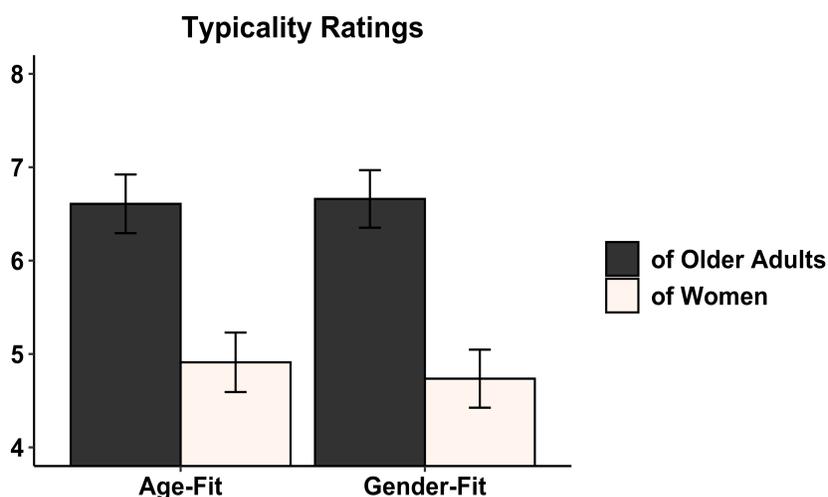
**Checklist stereotypes: Femininity.** According to ICT, participants’ impressions of older women—that is, the words they use to characterize these women—should also be rated as “more

feminine” in the gender-fit condition than in the age-fit condition. To examine whether this was the case, the average stereotypic femininity of participants’ checklist nominations were, as above, subjected to a one-way ANCOVA (fit: age, gender) that controlled for participants’ social desirability concerns. This analyses revealed that participants did not choose “more feminine” attributes for older women in the gender-fit condition ( $M = 4.08$ ,  $SD = 0.47$ ) than in the age-fit condition ( $M = 4.10$ ,  $SD = 0.45$ ):  $M_{\text{diff}} = 0.02$ , 95% CI[-0.11, 0.14],  $\beta = 0.04$ ,  $F(1, 205) = 0.08$ ,  $p = .77$ . Thus, there was no evidence that participants applied “more feminine” attributes to older women in the gender-fit condition as compared with the age-fit condition (see Figure 11)—even in this experimental paradigm, which was designed to be an improvement over the paradigm used in Experiments 1a and 1b.

**Face ratings: Typical of older adults.** According to ICT, participants should rate the faces of older women as seeming more “typical of older adults” in the age-fit condition relative to the gender-fit condition. To investigate whether this hypothesis was supported, typicality ratings were subjected to a one-way ANCOVA (fit: age, gender) that controlled for participants’ social desirability levels. In contrast to the predicted pattern, older women’s faces were not rated as seeming any “more typical of older adults” in the age-fit condition ( $M = 6.61$ ,  $SD = 1.84$ ) than in the gender-fit condition ( $M = 6.66$ ,  $SD = 1.37$ ),  $M_{\text{diff}} = -0.05$ , 95% CI[-0.49, 0.39],  $\beta = -0.03$ ,  $F(1, 206) = 0.05$ ,  $p = .82$  (see Figure 12).

In addition, ICT proposes that older women should be perceptually assimilated to old men in the age-fit condition, and perceptually contrasted away from young women in this condition as well. To examine this possibility, ratings of how “typical of older adults” the faces seemed was subjected to a 2 (fit: age, gender)  $\times$  3 (target face: older woman, older man, young

woman) mixed ANCOVA with repeated measures on the final factor. The covariate in this analysis was participants' social desirability levels. In contrast to the predictions of ICT, the extent to which older women were rated similarly to older men [ $M_{\text{diff}} = -0.02$ , 95% CI[-0.34, 0.30],  $\beta = -0.01$ ,  $F(1, 621) = 0.02$ ,  $p = .88$ ] was not moderated by whether participants were in the age-fit condition vs. the gender-fit condition:  $F(1, 621) = 1.85$ ,  $p = .18$ ,  $\omega_p^2 < .01$ . In addition, differences in how older people vs. young women were rated [ $M_{\text{diff}} = 2.96$ , 95% CI[2.69, 3.24],  $\beta = 1.36$ ,  $F(1, 623) = 438.90$ ,  $p < .001$ ] did not increase in the age-fit condition vs. the gender fit condition:  $F(1, 621) = 2.48$ ,  $p = .12$ ,  $\omega_p^2 < .01$ . In short, there was no evidence that older women were perceptually assimilated to other older people or that older people were perceptually contrasted away from young people in the age-fit (vs. gender-fit) condition.



*Figure 12.* Typicality ratings of old women's faces (Exp. 2), broken down by whether participants were in conditions that emphasized the fit of age (left), or gender (right) categories. Means are encompassed by 95% confidence intervals.

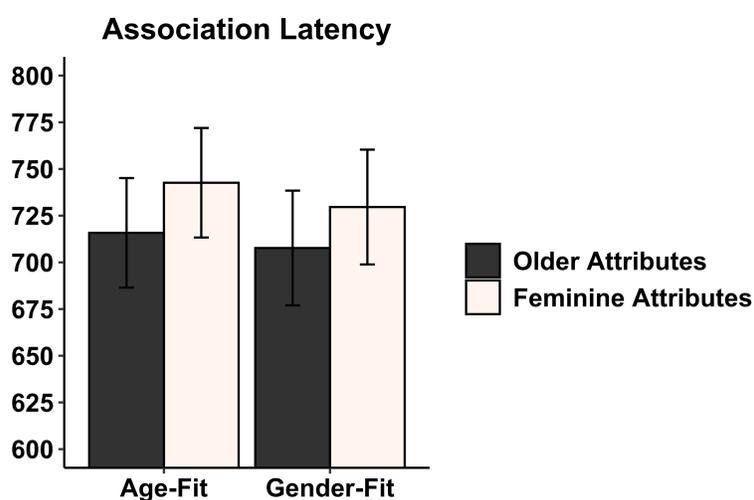
**Face ratings: Typical of women.** According to ICT, participants should rate the faces of older women as being more typical of women in the gender-fit condition relative to the age-fit condition. To investigate this possibility, typicality scores were subjected to a one-way

ANCOVA (fit: age, gender) that controlled for participants' social desirability levels. This analysis revealed that in contrast to predictions, older women were not rated as seeming more “typical of women” in the gender-fit condition ( $M = 4.74$ ,  $SD = 1.54$ ) than in the age-fit condition ( $M = 4.91$ ,  $SD = 1.73$ ):  $M_{\text{diff}} = -0.02$ , 95% CI[-0.34, 0.30],  $\beta = -0.01$ ,  $F(1, 621) = 0.02$ ,  $p = .88$  (see Figure 12).

In addition, ICT proposes that older women should be perceptually assimilated to young women in the gender-fit condition, and perceptually contrasted away from older men in the gender-fit condition as well. To examine this possibility, ratings of how “typical of women” the faces seem was subjected to a 2 (fit: age, gender)  $\times$  3 (target face: older woman, older man, young woman) mixed ANCOVA with repeated measures on the final factor. The covariate in this analysis was participants' social desirability levels. In contrast to predictions, perceptual differences between older women and young women [ $M_{\text{diff}} = -1.37$ , 95% CI[-1.92, -0.82],  $\beta = -0.67$ ,  $F(1, 412) = 24.09$ ,  $p < .001$ ] were not in any way moderated by whether participants were in gender-fit vs. age-fit condition:  $F(1, 412) = 0.60$ ,  $p = .44$ ,  $\omega_p^2 < .01$ . This is to say that older women were not perceptually assimilated to young women in the gender-fit condition. Moreover, there was no evidence that differences between ratings of women targets and older men [ $M_{\text{diff}} = 3.69$ , 95% CI[3.22, 4.16],  $\beta = 1.81$ ,  $F(1, 411) = 233.42$ ,  $p < .001$ ] became starker in the gender-fit condition than in the age-fit condition:  $F(1, 411) = 0.12$ ,  $p = .72$ ,  $\omega_p^2 < .01$ . Thus, there was no evidence that older women were perceptually assimilated to other women—or that these women were perceptually contrasted away from older men—in the gender-fit (vs. age-fit) condition.

**Brief IAT.** ICT proposes that when social contexts comparatively fit the lens of *gender*, participants should stereotypically associate older women with traits related to these women's

gender groups to a faster degree than with traits related to these women’s age groups. In contrast, when social contexts comparatively fit the lens of *age*, participants should do the reverse. To test this prediction, participants’ response latencies on the brief-IAT were subjected to a 2(fit: age, gender)  $\times$  2 (association type: older women + gender, older women + age) mixed ANOVA with repeated measures on the second factor. This analyses revealed only a main effect of association type, such that older women were associated more quickly with “old-people” concepts than with “women” concepts:  $M_{diff} = -77.15\text{ms}$ , 95% CI[-104.52, -49.78],  $\beta = -0.15$ ,  $F(1, 4888) = 30.53$ ,  $p < .001$ . In contrast to predictions, the speed with which older women were associated with “old people” vs. “women” concepts was not moderated by whether participants had been in the condition that comparatively emphasized these women’s age groups vs. these women’s gender groups:  $F(1, 4888) = 0.09$ ,  $p = .77$ ,  $\omega_p^2 < .01$ . Thus, the hypothesis that comparative fit manipulations would influence the traits that participants spontaneously associate with older women from the conversation they had seen was not supported (see Figure 13).



*Figure 13.* Speed (in milliseconds; Exp 2) with which participants in associated older women from the conversation they had seen with either “old people” attributes or “women” attributes on a brief-IAT. Results are broken down by which comparative fit condition participants had encountered prior to the brief-IAT. Means are encompassed by 95% confidence intervals.

## Discussion

In Experiments 1a and 1b, there was strong evidence that participants attended to social identities in a one-at-a-time fashion. That is, participants' memory errors in these experiments showed clear evidence that memory for one social identity (e.g., race), tended to trade off with memory for a separate social identity (e.g., gender). Yet these experiments revealed no evidence that which identity participants were attending to shaped how they stereotyped social targets at a later point in time. The goal of Experiment 2 was to examine whether changing the dependent variables—by asking participants to report on what they personally believe about targets (vs. what society more generally would believe about targets; Devine & Elliot, 1995), as well as by including a reaction time measure (the brief-IAT)—might reveal a tendency for participants to apply stereotypes to targets in a manner that was consistent with their lens use. However, Experiment 2 revealed more or less the same pattern of findings as Experiments 1a and 1b. That is, while the participants in Experiment 2 *did* exhibit a tendency to pay attention to a lens-relevant identity more than to a lens-irrelevant identity, this lens usage did not translate into a tendency for participants to stereotype targets in lens-specific ways. This issue will be elaborated upon in the general discussion.

### Chapter III: Initial Evidence of Lens-Dependent Stereotyping

### Experiment 3

One reason why participants in Experiments 1a-2 may have refrained from stereotyping targets in lens-specific ways is that which lens perceivers use may be highly temporally flexible. Thus, it may have been the case in these experiments that participants *were* using the context-afforded lens during the conversation that unfolded (i.e., when the comparative fit manipulation was present), but may have stopped using the context-afforded lens after the conversation was over (that is, when there was no longer a comparative fit manipulation that invited participants to use a particular lens over alternatives). Experiment 3 was designed to address this possibility by ensuring that participants' stereotype usage was measured *while* the social context was inviting them to think of targets through the vantage of one lens vs. another.

In Experiment 3, participants completed a variant of the IAT (Greenwald, Poehlman, Uhlmann, & Banaji, 2009) that assessed their tendency to stereotypically associate science concepts (e.g., Physics, Chemistry) with male targets more readily than with female targets. This bias (often termed a 'male-science' bias) has been documented in every nation in which the IAT has been administered, and it has been implicated in several consequential phenomena (for a recent meta-analysis, see Miller, Eagly, & Linn, 2014). As an example, nation-level variation in male-science biases predicts the magnitude of that nation's gender gap in math achievement between its 8<sup>th</sup> grade boys and girls (Nosek et al., 2009). The aim of Experiment 3 was to examine whether a male-science bias would emerge when the social context invited participants to view targets through the lens of *gender*, but whether this bias would disappear completely when the social context instead invited participants to view targets through the lens of *age*.

#### Method

In Experiment 3, participants completed two IATs in a within-person experimental design. One of the IATs was the traditional gender-science IAT, wherein participants had to sort faces by their gender groups and words by whether they were science-related or liberal-arts-related. The other IAT featured the exact same stimulus face and words, but in this variant, participants sorted faces by their *age* groups rather than by their gender groups. Across both IATs, the stimulus faces featured old and young men and women. The experimental design, hypotheses, and analyses plan for Experiment 3 were pre-registered.

**Participants.** A total of  $N = 114$  participants completed Experiment 3, of whom  $n = 21$  (18.42%) were excluded either for: a) failing to respond “yes” to the question, “Did you take this study seriously?”; b) reporting a technical issues that prevented the IAT from functioning properly; or c) having a mean response latency that was more than three median absolute deviations from the median of participants’ average response latencies (Leys, Ley, Klein, Bernard, & Licata, 2013). All three of these exclusion criteria were pre-registered. The 93 participants who remained were mostly male (54 male, 37 female, 2 non-specified), mostly White (70 White, 6 Black, 7 Latinx, 7 Asian, 2 American Indian, 1 multiracial), and the majority of them (51.6%) held a bachelor’s degree or higher. Their ages spanned from 21 to 71 ( $M_{\text{age}} = 36.22$ ,  $SD = 10.22$ ), and they skewed slightly toward political liberalism ( $M = 3.47$ ,  $SD = 3.00$ , on the same 11-point scale as in previous experiments).

