University of Denver

Digital Commons @ DU

Electronic Theses and Dissertations

Graduate Studies

2021

Anger Bias in the Evaluation of Crowds

Diana Mihalache
University of Denver

Follow this and additional works at: <https://digitalcommons.du.edu/etd>

Part of the [Other Psychology Commons](#), and the [Social Psychology Commons](#)

Recommended Citation

Mihalache, Diana, "Anger Bias in the Evaluation of Crowds" (2021). *Electronic Theses and Dissertations*. 1961.

<https://digitalcommons.du.edu/etd/1961>

This Dissertation is brought to you for free and open access by the Graduate Studies at Digital Commons @ DU. It has been accepted for inclusion in Electronic Theses and Dissertations by an authorized administrator of Digital Commons @ DU. For more information, please contact jennifer.cox@du.edu, dig-commons@du.edu.

ANGER BIAS IN THE EVALUATION OF CROWDS

A Dissertation

Presented to

the Faculty of the College of Arts, Humanities and Social Sciences

University of Denver

In Partial Fulfillment

of the Requirements for the Degree

Doctor of Philosophy

by

Diana Mihalache

August 2021

Advisor: Timothy Sweeny, Ph.D.

Author: Diana Mihalache
Title: ANGER BIAS IN THE EVALUATION OF CROWDS
Advisor: Timothy Sweeny, Ph.D.
Degree Date: August 2021

Abstract

Although people are good at classifying emotions, they also make mistakes. These errors tend to be negatively biased and potentially serve a protective function. Research on biases in emotion perception has largely focused on perception of individual faces and little is known about biases in evaluations of crowds. In the first investigation, I conducted six experiments, evaluating anger bias—a tendency to judge facial expressions as angry—in the context of single faces and emotionally homogenous crowds. I found that observers were biased to classify faces as angry, especially when evaluating crowds. This amplified bias emerged in the context of perceptual uncertainty and reached peak intensity for crowds with four members. Observers endorsed anger bias regardless of whether angry faces were discriminated against positive (happy) or negative (fearful) expressions. Anger bias persisted despite variability in identity and gender but was strongest for evaluations of male faces. In the second investigation, I conducted two experiments evaluating anger bias in the context of emotionally heterogenous crowds. Observers endorsed anger bias in the context of lower intensities of expression. Although, observers showed difficulty in accurately classifying crowds consisting of relatively balanced number of angry and happy faces displaying higher intensities of expression, these errors were not biased. In other words, anger bias emerged when judgments were difficult in the context of low perceptual information, but not in the context of clear and yet contradictory information. This series of studies provide insight into sensitivity and

bias in crowd perception and suggest that bias is amplified in crowds in the context of scarce diagnostic information.

Acknowledgments

I would like to thank Dr. Timothy Sweeny for his guidance throughout the dissertation and graduate school process. I truly appreciate the dedication, support, and humor in his approach to mentorship. I would like to express gratitude to my dissertation committee members for their insights and willingness to be a part of this process. I would also like to thank my collaborator, Dr. Stefan Uddenberg, for being there to discuss ideas and tackle challenges. Additionally, I would like to extend my appreciation to all the members of the Visual Perception, Emotion, and Cognition Lab at the University of Denver.

Lastly, I would like to thank my friends and family for their support, as well as Nietzsche, the dog, who started this journey with me in the first month of my undergraduate degree and sat by my side through every obstacle and celebration along the way.

Table of Contents

CHAPTER 1: GENERAL INTRODUCTION	1
A brief overview of facial expressions	1
Biases in perception of single faces	3
Negativity bias beyond face perception	7
Biases may be adaptive	8
Crowds	10
Present Research	15
CHAPTER 2: ANGER BIAS IN THE EVALUATION OF CROWDS	19
INTRODUCTION	19
Experiment 1: Are crowds particularly susceptible to being misperceived as angry?	22
Experiment 2: Is amplified evaluative bias for crowds especially strong for categorizations of anger?	29
Experiment 3: Is anger bias for crowds amplified by other kinds of perceptual uncertainty, or is it specific to emotional ambiguity?	32
Experiment 4: How does the strength of anger bias relate to the size of a crowd? ...	36
Experiment 5: Does anger bias occur for crowds with additional variability in identity?	38
Experiment 6: Does anger bias occur for crowds with additional variability in gender?	42
DISCUSSION	45
CHAPTER 3: IMPACT OF RATIO AND INTENSITY ON ANGER BIAS IN MIXED-EMOTION CROWDS	51
INTRODUCTION	51
Experiment 1: How is anger bias impacted by variability in proportion of angry to happy faces within a crowd?	55
Experiment 2: How does the anger bias interact with differing intensities of emotions within emotionally heterogeneous crowds?	65
DISCUSSION	72
CHAPTER 4: SUMMARY AND LIMITATIONS	78
REFERENCES	84
APPENDICES	103
APPENDIX A: INVESTIGATION 1	103
APPENDIX B: INVESTIGATION 2	116

List of Figures

CHAPTER 2

Figure 1.....	25
Figure 2.....	28
Figure 3.....	31
Figure 4.....	34
Figure 5.....	35
Figure 6.....	38
Figure 7.....	40
Figure 8.....	44

CHAPTER 3

Figure 1.....	60
Figure 2.....	62
Figure 3.....	64
Figure 4.....	70
Figure 5.....	71

CHAPTER 1: General Introduction

A brief overview of facial expressions

Facial expressions reflect morphological changes in the face and often communicate information about a person's internal affect (Ekman, 1992; Ekman, Frisen, & Hager, 2002). Researchers have outlined two dominant conceptualizations of facial expressions: categorical and dimensional. Within the categorical framework, there are discrete facial expressions, with happiness, anger, sadness, fear, disgust, and surprise as the main six that are often reported to have unique adaptive functions and are consistent across many cultures (Ekman & Cordaro, 2011; Izard, 1994; Panksepp & Watt, 2011). These expressions of emotion are typically identified based on particular changes in facial musculature, such as an upward or downward mouth, widening or narrowing of the eyes, and brow lowering or raising. Within the dimensional and constructionist frameworks, perceivers process information along the valence and arousal dimensions, which encompass the negative–positive and excitatory–calming continuums (Feldman Barrett, 2006; Russel, 1994). The observer then reconstructs the category based on context, with encoding of affect occurring prior to explicit categorization of expression.

Research suggests that perceivers are reliably able to identify facial expressions and discriminate them from each other (Calvo & Lundqvist, 2008; Elfenbein & Ambady, 2002; Tottenham et al., 2009). People are even skilled at emotion recognition when faces are scrambled and obscured (Bombari et al., 2013) or are seen in the periphery for

extremely short periods of time (Bayle, Schoendorff, Henaff, & Krolak-Salmon, 2011; Calvo & Lundqvist, 2008; Goren & Wilson, 2006; Milders, Sahraie, & Logan, 2008). People nonetheless make mistakes when classifying facial expressions, and some studies have focused on the types of confusions or systematic errors perceivers tend to make. Fear and surprise are frequently confused, (Palermo & Coltheart, 2004; Tottenham et al., 2009), sadness is occasionally confused with disgust, fear, or neutral facial expressions, and disgust is occasionally confused with anger (Palermo & Coltheart, 2004; Recio, Schacht, & Sommer, 2013; Tottenham et al., 2009).

The majority of research on classification of facial expressions has used posed and prototypical or full-blown representations of facial expressions, as well as forced-choice response tasks with predetermined categorical labels (Calvo & Nummenmaa, 2016). When perceivers are given the opportunity to freely evaluate spontaneous facial expressions, they display more difficulty with recognition and often assign more than one label (Kayyal & Russel, 2013; Motley & Camden, 1988; Naab & Russell, 2007). Perceivers also make more errors when evaluating subtle facial expressions, which are often more representative of real-world interactions (e.g., Calder, Rowland et al., 2000; Matsumoto & Hwang, 2014). Nonetheless, facial expressions are reliably discriminated, with happiness being identified more accurately and quickly compared to other emotions (Calvo & Lundqvist, 2008).

Not only are perceivers able to identify facial expressions, they may also use them to infer other information, which in turn may impact their behavior. For example, facial expressions depicting anger can signal the potential for threat as it frequently precedes violence and may influence the observer toward avoidance (Knutson, 1996; Yik &

Russell, 1999). Happiness, on the other hand, may signal warmth and potential for affiliation and may influence the observer to approach (Hess, Blairy, & Kleck, 2000; Knutson, 1996). In short, facial expressions are essential to social behavior as they provide information about other people's emotions and intentions and help people to guide their own behavior.

Biases in perception of single faces

Despite being adept at recognizing facial expressions, research suggests that people are nonetheless often biased toward rating individual faces negatively. For example, neutral faces with certain craniofacial characteristics (e.g., a low eyebrow ridge and wide nostrils) are particularly susceptible to anger misattribution (Shasteen, Sasson, & Pinkhsm, 2015). Importantly, not only are these features associated with temporary emotion states like anger, they are also associated with traits like aggression (Carré, McCormick, & Mondloch, 2009; Said, Sebe, & Todorov, 2009; Stillman, Maner, & Baumeister, 2010). For example, some faces decomposed into their low-spatial-frequency (LSF) components are rated as neutral and yet threatening (Bar, Neta, & Linz, 2006). Thus, perceivers seem to be particularly biased toward detection of threat and not just emotional states. These biases seem to frequently emerge under conditions of uncertainty, such as in the context of ambiguous or subtle facial expressions, rushed judgments, and brief and unpredictable presentation times. For example, an ambiguous facial expression like surprise can be judged as either positive or negative depending on the context, and when perceivers are required to make quick judgments, they exhibit a bias toward negative interpretations of surprise (Neta & Tong, 2016). However, this bias diminishes

when perceivers are required to delay their evaluations and engage in reappraisal strategies.

Thus, this negative valence bias may reflect early interpretations and result from an initial automatic process in the context of limited time and ambiguous or incomplete information (Davis et al., 2016; Neta & Tong, 2016). For example, perceivers display a tendency to label neutral facial expressions as threatening after quick 39-msec exposures (Bar, Neta, & Linz, 2006). LSF information is typically processed quicker via rapid magnocellular projections compared to detailed information (i.e., high spatial frequencies), and LSF surprise faces are also more likely to be interpreted negatively (Neta & Whalen, 2010). Negativity bias diminishes with full intensity LSF faces, suggesting that subtle or low intensity expressions may be especially vulnerable. Even children as young as 7 months of age seem to exhibit a negativity bias (Vaish, Grossmann, & Woodward, 2008). Relative to adolescents, they are more likely to rate surprised and neutral faces negatively, and their responses are marked by fast reaction times and increased corrugator activity (Tottenham et al., 2013).

Perception of facial expressions and gender. Gender has been argued to be differentially associated with emotion and affordance of threat (Becker et al., 2007). A number of studies have demonstrated that men tend to be associated with angry expressions, whereas women are more likely to be associated with happy or fearful expressions (Hess, Blairy, & Kleck, 2000; Hess, Adams, & Kleck, 2004). Perceivers are quicker and more accurate at identifying angry male faces compared to angry female faces and happy female faces compared to angry male faces (Becker et al., 2007). Neutral male faces are also more likely to be misclassified as angry whereas neutral female faces

are more likely to be misclassified as happy (Ashby & Townsed, 1986; Becker et al., 2007). Furthermore, features typically associated with angry expressions increase the likelihood of a face to be categorized as male, whereas features associated with happy expressions increase the likelihood of a face to be judged as female (Hess et al., 2009).

Some of these findings could be due to structural similarities between male or female faces and certain emotions (e.g., lowered eyebrows being associated more with both anger and more typically male faces), biasing certain androgynous faces toward anger, fear, or happiness depending on those structural markings (Hess, 2009). There are several possible reasons for the association between male faces and anger. It has been argued that men are more likely to exhibit aggression and that certain cultures impose expectations for men to be more aggressive (Biertnat & Manis, 1994; Fabes & Martin, 1991; Fisher 1993). Alternatively, there is also evidence that these biases may be driven by perception of dominance and affiliation rather than gender per se (Becket et al., 2007; Hess et al., 2004; 2009). Thus, certain facial structures, such as distance between eyes and eyebrows, roundness of face, and shape of mouth region, may be markers for dominance, and in turn, underlie gender bias (Becket et al., 2007; Senior et al., 1999). Thus, men may be perceived as more aggressive because gender markers and dominance markers are confounded within the structure of the face, which in turn may also interact with experiences and social expectations.

Internal processes in the perceiver may impact perception. In addition to the external characteristics of the observed face, internal traits and states within the perceiver may also contribute to misperception of facial expressions. For example, individuals with aggressive traits, including juvenile and adult criminal offenders as well as children and

adults with subclinical aggressiveness, display a “hostile attribution bias” in which they are more likely to misinterpret ambiguous actions of others in a hostile manner (Dodge, 2006; Knyazev, 2008; Schönenberg & Jusyte, 2014). Individuals with social anxiety tend to be more sensitive or accurate in their detection of threats as well as more biased to judge ambiguous, or subtle, facial expressions as threatening (Yoon & Zinbarg, 2008). Relative to non-anxious individuals, they may also be quicker at generating angry judgments and slower at generating happy judgments (Maoz et al., 2015). Bell and colleagues (2011) found that anxious and non-anxious individuals were similarly accurate in their classifications of emotions, but differed in the types of errors they made, with anxious individuals miscategorizing faces as angry and non-anxious individuals miscategorizing faces sad.

Attentional biases and processes in individuals with anxiety are complex. Specifically, individuals may show facilitated attention toward threatening stimuli, difficulty disengaging from it, as well as avoidance, with some of these processes being more automatic and others more strategic (Bar-Haim et al., 2007; Cisler & Koster, 2009). Of note, most of the studies on perceptual biases in anxiety have been done with single faces. Research into perception of crowds in individuals with anxiety is scarce and has yielded inconsistent results. Some studies found that individuals with social anxiety perceive crowds as more disapproving and emotionally costly (Douilliez et al., 2011; Gilboa-Schechtman et al., 2015). However, other studies noted that despite longer fixations on angry faces, individuals with anxiety do not rate faces more negatively relative to non-anxious individuals (Lange et al., 2008; 2011).

Negativity bias beyond face perception

These biases in face perception with facial expressions being classified as angry or threatening are consistent with other areas of research that have documented a negativity bias. Several researchers have described a fundamental negativity bias or a positive-negative asymmetry effect, where negative information is weighted more heavily than positive information, especially when contextual information is limited or ambiguous (Baumeister et al., 2001; Cacioppo & Berntson, 1994; Ito et al., 1998; Rozin & Royzman, 2001). For example, learning negative information about an individual is much more likely to impact impression formation than learning something positive about them (Peeters & Czapinski, 1990). Individuals also form these impressions with less information and report feeling more confident in their evaluations (Yzerbyt & Leyens, 1991; Hamilton & Zanna, 1972). In other words, people tend to form bad impressions of others quickly and confidently based on very little information.

Although people occasionally show a positivity bias, they seem to nonetheless rely heavily on a negativity bias. For example, infants and toddlers may exhibit an initial positivity bias—in the absence of information, they display a drive to approach and explore their environment. However, in the presence of both positive and negative information, they tend to respond more heavily to the negative information (Vaish, Grossmann, & Woodward, 2008). Adults tend to overestimate the likelihood of positive events in the future (Weinstein, 1980), but forecast that negative experiences will have a more powerful impact on their lives than positive events (Gilbert et al., 1998). People also tend to generate more negative than positive labels for emotions, they remember negative events more saliently, and they exert greater amounts of energy to escape from

negative moods than to generate positive ones (Baumeister et al., 2001; Finkenauer & Rimé, 1998).

Negative information also impacts decision making and information processing. For example, individuals tend to weigh potential losses more heavily than gains, especially in the context of uncertainty (Kahneman & Tversky, 1979, 1984). Negative information may also elicit greater information processing, as people tend to allocate more automatic and intentional attention to negative information (Fiske, 1980; Schwarz, 2001; Sinclair & Mark, 1992). For example, people's performance on a Stroop task suggests increased interference from negative versus positive words (Pratto & John, 1991), and people require more time to make inferences about others' bad behaviors compared to good behaviors (Krull & Dill, 1998). Additionally, people experiencing sadness tend to be more thorough in certain evaluations, although it is important to note that individuals experiencing high arousal negative emotions tend to make quicker and riskier decisions (Leither & Baumeister, 1996).

Biases may be adaptive

Not only are biases not necessarily arbitrary, they may serve a protective purpose (Haselton et al., 2009; Nesse, 2005). People that are attuned to detecting threat in their environment are more likely to survive, and they may rely on error management strategies to mitigate risk (Haselton et al., 2009; Holbrook et al., 2014; Johnson et al., 2013). When making a judgment in the context of uncertainty, one can make two types of errors: a false positive (assuming something is true or present when it is not) or a false negative (assuming something is false or absent when it is actually true). When the costs of false positives versus false negatives are weighted differently, people are likely to bias

their decision toward the less costly error (Johnson et al., 2013). These types of biases have been documented across several domains of research, including perception of auditory looming (Neuhoff, 1998), steepness (Profit, 2006), heights (Jackson & Cormack, 2007), avoidance of potentially diseased food (Garcia, Hankins, & Rusiniak, 1976) or individuals (Kurzban & Leary, 2001), and mating (Garver-Apgar, Gangestad, & Simpson, 2007). In these examples, the costs of false positives (e.g., object arrives slower, extra caution around cliffs, or missing out on potential sources of food or social partners) are relatively lower compared to costs of false negatives (e.g., object arrives quicker, recklessness around cliffs, serious sickness or potential death, and pregnancy and loss of agency). In short, this “better safe than sorry” approach can lead to assumptions of threat even in the absence of threat.

Neuberg and colleagues (2011, 2016) proposed that this self-protection system is composed of internally consistent processes, including perception (e.g., angry facial expressions), emotional reactions (fear), cognition (e.g., inference about intent), and behavior (e.g., escape or fight). These processes may incorporate multiple neural systems including the amygdala, which is often identified in threat-detection (Johansen et al., 2010; LeDoux, 2000), the primary motor cortex and dorsal basal ganglia, which are important in motor-preparation (Butler et al., 2007), the cingulate cortex and insula which are implicated in less immediate threat responses (Fiddick, 2011), as well as a cascade of cortico–striato–pallido–thalamo–cortical loops related to security motivation (Woody and Szechtman, 2011). Biases can be especially adaptive to employ because they are quick to engage, which can be necessary in dangerous situations when there is no time for thorough decision making (Johnson et al., 2013). In short, these precautionary biases are

most likely to emerge in the context of uncertainty, asymmetric costs, and pressures to make a quick decision.

Crowds

Collectives of people are socially important, as they facilitate group behavior, drive social organization, and amplify emotional experiences (Phillips, Weisbuch, & Ambady, 2014). Social interaction at the group level is a critical element in human survival, allowing for greater collective access to protection and resources (Van Vugt & Kameda, 2012; Toob & Cosmides, 2005). However, they can also pose threats, such as in group punishment and exclusion from the group, as well as interpersonal violence and exposure to infectious disease (Cosmides et al., 2005; Kurzban & Leary, 2001; Neuberg, Kenrick, Schaller, 2011; Sidanius & Pratto, 1999). Although there was a concentrated effort toward understanding social behavior in groups in the first half of the 20th century (e.g., Triplet, 1898; Lindzey, 1954), the cognitive revolution shifted the focus of many research studies away from groups and toward cognitive and emotional processes within the individual (Hogg & Tindale, 2008). However, some social psychologists have proposed that social behaviors in crowds are sometimes guided by a “collective consciousness” or a “group mind” and, therefore, studying them at an individual level is insufficient (Le Bon, 1960; McDougall, 1920; Resnik, Levine, & Teasley, 1991; Thompson & Fine, 1999). There has since been a significant increase in understanding group cognition; however, affect and emotion remain a relatively understudied area of research.

Emotions in crowds may be amplified due to “emotion contagion,” which refers to “catching feelings” observed in others—an automatic process of mimicking other

people's facial expressions, postures, and movements (Le Bon, 1896; Hatfield, Cacioppo & Rapson, 1993, 1994). Socially induced affect is another term used to describe how one's own affect can change depending on that of another person (McIntosh, Druckman, & Zajonc, 1994). Some theories of crowd behavior have pathologized the crowd attributing rises in conflict to "group mind" and "emotional contagion" (Allport, 1924; Baumeister, Ainsworth, & Vohs, 2016; Le Bon, 1896). For example, individuals in a highly aroused and cohesiveness crowd may experience a loss of self due to deindividuation; an increased sense of anonymity and diffused responsibility can then lead to members of the crowd "losing control" and becoming violent (Prati & Pietrantonio, 2009; Stott & Reicher, 1998). The elaborated social identity model (ESIM) describes crowd behavior in the context intergroup relations—not only does the crowd influence its internal members, but the way an outsider views and acts toward the crowd may also impact that crowd's behavior (Drury & Reicher, 2000; Reicher, 1996; Stott & Reicher, 1998).