**Procedure.** Participants in Experiment 3 were randomly assigned either to take the gender-science IAT prior to the age-science IAT, or to take the age-science IAT prior to the gender-science IAT. Both IATs featured the same stimulus faces and stimulus words. The stimulus faces included 6 older men, 6 young men, 6 older women, and 6 young women (whose

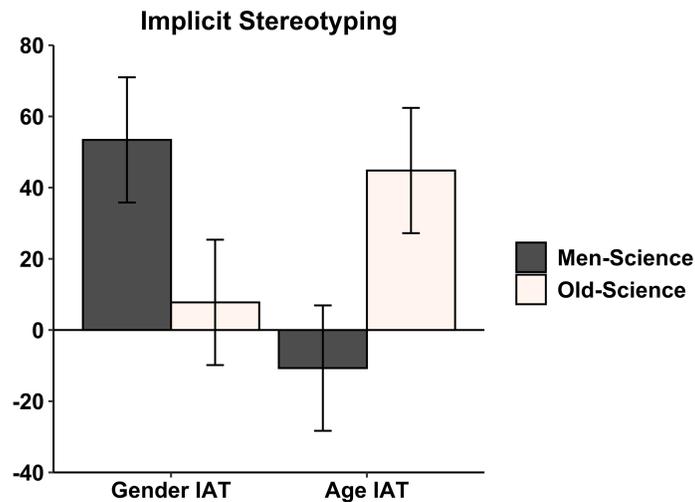
faces were drawn from the FACES database: Ebner et al., 2010). The words were either science-related (e.g., Physics, Chemistry, Engineering) or liberal-arts-related (e.g., History, English, Humanities). In the gender-science variant of this IAT, participants completed blocks in which they either associated male faces with science concepts—and female faces with liberal arts concepts—or in which they completed the reverse of these associations (that is, male faces with liberal arts concepts and female faces with science concepts). All participants completed 60 trials of each pairing type. In the age-science variation of this IAT, the stimuli were exactly the same, however the social category of focus was age groups (old, young) rather than gender groups. On this version of the IAT, participants completed blocks in which they either associated older faces with science concepts—and young faces with liberal arts concepts—or in which they completed the reverse of these associations (that is, old faces with liberal arts concepts and young faces with science concepts). Again, participants completed 60 trials of each pairing type. Within all IATs, block ordering was randomized, as was which key participants had to press for each of these respective pairing types.

**Results.** Response latencies (in milliseconds) for facial stimuli were regressed onto a mixed linear model that was analogous to a 2 (IAT type: gender-science; age-science)  $\times$  2 (gender pairing: male-science + female-arts; male-arts + female-science)  $\times$  2 (age pairing: old-science + young-arts; old-arts + young-science) repeated measures ANOVA. This model included estimates of three random effects: a random effect of IAT block intercept, which adjusts for any variation in mean response latency that would be attributable to the fact that some blocks come earlier in the experiment than others; a random effect of participant intercept, which adjusts for the fact that the full factorial design of this experiment was nested within person; and

a random effect of stimulus intercept, which adjusts for the fact that observations were also nested within particular stimulus faces. This multilevel modeling approach is preferable to the computation of *D*-scores because *D*-scores do not adjust for random variation in stimuli (which, when not adjusted for, increases the likelihood of making Type-I errors: Judd, Westfall, & Kenny, 2012; Skinner & Rae, 2019).

The main hypotheses in Experiment 3 were that a) participants would exhibit a tendency to associate men with science (and women with liberal arts) more quickly than the reverse; but that b) this bias would be eliminated when participants were categorizing targets by *age* rather than by gender. Supporting the first of these hypotheses, the 2 (IAT Type)  $\times$  2 (gender pairing)  $\times$  2 (age pairing) analysis described above revealed a main effect of gender pairing. The nature of this main effect was that participants were indeed faster to associate men with science (and women with liberal arts:  $M = 747.40$ ,  $SE = 14.47$ ) than the reverse ( $M = 768.75$ ,  $SE = 14.48$ ),  $M_{\text{diff}} = -21.36\text{ms}$ , 95% CI[-33.80, -8.91],  $\beta = -0.06$ ,  $F(1, 11050) = 11.31$ ,  $p < .001$ . Thus, participants did indeed exhibit a tendency to implicitly associate men (more than women) with science concepts. Moreover, and in support of the second hypothesis, this tendency was moderated by IAT type:  $F(1, 11054) = 25.47$ ,  $p < .001$ ,  $\omega_p^2 < .01$ . The nature of this interaction was that the tendency to implicitly associate men (more than women) with science was present when faces were being categorized by their gender groups [ $M_{\text{diff}} = -53.41\text{ms}$ , 95% CI[-71.00, -35.82],  $\beta = -0.15$ ,  $F(1, 11050) = 35.41$ ,  $p < .001$ ], but *not* when faces were being categorized by their age groups [ $M_{\text{diff}} = 10.70$ , 95% CI[-6.92, 28.31],  $\beta = 0.03$ ,  $F(1, 11050) = 1.42$ ,  $p = .23$ ]. Thus, this experiment yielded very strong support for ICT. When participants were given the processing goal of attending to the lens of *gender*, they stereotypically associated men more than

women with science. Yet when participants were instead given the processing goal of attending to the lens of *age*, they ceased to exhibit any gender stereotyping at all (see Figure 14).



*Figure 14.* Average male-science (dark gray) and old-science (light pink) implicit associations (Exp. 3), broken down by whether participants were completing a gender-lens IAT (left) or an age-lens IAT (right). Higher scores indicate stronger implicit associations (in milliseconds), and means are encompassed by 95% confidence intervals.

Although there were no pre-registered predictions related to age stereotyping, the 2 (IAT Type)  $\times$  2 (gender pairing)  $\times$  2 (age pairing) analysis described above also revealed a main effect of age pairing, such that participants tended to associate older adults with science (and young adults with liberal arts:  $M = 744.93$ ,  $SE = 14.47$ ) more quickly than the reverse of these pairings ( $M = 771.22$ ,  $SE = 14.47$ ),  $M_{\text{diff}} = -26.29\text{ms}$ , 95% CI[-38.73, -13.84],  $\beta = -0.07$ ,  $F(1, 11050) = 17.13$ ,  $p < .001$ . Interestingly, magnitude of this effect was also moderated by IAT type in a way that was highly sensible from the perspective of ICT:  $F(1, 10833) = 8.50$ ,  $p = .004$ ,  $\omega_p^2 < .01$ . That is, the tendency to associate older adults with science more quickly than young adults *only* emerged when participants were categorizing targets by age:  $M_{\text{diff}} = -44.81\text{ms}$ , 95% CI[-62.40, -27.21],  $\beta = -0.12$ ,  $F(1, 10275) = 24.21$ ,  $p < .001$ . When participants were categorizing targets by gender, this tendency, too, disappeared:  $M_{\text{diff}} = -7.76\text{ms}$ , 95% CI[-25.38, 9.85],  $\beta = -0.02$ ,

$F(1, 11050) = 0.77, p = .39$ . Thus, age stereotyping only emerged when participants were instructed to use the lens of *age*—it did not emerge when participants were instructed to use the lens of *gender*.

## **Discussion**

The findings of Experiment 3 go beyond those of Experiments 1a-2 in that these findings suggest that patterns of stereotyping—in addition to patterns of social categorization—can be highly compartmentalized. Specifically, the findings of Experiment 3 suggest that when participants are using the lens of gender to think about targets, they exhibit gender stereotyping against targets *but not age stereotyping*. Conversely, when participants are using the lens of age to think about targets, they exhibit age stereotyping *but not gender stereotyping*. Moreover, Experiment 3 suggests that the same target can face sweepingly different patterns of stereotyping depending on which lens is made salient to perceivers. Older women, in these data, are readily associated with science concepts when viewed through the lens of their age, but not when viewed through the lens of their gender.

Chapter IV: Evidence of Intersectional Lens Usage and Stereotyping

## Experiment 4

One well-known bias in the social perception literature is that when presented with an array of faces, people tend to recognize threatening facial expressions more quickly than non-threatening facial expressions. That is, perceivers are faster to discriminate angry faces from a set of neutral distractor faces than they are happy faces from a set of neutral distractor faces (Öhman, Lundqvist, & Esteves, 2001; Shasteen, Sasson, & Pinkham, 2014; 2015). Moreover, anger-detection biases such as these have been shown to be exacerbated (among White perceivers) when the target faces are Black rather than White (Ackerman et al., 2006; Hugenberg, 2005), as well as when the target faces are male rather than female (Öhman, Juth, & Lundqvist, 2010; Williams & Mattingly, 2006). These patterns of moderation are presumed to be attributable to stereotypes in perceivers' minds that conflate the Blackness (vs. Whiteness) with threat, and that likewise conflate men (vs. women) with threat. The purpose of Experiment 4 was to examine whether each of these biases trade off not only with each other, but also with an *intersectional* bias that perceivers may have in which they associate anger with Black men's faces more than readily than with the faces of other intersectional groupings (i.e., Black women, White men, White women). That is to say that Experiment 4 was designed to examine whether intersectional perceptual biases trade off with more simplistic, non-intersectional biases.

### Method

In Experiment 4, participants completed a speeded anger detection task (adapted from Öhman et al., 2001). In contrast to prior anger-detection experiments, which featured either male targets only (whose race was manipulated: e.g., Hugenberg, 2005), or White targets only (whose gender was manipulated: e.g., Williams & Mattingly, 2006), Experiment 4 featured targets

whose race and gender was simultaneously manipulated in a 2 (target race: Black, White)  $\times$  2 (target gender: male, female) within-person design. In addition, participants completed the anger-detection task three times: once in which no lens was made salient (control condition), once in which the lens of gender was made salient (gender-lens condition), and once in which the lens of race was made salient (race-lens condition). Thus, the full design of Experiment 4 was a 2 (target race)  $\times$  2 (target gender)  $\times$  3 (condition) within-person experiment.

**Participants.** A total of 223 undergraduates at Northwestern University participated in exchange for course credit. Of these,  $n = 19$  (8.5%) were excluded either because they (a) had a mean response latency that was more than three median absolute deviations from the median of participants' mean response latencies (Leys et al., 2013), or because they (b) correctly responded on fewer than 78.6% of Experiment 3's trials (78.6% was determined as the cutoff because it is three median absolute deviations below the median of the participants' mean rates of correct responding). The final sample comprised  $N = 204$  participants, of whom the majority were female (113 female, 87 male, 4 non-specified), the majority were White (94 White, 63 Asian, 16 Black, 14 Latinx, 3 Pacific Islander, 9 multiracial, 5 non-specified), and of whose ages ranged from 18 to 22 ( $M_{\text{age}} = 18.85$ ,  $SD = 0.97$ ). Finally, this sample skewed toward political liberalism ( $M = 3.75$ ,  $SD = 1.78$ , on a scale from 0 = *extremely liberal* to 10 = *extremely conservative*).

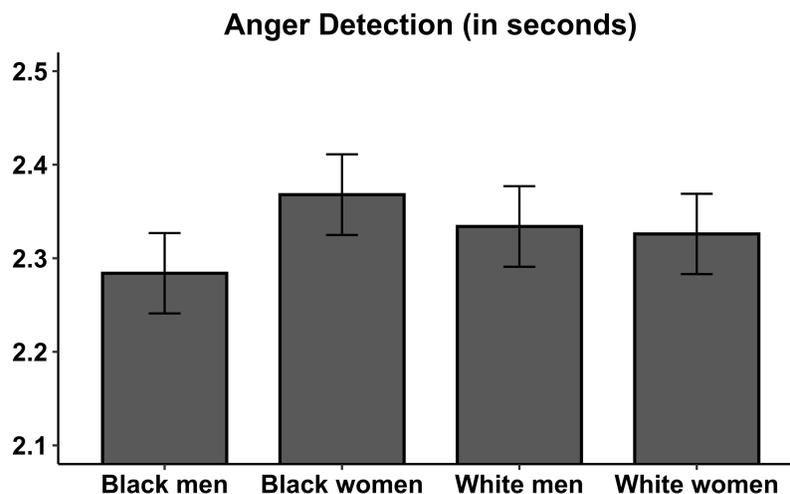
**Procedure.** Participants were told that they would be presented with groups of people on a computer screen. Their task would be to indicate as quickly as possible whether everyone in the group had a neutral facial expression, or whether one person in the group was angry. Participants then completed three sets of 72 trials: a control set, a gender-lens set, and a race-lens set. The control set was always presented to participants first; the gender-lens and race-lens sets

were always presented to participants second and third, but their ordering was counterbalanced. On every trial (regardless of set), participants saw 8 faces (2 Black men, 2 Black women, 2 White men, 2 White women, drawn randomly from a broader set of 48 faces; Ma et al., 2015). On a random third of each set's 72 trials, all the faces were neutral; on the other two thirds of each set's 72 trials, one of the 8 faces was angry. The key dependent variable was how long (in milliseconds) it took participants to notice that a face was angry on the trials in which one of the faces was indeed angry. During the control set of 72 trials, there were no additional instructions. During the gender-lens set, each trial was interspliced with a brief gender-categorization task in which participants saw 4 randomly-selected target faces and had to categorize them as 'male' or 'female.' During the race-lens set, each trial was interspliced instead with a brief race-categorization task in which participants saw 4 randomly-selected target faces and had to categorize them as "Black" or "White." The purpose of intersplicing each of the trials with either a gender-categorization task or a race-categorization task, respectively, was to make the lenses of *gender* or *race* highly accessible to participants.

**Results.** Participants in Experiment 4 were expected to exhibit a gender bias in the gender-lens condition, a racial bias in the race-lens condition, and an intersectional bias in the control condition. These three conditions were expected to reveal, respectively, a main effect of gender (but not a main effect of race); a main effect of race (but not a main effect of gender); and an interaction between race and gender. In order to examine whether this was indeed the case, participants' reaction times were subjected to a multilevel model that was equivalent to a 2 (target race)  $\times$  2 (target gender)  $\times$  3 (condition) within-person ANOVA. This model included a

random effect of participant intercept, which adjusts for the fact that the full factorial design was nested within person.

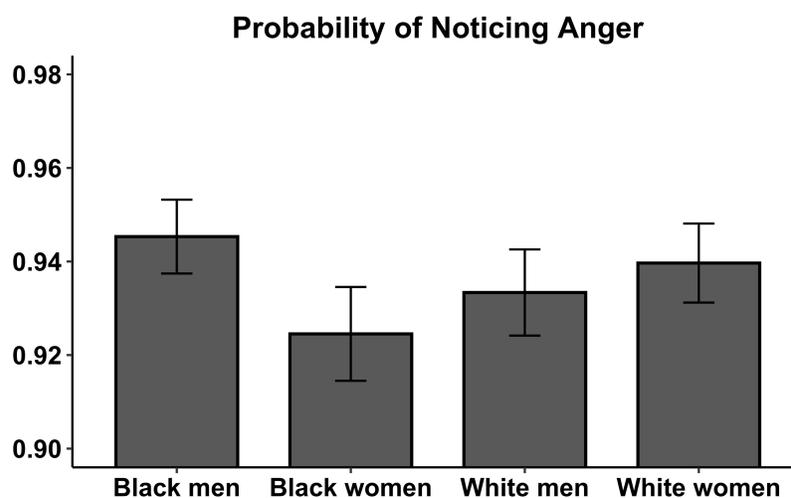
Subjecting reaction times to the analysis described above revealed a main effect of target gender, such that participants were faster to notice anger on the faces of male targets ( $M = 2.31$ ,  $SE = 0.02$ ), than on the faces of female targets ( $M = 2.35$ ,  $SE = 0.02$ ),  $M_{\text{diff}} = -38.24\text{ms}$ , 95% CI $[-59.78, -16.70]$ ,  $\beta = -0.04$ ,  $F(1, 25820) = 12.11$ ,  $p < .001$ . In addition, the main effect of target gender interacted with target race [ $F(1, 25820) = 17.46$ ,  $p < .001$ ,  $\omega_p^2 < .01$ ], such that the main effect of gender *only held* when targets were Black:  $M_{\text{diff}} = -84.16\text{ms}$ , 95% CI $[-114.68, -53.64]$ ,  $\beta = -0.09$ ,  $F(1, 25820) = 29.21$ ,  $p < .001$ . In contrast, when the targets were White, participants were no longer faster to identify anger on men's faces than on women's faces:  $M_{\text{diff}} = 7.68\text{ms}$ , 95% CI $[-22.72, 38.08]$ ,  $\beta = 0.01$ ,  $F(1, 25820) = 0.25$ ,  $p = .62$ . Thus, there was evidence that in general, participants exhibited an *intersectional* perceptual bias in which they noticed anger most quickly on the faces of Black men specifically (see Figure 15). However, there was no evidence that this intersectional perceptual bias was moderated by condition. That is, the target race  $\times$  target gender interaction was not significantly dependent on condition: all  $F$ s  $< 0.23$ , all  $p$ s  $> .63$ . Moreover, there was no evidence that the main effect of target gender varied as a function of whether the condition emphasized the lens of gender vs. the lens of race [ $F(1, 25820) = 0.53$ ,  $p = .47$ ], and there was only marginal evidence that the main effect of race varied as a function of whether the condition emphasized the lens of race vs. gender [ $F(1, 25820) = 2.91$ ,  $p = .09$ ,  $\omega_p^2 < .01$ ]. Taken together, while there was evidence of an intersectional perceptual bias (suggesting participants' use of an intersectional lens) there was not any evidence that this lens could be supplanted by the more simplistic lenses of *gender* or *race*, respectively.



*Figure 15.* Speed (in seconds) with which participants in Experiment 4 noticed expressions of anger on the faces of Black men, Black women, White men, and White women. Means are encompassed by 95% confidence intervals.

For the most part, participants tended to notice expressions of anger when they were present (anger was correctly identified on 92% of the trials in which it was present). However, another analysis that can be done on these data is to examine whether the probability of noticing anger (coded as  $1 = \text{anger detected}$ ,  $0 = \text{anger not detected}$ ) varied systematically as a function of the  $2$  (target race)  $\times$   $2$  (target gender)  $\times$   $3$  (condition) model described above. To examine whether this was the case, the probability of noticing anger was subjected to a (logit) multilevel model that was statistically analogous to a  $2 \times 2 \times 3$  within-person logistic ANOVA. This model had just one random effect: a random effect of participant intercept. This analyses yielded a highly similar pattern to that of participants' response times. That is, there was a main effect of target gender, such that the probability of noticing anger was significantly higher for male faces than for female faces (13% higher; odds ratio = 1.13,  $z = 2.62$ ,  $p = .009$ ), but this was moderated by targets' race ( $z = 4.94$ ,  $p < .001$ ). The nature of this interaction was that this main effect of gender held only for Black targets (such that people were 41% more likely to notice anger on the

faces of Black men than on Black women; odds ratio = 1.41,  $z = 5.36$ ,  $p < .001$ ). When the targets were White, in contrast, participants were not more likely to detect anger on the faces of men vs. women (odds ratio: 0.90,  $z = -1.64$ ,  $p = .10$ ; see Figure 16). There was some evidence that this race-by-gender interaction was stronger on control trials than on race- and gender-lens trials:  $z = 3.19$ ,  $p = .001$ . No other meaningful interactions with condition emerged (all  $ps > .35$ ). Thus, while anger-detection rates revealed evidence of an intersectional perceptual bias, there was no evidence that this bias was replaced with a gendered perceptual bias in the gender-lens condition, nor was there evidence that this bias was replaced with a racialized perceptual bias in the race-lens condition.



*Figure 16.* Probability that participants in Experiment 4 noticed (vs. did not notice) expressions of anger on the faces of Black men, Black women, White men, and White women. Probabilities are encompassed by 95% confidence intervals.

## Discussion

Experiments 1a-3 all investigated whether singular, simplistic lenses (e.g., gender) can come into focus at the expense of other singular, simplistic lenses (e.g., age). Generally speaking, the results from these studies accorded with the possibility that they do. However, ICT also

argues that lenses can be intersectional and complex (e.g., a lens for *older woman*, specifically) in addition to singular and simplistic, and that intersectional lenses can trade off in perceivers' minds with singular, simplistic lenses. Experiment 4 was designed to empirically test this possibility. Specifically, Experiment 4 was designed to examine whether an intersectional stereotypic bias (a bias directed toward Black men) would be attenuated and replaced with a gender bias when the lens of *gender* was made accessible to participants, and whether it would likewise be replaced by a race bias when the lens of *race* was made accessible to participants. Although Experiment 4 did provide evidence of an intersectional perceptual bias—a bias to notice anger most quickly and reliably on the faces of Black men, and to notice anger least quickly and least reliably on the faces of Black women—it did not provide evidence that this bias could be supplanted with more simplistic gender biases or racial biases, respectively. One possible reason for why this may have been the case is that, as noted previously, the process of lens selection (and in particular the process of switching from one lens to another) may be quite quick and flexible over time. This is to say that participants may have been using the lens of *gender*, for example, when categorizing faces by their gender groups, but may have quickly reverted to using an *intersectional* lens the minute they were presented with a 2 (Black, White) × 2 (male, female) array of faces on the computer screen. Such a possibility is reasonable from the perspective of ICT. However, even if lens switching *can* be quick and flexible over time, it may be the case that such quick and flexible lens switching is relatively rare in daily life (that is, in contexts outside of the laboratory). Considerations related to external validity will be discussed further in the general discussion.

## Experiments 5a and 5b

Experiment 4 revealed that perceivers do occasionally view targets through intersectional lenses. However, Experiment 4 did not provide evidence that perceivers will use simplistic lenses in lieu of intersectional lenses when the situation calls for it. Experiments 5a and 5b were designed to continue investigating to what extent intersectional lenses trade off in perceivers' minds with more simplistic lenses. Experiment 5a investigated whether perceivers would associate Black women with weapons when viewing these women through the lens of *race*, but whether they would cease to do so when viewing these women through an intersectional lens or the a gender lens. Experiment 5b investigated whether perceivers would associate older women with church concepts when viewing these women through an intersectional lens (i.e., as *old women*), but whether this bias would likewise become attenuated when viewing these women through the lenses of either *gender* or *age*.

### Experiment 5a

Previous research has documented that perceivers in the U.S. associate crime with Black people more than with White people (Eberhardt et al., 2004; 2006). One way by which this bias has been indexed is through a variant of the IAT called the race-weapons IAT (Glaser & Knowles, 2008). In this version of the IAT, participants see Black and White faces along with words that can be categorized as weapons (e.g., pistol, shotgun) or harmless objects (e.g., laptop, sweater). To the extent that perceivers are faster to associate Black faces with weapons—and White faces with harmless objects—than the reverse of these pairings, researchers can infer that participants harbor an implicit stereotype linking Blackness and crime.

As was the case with the male-science bias discussed previously, implicit associations on the race-weapons IAT have been shown to be consequential. For example, individuals who have stronger associations between Black faces and weapons exhibit an increased “shooter bias” in which they are faster to shoot armed Black individuals vs. armed White individuals in a video game task (Glaser & Knowles, 2008; Correll, Park, Judd, & Wittenbrink, 2002). Experiment 5a was designed to examine how pliable this (consequential) association between race and weapons might be. Specifically, Experiment 5a was designed to examine whether Black women would be associated with weapons more quickly than White men when viewed through the lens of *race*, whether this bias would be attenuated when these women were viewed through an *intersectional* lens, and whether this bias would reverse entirely when these women were viewed through the lens of *gender*.

**Method.** Participants in Experiment 5a all completed a race-weapons IAT that featured targets who were Black women or White men, but participants were randomly assigned to one of three between-person conditions (race lens, gender lens, intersectional lens). In the *race-lens* condition, participants completed the race-weapons IAT as usual. In the *gender-lens* condition, participants were instructed to categorize targets by gender rather than by race. Finally, in the *intersectional-lens* condition, participants were instructed to categorize targets by their intersection (that is, as “Black women” or “White men”). The predictions of Experiment 5a were a) that participants would be faster to associate Black women with weapons and White men with harmless objects (than the reverse of these associations) when categorizing targets by their race; b) that this bias would be reversed when participants were categorizing targets by their gender; and c) that this bias would disappear (that is, be no different from zero) when participants were

categorizing targets by their intersection. All predictions, exclusions, and analytic choices for Experiment 5a were pre-registered.

**Participants.** A total of 321 U.S. citizens were recruited from MTurk to complete Experiment 5a. Of these,  $n = 27$  (8.4%) were excluded either for a) not responding “yes” to the question, “Did you take this study seriously?”; b) reporting technical difficulties during the IAT itself; or c) being more than three median absolute deviations from the median of the sample’s average response latency on the IAT (that is, for being either too fast or too slow: Leys et al., 2013). The remaining participants were mostly male (183 male, 111 female), mostly White (223 White, 29 Black, 19 Asian, 15 Latinx, 2 American Indian, 2 Pacific Islander, 2 multi-racial, 2 non-specified), and had ages spanning from 20 to 72 ( $M = 36.95$ ,  $SD = 11.15$ ). In addition, 50.34% of the sample held at least a bachelor’s degree, and they skewed toward political liberalism ( $M = 3.74$ ,  $SD = 3.00$ , on an 11-point scale from  $0 = extremely liberal$  to  $10 = extremely conservative$ ).

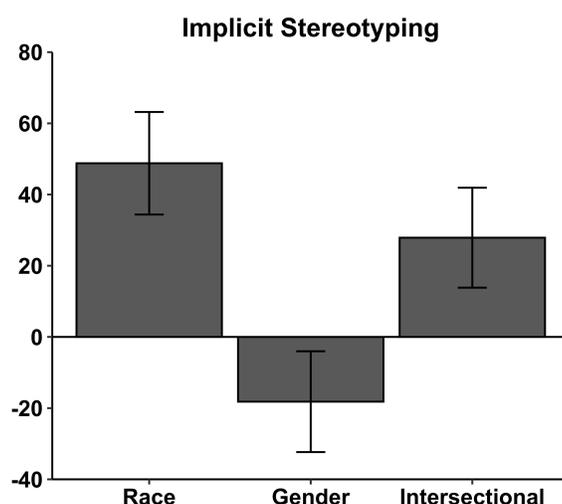
**Procedure.** Participants in Experiment 5a were randomly assigned to one of three between-person conditions: a race-lens condition, a gender-lens condition, or an intersectional-lens condition. All three IATs featured the same stimulus faces and stimulus words. The stimulus faces included 11 Black women’s faces and 11 White men’s faces (all of whom appeared to be in their 20s; faces taken from the MR2 database: Strohminger, Gray, Chituc, Heffner, Schein, & Heagins, 2016). The words were either weapons (e.g., gun, knife, blade) or harmless objects (e.g., camera, soda, wallet). In all three conditions, participants completed blocks in which they either associated Black women’s faces with weapons—and White men’s faces with harmless objects—or in which they completed the reverse of these associations (that is, Black women’s

faces with harmless objects and White male faces with weapons). All participants completed 60 trials of each pairing type. In the race-lens condition, participants were instructed to categorize targets by race (Black, White); in the gender-lens condition, participants were instructed to categorize targets by sex (female, male); and in the intersectional-lens condition, participants were instructed to categorize targets by their intersectional identities (Black women; White men). Within all IATs, block ordering was randomized, as was which key participants had to press for each of these respective pairing types.

**Results.** Response latencies (in milliseconds) were regressed onto a mixed linear model that was statistically analogous to a 3 (condition: race lens, gender lens, intersectional lens)  $\times$  2 (trial type: Black-weapons + White-objects; Black-objects, White-weapons) mixed ANOVA with repeated measures on the second factor. This model included estimates of two random effects: a random effect of participant intercept, which adjusts for the fact that trials were nested within person; and a random effect of stimulus intercept, which adjusts for the fact that each trial featured one of several possible stimuli that were drawn randomly from the stimulus pool (Judd et al., 2012).

According to ICT, racial stereotypes that cause perceivers to associate Blackness with weapons should manifest most strongly when perceivers are using the lens of race. To examine whether this was indeed the case, response latencies on the race-weapons IAT were subjected to the 3  $\times$  2 analysis described above. This analyses revealed that as anticipated, there was a main effect of trial type: participants associated Black women with weapons—and White men with harmless objects ( $M = 759.08$ ,  $SE = 15.03$ )—more quickly than they completed the reverse of these associations ( $M = 778.57$ ,  $SE = 15.02$ ),  $M_{diff} = -19.49\text{ms}$ , 95% CI[-27.69, -11.30],  $\beta = -$

0.05,  $F(1, 34890) = 21.72, p < .001$ . Moreover, and in line with the prediction outlined above, this Black-weapons association was moderated by experimental condition [interaction:  $F(1, 34890) = 24.18, p < .001, \omega_p^2 < .01$ ], such that this bias was stronger in the race-lens condition [ $M_{\text{diff}} = -48.79\text{ms}$ , 95% CI[-63.19, -34.38],  $\beta = -0.12, F(1, 34880) = 44.08, p < .001$ ] than it was in the other two conditions [ $M_{\text{diff}} = -4.85$ , 95% CI[-14.82, 5.12],  $\beta = -0.01, F(1, 34890) = 0.91, p = .34$ ]. Thus, the hypothesis that Black women would be most strongly associated with weapons in the race lens condition was supported (see Figure 17).