How good are people at perceiving crowds? Whether in the context of a classroom, a protest, or a sporting event, individuals observing a crowd (as well as those within the crowd) require an ability to efficiently examine the group and make judgments about it. Perception of crowds is aided by an efficient mechanism—ensemble coding. Instead of recruiting focused attention to inspect individual constituents one at a time, ensemble coding pools and compresses information about a crowd into a summary representation. (Alvarez, 2011; Sweeny & Whitney, 2014). Thus, the visual system is able to represent groups rapidly and accurately (Whitney, Haberman, Sweeny, 2014). Ensemble coding has been extensively evaluated as a computational process for

evaluating low-level visual features including position (Alvarez & Oliva, 2008), brightness (Bauer, 2009), orientation (Ariely, 2001), and size (Chong & Treisman, 2003). Perceivers are able to make quick and strikingly accurate judgments about the average of each of these visual features as they occur in groups (e.g., what is the average brightness of a set of circles? What is the average size of oranges arranged in grocery display?). More recently, this computational process has also been investigated with more complex visual features processed in higher-level stages of visual processing, which play a role in social interactions. In less time than it takes to make a saccade (e.g., approximately 200-msec), people are able to extract the gist of a group's identity (Yamanashi Leib et al., 2012), gender (Haberman & Whitney, 2007), gaze direction (Sweeny & Whitney, 2014), and emotion (Elias, Dyer, & Sweeny, 2017).

In some cases, people can extract the averaged characteristics of a group more accurately than that of an individual (Alvarez, 2011; Sweeny, Haroz, & Whitney, 2013; Elias & Sweeny, 2017). For example, Sweeny et al. (2013) demonstrated that perceivers are better at estimating the direction of a crowd in motion than the direction of a single individual. A single neural representation can be randomly perturbed by internal noise (e.g., Sweeny, Grabowecky, Kim, & Suzuki, 2011), but when several noisy representations are averaged, the variability in each cancels out, leaving only a sharpened representation of the group mean (Sweeny, Haroz, & Whitney, 2013). Ensemble coding may also account for potential variability in a set (Haberman, Lee, & Whitney, 2015). Perceivers are able to discount outliers (Haberman & Whitney, 2010), and there is some evidence to suggest that not all items of the set contribute to the mean equally (Alvarez, 2011). Instead the visual system engages in precision-weighted averaging. For example,

probability, location, and size of certain items may contribute to how heavily they are weighted in the average (Haberman & Whitney, 2010; Kanaya, Hayashi, & Whitney, 2018; Mareschal, Morgan, & Solomon, 2010). This can be especially important in crowds, where there are multiple sources of heterogeneity, including gender, race, and emotion. The ability to extract a sense of a group's variance in addition to their average can provide essential information about emotional crowds (Haberman, 2015). Crowds are rarely unanimous in their expressions as members may display different emotional categories or different intensities of emotion within a category. For example, a group of faces all expressing moderate anger might have the same average as a group of faces expressing mild and intense anger. Variance thus can provide information about the reliability of the average or the coherence of a particular crowd.

In summary, compared to inspecting a group one member at a time, ensemble coding provides the ability to extrapolate a concise summary representation of the crowd's overall characteristics. This mechanism can be especially useful in social contexts in which it is critical to understand emotion or behavior at a group rather than an individual level. For example, a panicked crowd versus a single panicked individual within a crowd provides important information about a potential threat, and there is a need to extract this information quickly and automatically. When making these types of evaluations, it is important to understand the observer's sensitivity—their ability to accurately discriminate between categories, and the observer's response bias—their tendency to endorse a particular category. Research on ensemble coding, however, has largely focused on sensitivity, and little is known about biases in evaluations of crowds.

Biases in perception of emotions in crowds. Although research on biases in crowd perception is limited, some studies have focused on processing biases in emotion perception in crowds. For example, Hansen & Hansen (1988) introduced the face-in-the-crowd effect when they found that perceivers were quicker at detecting an angry face in a neutral or a happy crowd relative to detecting a happy face among neutral or angry faces. Because larger crowds did not impact perceivers' ability to detect the angry target faces, it was concluded that angry faces "pop out" in a crowd. Although some subsequent studies (Horstmann & Bauland, 2006; Öhman, Lundqvist, & Esteves, 2001) found evidence to support this Angry Superiority Effect (ASE), others have disputed it. For example, Becker and colleagues (2011; 2014) found that when using photographic or hyper-realistic faces rather than schematic ones, perceivers were faster and more accurate at detecting happy rather than angry faces. Other researchers have also found this Happy Superiority Effect (HSE; Hunt et al., 2007; Purcell, Stewart & Skov, 1996), particularly when stimuli were matched on low-level visual features, such as luminance and contrast.

It is questionable whether ASE and HSE reflect "pop out" effects, as evidence suggests that perceivers engage in serial rather than parallel processing of these faces. For example, duplicates of faces do not speed up the search and the position of faces impacts reaction times (Becker, 2011; Hampton et al., 1989). Additionally, it is unclear whether the face-in-the-crowd effect is due to the saliency of the target face (e.g., angry) or efficiency in processing of the distractor faces (e.g., happy faces) in the case of ASE, and vice versa for HSE. To assess some of these biases in processing angry versus happy faces within a crowd, Bucher and Voss (2019) took a different approach. Instead of asking perceivers to find a particular face in the crowd, they asked them to report the

overall mood of the crowd. They found that perceivers were more likely to attend to happy faces, focus on them for longer durations, and evaluate them more accurately. It is important to note that in all of these studies, perceivers had the time to serially search the crowds. Little is known about biases in perceivers who are only able to view crowds for a brief moment. However, one study that used quick presentation times and assessed biases in anxious and non-anxious individuals also found a positivity bias in healthy perceivers (Yang et al., 2013). Of note, these crowds consisted of both female and male faces and displayed “full-blown” prototypical facial expressions. In short, there is a huge gap in understanding the extent to which perceivers are susceptible to biases in perception of emotions in crowds, especially under conditions that present uncertainty.

Present Research

If perceivers are more likely to attribute threat to people with greater capacity for harm (Neel et al., 2012), then based on numerosity alone, crowds should be especially vulnerable to biased evaluations. Negative biases are also more likely to be engaged in circumstances where uncertainty or ambiguity is high, such as when perception is fleeting, uncertain, or unpredictable (Davis et al., 2016; Neta & Tong, 2016). Crowd perception often relies on ensemble coding, extracting a quick gist summary, with much of that information from the visual periphery. Thus, crowds impose perceptual and computational constraints on perceivers that may encourage the use of heuristic processes and exacerbate negative biases.

In the first investigation, I evaluated anger bias in response to single faces and emotionally unanimous crowds. In Experiment 1, I predicted that observers would engage a bias to categorize faces as “angry” (relative to happy), particularly in response

to crowds compared to individuals. I hypothesized that these biased evaluations would emerge under conditions of increased perceptual uncertainty or ambiguity, such as when faces display subtle facial expressions. Particularly, I expected anger bias to emerge when sensitivity, or ability to discriminate between facial expressions, is low. If bias is related to a target's affordance to inflict harm, it should be specific for judgments of anger. However, if anger bias only emerges in the context of evaluating angry versus happy faces, it may reflect a more general negativity bias. In Experiment 2, I thus evaluated whether anger bias would persist even in a design context in which only angry and fearful faces were discriminated. Here, I predicted a liberal bias specific to the classification of anger. If anger bias is engaged most strongly when diagnostic information relevant to discriminating expressions is limited, then ambiguity from very basic perceptual dimensions, like visibility, may also trigger bias. In Experiment 3, I thus predicted that occluding the faces, and thereby limiting visibility and increasing uncertainty, should result in similarly-biased evaluations of anger in single faces and crowds. The purpose of Experiment 4 was to evaluate whether crowd size impacts the magnitude of anger bias. I predicted that anger bias would have a non-linear relationship with crowd size, where it initially increases with each member, but then plateaus as the crowd grows larger. The crowds in the first four experiments were restricted to the same identity in order to isolate the effect of numerosity. That is, all the faces in a crowd were of the same individual. The purpose of Experiment 5 and Experiment 6 was to introduce additional variability in terms of identity and gender. In Experiment 5, I predicted that observers would continue to adopt an amplified anger bias in response to crowds whose members consisted of different identities. In Experiment 6, I evaluated anger bias in response to single male and

female faces as well as male and female crowds. I predicted that bias with low-intensity female faces would follow the same pattern as with male faces, albeit with a reduced magnitude, and possibly neutralize or shift to a happy bias for high-intensity expressions.

In the second investigation, I evaluated the persistence of anger bias under additional conditions of uncertainty within the context of mixed-emotion crowds, consisting of varying ratios of angry to happy faces and varying intensities of emotional expression. In the first investigation, all of the crowds presented were homogeneous in terms of facial expression (i.e., all angry or all happy). However, outside the lab, the individuals that constitute a crowd are unlikely to unanimously express the same emotion at the same time. Thus, the purpose of the second investigation was to evaluate how observers navigate contradictory information when a crowd's constituents express divergent emotions at different intensities. In Experiment 1, I evaluated anger bias in response to crowds with varying ratios of angry to happy faces. I predicted that, consistent with the first investigation, observers would exhibit more anger bias under conditions of uncertainty. Specifically, I predicted that anger bias would emerge in response to more balanced ratios of angry to happy faces and diminish in response to highly imbalanced ratios of angry to happy faces. Additionally, I predicted that bias would interact with emotional intensity, such that observers would show more bias in response to low intensity crowds relative to high intensity crowds. In Experiment 2, I evaluated the extent to which heterogeneity in emotion and intensity impacts anger bias. Specifically, in these crowds, half of the faces were happy, and half of the faces were angry, but the intensities between the two emotions were varied. I predicted that anger bias would be stronger when the two emotions were closer in intensity (e.g., half of the

crowd expressing 50% intensity and the other half of the crowd expressing 60% intensity) than when they were further apart in intensity (e.g., half of the crowd expressing 50% intensity and the other half of the crowd expressing 100% intensity). I also predicted that bias would be stronger when the two emotions were in the lower intensity range (low-medium) and weaker in the higher intensity range (medium-high). The second investigation, thus, provides additional insight into how other sources of uncertainty (i.e., variability in emotion and intensity within the crowd) impact anger bias.

CHAPTER 2: Anger bias in the evaluation of crowds

Introduction

The ability to discriminate facial expressions allows perceivers to gather information about others' internal affective states and intentions, influencing how people interact and react to each other (Frijda & Mesquita, 1994). Indeed, people are quite skillful at categorizing prototypical displays of emotion (Calvo & Nummenmaa, 2016). Yet perceivers sometimes evaluate subtle, difficult to see, or obscured expressions and, in some instances, these judgments must be made quickly. This is a particular challenge in the context of crowds where there is an abundance of information— much of it in the visual periphery— and the capacity to carefully evaluate each individual is limited (Awh, Barton, & Vogel, 2007; Chong & Treisman, 2005). Under these circumstances, errors in emotion classification are inevitable. These errors may not necessarily be arbitrary but instead may reflect affectively-adaptive biases that are potentially modified by crowd context. People tend to misinterpret single faces as expressing anger or hostility (Holbrook et al., 2014). Yet surprisingly little is known about negativity bias in perception of crowds. Here, we begin to address this gap by evaluating anger bias—a tendency to judge facial expressions as angry—in the context of single faces and crowds.

Biased evaluations are often portrayed as being adaptive. For example, error management theory is predicated on the idea that overestimating the presence of threat is less costly than underestimating its presence, and it proposes that people avoid high-risk

mistakes to minimize potentially negative outcomes (Haselton et al., 2009; Johnson, Blumstein, Fowler, & Haselton, 2013; Nesse, 2005). Accordingly, people may more liberally endorse the presence of anger, especially when their judgments are made with low confidence (Holbrook et al., 2014). This adaptive mechanism is not limited to facial expressions—it similarly cultivates protective judgments about loudness or proximity of sounds (Neuhoff, 1998) and it encourages loss-aversion during decision making (Kahneman & Tversky, 1984). Of course there are special circumstances which can lead perceivers to adopt positive biases, like evaluations of in-group members (Lazerus, Ingbreetsen, Stoller, Freeman, & Cikara, 2016) or prolonged decision time, but initial judgments of affect tend to be negative (Neta & Tong, 2016). Perceivers may be inclined to judge individuals as expressing threat (or anger) presumably to avoid danger (Becker, Kenrick, Neuberg, Blackwell, & Smith, 2007; Gibson, 1979; Holbrook et al., 2014). Importantly, perceived threat depends on an agent's capacity to inflict harm on the viewer. If a single individual is assumed to pose potential harm, then to what extent might many individuals, or a crowd, carry an increased affordance to inflict harm and, in turn, elicit a more strongly biased evaluation?

Crowds and groups are socially important—they allow collective behavior, drive social organization, and amplify emotional experiences (Phillips, Weisbuch, & Ambady, 2014; Shteynberg et al., 2014). To evaluate collectives, humans rely on ensemble coding—a visual mechanism that enables rapid extraction of summary or gist information, including average orientation (Parkes et al., 2001), size (Ariely, 2001), location (Alvarez & Oliva, 2008), gaze (Sweeny & Whitney, 2014), identity (Yamanashi Leib et al., 2012), facial expression (Haberman & Whitney, 2007), and biological motion

(Sweeny, Haroz, & Whitney, 2013). Despite an abundance of research on sensitivity, or the ability to correctly discriminate between the actual attributes of crowds, little research has been conducted on biases that occur when perceivers make rapid judgments about crowds.

Crowds should be especially vulnerable to biased evaluations. Logically, based on numerosity alone, a crowd has greater capacity to inflict harm on a perceiver than an individual. Negative biases are also more likely to be engaged in circumstances where uncertainty or ambiguity in perception or decision making is high, such as when perception is fleeting, uncertain, or unpredictable (Davis, Neta, Kim, Moran, & Whalen, 2016; Neta & Tong, 2016). Crowd perception is often made in just a glance, based on gist information, with much of that information from the visual periphery. Thus, crowds impose perceptual and computational constraints on perceivers that may encourage the use of heuristic processes and exacerbate negative biases. We predicted that observers would engage a bias to categorize faces as “angry,” particularly in response to crowds compared to individuals. We hypothesized that these biased evaluations would emerge under conditions of increased perceptual uncertainty or ambiguity, such as when faces sport subtle facial expressions or are difficult to see due to occlusion. Particularly, we expected anger bias to emerge when sensitivity, or ability to discriminate between facial expressions, is low. We tested and found support for these predictions in the six experiments described below.

Experiment 1: Are crowds particularly susceptible to being misperceived as angry?

Methods

Observers. Eighty-three undergraduate students (58 female, 27 male) from the University of Denver participated in Experiment 1. This experiment was run as part of a larger study on emotion perception and affective reactions. The number of participants in this sample reflected an attempt to capture potentially weak physiological effects associated with this larger study that are not described in detail here. Our post-hoc observed-power for examining the main effect of crowd size in this experiment was 0.71 (with alpha set at .05, captured by a matched-pairs t-test between bias in the crowd and single conditions, accounting for the correlation between these conditions). All participants provided informed consent and received course credit for their participation. Each observer had normal or corrected-to-normal visual acuity and completed the experiment in a dimly lit room. This study was approved by the Institutional Review Board at the University of Denver and the research was carried out in accordance with the provisions of the World Medical Association Declaration of Helsinki.

Stimuli. Our face set consisted of four male actors from the NimStim face set (Tottenham, Borscheid, Ellertsen, Marcus, & Nelson, 2002). We were interested in examining anger bias across a range of facial expression intensities. Consistent with our theorizing in the Introduction, we predicted that anger bias would emerge most strongly for judgments of weaker intensity facial expressions, as they are more difficult to discriminate and therefore pose more uncertainty than high-intensity facial expressions. To create a range of facial expression intensities, we morphed full-intensity exemplar

expressions from each of the four actors with their own neutral faces. Specifically, we used Fantamorph software (Version 5) to create linear interpolations (i.e., morphs) between a neutral expression and two emotional expressions—angry or happy—for each of the actors. A norming experiment run prior to this investigation confirmed that these interpolations did indeed produce linear changes in perceived intensity for all the stimuli in our set (see Appendices). Note that the appropriateness of our stimulus set (and our ability to measure changes in the magnitude of anger bias) did not depend on the lowest-intensity expressions being perceived as perfectly neutral in valence—a bias can be measured even with stimuli that contain diagnostic visual information (e.g., a moderate level of expressiveness).

In our first experiment, we presented faces with five levels of emotional intensity (20%, 40%, 60%, 80%, and 100%) from these two expression ranges. We used a Gaussian edge-blurring tool (Adobe Photoshop CS6, Version 13.0) to smooth the external contours of the face and hair and to diminish the rough edges of the cropped faces. The background of each image was then replaced with a uniform gray color (RGB = 170,170,170). The stimuli were presented in MATLAB (R2014b) on a uniform background (RGB = 170,170,170; luminance = 27.5 cd/m²) on a CRT monitor with a screen size of 27.3cm x 36.5cm, a resolution of 1024 x 768. Each face subtended a visual angle of 3.75° x 4.53°.

Procedure. We used a within-subjects design in which each observer viewed and evaluated emotional expressions on a single face (the single condition) and from a collection of 12 faces (the crowd condition). The task was to categorize a single person's or a crowd's facial expression(s) as happy or angry by pressing the right or left arrow key

(counterbalanced across observers). Observers were limited to a binary decision to ensure that they could not opt out of making an emotional classification. This was particularly important as we predicted enhanced bias on trials in which emotion classification was difficult and uncertainty was high. There was no time limit put on observers to provide a response. On each trial faces were randomly distributed around a centrally presented fixation point. In the single condition, one face appeared at one of the four central positions (randomly selected on each trial) to ensure that any potential differences in emotion discrimination (i.e., sensitivity) between single and crowd conditions would not be due to differences in visual acuity (that is, both conditions included diagnostic information near foveal vision). In the crowd condition, the faces appeared in twelve positions scattered across the screen. Faces within each crowd were identical, always with the same actor and intensity of facial expression. For example, a crowd trial could feature 12 images of the same person depicting a 40% happy expression (Fig. 1). We introduced this redundancy intentionally, accepting that it degraded our crowds in terms of real world sources of variability, such as gender, expressiveness, and identity, because it provided a more important upside for our initial investigation. That is, because the crowds did not contain any additional information for categorizing the expression, any difference in bias in the crowd condition relative to the single condition could only be attributable to the increase in number of faces. The identity and intensity of the face, however, varied across trials.

Figure 1. (A) Morphs between a neutral expression and happy (top row), angry (middle row), and fearful (bottom row) expressions from one actor in our stimulus set. Intensities depicted here include 0%, 20%, 40%, 60%, 80%, and 100%. (B) A crowd of 12 faces from Experiment 1. Face identities were identical within each crowd on a given trial for most of the experiments in this investigation. This ensured that our predicted result of increased bias for evaluations of crowds (relative to individuals), were it to occur, could be accounted for by numerosity alone.

To prevent observers from focusing on specific locations, we randomly and independently jittered the location of each face by 1-to-15 pixels both horizontally and vertically on each trial. Single faces and crowds were presented for 100-ms, a presentation time brief enough to prevent observers from making deliberate saccades to individual faces (Findlay & Walker, 1999). Each face was followed by a pattern mask (70 rectangular pieces derived from the preceding image, reshuffled into a new image) shown for 250-msec. This approach ensured that the emotional faces and scrambled masks matched in terms of low-level image characteristics, thus decreasing the visibility of the masked image (Enns & Oriet, 2007) and preventing residual visual processing (Rolls, Tovee, & Panzeri, 1999). Upon viewing each single face or crowd, observers indicated whether that person's or crowd's emotional expression was happy or angry by pressing the right or the left arrow on a keyboard. Identity, emotion category, and emotion

intensity were randomly determined on each trial. Trials were randomized for each block and observer. Each observer completed six blocks for a total of 480 trials. Each condition (e.g., a crowd of happy faces at 20% intensity) was repeated for 24 trials across all six blocks.

Analyses. We used Signal Detection Theory (SDT) for our primary analyses, which allowed us to separately measure both sensitivity and bias in perceptual decision making (e.g., whether a signal— in this case anger— is present) in the context of uncertainty. In this experiment, there were two possible responses (angry or happy), resulting in four possible SDT outcomes (see table below): “hit” (angry face(s) identified as angry), “miss” (angry face(s) identified as happy), “false alarm” (happy face(s) identified as angry), and “correct rejection” (happy face(s) identified as happy).

Response	Signal (Angry)	Noise (Happy)
Yes: “Angry”	Hit	False Alarm
No: “Happy”	Miss	Correct Rejection

Table 1. Signal Detection Theory analysis parameters for crowds of angry and happy faces.

Based on these four response outcomes, we derived d' , an index of sensitivity which captures an observer’s ability to correctly discriminate between angry and happy expressions (Dorfman & Alf, 1968; Macmillan, 1993). We also calculated criterion, or c , which measures an observer’s bias to respond a certain way (i.e., “angry”) independent of whether the signal (in this case, anger) was present or not. For the single and crowd conditions, for each observer, we calculated d' and criterion for each level of expression intensity using z -transformed hit rates (h) and false alarm rates (fa), where hit rate =

$$\frac{(\text{hit})}{(\text{hit}+\text{miss})}, \text{ false alarm rate} = \frac{(\text{f.a.})}{(\text{f.a.}+\text{correct rejection})}, d' = Z_h - Z_{fa}, \text{ and } c = -0.5(Z_h - Z_{fa}).$$

We replaced hit and false alarms rates of 1.0 with a value determined as $1 - \frac{1}{2*n}$, with n reflecting the number of possible hits of false alarms. We replaced hit and false alarms rates of 0.0 with a value determined as $\frac{1}{2*n}$.