*Figure 17.* Average speed (in milliseconds) with which participants in Experiment 5a associated Black women (vs. White men) with weapons (vs. harmless objects). Values above zero indicate a greater association between Black women and weapons; scores below zero indicate a greater association between White men and weapons. Results are broken down by whether participants were in the race-lens (left), gender-lens (middle), or intersectional-lens (right) condition. Means are encompassed by 95% confidence intervals.

In addition, the magnitude of the race-weapons bias significantly differed between the gender-lens condition and the intersectional-lens condition: interaction  $F(1, 34888) = 20.52, p < .001, \omega_p^2 < .01$ . In the intersectional-lens condition, participants exhibited a Black-weapons bias (albeit to a weaker degree than did participants in the race-lens condition:  $M_{\text{diff}} = -27.88\text{ms}$ , 95% CI[-41.91, -13.84],  $\beta = -0.07, F(1, 34880) = 15.15, p < .001$ ). However, in the gender-lens

condition, participants exhibited a *reversal* of the Black-weapons bias. In this condition, participants were faster to associate White *men* with weapons than they were to associate Black women with weapons:  $M_{\text{diff}} = 18.18\text{ms}$ , 95% CI[4.02, 32.34],  $\beta = 0.04$ ,  $F(1, 34890) = 6.34$ ,  $p = .012$ . Thus, there was strong evidence that patterns of racial stereotyping that are directed toward Black women manifest most strongly these women are viewed through the lens of *race*, and that they can become attenuated or even reversed when these women are viewed intersectional or gender lenses.

### **Experiment 5b**

Experiment 5a provided initial evidence that viewing targets through an intersectional lens can bring stereotypes to mind that are distinct from those that come to mind by way of singular lenses. For example, when viewed through intersectional lenses, Black women were associated with weapons more strongly than White men were associated with weapons. However, when viewed through the lens of *gender*, this bias reversed such that White men were associated with weapons more than Black women were. An issue with Experiment 5a, however, is that it remains unclear what was causing the pattern of stereotyping observed in the intersectional lens condition. On the one hand—as ICT predicts—it may have been the case that the intersection of *Black womanness* brought its own stereotypic associations to perceivers' minds. On the other hand, it could have been the case that perceivers in the intersectional lens condition were not thinking of Black womanness in and of itself, but were instead averaging together the stereotypes that came to mind when thinking of *Blackness* by itself or *womanness* by itself. The purpose of Experiment 5b was to examine whether intersectional lenses can indeed bring their own stereotypes to perceivers minds—stereotypes that are something greater than the

average of what comes to mind when perceivers are using singular lenses. In particular, Experiment 5b tested whether a) older women would be associated with church concepts (Payne & Whittington, 1975)—and whether young men would be associated with fraternity concepts (Ashmore, Del Boca, & Beebe, 2002)—more quickly than the reverse of these pairings, and it tested whether b) this bias would manifest *most strongly* when targets are being viewed through intersectional lenses rather than simplistic lenses.

**Method.** Participants in Experiment 5b all completed a ‘church-lady’ IAT that featured targets who were older White women and young White men, and that featured words that were either church-related (e.g., God) or fraternity-related (e.g., Beer). As in Experiment 5a, participants were randomly assigned to one of three between-person conditions. In the *age-lens* condition, participants completed the ‘church-lady’ IAT while categorizing faces by their age groups. In the *gender-lens* condition, participants were instructed to categorize targets by their gender groups rather than by their age groups. Finally, in the *intersectional-lens* condition, participants were instructed to categorize targets by their intersection (that is, as either “old women” or “young men,” specifically). The predictions of Experiment 5b were that a) participants would be faster to associate older women with church concepts—and young men with fraternity concepts—than the reverse, and that b) this bias would be significantly stronger in the intersectional lens condition than in the other two conditions. All predictions, exclusions, and analyses for Experiment 5b were registered.

**Participants.** A total of 374 U.S. citizens were recruited from MTurk to complete Experiment 5b. Of these,  $n = 36$  (9.6%) were excluded either for a) not responding “yes” to the question, “Did you take this study seriously?”; b) reporting technical difficulties during the IAT

itself; or c) being more than three median absolute deviations from the median of participants' average response latencies on the IAT (that is, for being either too fast or too slow: Leys et al., 2013). The remaining participants were mostly male (202 male, 136 female), mostly White (249 White, 38 Black, 25 Asian, 18 Latinx, 1 American Indian, 6 multi-racial, 1 non-specified), and had ages spanning from 21 to 75 ( $M = 37.83$ ,  $SD = 11.94$ ). In addition, 49.41% of the sample held at least a bachelor's degree, and they skewed toward political liberalism ( $M = 3.64$ ,  $SD = 2.87$ , on an 11-point scale from  $0 = \textit{extremely liberal}$  to  $10 = \textit{extremely conservative}$ ).

**Procedure.** Participants in Experiment 5b were randomly assigned to one of three between-person conditions: an age-lens condition, a gender-lens condition, or an intersectional-lens condition. All three IATs featured the same stimulus faces and stimulus words. The stimulus faces included 15 older women's faces and 15 young men's faces (all of whom were White; faces taken from the FACES database: Ebner et al., 2010). The words were either church-related (e.g., faith, church, prayer, rosary) or fraternity-related (e.g., college, party, drinking, jock). In all three conditions, participants completed blocks in which they either associated older women's faces with church-related words—and young men's faces with fraternity-related words—or in which they completed the reverse of these associations (that is, older women's faces with fraternity-related words and young men's faces with church-related words). All participants completed 60 trials of each pairing type. In the age-lens condition, participants were instructed to categorize targets by age (old, young); in the gender-lens condition, participants were instructed to categorize targets by sex (female, male); and in the intersectional-lens condition, participants were instructed to categorize targets by their intersectional identities (as old women and young

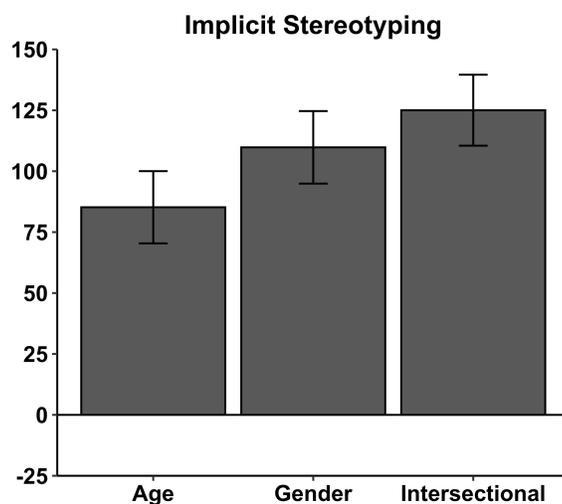
men, respectively). Within all IATs, block ordering was randomized, as was which key participants had to press for each of these respective pairing types.

**Results.** Response latencies (in milliseconds) were regressed onto a mixed linear model that was statistically analogous to a 3 (condition: age lens, gender lens, intersectional lens)  $\times$  2 (trial type: church-lady + frat-boy; church-boy + frat-lady) mixed ANOVA with repeated measures on the second factor. This model included estimates of two random effects: a random effect of participant intercept, which adjusts for the fact that trials were nested within person; and a random effect of stimulus intercept, which adjusts for the fact that each trial featured one of several possible stimuli drawn randomly from the stimulus pool (Judd et al., 2012).

According to ICT, any tendency there might be for participants to associate older women with church—and young men with fraternities—should be most pronounced when participants are using an intersectional lens for viewing targets. Subjecting participants' response latencies to the analysis described above revealed evidence that is highly consistent with this reasoning. That is, participants were faster to respond on trials that were congruent with the 'church-lady' stereotype ( $M = 791.43$ ,  $SE = 15.12$ ) than they were on trials that were incongruent with this stereotype ( $M = 898.17$ ,  $SE = 15.12$ ):  $M_{\text{diff}} = -106.68\text{ms}$ , 95% CI $[-115.20, -98.16]$ ,  $\beta = -0.22$ ,  $F(1, 40295) = 602.36$ ,  $p < .001$ . More critically, this bias was significantly interacted with condition [ $F(1, 40295) = 9.04$ ,  $p = .003$ ,  $\omega_p^2 < .01$ ], such that it was significantly more pronounced in the intersectional-lens condition [ $M_{\text{diff}} = -125.04\text{ms}$ , 95% CI $[-139.60, -110.48]$ ,  $\beta = -0.26$ ,  $F(1, 40295) = 283.28$ ,  $p < .001$ ] than it was in the singular-lens conditions [ $M_{\text{diff}} = -97.50\text{ms}$ , 95% CI $[-108.01, -87.00]$ ,  $\beta = -0.20$ ,  $F(1, 40295) = 331.02$ ,  $p < .001$ ; see Figure 18]. This is to say that the intersectional stereotypes that characterize older women and religious were

most strongly associated with these women when they were viewed through an intersectional lens.

Finally, there was evidence that the strength of the ‘church-lady’ bias was moderated by which singular-lens condition participants were assigned:  $F(1, 40295) = 5.26, p = .022, \omega_p^2 < .01$ . The nature of this interaction was that the ‘church-lady’ bias was significantly stronger in the gender-lens condition [ $M_{\text{diff}} = -109.80\text{ms}$ , 95% CI[-124.69, -94.91],  $\beta = -0.23, F(1, 40295) = 208.95, p < .001$ ], than it was in the age-lens condition [ $M_{\text{diff}} = -85.21\text{ms}$ , 95% CI[-100.03, -70.39],  $\beta = -0.18, F(1, 40295) = 126.99, p < .001$ ]. Although, this latter difference in conditions was not anticipated, it is consistent with the idea that different lenses may have different implications for how strongly a target is associated with particular stereotypic concepts.



*Figure 18.* Average speed (in milliseconds) with which participants in Experiment 5b associated older women (vs. young men) with church (vs. fraternity) concepts. Values above zero indicate a greater association between older women and church (and between young men and fraternities). Results are broken down by whether participants were in the age-lens (left), gender-lens (middle), or intersectional-lens (right) condition. Means are encompassed by 95% confidence intervals.

## Discussion

Experiments 5a and 5b were designed to investigate whether the same targets (e.g., Black women) might be stereotyped differently depending on whether they are stereotyped through the inflection of one singular lens (e.g., *gender*), another singular lens (e.g., *race*) or through the inflection of an intersectional lens (e.g., *race-by-gender*). Experiment 5a revealed that Black women are indeed associated most strongly with racial stereotypes when viewed through a *race* lens as compared with other lenses. Furthermore, Experiment 5a suggested that stereotypes disadvantaging Black women relative to White men can be supplanted by stereotypes that advantage Black women relative to White men. When perceivers viewed these women through the lens of *gender*, they associated White men with weapons more quickly than they associated Black women with weapons—a pattern that contradicts the well-established Black-weapons bias that this task typically unveils (e.g., Glaser & Knowles, 2008).

Experiment 5b was designed to investigate whether intersectional lenses can bring certain stereotypes to mind more strongly than would non-intersectional lenses. In particular, Experiment 5b investigated whether ‘church-lady’ and ‘frat-boy’ stereotypes might be most strongly applied to older women and young men, respectively (Ashmore et al., 2002; Payne & Whittington, 1975), when these targets were viewed through an intersectional lens. Consistent with this possibility, Experiment 5b showed that the tendency to associate older women with church-concepts—and to associate young men with fraternity concepts—manifested to a significantly stronger degree when these targets were viewed through intersectional lenses than when these targets were viewed through the simplistic lenses of *age* or *gender*, respectively. This suggests not only that intersectional lenses can trade off in perceivers minds with more simplistic

lenses, but that intersectional lenses can themselves bring stereotypes to mind that would not come to mind as strongly when non-intersectional lenses are in use.

## Chapter V: General Discussion

## General Discussion

Prevalent theories in social psychology tend to argue for one of two perspectives when explaining how perceivers engage in intersectional stereotyping: either a) that perceivers invariably attend to certain social identities more than others (termed dominance theories: Pietraszewski et al., 2015), or b) that perceivers invariably attend to all of targets' social identities at once (termed integration theories: e.g., Freeman et al., 2020). The main issue with the former perspective is that it has been falsified many times over. The main issue with the latter perspective is that it is virtually unfalsifiable. In contrast to these perspectives, this dissertation advocates for a “middle path”—a theoretical perspective that is flexible enough to account for contradictory findings in the literature on intersectional stereotyping, but not so flexible as to be virtually unfalsifiable (Petsko & Bodenhausen, 2020). The name of the middle-path perspective developed here is intersectional categorization theory (ICT), which argues that perceivers sharpen their focus on just one identity—or one intersection of identities—at a time, as a function of the social context. This is to say that when perceivers are thinking of gay Black men as *Black*, they should not, at least in these contexts, be inclined to attend to these men's sexual orientation groups (Petsko & Bodenhausen, 2019a, Experiment 1). Likewise, when perceivers are thinking of gay Black men as *gay*, they should not, at least in these contexts, be inclined to attend to these men's racial groups (Petsko & Bodenhausen, 2019b). Finally, ICT argues that when perceivers are thinking gay Black men not as *gay* or as *Black*, but as *gay Black men* specifically, they should be likely to exhibit patterns of stereotyping against these men that are intersectional in nature, and that are not necessarily the sum of how perceivers think of “gayness” by itself or “Blackness” by itself (Calabrese et al., 2017).