We provided the raw data as well as the calculated Hits, False Alarms, Misses, Correct Rejections, Hit Rates, and False Alarm Rates for each experiment at <https://osf.io/8npx6/>. Additionally, we included plots visualizing Hit Rates and False Alarm Rates for all six experiments. We also conducted additional exploratory analyses with reaction time as our dependent variable. We did this for each experiment in our investigation. However, because RT was not our primary interest, we share these analyses and results in the Appendices.

Results. We began with a repeated-measures ANOVA using d' as our dependent variable with factors of Numerosity (single, crowd) and Intensity (20%, 40%, 60%, 80%, and 100%). This analysis yielded main effects of Numerosity $F(1, 82) = 119.2, p < .001, \eta_p^2 = 0.402$, and Intensity $F(4, 328) = 178.4, p < .001, \eta_p^2 = 0.824$, and an interaction between Numerosity and Intensity $F(4, 328) = 9.51, p < .001, \eta_p^2 = 0.104$ (Fig. 2a). These results are not surprising—emotion categorization should be (and was) better when faces portrayed intense expressions, especially when many faces were present to carry this information. This d' analysis is nevertheless an important first step. We predicted an increase in anger bias when uncertainty was high, and this initial analysis confirmed that sensitivity was indeed poor (and that uncertainty was high) at low intensities of facial expression.

Figure 2. Sensitivity (A) and bias (B) for classification of angry and happy expressions in Experiment 1. Each panel depicts performance across changes in expression intensity, separately for crowds and single faces. Increasingly negative criterion values are plotted up the y-axis in panel B, reflecting a liberal bias for classifying anger. Error bars in each panel represent 95% confidence intervals.

We then turned to our main analysis, evaluating criterion first relative to a null-value of zero and then as a function of numerosity and emotional intensity. We found an overall bias to over-report facial expressions as angry in response to both single faces as well as crowds of faces. Note that a criterion (c) value of zero indicates no bias, a negative value reflects a liberal bias to report the presence of anger, and a positive value reflects conservative use of the anger response label. When collapsed across all intensities of facial expression, average criterion was consistently negative: -0.171 (CI: -0.25, -0.091) in the single face condition (one-sample t-test: $t[82] = -4.284$, $p < 0.001$, $d = 0.47$) and -0.274 (CI: -0.35, -0.20) in the crowd condition (one-sample t-test: $t[82] = -7.709$, $p < 0.001$, $d = 0.85$). Across observers, the strength of anger bias for evaluations of single faces and crowds was positively correlated ($R^2 = 0.60$), suggesting a shared mechanism.

A repeated-measures ANOVA with criterion as our dependent variable revealed significant main effects of Numerosity, $F(1, 82) = 16.57$, $p < .001$, $\eta_p^2 = 0.13$ and Intensity, $F(4, 328) = 25.84$, $p < .001$, $\eta_p^2 = 0.372$, as well as an interaction between Numerosity and Intensity $F(4, 328) = 3.822$, $p < .01$, $\eta_p^2 = 0.045$. Observers were more likely to misclassify crowds as being angry compared to single faces, and this amplification of anger bias in response to crowds was particularly evident at lower intensities of facial expressions. Anger bias, in general, was present at lower intensities for both crowds and single faces, dissipating at higher intensities of facial expression (Fig. 2b). Thus, these data suggest that high perceptual uncertainty (e.g., low intensity facial expressions) and high threat potential (e.g., crowds) may independently and interactively engage anger bias.

Experiment 2: Is amplified evaluative bias for crowds especially strong for categorizations of anger?

Angry and fearful expressions can signal potential threat, but unlike anger, which is typically associated with approach behavior from another person, fear is associated with avoidant behavior (Adams, Ambady, Macrae, & Kleck, 2006). If the bias we measured in Experiment 1 is related to a target's affordance to inflict harm, it should be specific for judgments of anger, and/or stronger for evaluations of anger compared with fearful expressions. Experimentally speaking, such an anger bias would persist even in a design context in which only angry and fearful faces were discriminated. If, on the other hand, the bias reflects a more generic heuristic to report negative emotion relative to positive emotion, then an anger-specific bias should not be present when perceivers

discriminate angry and fearful expressions. Experiment 2 disambiguated these hypotheses. We predicted the former, indicating a liberal bias specific to the classification of anger.

Method

Observers. In order to obtain the same power as in Experiment 1 (with alpha set at .05, based on a simple two-tailed matched-pairs t-test between bias in the crowd and single conditions), we would have, of course, needed to once again collect data from 83 observers. But unlike Experiment 1, we did not run Experiment 2 while simultaneously collecting physiological data. In Experiment 1, observers were required to wait four seconds between image presentations, so as not to interfere with the physiological recording. This was not a concern in Experiment 2, which allowed us to nearly double the trial count for each observer. We thus decided to run only thirty observers. Thirty new observers (24 female, 6 male) provided consent and participated in Experiment 2. Each observer had normal or corrected-to-normal visual acuity.

Stimuli and Procedure. The stimuli and procedures were identical to those in Experiment 1 with the exception that observers were required to discriminate between two negatively-valenced facial expressions—angry and fearful. We used the same morphing procedure as in Experiment 1 to create faces with six intensities of fearful and angry expressions (2%, 20%, 40%, 60%, 80%, and 100%). Because we predicted that anger bias would be amplified under conditions of uncertainty, we included this extremely low intensity (2%) of facial expression in each of our subsequent experiments. Observers discriminated between fearful faces and the angry faces from the previous experiment. Observers completed 10 blocks for a total of 960 trials, with the exception of

three observers who completed nine blocks for a total of 864 trials. Each condition (e.g., a crowd of fearful faces at 20% intensity) was repeated for 40 trials across all 10 blocks.

Results. In this experiment, each observer's d' (sensitivity) and criterion (bias) were calculated for responding "angry" or "afraid." A repeated-measures ANOVA with d' as the dependent variable yielded main effects of Numerosity, $F(1, 29) = 54.61$, $p < .001$, $\eta_p^2 = 0.343$, and Intensity, $F(5, 145) = 97.69$, $p < .001$, $\eta_p^2 = 0.906$, as well as an interaction between Numerosity and Intensity $F(5, 145) = 10.12$, $p < .001$, $\eta_p^2 = 0.259$ (Fig. 3a). Consistent with Experiment 1, sensitivity, or the ability to discriminate between facial expressions, improved with increases in the intensity of expression, especially for judgments of crowds.

Figure 3. Sensitivity (A) and bias (B) for classification of angry and fearful expressions in Experiment 2. Each panel depicts performance across changes in expression intensity, separately for crowds and single faces. Error bars in each panel represent 95% confidence intervals.

Criterion analyses revealed that anger bias persisted in this experiment with observers misclassifying single faces and crowds as "angry" relative to "afraid." Note

again that a negative value of criterion (c) reflects a liberal bias to report the presence of anger. When collapsed across all intensities of facial expression, average criterion was -0.27 (CI: -0.38, -0.16) in the single-face condition (one-sample t-test against a null-value of zero: $t[29] = -4.96$, $p < .001$, $d = 0.90$) and -0.43 (CI: -0.57, -0.29) in the crowd condition (one-sample t-test: $t[29] = -6.36$, $p < 0.001$, $d = 1.16$). Across observers, the strength of anger bias with single faces and crowds was positively correlated ($R^2 = 0.77$). A repeated-measures ANOVA with criterion as the dependent variable yielded main effects of Numerosity $F(1, 29) = 23.46$, $p < .001$, $\eta_p^2 = 0.384$, and Intensity $F(5, 145) = 11.25$, $p < .001$, $\eta_p^2 = 0.555$, and an interaction between Numerosity and Intensity $F(5, 145) = 2.332$, $p < .05$, $\eta_p^2 = 0.074$. These results are consistent with Experiment 1, showing that observers are more likely to evaluate faces as angry when facial expressions are weak and that this bias is amplified for judgments of crowds compared to single individuals (Fig. 3b). Importantly, this experiment provided evidence that observers did not merely adopt a generic negative bias when required to decide between two oppositely-valenced emotions, such as happy and angry. Instead, high uncertainty due to low intensity of expression interacted with numerosity to elicit a potent evaluative bias specific to anger.

Experiment 3: Is anger bias for crowds amplified by other kinds of perceptual uncertainty, or is it specific to emotional ambiguity?

If anger bias is engaged most strongly when diagnostic information relevant to discriminating expressions is limited, then ambiguity from very basic perceptual dimensions, like visibility, may trigger bias similarly to the weak expressionality in

Experiments 1 and 2. For example, observers have been shown to sometimes misclassify masked faces as angry (Nikitin & Freund, 2015). We predicted that occluding the faces, and thereby limiting visibility and increasing uncertainty, should result in similarly-biased evaluations of anger in single faces and crowds. Such a result would indicate that anger bias may be a heuristic that operates when perceptual decisions are compromised via uncertainty from a variety of sources.

Method

Observers. Thirty new observers (27 female, 3 male) provided consent and participated in Experiment 3. We selected this sample size based on the large effect of set size in Experiment 2, which produced observed power of .99 with an N of 30. Each observer had normal or corrected-to-normal visual acuity.

Stimuli and Procedure. The stimuli and procedures were identical to those in Experiment 1 with some adjustments that allowed us to test whether perceptual uncertainty, independent of expressive intensity, enhanced biased evaluations of anger. We kept the intensity of facial expressions constant at 40% across all trials because this mid-range value elicited a consistently stronger bias for the crowd condition compared with the single condition in our previous two experiments. We instead varied the visibility of the faces by introducing visual noise. On a given trial, 0%, 15%, 30%, 45%, 60%, 75%, or 90% of the image pixels in each face were occluded by grey pixels (Fig. 4). Each face on crowd trials was occluded by the same amount of visual noise. Each observer completed 10 blocks for a total of 1120 trials, with the exception of two observers who completed nine blocks (1008 trials total) and 2 observers who completed

12 blocks (1344 trials total). Each condition (e.g., a crowd of happy faces with 30% occlusion) was repeated for 40 trials across all 10 blocks.

Figure 4. The range of occlusion levels (90%, 75%, 60%, 45%, 30%, 15%, and 0%) for one actor and expression in our stimulus set.

Results. We conducted a repeated-measures ANOVA with d' as the dependent variable and factors of Numerosity (single, crowd), and Occlusion (0%, 15%, 30%, 45%, 60%, 75%, or 90%). This analysis yielded main effects of Numerosity, $F(1, 29) = 92.63$, $p < .001$, $\eta_p^2 = 0.707$, and Occlusion, $F(6, 174) = 130$, $p < .001$, $\eta_p^2 = 0.921$, as well as an interaction between Numerosity and Occlusion, $F(6, 174) = 9.879$, $p < .001$, $\eta_p^2 = 0.254$ (Fig. 5a). Sensitivity increased as occlusion decreased, with crowds being easier to discriminate relative to single faces, especially when the faces were easier to see.

Figure 5. Sensitivity (A) and bias (B) for classification of angry and happy expressions in Experiment 3. Each panel depicts performance across changes in the amount of face occlusion, separately for crowds and single faces. Error bars in each panel represent 95% confidence intervals.

Consistent with Experiments 1 and 2, observers demonstrated a tendency to classify facial expressions as angry. When collapsed across all magnitudes of occlusion, average criterion was -0.24 (CI: -0.37, -0.12) in the single-face condition (one-sample t -test against a null-value of zero: $t[29] = -3.89$, $p < .001$, $d = 0.71$) and -0.28 (CI: -0.40, -0.15) in the crowd condition (one-sample t -test: $t[29] = -4.55$, $p < 0.001$, $d = 0.83$).

Across observers, the strength of angry bias for evaluations of single faces and crowds was positively correlated ($R^2 = 0.67$). A repeated-measures ANOVA with criterion as the dependent variable yielded a main effect of Occlusion, $F(6, 174) = 11.73$, $p < .001$, $\eta_p^2 = 0.483$, no main effect of Numerosity, $F(1, 29) = 0.745$, $p = 0.395$ (Fig. 5b), and an interaction between Numerosity and Occlusion, $F(6, 174) = 3.099$, $p < .01$, $\eta_p^2 = 0.097$.

When considered with the results of Experiments 1 and 2, these results suggest that people produce biased evaluations of anger when perceptual uncertainty, and not just emotional uncertainty, is high. Of note, at the lowest levels of occlusion (e.g., 0% and 15%), anger bias was amplified in the single condition relative to the crowd condition.

This may reflect a scaling effect, since these conditions featured lower perceptual uncertainty relative to the other occlusion conditions in this experiment. We are thus cautious about comparing the 0% occlusion condition in this experiment with the 40% intensity condition in Experiment 1, for example.

Experiment 4: How does the strength of anger bias relate to the size of a crowd?

Crowds are powerful visual signals. For example, people are more likely to orient their attention toward a crowd's point of gaze than that of an individual (Gallup et al., 2012). Importantly, this amplified influence of crowds on perceiver behavior is non-linear, increasing from 1 to 5 members but plateauing with increases in size up to 15 members (Milgram, Bickman, & Berkowitz, 1969). We predicted anger bias would have a similar non-linear relationship with crowd size.

Method

Observers. Thirty new observers (24 female, 6 male) provided consent and participated in Experiment 4. We selected this sample size based on the results of Experiments 2 and 3. Each observer had normal or corrected-to-normal visual acuity.

Stimuli and Procedure. The stimuli and procedures were identical to those in Experiment 1 with the exception that the number of faces displayed on each trial was variable, and featured 1, 2, 4, 6, 12, or 24 faces. We also included an additional intensity of facial expression (2%). We included more incremental increases in size for the smaller crowds (e.g., 2, 4, 6) because we expected that each additional member may have a more substantial perceived impact in smaller sets, whereas the magnitude of anger bias may already peak by the time a crowd reaches 12 or 24 members. On each trial, observers

indicated whether the face or faces depicted a happy or an angry expression. Each observer completed three blocks for a total of 864 trials. Each condition (e.g., a crowd of four happy faces with intensities collapsed) was repeated for 72 trials across all three blocks.

Results. This experiment included several combinations of Numerosity (six levels) and Intensity (six levels). Because evaluating each level of intensity for each crowd size would yield 36 values of d' (and criterion) and the goal of this experiment was to examine the impact of crowd size, we simplified our analyses and increased power by conducting our signal detection analyses only across the six levels of numerosity (collapsing across intensities). A repeated-measures ANOVA with d' as the dependent variable yielded a main effect of Numerosity, $F(5, 145) = 27.54$, $p < .001$, $\eta_p^2 = 0.487$. This indicates that sensitivity for discriminating happy and angry expressions increases with the number of faces in the crowd until plateauing around four to six members (Fig. 6a).

A repeated-measures ANOVA with criterion as the dependent variable yielded a main effect of Numerosity, $F(5, 145) = 3.953$, $p < .01$, $\eta_p^2 = 0.12$. Anger bias increased non-linearly with the number of faces in the crowd, reaching peak potency at four faces (Fig. 6b).

Figure 6. Sensitivity (A) and bias (B) for classification of angry and happy expressions in Experiment 4. Each panel depicts performance across changes in the number of faces visible to observers. Error bars in each panel represent 95% confidence intervals.

Experiment 5: Does anger bias occur for crowds with additional variability in identity?

Real world crowds contain variability in terms of facial expression, gender, and of course, identity. We restricted these sources of variability in Experiments 1 through 4 to isolate the effect of numerosity. In Experiment 5, we began to evaluate anger bias while re-introducing natural sources of variability, starting with identity. We predicted that observers would continue to adopt an amplified anger bias in response to crowds whose members consisted of different identities, especially under conditions of higher uncertainty.

Method

Observers. Thirty new observers (19 female, 11 male) provided consent and participated in Experiment 5. We selected this sample size based on the results of Experiments 2, 3, and 4. Each observer had normal or corrected-to-normal visual acuity.

Stimuli and Procedure. The stimuli and procedures were identical to those in Experiment 1 with a few notable exceptions. Unlike the previous experiments in which each crowd consisted of members with the same identity, Experiment 5 included two types of crowds—one in which each identity within the crowd was identical (homogenous) and another in which each identity was different¹ (heterogenous). The crowd condition also consisted of 4 faces rather than 12, since Experiment 4 showed that sensitivity and bias plateau around this crowd size. We also included a weaker intensity (2%) of facial expression. The faces in the single and crowd conditions were restricted to the four central positions around the fixation point. Observers discriminated between the happy and angry faces from the previous experiments. Each observer completed eight blocks for a total of 1,152 trials. Each condition (e.g., a crowd homogenous happy faces at 20% intensity) was repeated for 40 trials across all eight blocks.

Results. We conducted a repeated-measures ANOVA with d' as the dependent variable and factors of Crowd Type (single, homogenous crowd, heterogenous crowd) and Intensity (2%, 20%, 40%, 60%, 80%, 100%). This analysis yielded main effects of Crowd Type $F(2, 58) = 59.81, p < .001, \eta_p^2 = 0.562$, Intensity, $F(5, 145) = 157.1, p < .001, \eta_p^2 = 0.915$, as well as an interaction between Crowd Type and Intensity, $F(10, 290) = 13.96, p < .001, \eta_p^2 = 0.325$ (Fig. 7a). Sensitivity for discriminating expressions improved as emotional intensity increased across all conditions. The interaction shows that sensitivity in the single-face condition was lower than sensitivity in the homogenous and heterogenous crowd conditions at higher intensities of emotional expression. Of note, there was no difference in sensitivity between crowds composed of identical versus

¹ We are unable to show an example of the mixed identity condition due to copyright restrictions.

different identities for five out of the six comparisons; we conducted six paired samples t-tests comparing the two types of crowds at each level of Intensity: 2% $p = 0.7296$, 20% $p = 0.3439$, 40% $p = 0.0046$, 60% $p = 0.4871$, 80% $p = 0.497$, 100% $p = 0.3921$ ($\alpha = 0.008333$ with a Bonferroni correction for multiple comparisons).

Figure 7. Sensitivity (A) and bias (B) for classification of angry and happy expressions in Experiment 5. Each panel depicts performance across changes in expression intensity, separately for homogeneous crowds (same identity), heterogeneous crowds (unique identities), and single faces. Error bars in each panel represent 95% confidence intervals.

Consistent with the previous experiments, observers demonstrated an overall tendency to classify facial expressions as angry. When collapsed across all intensities of facial expression, the average criterion was -0.44 (CI: -0.60, -0.29) in the single-face condition (one-sample t-test against a null-value of zero: $t[29] = -5.78$, $p < .001$, $d = 1.05$), -0.44 (CI: -0.57, -0.30) in the homogenous crowd condition (one-sample t-test: $t[29] = -6.73$, $p < 0.001$, $d = 1.23$), and -0.44 (CI: -0.57, -0.30) in the heterogenous crowd condition (one-sample t-test: $t[29] = -6.61$, $p < 0.001$, $d = 1.21$). A repeated measures ANOVA with criterion as the dependent variable yielded a main effect of Intensity, $F(5, 145) = 50.47$, $p < .001$, $\eta_p^2 = 0.812$, no main effect of Crowd Type, $F(2, 58) = 0.044$, $p =$

0.957, and an interaction between Intensity and Crowd Type, $F(10, 290) = 10.79$, $p < .001$, $\eta_p^2 = 0.271$, but (Fig. 7b).

Consistent with our predictions, anger bias was strongest when observers evaluated weakly expressive faces. Weakly expressive faces were also discriminated with low sensitivity. Our data suggest that increased bias does not depend solely on the presence of a crowd or the boost in discrimination that comes from seeing a collective of strongly expressive faces. Rather, there are two factors that interact to elicit anger bias—uncertainty (which increases the likelihood of making an error in judgment) and potential threat (which increases the cost of that error). When uncertainty is equally high for discriminating single faces and crowds (i.e., at low intensities of facial expression), observers appear to adopt an amplified anger bias in response to crowds, which carry higher potential for harm. When discrimination sensitivity for crowds is improved relative to an individual face (i.e., at high intensities of facial expression), amplified anger bias for crowd judgments is either diminished (as in Experiments 1 and 2) or even weaker than for single faces (as in Experiment 3 and 5), which carry higher levels of uncertainty. Of note, there was no difference in criterion (i.e., bias between crowds composed of identical versus different faces; we conducted six paired samples t-tests comparing the two types of crowds at each level of Intensity: 2% $p = 0.9685$, 20% $p = 0.3848$, 40% $p = 0.9129$, 60% $p = 0.8607$, 80% $p = 0.8169$, 100% $p = 0.3184$ ($a = 0.008333$ with a Bonferroni correction for multiple comparisons).

Experiment 6: Does anger bias occur for crowds with additional variability in gender?

Gender is differentially associated with emotion and affordance of threat (Becker et al., 2007; Neel, Becker, Neuberg, & Kenrick, 2012). Male faces tend to be associated with angry expressions whereas female faces are more likely to be associated with happy or fearful expressions (Hess, Adams, & Kleck, 2004). We thus included only male faces in the preceding experiments to isolate a stable effect of anger bias. In Experiment 6, we examined anger bias during evaluations of crowds of female faces, male faces, and crowds inclusive of both. We expected that bias with low-intensity female faces would follow the same pattern as with male faces, albeit with a reduced magnitude, and possibly neutralize or shift to a happy bias for high-intensity expressions.