## Summary of Findings and Contributions

Seven experiments were designed to provide preliminary tests of ICT's assumptions. Experiments 1a, 1b, and 2 examined whether sharpening perceivers' focus on the lens of *gender* could cause perceivers to regard targets as interchangeable with other members of their gender groups, but cause perceivers to no longer regard targets—in these moments—as interchangeable with other members of their age groups (Experiment 1a, Experiment 2) or racial groups (Experiment 1b), respectively. Generally speaking, these experiments provided strong support for this idea. When a lens provided comparative fit to the social context in these experiments, perceivers exhibited strong memory for lens-relevant identities, but not for lens-irrelevant identities. For example, when the experimental context comparatively fit the lens of gender, perceivers confused Black women with other *women*, but they exhibited no systematic tendency to confuse these Black women with other Black individuals. Experiment 3 extended these findings by suggesting that as one lens comes into focus, stereotypes relating to the lens-associated identity—but not to the lens-irrelevant identity—become associated with targets in the minds of perceivers. Specifically, Experiment 3 revealed that when the lens of *gender* is made highly accessible to perceivers (by way of giving perceivers the goal of attending to gender during an IAT), perceivers exhibit a pattern of implicit gender stereotyping such that they associate science with men more than with women (i.e., a male-science bias: Nosek et al., 2009). However, Experiment 3 also revealed that when the lens of *age* is made accessible to perceivers instead, perceivers *stop* exhibiting evidence of implicit gender stereotyping and they *start* exhibiting evidence of implicit age stereotyping (such that they associate older people with science more readily than they associate young people with science). Experiment 4 began to

grapple with the question of whether there would be moments in which perceivers attend to intersectional lenses specifically. This experiment revealed consistent evidence that on an anger-recognition task (Öhman et al., 2001; Shasteen et al., 2014; 2015), perceivers were most accurate (and quickest) at noticing anger on the faces of Black men, and they were least accurate (and slowest) at noticing anger on the faces of Black women. Thus, Experiment 4 provided evidence that there are social contexts in which perceivers attend to targets' intersecting race and gender identities. Experiments 5a and 5b expanded these findings by showing, first, that when perceivers are attending to targets' intersectional identities they exhibit patterns of implicit stereotyping that are of categorically different magnitudes than those they exhibit when attending to targets' non-intersectional identities. For example, the magnitude of an implicit Black-weapons bias on an IAT (Glaser & Knowles, 2008) differed significantly across each of three conditions: a race-lens condition (where it was strongest), and intersectional-lens condition (where it was significantly weaker), and a gender-lens condition (where it was not just weaker, but significantly reversed such that there was evidence a White-weapons bias). The second way that these experiments expanded the findings of Experiments 1-4 was by showing that in some contexts, intersectional lenses can bring stereotypic associations to perceivers' minds more strongly than would singular lenses. In particular, Experiment 5b revealed that when perceives were attending to older women's (and young men's) *intersectional* identities, they exhibited a stronger tendency to associate older women with church (Payne & Whittington, 1975)—and young men with fraternities (Ashmore et al., 2002)—than they did when attending to these targets' non-intersectional identities (i.e., their gender groups alone or their age groups alone). Collectively, these experiments provide strong evidence for the fundamental tenets of ICT: (a)

that perceivers use one lens at a time for making sense of other people; (b) that the lenses perceivers use can be singular and simplistic, or intersectional and complex; and (c) that different lenses can prescribe categorically distinct sets of stereotypes that perceivers use as frameworks for thinking about targets.

### **Limitations**

Of note, the experiments described above are not without their limitations. Experiments 1a, 1b, and 2, for example, *did* reveal evidence that perceivers' attention to lens-associated identities can be compartmentalized, but they did not reveal that perceivers' attentional patterns altered the social stereotypes that they applied to targets. That is, while perceivers did exhibit evidence of using social categories in a compartmentalized way on the who-said-what task (Taylor et al., 1978), they did not exhibit a tendency to stereotype older women, for example, as seeming "older" when these women's age groups comparatively fit the social context than when these women's gender groups comparatively fit the social context. Conceivable reasons why this may have been the case are that comparative fit manipulations either a) only sharpen perceivers' attention on lens-associated identities for a short duration of time (that is, during the comparative fit manipulation itself, but not after this manipulation concludes), or that b) comparative fit manipulations do not activate identity-relevant stereotypes in the minds of perceivers. The former possibility implies that lens usage may be highly flexible, with lenses coming in and out of focus as the immediate social environment demands. The latter possibility poses problems for ICT, which currently argues that comparative fit manipulations should influence not just which identities perceivers attend to, but also what stereotypes come to perceivers' minds when thinking about the targets of their perceptions (Oakes, 1987; 1994). Future work should seek to adjudicate

between these two possibilities, as this would be informative for understanding the role that comparative fit plays in shaping (or not shaping) person perception. Another notable limitation of these findings was that while Experiment 4 *did* reveal that perceivers occasionally attend to intersectional identities themselves, it did not reveal that intersectional lens usage—on the part of perceivers—could be supplanted by singular lens usage by making singular lenses highly accessible. That is, an intersectional tendency for perceivers to notice anger most quickly on Black men’s faces could not be supplanted by a tendency for perceivers to notice anger most quickly on men’s faces regardless of these men’s race (in the gender-lens condition), or on Black faces regardless of these individuals’ gender (in the race-lens condition). This finding, too, poses issues for ICT. Specifically, this finding poses issues for ICT’s claim that making the lens of gender more accessible, for example, should bring targets’ gender into perceivers’ focus *at the expense* of targets’ intersectional identities. Although all seven experiments are collectively supportive of ICT’s main arguments, more work is needed to understand why some accessibility manipulations—for example, those used in Experiments 3, 5a, and 5b—work well, whereas others (e.g., those used in Experiment 4) do not.

Another limitation of the present experiments is that although they are high in internal validity, they are relatively low in external validity (for a critique related to this issue, see Baumeister, Vohs, & Funder, 2007). For example, the factor of comparative fit was manipulated by showing participants artificial conversations—on a computer screen—in which members of two different groups (e.g., older people, young people) debated one another on a topic that was *specifically normed* to be unrelated to their group memberships (Experiments 1a, 1b, and 2). Real-world social exchanges—exchanges that occur in naturalistic settings—are rarely, if ever,

as clear-cut or heavy handed as this. Given that the purpose of ICT is to explain how it is that intersectional stereotyping manifests across different *real-world* contexts, the fact that these experiments examine computer-simulated contexts alone provides only limited support for ICT. However, it is worth noting that many of the factors ICT highlights (e.g., distinctiveness, fit) can indeed be studied in naturalistic settings. For example, settings vary in the extent to which they normatively fit different lenses. To harken back to an example from earlier, the lens of *race* may fit social reality to a greater extent at a Black Lives Matter rally than at the post office. It would be useful for future research to consider whether basic patterns of intergroup stereotyping readily shift across naturalistic environments such as these.

### **Future Directions**

Much of ICT remains to be tested. As one example, ICT argues that rare and distinctive identities (e.g., McGuire, McGuire, Child, & Fujioka, 1978), in addition to those that are accessible and fitting of social reality (Bruner, 1957; Oakes et al., 1991), ought to be attention-grabbing for perceivers. However, the experiments described here did not manipulate how rare vs. common particular social identities were in a given social environment. It would be informative to investigate whether gay Black men, for example, are indeed more likely to be viewed through the lens of their sexual orientation when surrounded by heterosexual people, yet more likely to be viewed through the lens of their racial groups when surrounded by White people. As a second example, ICT proposes that once a lens is selected, perceivers' come to accentuate the stereotypes implied by that lens in ways that allow them to see themselves favorably (Tajfel & Turner, 1979), and that also maximize the perceptual contrast (i.e., the meta-contrast: Turner et al., 1987) between groups of people that the perceiver is viewing through that

lens. The former might imply, for example, that White women stereotype Black women more positively when viewing Black women through the lens of *gender* than when viewing Black women through the lens of *race*. The latter might imply that the meaning that White women make of “Black women” depends on whether Black women are being compared against Latina women, for example, versus East-Asian women. Both of these possibilities have not yet been tested, but, if supported, would be hugely informative for the study of intersectional stereotyping. Finally, the experiments described here look only at how lens-usage influences social categorization and stereotype activation. Future research would benefit immensely from examining how compartmentalized lens-usage influences perceivers’ *behaviors*. As noted previously, behaviors are assumed to follow from lens-inflected stereotyping in ways that are relatively automatic. This is to say that viewing someone through the lens of *gender* ought to bring about gender-based discrimination, whereas viewing someone through the lens of *age* ought to bring about age-based discrimination, instead. However, discrimination on the basis of *certain* identities can often be more proscribed than discrimination on the basis of other identities (e.g., Krieger & Fiske, 2006). To harken back to an example from earlier, suppose that a perceiver is prejudiced against both older people and women, but lives in a culture that proscribes discriminatory treatment on the basis of gender. These conditions would imply that perceivers *would* exhibit discrimination against older women when viewing these women through the lens of age, but *would not* exhibit discrimination against older women when viewing these women through the lens of gender—even if the perceiver feels equally negatively toward both social identities. Such a possibility, if supported, would be informative about when psychologists can expect particular targets to face various forms of discrimination.

## Concluding Remarks

In summary, scientific understanding of intersectional stereotyping stands much to gain from considering the possibility that person perception may be a compartmentalized process. It may be the case that rather than focusing on some social identities more than others—and that rather than focusing on all social identities at once—perceivers instead focus their attention on *just one social identity at a time*. Systematic tests of this possibility suggest that perceivers do indeed sharpen their attention on some social identities at the expense of attending to others (Experiments 1a, 1b, 2, 3, 5a, 5b), that these identities can be specified in intersectional ways (Experiments 4, 5a, and 5b), and that as perceivers' lenses for viewing social targets change, so too do the stereotypic attributes that they come to associate with targets change (Experiments 3, 5a, and 5b). This framework implies that even the same social targets can be stereotyped in fundamentally different ways depending on the lens that their surrounding context affords to perceivers. It implies that just as there may be moments in which older Black men and young Black boys are stereotyped in as highly similar to each other (e.g., Todd et al., 2016), so too are there likely to be moments in which older Black men and young Black boys are stereotyped in opposition to each other. A broader implication of this reasoning is that although there are not one-size-fits-all prescriptions that scientists can make about how Black women, for example, are likely to be stereotyped *across* contexts, there can be meaningful prescriptions that scientists can make about how these women are likely to be stereotyped *within* contexts (see also, Bodenhausen & Petsko, in press). Indeed, contexts that invite perceivers to use the lenses of *gender* by itself, *race* by itself—or intersections of the two—ought to unveil the very patterns of behavior that Crenshaw (1989) observed in the legal system. Namely, that:

Black women sometimes experience discrimination in ways [that are] similar to white women's experiences; sometimes they share very similar experiences with Black men ...

And sometimes, they experience discrimination as *Black women*. (italics added; p. 149)

The findings of this dissertation lend credence to Crenshaw's observations, and they underscore the utility of ICT for explaining the complicated business of intersectional person perception.

## References

- Abrams, D., & Hogg, M. A. (2004). Metatheory: Lessons from social identity research. *Personality and Social Psychology Review, 8*, 98-106.
- Abrams, D., & Hogg, M. A. (1988). *Social identifications: A social psychology of intergroup relations and group processes*. New York, NY: Routledge.
- Abrams, D., & Hogg, M. A. (2010). Social identity and self-categorization. In J. F. Dovidio, M. Hewstone, & V. M. Esses (Eds.), *The SAGE handbook of prejudice, stereotyping and discrimination* (pp. 179-193). London: Sage.
- Ackerman, J., Shapiro, J. R., Neuberg, S. L., Kenrick, D. T., Becker, D. V., Griskevicius, V., ... Schaller, M. (2006). They all look the same to me (unless they're angry): From out-group homogeneity to out-group heterogeneity. *Psychological Science, 17*, 836-840.
- Adorno, T. W., Frenkel-Brunswik, E., Levinson, D. J., & Sanford, R. M. (1950). *The authoritarian personality*. New York: Harper.
- Altemeyer, B. (1988). *Enemies of freedom: Understanding right-wing authoritarianism*. San Francisco, CA, US: Jossey-Bass.
- Ambady, N. (2010). The perils of pondering: Intuition and thin slice judgments. *Psychological Inquiry, 21*, 271-278.
- Anderson, N. H. (1971). Integration theory and attitude change. *Psychological Review, 78*, 171-206.
- Asch, S. E. (1946). Forming impressions of personality. *The Journal of Abnormal and Social Psychology, 41*, 258-290.

- Asch, S. E., & Zukier, H. (1984). Thinking about persons. *Journal of Personality and Social Psychology, 46*, 1230-1240.
- Ashmore, R. D., Del Boca, F. K., & Beebe, M. (2002). "Alkie," "frat brother," and "jock": Perceived types of college students and stereotypes about drinking. *Journal of Applied Social Psychology, 32*, 885-907.
- Bandura, A. (1977). Self-efficacy: Toward a unifying theory of behavioral change. *Psychological Review, 84*, 191-215.
- Barnum, P., Liden, R. C., & DiTomaso, N. (1995). Double jeopardy for women and minorities: pay differences with age. *The Academy of Management Journal, 38*, 863-880.
- Bates, D., Mächler, M., Bolker, B. M., & Walker, S. C. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software, 67*, 1-48.
- Baumeister, R. F., Vohs, K. D., & Funder, D. C. (2007). Psychology as the science of self-reports and finger movements: Whatever happened to actual behavior? *Perspectives on Psychological Science, 4*, 396-403.
- Becker, D. V., Neel, R., & Anderson, U. S. (2010). Illusory conjunctions of angry facial expressions follow intergroup biases. *Psychological Science, 21*, 938-940.
- Biernat, M., & Vescio, T. K. (1993). Categorization and stereotyping: Effects of group context on memory and social judgment. *Journal of Experimental Social Psychology, 29*, 166-202.
- Billig, M., & Tajfel, H. (1973). Social categorization and similarity in intergroup behaviour. *European Journal of Social Psychology, 3*, 27-52.
- Blair, I. V., Judd, C. M., & Chapleau, K. M. (2004). The influence of afrocentric facial features in criminal sentencing. *Psychological Science, 15*, 674-679.

- Bodenhausen, G. V. (2010). Diversity in the person, diversity in the group: Challenges of identity complexity for social perception and social interaction. *European Journal of Social Psychology, 40*, 1-16.
- Bodenhausen, G. V., & Macrae, C. N. (1998). Stereotype activation and inhibition: Advances in Social Cognition, Volume XI. In R. S. Wyer, Jr. (Ed.), *Stereotype Activation and Inhibition: Advances in Social Cognition, Volume XI* (pp. 1-52).
- Bodenhausen, G. V., & Petsko, C. D. (in press). Complications in predicting intergroup behavior from implicit biases: One size does not fit all. In J. A. Krosnick, T. H. Stark, & A. L. Scott (Eds.), *The Cambridge Handbook of Implicit Bias and Racism*. Cambridge, UK: Cambridge University Press.
- Brauer, M., & Chaurand, N. (2010). Descriptive norms, prescriptive norms, and social control: An intercultural comparison of people's reactions to uncivil behaviors. *European Journal of Social Psychology, 40*, 490-499.
- Brewer, M. B. (1988). A dual process model of impression formation. In *A dual process model of impression formation* (pp. 1-36). Hillsdale, NJ, US: Lawrence Erlbaum Associates, Inc.
- Brewer, M. B., & Gardner, W. (1996). Who is this "we"? Levels of collective identity and self representations. *Journal of Personality and Social Psychology, 71*, 83-93.
- Brown-Iannuzzi, J. L., Dotsch, R., Cooley, E., & Payne, B. K. (2017). The relationship between mental representations of welfare recipients and attitudes toward welfare. *Psychological Science, 28*, 92-103.
- Bruner, J. S. (1957). On perceptual readiness. *Psychological Review, 64*, 123-152.