Method

Observers. Thirty-eight new observers (30 female, 7 male, 1 identified as other) provided consent and participated in Experiment 6. We increased the sample size due to a decrease in number of trials per condition compared to the previous five experiments. Each observer had normal or corrected-to-normal visual acuity.

Stimuli and Procedure. Stimuli and procedure were similar to the previous experiments with a few notable differences. Four identities from the NimStim face set were used: two male faces from the previous experiments and two new female faces. Interpolations between a neutral expression and two emotional expressions—happy and angry (2%, 20%, 40%, 60%, 80%, and 100% intensity)—were created for each of the female faces. This experiment featured five conditions. The single-male condition

featured one of the two male faces displayed at one of the four central locations on the screen. The single-female condition used a similar configuration. The male-crowd condition featured 12 faces, always just one of the two identities on the set. The female-crowd condition used a similar configuration. The mixed-crowd condition also featured 12 faces, but with two male and two female identities, each repeated 3 times. Observers discriminated between happy and angry expressions. Each observer completed 10 blocks for a total of 1,440 trials. Each condition (e.g., a male crowd of happy faces at 20% intensity) was repeated for 20 trials, with the exception of the mixed-crowd, which was repeated for 40 trials, across all 10 blocks.

Results. We conducted repeated-measures ANOVAs with d' and criterion as the dependent variables. We evaluated factors of Crowd Type (single-male, single-female, male-crowd, female-crowd, and mixed-crowd) and Intensity Level. We collapsed data from the 2%, 20%, 40% intensities into a low-intensity condition, and we collapsed data from the 60%, 80%, 100% conditions into a high-intensity condition. We did this to make evaluation across the five different crowd types more manageable (reducing the number of comparisons), and because our predictions about intensity were not specific to any of the six levels in particular. The analysis of d' revealed a main effect of Crowd Type, $F(4, 148) = 74.33, p < .001, \eta_p^2 = 0.756$, a main effect of Intensity Level, $F(1, 37) = 469.5, p < .001, \eta_p^2 = 0.949$, as well as an interaction between Crowd Type and Intensity Level, $F(4, 148) = 24.78, p < .001, \eta_p^2 = 0.401$ (Fig. 8a). Figure 8a illustrates that emotion discrimination was better for faces with intense expressions compared to faces with weak expressions. Sensitivity was nonetheless better for evaluations of crowd emotion, regardless of gender, compared to evaluations of single faces, particularly at high

intensity of facial expression. Of note, although data were collapsed into low- and high-intensity conditions in the analyses above, the figure presents the data across all six conditions to be consistent with the previous experiments.

Figure 8. Sensitivity (A) and bias (B) for classification of angry and happy expressions in Experiment 6. Each panel depicts performance across changes in expression intensity, separately for male faces, female faces, both in crowds and seen as single faces, and for crowds of male and female faces. Error bars in each panel represent 95% confidence intervals.

A repeated measures ANOVA of criterion yielded a main effect of Crowd Type $F(4, 148) = 49.52, p < .001, \eta_p^2 = 0.802$, a main effect of Intensity Level, $F(1, 37) = 46.5, p < .001, \eta_p^2 = 0.819$, as well as an interaction between Crowd Type and Intensity Level, $F(4, 148) = 39.75, p < .001, \eta_p^2 = 0.518$ (Fig. 8b). Across all conditions, anger bias was strongest for judgments of faces with low-intensity expressions. These results again indicate that poor sensitivity and high uncertainty are key factors in amplifying anger bias. The relationship between anger bias, gender, and crowd type was also clear for evaluations of weakly expressive faces. Paired samples t-tests revealed that bias was stronger in the single-male condition compared with the single-female condition ($t[37] =$

-9.995, $p < 0.001$, $d = 1.09$). Bias was also stronger in the male-crowd condition compared with the female-crowd condition ($t[37] = -5.91$, $p < 0.001$, $d = 0.71$). Anger bias in the mixed-crowd condition fell between these two conditions, as predicted.

The strength of anger bias was only partly consistent with our hypotheses at higher levels of intensity, when uncertainty was less of a factor. Observers adopted a happy bias when evaluating female faces with high intensity of facial expression. Average criterion was +0.12 (CI: -0.001, 0.235) in the single-female condition (compared against a null value of zero with a one-sample t-test: $t[37] = 2.05$, $p < .05$, $d =$) and +0.18 (CI: 0.10, 0.27) in female-crowd condition (one-sample t-test: $t[37] = 4.26$, $p < 0.001$, $d = 0.69$). Yet, surprisingly, anger bias was actually stronger in the single-male condition ($M = -0.53$, CI: -0.65, -0.41) compared with the male-crowd condition ($M = -0.29$, CI: -0.39, -0.19), $t(37) = -7.62$, $p < 0.001$, $d = 1.01$. This pattern is consistent with Experiment 3 and 5, indicating that the amplification of anger bias for evaluations of a crowd does not happen consistently at high-intensities of expression.

Discussion

This series of studies demonstrated that people are biased to classify facial expressions as angry, especially when viewing crowds of faces. We isolated a few factors that elicit and escalate these erroneous evaluations, including numerosity and gender, and especially uncertainty (in terms of subtlety or visibility of facial expressions). When uncertainty was high, observers showed a tendency to over-endorse anger for single faces and especially for crowds. This additional bias was non-linearly related to crowd size, peaking at approximately four constituents. Anger bias presented itself independent of gender, although it was strongest for evaluations of male faces, and it persisted in the

context of variable identities. The bias to report anger occurred regardless of whether angry faces were discriminated against positive (happy) or negative (fearful) expressions. Thus, it did not simply reflect a generic negativity bias, although one may very well exist.

The relationship between bias and crowds was less clear when perceptual information was clear or unambiguous, in which case anger bias either dissipated or even reversed. For example, anger bias shifted to a happy bias for female faces, which is consistent with the idea that in some circumstances female faces may be associated with positive affordances and emotions (Becker et al., 2007). When expressions were easier to discriminate, anger bias was also sometimes stronger for evaluations of single male faces than crowds of male faces. Our sensitivity (d') analyses point to one plausible explanation. When faces are weakly expressive, perceptual uncertainty is high for both crowds and single faces (because diagnostic visual information is scarce), and in this context the increased threat-potential of crowds drives the anger bias. In contrast, the additional numerosity of a crowd may be less impactful for driving bias when its constituents are easy to discriminate. Our data suggest that high perceptual uncertainty and high threat potential may independently and interactively engage anger bias.

We framed our predictions in the context of error-management theory (Haselton et al., 2009; Johnson et al., 2013) and the theory of affordances (Becker et al., 2007; Gibson, 1979; Holbrook et al., 2014). Specifically, we theorized that people overestimate a person's (or crowd's) anger according to their potential to inflict harm, and we proposed that crowds have higher affordance to inflict harm compared to individuals based on their numerosity. Although we did not directly assess theories of error-

management or affordances, our results are consistent with these theoretical frameworks and provide a blueprint for motivating future research.

Real world events are often fraught with fleeting, chaotic, and unpredictable information, requiring quick action. Anger bias may be a heuristic for guiding decision making in the face of these sorts of pressures (Fowler & Johnson, 2011) that can be rewarded and strengthened over time (Levine, 1959). Indeed, negatively biased evaluations of single faces are amplified in the context of unpredictable exposure (Davis et al., 2016) and rapid evaluation, and are diminished when observers are required to delay their judgments (Neta & Tong, 2016). Future work should examine if angry bias with crowds is similarly malleable. Additionally, anger bias may be a heuristic that emerges specifically when perceivers are forced into categorical decisions. This is important to consider given evidence that, when given the chance, people tend to provide multi-dimensional evaluations of facial expressions, attributing a mixture of emotions (Hall & Matsumoto, 2004; Riediger, Voelkle, Ebner, & Lindenberger, 2011). Thus, it is unclear whether a similarly strong anger bias would emerge, especially in crowds, if observers in our investigation had not been forced into a binary decision. There is, nonetheless, value in investigating binary choices as many actions or decisions are ultimately based on coarse categorical evaluations.

Most work on bias in emotion perception has focused on single faces. Conversely, the study of crowd perception, also known as ensemble coding, has been dominated by examination of sensitivity (e.g., Whitney, Haberman, & Sweeny, 2014). Recent work shows that extreme members of a set tend to dominate evaluations of crowds. This phenomenon (for the perception of a more basic visual computation—object size)

occurred presumably because larger objects “pop out” (Kanaya, Hayashi, & Whitney, 2018). Our crowds were homogeneous, and thus could not have depended on this same mechanism. Additionally, the unique pattern of increased sensitivity (i.e., higher d' values) for high-intensity expressions hints that the crowd judgments in our experiments may even reflect an alternative mechanism—probability summation—rather than ensemble integration (Blake, Martens, Garrett, & Westendorf, 1980). Our data also do not allow us to determine that the process of sampling, integrating, and then averaging (or pooling) information across the faces in crowds even occurred in our experiments. Thus, we cannot be certain that ensemble coding occurred in our experiments, nor was this our objective. For these reasons, comparing our results with outcomes from studies that directly examine ensemble coding should be done cautiously. Our findings nonetheless add an important complementary piece to the growing crowd-perception literature, which has largely focused on discrimination rather than bias in perception of facial expressions.

The affective bias we have shown here could potentially be moderated by a perceiver’s internal characteristics, like risk-aversion or anxiety. People with social anxiety tend to attribute a higher emotional cost to interacting with groups than individuals (Douilliez, Yzerbyt, Gilboa-Schechtman, & Philippot, 2012), and tend to interpret ambiguous information as hostile (Bell et al., 2011; Yoon & Zinbarg, 2008). To our knowledge, just one study has examined biased evaluations of crowds among people with social anxiety (Lira, Sang, & Chong, 2013), but the authors used emotionally heterogeneous crowds of full-intensity expressions of happiness and anger, which makes comparison with our results difficult. New research examining how biases may contribute to the maintenance of clinical anxiety is thus important to consider. The strength of anger

bias may also depend on an interaction between an observer's gender identity and the perceived gender of faces in a crowd. We hope that future work will examine this relationship in more detail, taking into account each observer's attitudes about gender and more systematically manipulating gender cues on individual faces and crowds.

Our results are unlikely to be due to differences in the deployment of visual attention, different patterns of fixation across our conditions, or engagement of a visual search process. We acknowledge that because we did not use an eye tracker, we cannot be sure that observers always maintained fixation at the center of the screen. Given the short presentation time we used in each of our experiments, this should not be a concern. Single faces and crowds were always displayed for 100-msec, a duration too brief for observers to have made deliberate saccades to individual faces (Findlay & Walker, 1999) in either condition. Moreover, the faces within any given crowd were always identical to one another, so attending to (or searching for) different faces, even randomly or unintentionally, would not have imparted any benefit to observers. Finally, faces in the single condition were always placed near fixation, so initiation of rapid, longer-distance saccades on these trials was not necessary. The amplified liberal bias in the crowd condition thus appears to be rooted in the effect of numerosity.