- Buchanan, N. T., & Fitzgerald, L. F. (2008). Effects of racial and sexual harassment on work and the psychological well-being of African-American women. *Journal of Occupational Health Psychology, 13*, 137-151.
- Calabrese, S. K., Earnshaw, V. A., Magnus, M., Hansen, N. B., Krakower, D. S., Underhill, K., ... Dovidio, J. F. (2018). Sexual stereotypes ascribed to Black men who have sex with men: An intersectional analysis. *Archives of Sexual Behavior, 47*, 143-156.
- Carpinella, C. M., Chen, J. M., Hamilton, D. L., & Johnson, K. L. (2015). Gendered facial cues influence race categorizations. *Personality & Social Psychology Bulletin, 41*, 405-419.  
<https://doi.org/10.1177/0146167214567153>
- Chapman, L. J. (1967). Illusory correlation in observational report. *Journal of Verbal Learning & Verbal Behavior, 6*, 151-155.
- Cialdini, R. B. (2012). The focus theory of normative feedback conduct. In P. A. M. Van Lange, A. W. Kruglanski, & E. T. Higgins, *Handbook of theories of social psychology* (pp. 295-312). London, UK: Sage.
- Cole, E. R. (2009). Intersectionality and research in psychology. *American Psychologist, 64*, 170-180.
- Correll, J., Park, B., Judd, C. M., & Wittenbrink, B. (2002). The police officer's dilemma: Using ethnicity to disambiguate potentially threatening individuals. *Journal of Personality and Social Psychology, 83*, 1314-1329.
- Craig, M. A., & Bodenhausen, G. V. (2018). Category (non)fit modulates extrapolative stereotyping of multiply categorizable social targets. *Social Cognition, 35*, 559-588.

- Crandall, C. S., & Eshleman, A. (2003). A justification-suppression model of the expression and experience of prejudice. *Psychological Bulletin, 129*, 414-446.
- Crenshaw, K. (1989). Demarginalizing the intersection of race and sex: A Black feminist critique of antidiscrimination doctrine, feminist theory and antiracist politics. *The University of Chicago Legal Forum, 140*, 139-167.
- Crenshaw, K. (1991). Mapping the margins: Intersectionality, identity politics, and violence against women of color. *Stanford Law Review, 43*, 1241-1299.
- Crisp, R. J., & Beck, S. R. (2005). Reducing intergroup bias: The moderating role of ingroup identification. *Group Processes & Intergroup Relations, 8*, 173-185.
- Crisp, R. J., & Hewstone, M. (1999). Differential evaluation of crossed category groups: Patterns, processes, and reducing intergroup bias. *Group Processes & Intergroup Relations, 2*, 307-333.
- Crisp, R. J., & Hewstone, M. (2007). Multiple social categorization. *Advances in Experimental Social Psychology, 39*, 163-254.
- Dasgupta, N. (2013). Implicit attitudes and beliefs adapt to situations: A decade of research on the malleability of implicit prejudice, stereotypes, and the self-concept. *Advances in Experimental Social Psychology, 47*, 233-279.
- Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology, 56*(1), 5-18.
- Devine, P. G., & Elliot, A. J. (1995). Are racial stereotypes really fading? The Princeton trilogy revisited. *Personality and Social Psychology Bulletin, 21*, 1139-1150.

- Diener, E., Fraser, S. C., Beaman, A. L., & Kelem, R. T. (1976). Effects of deindividuation variables on stealing among Halloween trick-or-treaters. *Journal of Personality and Social Psychology, 33*, 178-183.
- Doise, W., Deschamps, J. C., & Meyer, G. (1978). The accentuation of intra-category similarities. In *Differentiation between social groups. Studies in the social psychology of intergroup relations* (pp. 159-168). London: Academic Press.
- Dollard, J., Miller, N. E., Doob, L. W., Mowrer, O. H., & Sears, R. R. (1939). *Frustration and aggression*. New Haven, CT, US: Yale University Press. <https://doi.org/10.1037/10022-000>
- Dunton, B. C., & Fazio, R. H. (1997). An individual difference measure of motivation to control prejudiced reactions. *Personality and Social Psychology Bulletin, 23*, 316-326.
- Eberhardt, J. L., Davies, P. G., Purdie-Vaughns, V. J., & Johnson, S. L. (2006). Looking deathworthy: Perceived stereotypicality of Black defendants predicts capital-sentencing outcomes. *Psychological Science, 17*, 383-386.
- Eberhardt, J. L., Goff, P. A., Purdie, V. J., & Davies, P. G. (2004). Seeing Black: Race, crime, and visual processing. *Journal of Personality and Social Psychology, 87*, 876-893.
- Ebner, N., Riediger, M., & Lindenberger, U. (2010). FACES—A database of facial expressions in young, middle-aged, and older women and men: Development and validation. *Behavior Research Methods, 42*, 351-362.
- Fazio, R. H. (1990). Multiple processes by which attitudes guide behavior: The MODE model as an integrative framework. *Advances in Experimental Social Psychology, 23*, 75-109.

- Fazio, R. H., & Olson, M. A. (2014). The MODE model: Attitude—behavior processes as a function of motivation and opportunity. In J. W. Sherman, B. Gawronski, & Y. Trope (Eds.), *Dual-process theories of the social mind* (pp. 155-171). New York: Guilford.
- Fein, S., & Spencer, S. J. (1997). Prejudice as self-image maintenance: Affirming the self through derogating others. *Journal of Personality and Social Psychology*, *73*, 31-44.
- Fiske, S. T., & Neuberg, S. L. (1990). A continuum of impression formation, from category-based to individuating processes: Influences of information and motivation on attention and interpretation. In M. P. Zanna (Ed.), *Advances in Experimental Social Psychology* (Vol. 23, pp. 1-74). Academic Press. [https://doi.org/10.1016/S0065-2601\(08\)60317-2](https://doi.org/10.1016/S0065-2601(08)60317-2)
- Forscher, P. S., Cox, W. T. L., Graetz, N., & Devine, P. G. (2015). The motivation to express prejudice. *Journal of Personality and Social Psychology*, *109*(5), 791-812.
- Freeman, J. B., & Ambady, N. (2011). A dynamic interactive theory of person construal. *Psychological Review*, *118*, 247-279.
- Freeman, J. B., Penner, A. M., Saperstein, A., Scheutz, M., & Ambady, N. (2011). Looking the part: Social status cues shape race perception. *PLOS ONE*, *6*, e25107.
- Freeman, J. B., Stoller, R. M., & Brooks, J. A. (2020). Dynamic interactive theory as a domain-general account of social perception. *Advances in Experimental Social Psychology*, *63*, 237-287.
- Freud, S. (1922). *Group psychology and the analysis of the ego*. London: Hogarth Press.
- Gaertner, S. L., & Dovidio, J. F. (2000). *Reducing intergroup bias: The common ingroup identity model*. New York, NY, US: Psychology Press.

- Galinsky, A. D., Hall, E. V., & Cuddy, A. J. C. (2013). Gendered races: Implications for interracial marriage, leadership selection, and athletic participation. *Psychological Science, 24*, 498-506.
- Gelfand, M. J., Nishii, L., & Raver, J. L. (2006). On the nature and importance of cultural tightness-looseness. *Journal of Applied Psychology, 91*, 1225-1244.
- Gelfand, M. J., Raver, J. L., Nishii, L., Leslie, L. M., Lun, J., Lim, B. C., ... Yamaguchi, S. (2011). Differences between tight and loose cultures: A 33-nation study. *Science, 6033*, 1100-1104.
- Ghavami, N., Katsiaficas, D., & Rogers, L. O. (2016). Toward an intersectional approach in developmental science. The role of race, gender, sexual orientation, and immigrant status. *Advances in Child Development and Behavior, 50*, 31-73.
- Ghavami, N., & Peplau, L. A. (2013). An intersectional analysis of gender and ethnic stereotypes: Testing three hypotheses. *Psychology of Women Quarterly, 37*, 113-127.
- Ghavami, N., & Peplau, L. A. (2018). Urban middle school students' stereotypes at the intersection of sexual orientation, ethnicity, and gender. *Child Development, 89*, 881-896.
- Glaser, J., & Knowles, E. D. (2008). Implicit motivation to control prejudice. *Journal of Experimental Social Psychology, 44*, 164-172.
- Goff, P. A., Thomas, M. A., & Jackson, M. C. (2008). "Ain't I a woman?": Towards an intersectional approach to person perception and group-based harms. *Sex Roles, 59*, 392-403.

- Greenwald, A. G., Poehlman, T.A., Uhlmann, E. L., & Banaji, M. R. (2009). Understanding and using the implicit association test: III. Meta-analysis of predictive validity. *Journal of Personality and Social Psychology, 97*, 17-41.
- Hall, E. V., Galinsky, A. D., & Phillips, K. W. (2015). Gender profiling: A gendered race perspective on person-position fit. *Personality and Social Psychology Bulletin, 41*, 853-868.
- Hall, E. V., Hall, A. V., Galinsky, A. D., & Phillips, K. W. (2019). MOSAIC: A model of stereotyping through associated and intersectional categories. *Academy of Management Review, 44*, 643-672.
- Hamilton, D. L., & Gifford, R. K. (1976). Illusory correlation in interpersonal perception: A cognitive basis of stereotypic judgments. *Journal of Experimental Social Psychology, 12*, 392-407.
- Haslam, S. A., Oakes, P. J., Reynolds, K. J., & Turner, J. C. (1999). Social identity salience and the emergence of stereotype consensus. *Personality and Social Psychology Bulletin, 25*, 809-818.
- Haslam, S. A., & Turner, J. C. (1992). Context-dependent variation in social stereotyping 2: The relationship between frame of reference, self-categorization and accentuation. *European Journal of Social Psychology, 22*, 251-277.
- Haslam, S. A., Turner, J. C., Oakes, P. J., McGarty, C., & Hayes, B. K. (1992). Context-dependent variation in social stereotyping 1: The effects of intergroup relations as mediated by social change and frame of reference. *European Journal of Social Psychology, 22*(1), 3-20.

- Hewstone, M. (1994). Revision and change of stereotypic beliefs: In search of the elusive subtyping model. *European Review of Social Psychology*, 5, 69-109.
- Higgins, E. T. (1996). Knowledge activation: Accessibility, applicability, and salience. In *Social psychology: Handbook of basic principles* (pp. 133-168). New York, NY, US: Guilford Press.
- Hogg, M. A. (2001). A social identity theory of leadership. *Personality and Social Psychology Review*, 5, 184-200.
- Hogg, M. A. (2007). Uncertainty-Identity Theory. *Advances in Experimental Social Psychology*, 39, 69-126. [https://doi.org/10.1016/S0065-2601\(06\)39002-8](https://doi.org/10.1016/S0065-2601(06)39002-8)
- Hogg, M. A., Abrams, D., Otten, S., & Hinkle, S. (2004). The social identity perspective: Intergroup relations, self-conception, and small groups. *Small Group Research*, 35, 246-276.
- Hong, Y.-Y., Chiu, C.-Y., & Kung, T. M. (1997). Bringing culture out in front: Effects of cultural meaning system activation on social cognition. In K. Leung, Y. Kashima, U. Kim, & S. Yamaguchi (Eds.), *Progress in Asian social psychology* (Vol. 1, pp. 135-146). Singapore: Wiley.
- hooks, bell. (1984). *Feminist theory: From margin to center*. Pluto Press.
- Hornsey, M. J. (2008). Social identity theory and self-categorization theory: A historical review. *Social and Personality Psychology Compass*, 2, 204-222.
- Hosoda, M., Stone, D. L., & Stone-Romero, E. F. (2003). The Interactive effects of race, gender, and job type on job suitability ratings and selection decisions. *Journal of Applied Social Psychology*, 33, 145-178.

- Hugenberg, K. (2005). Social categorization and the perception of facial affect: Target race moderates the response latency advantage for happy faces. *Emotion, 5*, 267-276.
- Hugenberg, K., & Bodenhausen, G. V. (2003). Facing prejudice: Implicit prejudice and the perception of facial threat. *Psychological Science, 14*, 640-643.
- Johnson, K. L., Freeman, J. B., & Pauker, K. (2012). Race is gendered: How covarying phenotypes and stereotypes bias sex categorization. *Journal of Personality and Social Psychology, 102*, 116-131.
- Johnson, K. L., & Ghavami, N. (2011). At the crossroads of conspicuous and concealable: What race categories communicate about sexual orientation. *PLOS ONE, 6*, e18025.
- Jones, C. S., & Kaplan, M. F. (2003). The effects of racially stereotypical crimes on juror decision-making and information-processing strategies. *Basic and Applied Social Psychology, 25*, 1-13.
- 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.
- Judd, C. M., Westfall, J., & Kenny, D. A. (2012). Treating stimuli as a random factor in social psychology: A new and comprehensive solution to a pervasive but largely ignored problem. *Journal of Personality and Social Psychology, 103*, 54-69.
- Kang, S. K., & Bodenhausen, G. V. (2015). Multiple identities in social perception and interaction: Challenges and opportunities. *Annual Review of Psychology, 66*, 547-574.

- Kang, S. K., & Chasteen, A. L. (2009). Beyond the double-jeopardy hypothesis: Assessing emotion on the faces of multiply-categorizable targets of prejudice. *Journal of Experimental Social Psychology, 45*, 1281-1285.
- Kang, S. K., Chasteen, A. L., Cadieux, J., Cary, L. A., & Syeda, M. (2014). Comparing young and older adults' perceptions of conflicting stereotypes and multiply-categorizable individuals. *Psychology and Aging, 29*, 469-481.
- Kanter, R. M. (1977). Some effects of proportions on group life: Skewed sex ratios and responses to token women. *American Journal of Sociology, 82*, 965-990.
- Katz, D., & Braly, K. (1933). Racial stereotypes of one hundred college students. *The Journal of Abnormal and Social Psychology, 28*, 280-290.
- Kawakami, K., Amodio, D. M., & Hugenberg, K. (2017). Intergroup perception and cognition: An integrative framework for understanding the causes and consequences of social categorization. *Advances in Experimental Social Psychology, 55*, 1-80.  
<https://doi.org/10.1016/bs.aesp.2016.10.001>
- King, D. K. (1988). Multiple jeopardy, multiple consciousness: The context of a Black feminist ideology. *Signs: Journal of Women in Culture and Society, 14*, 42-72.
- Klauer, K. C., Hölzenbein, F., Calanchini, J., & Sherman, J. W. (2014). How malleable is categorization by race? Evidence for competitive category use in social categorization. *Journal of Personality and Social Psychology, 107*, 21-40.
- Krieger, L. H., & Fiske, S. T. (2006). Behavioral realism in employment discrimination law: Implicit bias and disparate treatment. *California Law Review, 94*, 997-1062.

- Kruglanski, A. W. (2012). Lay epistemic theory. In P. A. M. Van Lange, A. W. Kruglanski, & E. T. Higgins (Eds.), *Handbook of theories of social psychology: Vol. 1* (pp. 201-223). London, UK: Sage Publications
- 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.
- Kunda, Z., Miller, D. T., & Claire, T. (1990). Combining social concepts: The role of causal reasoning. *Cognitive Science*, *14*, 551-577.
- Kunda, Z., & Thagard, P. (1996). Forming impressions from stereotypes, traits, and behaviors: A parallel-constraint-satisfaction theory. *Psychological Review*, *103*, 284-308.
- Kurzban, R., Tooby, J., & Cosmides, L. (2001). Can race be erased? Coalitional computation and social categorization. *Proceedings of the National Academy of Sciences of the United States of America*, *98*, 15387-15392.
- Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2016). lmerTest: Tests in linear mixed effects models. *R package version 2.0-33*. <https://CRAN.R-project.org/package=lmerTest>
- Lakens, D. (2017). Equivalence tests: A practical primer for *t*-tests, correlations, and meta-analyses. *Social Psychological and Personality Science*, *8*, 355-362.
- Le Bon, G. (1908). *The crowd: A study of the popular mind*. London: Unwin.
- Lei, R. F., & Bodenhausen, G. V. (2017). Racial assumptions color the mental representation of social class. *Frontiers in Psychology*, *8*, 1-7. <https://doi.org/10.3389/fpsyg.2017.00519>

- Levin, S., Sinclair, S., Veniegas, R. C., & Taylor, P. L. (2002). Perceived discrimination in the context of multiple group memberships. *Psychological Science, 13*, 557-560.
- Leyens, J-P., Yzerbyt, V. Y., & Schadron, G. (1992). The social judgeability approach to stereotypes. *European Review of Social Psychology, 3*, 91-120.
- Leys, C., Ley, C., Klein, O., Bernard, P., Licata, L. (2013). Do not use standard deviation around the mean, use absolute deviation around the median. *Journal of Experimental Social Psychology, 49*, 764-766.
- Livingston, R. W., & Pearce, N. A. (2009). The teddy-bear effect: Does having a baby face benefit Black chief executive officers? *Psychological Science, 20*, 1229-1236.
- Livingston, R. W., Rosette, A. S., & Washington, E. F. (2012). Can an agentic Black woman get ahead? The impact of race and interpersonal dominance on perceptions of female leaders. *Psychological Science, 23*, 354-358.
- Lundberg, G. J. W., Neel, R., Lassetter, B., & Todd, A. R. (2018). Racial bias in implicit danger associations generalizes to older male targets. *PLOS ONE, 13*, e0197398.
- 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.
- Macrae, C. N., Bodenhausen, G. V., & Milne, A. B. (1995). The dissection of selection in person perception: Inhibitory processes in social stereotyping. *Journal of Personality and Social Psychology, 69*, 397-407.
- McDonald, M. M., Sidanius, J., & Navarrete, C. D. (2011). Developing a theory of gendered prejudice: An evolutionary and social dominance perspective. In R. M. Kramer, G. J.

- Leonardelli, & R. W. Livingston (Eds.), *Social cognition, social identity, and intergroup relations* (pp. 189-220). New York: Psychology Press.
- McGuire, W. J., McGuire, C. V., Child, P., & Fujioka, T. (1978). Salience of ethnicity in the spontaneous self-concept as a function of one's ethnic distinctiveness in the social environment. *Journal of Personality and Social Psychology, 36*, 511-520.
- Milgram, S. (1963). Behavioral Study of obedience. *The Journal of Abnormal and Social Psychology, 67*, 371-378.
- Miller, D. I., Eagly, A. H., & Linn, M. C. (2015). Women's representation in science predicts national gender-science stereotypes: Evidence from 66 nations. *Journal of Educational Psychology, 107*, 631-644.
- Monteith, M. J. (1993). Self-regulation of prejudiced responses: Implications for progress in prejudice-reduction efforts. *Journal of Personality and Social Psychology, 65*, 469-485.
- Mussweiler, T., Gabriel, S., & Bodenhausen, G. V. (2000). Shifting social identities as a strategy for deflecting threatening social comparisons. *Journal of Personality and Social Psychology, 79*, 398-409.
- Nelson, L. J., & Miller, D. T. (1995). The distinctiveness effect in social categorization: You are what makes you unusual. *Psychological Science, 6*, 246-249.
- Nicolas, G., de la Fuente, M., & Fiske, S. T. (2017). Mind the overlap in multiple categorization: A review of crossed categorization, intersectionality, and multiracial perception. *Group Processes & Intergroup Relations, 20*, 621-631.

- Niemann, Y. F., Jennings, L., Rozelle, R. M., Baxter, J. C., & Sullivan, E. (1994). Use of free responses and cluster analysis to determine stereotypes of eight groups. *Personality and Social Psychology Bulletin*, *20*, 379-390.
- Nosek, B. A., Bar-Anan, Y., Sriram, N., Axt, J., & Greenwald, A. G. (2014). Understanding and using the brief implicit association test: Recommended scoring procedures. *PLoS One*, *9*, e110938.
- Nosek, B. A., Smyth, F. L., Sriram, N., Lindner, N. M., Devos, T., Ayala, A., ... Greenwald, A. G. (2009). National differences in gender-science stereotypes predict national sex differences in science and math achievement. *Proceedings of the National Academy of Sciences*, *106*, 10593-10597.
- Oakes, P. J. (1994). The effects of fit versus novelty on the salience of social categories: A response to Biernat and Vescio (1993). *Journal of Experimental Social Psychology*, *30*, 390-398.
- 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, *Rediscovering the social group* (pp. 117-141). Oxford, UK: Blackwell.
- Oakes, P. J., Turner, J. C., & Haslam, S. A. (1991). Perceiving people as group members: The role of fit in the salience of social categorizations. *British Journal of Social Psychology*, *30*, 125-144.
- Öhman, A., Juth, P., & Lundqvist, D. (2010). Finding the face in a crowd: Relationships between distractor redundancy, target emotion, and target gender. *Cognition and Emotion*, *24*, 1216-1228.

- Öhman, A., Lundqvist, D., & Esteves, F. (2001). The face and the crowd revisited: A threat advantage with schematic stimuli. *Journal of Personality and Social Psychology, 80*, 381-396.
- Palma, T. A., Garcia-Marques, L., Marques, P., Hagá, S., & Payne, B. K. (2019). Learning what to inhibit: The influence of repeated testing on the encoding of gender and age information. *Journal of Personality and Social Psychology, 116*, 899-918.
- Payne, B., & Whittington, F. (1976). Older women: An examination of popular stereotypes and research evidence. *Social Problems, 23*, 488-504.
- Payne, B. K. (2001). Prejudice and perception: The role of automatic and controlled processes in misperceiving a weapon. *Journal of Personality and Social Psychology, 81*, 181-192.
- Pedulla, D. S. (2014). The positive consequences of negative stereotypes: Race, sexual orientation, and the job application process. *Social Psychology Quarterly, 77*, 75-94.
- Penner, A. M., & Saperstein, A. (2008). How social status shapes race. *Proceedings of the National Academy of Sciences, 105*, 19628-19630.
- Perszyk, D. R., Lei, R. F., Bodenhausen, G. V., Richeson, J. A., & Waxman, S. R. (2019). Bias at the intersection of race and gender: Evidence from preschool-aged children. *Developmental Science*, e12788. <http://doi.org/10.1111/desc.12788>
- Petrie, M., & Roman, P. M. (2004). Race and gender differences in workplace autonomy: A research note. *Sociological Inquiry, 74*, 590-603.
- Petsko, C. D., & Bodenhausen, G. V. (2020). Multifarious person perception: How social perceivers manage the complexity of intersectional targets. *Social and Personality Psychology Compass, 14*, e12518.

- Petsko, C. D., & Bodenhausen, G. V. (2019). Race-crime congruency effects revisited: Do we take defendants' sexual orientation into account? *Social Psychological and Personality Science, 10*, 73-81.
- Petsko, C. D., & Bodenhausen, G. V. (2019). Racial stereotyping of gay men: Can a minority sexual orientation erase race? *Journal of Experimental Social Psychology, 83*, 37-54.
- Pietraszewski, D. (2018). A reanalysis of crossed-dimension "who said what?" paradigm studies, using a better error base-rate correction. *Evolution and Human Behavior, 39*, 479-489.
- Pietraszewski, D., Curry, O. S., Petersen, M. B., Cosmides, L., & Tooby, J. (2015). Constituents of political cognition: Race, party politics, and the alliance detection system. *Cognition, 140*, 24-39.
- Pittinsky, T. L., Shih, M. J., & Trahan, A. (2006). Identity cues: Evidence from and for intra-individual perspectives on positive and negative stereotyping. *Journal of Applied Social Psychology, 36*, 2215-2239.
- Plant, E. A., & Devine, P. G. (1998). Internal and external motivation to respond without prejudice. *Journal of Personality and Social Psychology, 75*, 811-832.
- Plaut, V. C. (2010). Diversity science: Why and how difference makes a difference. *Psychological Inquiry, 21*, 77-99.
- Popper, K. R. (1959). *The logic of scientific discovery*. Oxford, England: Basic Books.
- Purdie-Vaughns, V., & Eibach, R. P. (2008). Intersectional invisibility: The distinctive advantages and disadvantages of multiple subordinate-group identities. *Sex Roles: A Journal of Research, 59*, 377-391.

- Ramírez-Esparza, N., Gosling, S. D., Benet-Martínez, V., Potter, J. P., & Pennebaker, J. W. (2006). Do bilinguals have two personalities? A special case of cultural frame switching. *Journal of Research in Personality, 40*, 99-120.
- Rattan, A., Steele, J., & Ambady, N. (2019). Identical applicant but different outcomes: The impact of gender versus race salience in hiring. *Group Processes & Intergroup Relations, 22*, 80-97.
- Reid, A., & Deaux, K. (1996). Relationship between social and personal identities: Segregation or integration? *Journal of Personality and Social Psychology, 71*, 1084-1091.
- Remedios, J. D., Chasteen, A. L., Rule, N. O., & Plaks, J. E. (2011). Impressions at the intersection of ambiguous and obvious social categories: Does gay+Black=likable? *Journal of Experimental Social Psychology, 47*, 1312-1315.
- Reynolds, W. M. (1982). Development of reliable and valid short forms of the Marlowe-Crowne social desirability scale. *Journal of Clinical Psychology, 38*, 119-125.
- Roccas, S., & Brewer, M. B. (2002). Social identity complexity. *Personality and Social Psychology Review, 6*(2), 88-106.
- Rosette, A. S., & Livingston, R. W. (2012). Failure is not an option for Black women: Effects of organizational performance on leaders with single versus dual-subordinate identities. *Journal of Experimental Social Psychology, 48*, 1162-1167.
- Rosette, A. S., Ponce de Leon, R., Koval, C. Z., & Harrison, D. A. (2018). Intersectionality: Connecting experiences of gender with race at work. *Research in Organizational Behavior, 38*, 1-22.

- Rule, N. O., Garrett, J. V., & Ambady, N. (2010). On the perception of religious group membership from faces. *PLoS ONE*, *5*, e14241.
- Rule, N. O., Macrae, C. N., & Ambady, N. (2009). Ambiguous group membership is extracted automatically from faces. *Psychological Science*, *20*, 441-443.
- Schadron, G., & Yzerbyt, V. Y. (1991). Social judgeability: Another framework for the study of social inference. *European Bulletin of Cognitive Psychology*, *11*, 229-258.
- Schug, J., Alt, N. P., & Klauer, K. C. (2015). Gendered race prototypes: Evidence for the non-prototypicality of Asian men and Black women. *Journal of Experimental Social Psychology*, *56*, 121-125.
- Schwartz, S. H. (1992). Universals in the content and structure of values: Theoretical advances and empirical tests in 20 countries. *Advances in Experimental Social Psychology*, *25*, 1-65.
- Sesko, A. K., & Biernat, M. (2018). Invisibility of Black women: Drawing attention to individuality. *Group Processes & Intergroup Relations*, *21*, 141-158.
- Sesko, A. K., & Biernat, M. (2010). Prototypes of race and gender: The invisibility of Black women. *Journal of Experimental Social Psychology*, *46*, 356-360.
- Shasteen, J.R., Sasson, N. J., & Pinkham, A. E. (2015). A detection advantage for facial threat in the absence of anger. *Emotion*, *15*, 837-845.
- Shasteen, J. R., Sasson, N. J., & Pinkham, A. E. (2014). Eye tracking the face in the crowd task: Why are angry faces found more quickly? *PLoS One*, e93914.
- Sherif, M. (1966). *In common predicament: Social psychology of intergroup conflict and cooperation*. Boston: Houghton-Mifflin.
- Sherif, M. (1967). *Group conflict and cooperation*. London: Routledge & Kegan Paul.

- Sidanius, J., & Pratto, F. (2012). Social dominance theory. In P. A. M. Lange, A. W. Kruglanski, & E. T. Higgins (Eds.), *Handbook of theories of social psychology* (pp. 418-438). London: Sage Publications.
- Sinclair, L., & Kunda, Z. (1999). Reactions to a black professional: Motivated inhibition and activation of conflicting stereotypes. *Journal of Personality and Social Psychology, 77*, 885-904.
- Skinner, A. L., & Rae, J. R. (2019). A robust bias against interracial couples among White and Black respondents, relative to multiracial respondents. *Social Psychological and Personality Science, 10*, 823-831.
- Stangor, C., Lynch, L., Duan, C., & Glas, B. (1992). Categorization of individuals on the basis of multiple social features. *Journal of Personality and Social Psychology, 62*, 207-218.
- Stephens, N. M., Markus, H. R., & Phillips, L. T. (2014). Social class culture cycles: How three gateway contexts shape selves and fuel inequality. *Annual Review of Psychology, 65*, 611-634.
- Strohinger, N., Gray, K., Chituc, V., Heffner, J., Schein, C., & Heagins, T. B. (2015). The MR2: A multi-racial, mega-resolution database of facial stimuli. *Behavior Research Methods, 48*, 1197-1204.
- Tajfel, H. (1959). Quantitative judgement in social perception. *British Journal of Psychology, 50*, 16-29.
- Tajfel, H. (1969). Social and cultural factors in perception. In G. Lindzey & E. Aronson (Eds.), *Handbook of social psychology, Vol 3*. Reading, MA: Addison-Wesley.

- Tajfel, H., Billig, M. G., Bundy, R. P., & Flament, C. (1971). Social categorization and intergroup behaviour. *European Journal of Social Psychology, 1*, 149-178.
- Tajfel, H., Sheikh, A. A., & Gardner, R. C. (1964). Content of stereotypes and the inference of similarity between members of stereotyped groups. *Acta Psychologica, 22*, 191-201.
- Tajfel, H., & Turner, J. C. (1979). An integrative theory of intergroup conflict. In W. G. Austin & S. Worchel (Eds.), *The social psychology of intergroup relations*. Monterey, CA: Brooks-Cole.
- Tajfel, H., & Turner, J. C. (1986). The social identity theory of intergroup behavior. In S. Worchel & W. G. Austin (Eds.), *Psychology of intergroup relations*. Chicago: Nelson-Hall.
- Taylor, S. E., Fiske, S. T., Etcoff, N. L., & Ruderman, A. J. (1978). Categorical and contextual bases of person memory and stereotyping. *Journal of Personality and Social Psychology, 36*, 778-793.
- Todd, A. R., Simpson, A. J., Thiem, K. C., & Neel, R. (2016). The generalization of implicit racial bias to young Black boys: Automatic stereotyping or automatic prejudice? *Social Cognition, 34*, 306-323.
- Todd, A. R., Thiem, K. C., & Neel, R. (2016). Does seeing faces of young Black boys facilitate the identification of threatening stimuli? *Psychological Science, 27*, 384-393.
- Tooby, J., & Cosmides, L. (2010). Groups in mind: The coalitional roots of war and morality. In H. Høgh-Olesen (Ed.), *Human Morality and Sociality: Evolutionary and Comparative Perspectives* (pp. 91-234). Palgrave-Macmillan.
- Tulving, E., & Pearlstone, Z. (1996). Availability versus accessibility of information in memory for words. *Journal of Verbal Learning and Verbal Behavior, 5*, 381-391.

Turner, J. C., Hogg, M. A., Oakes, P. J., Reicher, S. D., & Wetherell, M. S. (1987).

*Rediscovering the social group: A self-categorization theory.* Cambridge, MA, US: Basil Blackwell.

Urban, L. M., & Miller, N. (1998). A theoretical analysis of crossed categorization effects: A meta-analysis. *Journal of Personality and Social Psychology*, *74*, 894-908.

van Rijswijk, W., & Ellemers, N. (2002). Context effects on the application of stereotype content to multiple categorizable targets. *Personality and Social Psychology Bulletin*, *28*, 90-101.

van Twuyver, M., & van Knippenberg, A. (1995). Social categorization as a function of priming. *European Journal of Social Psychology*, *25*, 695-701.

Volpert-Esmond, H. I., & Bartholow, B. (2019). Explicit categorization goals affect attention-related processing of race and gender during person construal. *Journal of Experimental Social Psychology*, *85*. <https://doi.org/10.1016/j.jesp.2019.103839>

Williams, M. A., & Mattingly, J. B. (2006). Do angry men get noticed? *Current Biology*, *16*, PR402-PR404.

Wilson, J. P., Remedios, J. D., & Rule, N. O. (2017). Interactive effects of obvious and ambiguous social categories on perceptions of leadership: When double-minority status may be beneficial. *Personality and Social Psychology Bulletin*, *43*, 888-900.

Yamaguchi, M., & Beattie, G. (2020). The role of explicit categorization in the Implicit Association Test. *Journal of Experimental Psychology: General*, *149*, 809-827

Yap, M., & Konrad, A. M. (2009). Gender and racial differentials in promotions: Is there a sticky floor, a mid-level bottleneck, or a glass ceiling? *Industrial Relations*, *64*, 593-619.

- Yoder, J. D. (1991). Rethinking tokenism: Looking beyond numbers. *Gender & Society*, 5, 178-192.
- Yoder, J. D., Adams, J., & Prince, H. T. (1983). The price of a token. *Journal of Political and Military Sociology*, 11, 325-337.
- Zarate, M. A., & Smith, E. R. (1990). Person categorization and stereotyping. *Social Cognition*, 8, 161-185.
- Zou, L. X., & Cheryan, S. (2017). Two axes of subordination: A new model of racial position. *Journal of Personality and Social Psychology*, 112, 696-717.

## Appendix A

## Conversation Topic Pretesting (Exp. 1a)

A sample MTurk workers ( $N = 56$ ) pre-tested a series of conversation stances (e.g., zoos are unethical, Americans drink too much caffeine, etc.)—all taken from a high-school debate team website—on whether men *vs.* women would agree with them more (from 1 = *men would agree more* to 7 = *women would agree more*), and on whether young *vs.* older adults would agree with them more (from 1 = *young adults would agree more* to 7 = *older adults would agree more*).

Of the stances that were pretested, two were chosen based on these ratings. According to these ratings, the stance “Introverts are more cerebral than extraverts” was no more associated with men than with women (according to a one-sample *t*-test:  $t(55) = 1.32, p = .19$ ), and it was associated no more with older adults than with young adults (according to a one-sample *t*-test:  $t(55) = -0.29, p = .77$ ). In addition, the stance “People who mirror whomever they’re talking to are inauthentic” was no more associated with men than with women (one-sample *t*-test:  $t(54) = 0.43, p = .67$ ), and it was no more associated with older adults than with young adults (one-sample *t*-test:  $t(55) = -0.49, p = .63$ ).

## Appendix B

### “Oldness” Ratings of Checklist Attributes

A sample of MTurk workers ( $N = 81$ ) rated all 99 checklist traits on the degree to which they seem stereotypic of old people. Of these participants, I excluded  $n = 4$  (4.94%) for not responding “yes” to the question, “Did you take this survey seriously?” Participants rated all 99 traits, in a randomized order on how stereotypic they seemed, from  $1 = not\ at\ all$  to  $7 = very\ much$ .

Average “oldness” scores were created for each attribute, which were then imputed into participants’ trait selections in Experiment 1a. For illustrative purposes the “oldest” 10 words in the checklist were, in order: *loyal to family ties, tradition-loving, stubborn, conservative, conventional, very religious, faithful, straightforward, honest, and practical*. The 10 “least old” words in the checklist were, from “least” to “more old”: *athletic, criminal, violent, sexually perverse, mercenary, sensual, progressive, radical, treacherous, gluttonous, and aggressive*.

## Appendix C

## Results by Conversation Topic (Exp. 1a)

In Exp. 1a, age categorization is assessed by analyzing participants' errors in a 2(error type: within-age-group, between-age-group)  $\times$  3 (condition: age fit, control gender fit) mixed ANOVA with repeated measures on the first factor. Here, I analyze these same errors as a 2(error type: within-age-group, between-age-group)  $\times$  3 (condition: age fit, control, gender fit)  $\times$  2 (conversation topic: introversion, phoniness) mixed ANOVA with repeated measures on the first factor.

This analysis reveals the exact same pattern of results reported in the manuscript: a main effect of error type suggesting that participants are engaging in age categorization, on average ( $M_{\text{diff}} = 4.67$ , 95% CI[4.07, 5.26],  $\beta = 0.93$ ,  $F(1, 295) = 235.04$ ,  $p < .001$ ); and an interaction between error type and fit condition ( $F(1, 295) = 235.93$ ,  $p < .001$ ,  $\omega_p^2 = 0.44$ ). The interaction was almost identical in magnitude as that reported in the manuscript and was not moderated by which conversation topic participants were assigned to see [ $F(1, 295) = 0.03$ ,  $p = .86$ ,  $\omega_p^2 < 0.01$ ]. The only influence of conversation topic on participants' responses was a main effect: participants made slightly more total errors in the who-said-what task in the introversion conversation condition than in the phoniness conversation condition:  $M_{\text{diff}} = 0.64$ , 95% CI[0.02, 1.25],  $\beta = 0.13$ ,  $F(1, 295) = 4.15$ ,  $p = .042$ .

In Exp. 1a, gender categorization is assessed by analyzing participants' errors in a 2(error type: within-gender-group, between-gender-group)  $\times$  3 (condition: age fit, control gender fit) mixed ANOVA with repeated measures on the first factor. Here, I analyze these same errors as a 2(error type: within-gender-group, between-gender-group)  $\times$  3 (condition: age fit, control, gender

fit)  $\times$  2 (conversation topic: introversion, phoniness) mixed ANOVA with repeated measures on the first factor.

This analysis reveals the exact same pattern of results reported in the manuscript: a main effect of error type suggesting that participants are engaging in gender categorization, on average ( $M_{\text{diff}} = 5.09$ , 95% CI[4.51, 5.67],  $\beta = 0.99$ ,  $F(1, 295) = 294.02$ ,  $p < .001$ ); and an interaction between error type and fit condition ( $F(1, 295) = 215.41$ ,  $p < .001$ ,  $\omega_p^2 = 0.42$ ). The nature of this interaction was that there was substantially greater gender categorization in the gender-fit condition relative to the other two conditions. This interaction, however, *was* different in magnitude depending on which conversation topic participants were assigned ( $F(1, 295) = 21.96$ ,  $p < .001$ ,  $\omega_p^2 = 0.07$ ). Deconstructing this interaction reveals that the influence of the gender-fit condition (*vs.* the other two conditions) was weaker when the conversation topic was about introversion [ $F(1, 295) = 48.61$ ,  $p < .001$ ,  $\omega_p^2 = 0.14$ ] than when it was about the phoniness of self-monitors [ $F(1, 295) = 192.60$ ,  $p < .001$ ,  $\omega_p^2 = 0.39$ ]. See the supplemental figure below for gender categorization strength by condition.

## Appendix D

## Conversation Topic Pretesting (Exp. 1b)

A sample MTurk workers ( $N = 54$ ) pre-tested a series of conversation stances (e.g., everyone should be vegetarian, obesity is a disease, etc.)—all taken from a high-school debate team website—on whether men *vs.* women would agree with them more (from 1 = *men would agree more* to 7 = *women would agree more*), and on whether White *vs.* Black Americans would agree with them more (from 1 = *White Americans would agree more* to 7 = *Black Americans would agree more*).

Of the stances that were pretested, two were chosen based on these ratings. According to these ratings, the stance “Committing suicide should be made illegal” was no more associated with men than with women (according to a one-sample  $t$ -test:  $t(53) = 0.89, p = .38$ ), and it was associated no more with White Americans than with Black Americans (according to a one-sample  $t$ -test:  $t(52) = 0.57, p = .57$ ). In addition, the stance “Celebrities earn too much money” was no more associated with men than with women (one-sample  $t$ -test:  $t(52) = 0.81, p = .42$ ), and it was no more associated with White Americans than with Black Americans (one-sample  $t$ -test:  $t(52) = -0.34, p = .74$ ).

## Appendix E

## Results by Conversation Topic (Exp. 1b)

In Exp. 1b, race categorization is assessed by analyzing participants' errors in a 2(error type: within-race-group, between-race-group)  $\times$  3 (condition: race fit, control gender fit) mixed ANOVA with repeated measures on the first factor. Here, I analyze these same errors as a 2(error type: within-race-group, between-race-group)  $\times$  3 (condition: race fit, control, gender fit)  $\times$  2 (conversation topic: celebrities, suicide) mixed ANOVA with repeated measures on the first factor.

This analysis yields the two same results that I report in the manuscript. Namely, across conditions there is a pronounced tendency for participants to engage in race categorization [ $M_{\text{diff}} = 4.00$ , 95% CI[3.39, 4.61],  $\beta = 0.78$ ,  $F(1, 568) = 168.01$ ,  $p < .001$ ], and the degree of race categorization depends on whether or not participants were in the race-fit condition [ $F(1, 568) = 348.94$ ,  $p < .001$ ,  $\omega_p^2 = 0.38$ ]. However, the magnitude of the amount by being in the race-fit condition amplified race categorization was contingent on which conversation topic participants were assigned [ $F(1, 568) = 8.29$ ,  $p = .004$ ,  $\omega_p^2 = 0.01$ ]. The nature of this interaction was that the influence of race-fit (vs. other conditions) on participants' race-categorization was slightly weaker in the celebrity topic condition [ $F(1, 568) = 121.33$ ,  $p < .001$ ,  $\omega_p^2 = 0.17$ ] than in the suicide topic condition [ $F(1, 568) = 239.32$ ,  $p < .001$ ,  $\omega_p^2 = 0.30$ ].

In Exp. 1b, gender categorization is assessed by analyzing participants' errors in a 2(error type: within-gender-group, between-gender-group)  $\times$  3 (condition: race fit, control gender fit) mixed ANOVA with repeated measures on the first factor. Here, I analyze these same errors as a 2(error type: within-gender-group, between-gender-group)  $\times$  3 (condition: race fit, control,

gender fit)  $\times$  2 (conversation topic: celebrities, suicide) mixed ANOVA with repeated measures on the first factor.

This analysis yields the two same results that I report in the manuscript. Across conditions there is a pronounced tendency for participants to engage in gender categorization [ $M_{\text{diff}} = 6.13$ , 95% CI[5.53, 6.74],  $\beta = 1.12$ ,  $F(1, 284) = 399.24$ ,  $p < .001$ ], and the degree of gender categorization depends on whether or not participants were in the gender-fit condition [ $F(1, 284) = 233.33$ ,  $p < .001$ ,  $\omega_p^2 = 0.45$ ]. However, the magnitude of the amount by being in the gender-fit condition amplified gender categorization was contingent on which conversation topic participants were assigned [ $F(1, 284) = 22.98$ ,  $p < .001$ ,  $\omega_p^2 = 0.07$ ]. The nature of this interaction was that the influence of gender-fit (vs. other conditions) on participants' gender-categorization was weaker in the celebrity topic condition [ $F(1, 284) = 55.16$ ,  $p < .001$ ,  $\omega_p^2 = 0.16$ ] than in the suicide topic condition [ $F(1, 284) = 200.57$ ,  $p < .001$ ,  $\omega_p^2 = 0.41$ ].