Our data do not articulate whether anger bias is rooted in perceptual or cognitive mechanisms. Indeed, changes in criterion can be produced by both response bias as well as perceptual changes (Morgan, Dillenburg, Raphael, & Solomon, 2012; Witt, Taylor, Sugovic, & Wixted, 2015). Future investigations could potentially untangle the contributions of sensory and cognitive mechanisms to anger bias by using methods like drift-diffusion modelling (Ratcliff & Childers, 2015). For example, the liberal bias to

report anger we have shown here could be understood in more depth by considering the resting biases observers bring to each trial, the rate at which they accumulate information about certain expressions in crowds, and the thresholds they use to report seeing anger. Whether faces within crowds actually look angrier than single faces remains an open question. We also hope that our basic research will inspire evaluations of bias in more complex contexts, including naturalistic images, virtual environments, and real social interactions.

We demonstrated that a bias to report anger on individual faces is amplified when people evaluate crowds, and that this bias is engaged most strongly when diagnostic visual information is scarce. Our findings complement the many examples of biased processing of threat when it is more conspicuous, for example in terms of prioritized attention (e.g., Dominguez-Borràs & Vuilleumier, 2013), access to visual awareness (e.g., Capitão et al., 2014), or during visual search (e.g., Huang, Chang, and Chen, 2011). More generally, our work illustrates the value of considering bias in addition to sensitivity, especially for understanding how people see and understand information at the crowd level.

CHAPTER 3: Impact of ratio and intensity on anger bias in mixed-emotion crowds

Introduction

Perceiving groups is an important aspect of daily life. Whether in a meeting, a sporting event, or at a protest, the perceived collective mood may impact the behavior of those within the group as well as those observing the group (Wiltemuth & Heath, 2009). Although people are generally very good at discriminating between emotional expressions, when they do make errors, they are often biased toward negative interpretations (Bar, Neta, & Linz, 2006; Neta & Whalen, 2010). These biases are particularly prevalent in situations of high uncertainty, such as when it is difficult to identify facial expressions due to lack of visibility or subtlety of expressive intensity. Under these conditions, perceivers have a tendency to misinterpret faces as angry or hostile. This negativity bias can serve as a protective mechanism as it biases the errors that perceivers make toward less costly outcomes. Yet, bias in response to crowds is a largely unexplored area of research. We examined the extent to which variability in emotional expression and intensity within a crowd impacts anger bias.

Research on emotion classification biases has primarily focused on evaluation of single faces (Holbrook et al., 2014; Neta & Tong, 2016; Neta & Whalen, 2010). For example, an ambiguous facial expression like surprise can be judged as either positive or negative depending on the context, and when observers are required to make quick judgments, there is a bias toward negative interpretations of surprise (Neta & Tong,

2016). This negative valence bias may reflect early interpretations and result from an initial automatic process in the context of limited time and ambiguous or incomplete information (Davis et al., 2016; Neta & Tong, 2016). For example, there is a tendency to label neutral facial expressions as threatening when observers are exposed to these faces for a mere fraction of a second (Bar et al., 2006). Neutral faces with certain craniofacial characteristics (e.g., a low eyebrow ridge and wide nostrils) are particularly susceptible to anger misattribution (Shasteen, Sasson, & Pinkham, 2015). Importantly, not only are these features associated with temporary emotion states like anger, they are also associated with traits like aggression (Carré, McCormick, & Mondloch, 2009; Said, Sebe, & Todorov, 2009; Stillman, Maner, & Baumeister, 2010). These biases in face perception are consistent with other areas of research that describe a fundamental negativity bias or a positive-negative asymmetry effect, where negative information, particularly in an ambiguous or scarce context, is weighted more heavily than positive information (Baumeister et al., 2001; Cacioppo & Berntson, 1994; Ito et al., 1998; Rozin & Royzman, 2001). For example, people tend to form impressions more quickly and confidently based on negative information about others (Peeters & Czapinski, 1990), overestimate the impact of negative events on their life relative to positive events (Gilbert et al., 1998), and weigh potential losses more heavily than gains (Kahneman & Tversky, 1979, 1984).

These biases may serve an adaptive function, as people that are attuned to detecting threat in their environment are more likely to survive. Neuberg and colleagues (2011; 2016) proposed a precautionary self-protection system which is sensitive to detection of potential harm from others (e.g., an angry facial expression) and responds consistently across perception, affective activation, cognition, and behavioral processes.

In the context of uncertainty, observers may rely on error management strategies, biasing their response toward a less costly error. (Haselton et al., 2009; Johnson et al., 2013).

This “better safe than sorry” approach can in turn lead to endorsements of threat even in its absence. Biases can be especially adaptive to employ because they are quick and easy to engage, which can be useful or even necessary in dangerous situations when there is little time for protracted or deliberate decision making (Johnson et al., 2013). In short, these precautionary biases are most likely to emerge in the context of uncertainty, asymmetric costs, and pressures to make a quick decision.

Interactions with groups can have diverse consequences, as they have the potential to offer security and protection to individuals, but also exposure to interpersonal violence or infectious disease (Neuberg et al., 2011, 2016). Thus, it is important to perceive crowds, and especially socially-relevant information from crowds, quickly and accurately. To accomplish this task, the visual system relies on ensemble coding, which pools and compresses information about the crowd into a summary representation. (Alvarez, 2011; Sweeny & Whitney, 2014). Ensemble coding has been extensively evaluated as a computational process that operates at low-levels of visual analysis to provide summary representations of features such as position (Alvarez & Oliva, 2008), brightness (Bauer, 2009), orientation (Ariely, 2001), and size (Chong & Treisman, 2005), as well as at high-levels of visual analysis to summarize social features, such as identity (Yamanashi Leib et al., 2012), gender (Haberman & Whitney, 2007), gaze direction (Sweeny & Whitney, 2014), and emotion (Elias, Dyer, & Sweeny, 2017). In some cases, judgments of a crowd’s average properties can be more precise than judgments of an individual. For example, Sweeny et al. (2013) demonstrated that observers are better at

estimating the direction of a crowd in motion than the direction of a single individual. When several noisy representations are averaged, the variability of each signal is canceled out, potentially providing a more accurate representation of the group mean. In addition to the mean, ensemble coding may also extract the variance of a crowd, which provides information about the reliability of the mean estimate (Haberman, Lee, & Whitney, 2015). There is some evidence to suggest that observers may discount outliers and engage in precision-weighted averaging, where not all items contribute to the mean equally (Alvarez, 2011). For example, probability, location, and size of certain items may contribute to how heavily they are weighted in the average (Haberman & Whitney, 2010; Kanaya, Hayashi, & Whitney, 2018; Mareschal, Morgan, & Solomon, 2010). Research on ensemble coding, however, has largely focused on sensitivity. Little is known about biases that occur during evaluations of crowds.

If observers are more likely to attribute threat to those with greater capacity for harm (Holbrook et al., 2014; Neel et al., 2012), then based on numerosity alone, crowds should be especially vulnerable to biased evaluations. There is also some evidence to suggest that people may become more violent in crowds due to emotional contagion, which may synchronize and amplify emotions, and deindividuation, which may contribute to a sense of increased anonymity and diffused responsibility (Allport, 1924; Baumeister, Ainsworth, & Vohs, 2016; Le Bon, 1896). Perceivers thus might be more likely to attribute threat to crowds. Indeed, previous work found that observers are biased to classify faces as angry, which is amplified in response to crowds (Mihalache, Lamer, Allen, Maher, & Sweeny, in revision). This anger bias in this previous work emerged in the context of uncertainty, particularly in response to subtle or weak expressions of

intensity. These crowds, however, were homogenous in their emotion and intensity (e.g., a crowd of all happy faces with the same intensity of expression). It is unclear whether emotionally heterogenous crowds would be similarly susceptible to anger bias.

Outside the lab, the individuals that constitute a crowd are unlikely to unanimously express the same emotion at the same intensity and at the same time. The purpose of the present study was to evaluate to what extent anger bias emerges when observers have to navigate contradictory information within the crowd. In Experiment 1, observers evaluated crowds of faces with varying ratios of angry to happy expressions. We predicted that anger bias would emerge in response to more balanced ratios of angry to happy expressions (e.g., 5 to 3) and diminish in response to highly imbalanced ratios (e.g., 6 to 2 or 8 to 0). Additionally, we predicted that the strength of anger bias would interact with emotional intensity, such that observers would show more bias in response to low intensity crowds relative to high intensity crowds.

Experiment 1: How is anger bias impacted by variability in proportion of angry to happy faces within a crowd?

Method

Observers. 47 undergraduate students (37 identified as female, nine identified as male, and one identified as other) were recruited from the University of Denver to participate in Experiment 1. This sample size was selected based on the results of a previous investigation (Mihalache, Lamer, Allen, Maher, & Sweeny, in revision). Based on the large effect size of criterion for crowds versus single faces (Cohen's $d = 1.18$), an N of 30 would provide us with power of .99 in this investigation, assuming an equally

large effect size and number of trials per observer. We predicted a smaller effect size as emotionally heterogeneous crowds would likely introduce more noise in visual processing; thus, we selected an N of 45. All participants provided informed consent and received course credit for their participation. Each observer had normal or corrected-to-normal visual acuity and completed the experiment in a dimly lit room. This study was approved by the Institutional Review Board at the University of Denver and the research was carried out in accordance with the provisions of the World Medical Association Declaration of Helsinki.

Stimuli. Our face set consisted of four White men from the NimStim face set (Tottenham et al., 2002). We used Fantamorph software (Version 5) to create linear interpolations (i.e., morphs) between a neutral expression and two emotional expressions—angry or happy—for each of the actors. There are 50 morphs of emotional expressiveness (ranging from 2% to 100% emotional intensity in increments of 2%) for both expressions. We then used a Gaussian edge-blurring tool (Adobe Photoshop CS6, Version 13.0) to smooth the external contours of the face and hair and to diminish the rough edges of the cropped faces. The background of each image was then replaced with a uniform gray color (RGB = 170,170,170). The stimuli were presented on a uniform background (RGB = 170,170,170; luminance = 27.5 cd/m²) on a CRT monitor with a screen size of 27.3cm x 36.5cm, a resolution of 1024 x 768.

Procedure. We used a within-subjects design in which each observer viewed and evaluated emotional expressions in crowds consisting of happy and angry faces. Faces within each trial consisted of the same identity, but the identity varied between trials. Each crowd consisted of eight faces. Varying ratios of emotions were randomly presented

with the following combinations of angry to happy faces: 8:0, 6:2, 5:3, 4:4, 3:5, 6:2, 0:8.

We excluded the 7:1 and 1:7 angry to happy ratios so that results could not be explained in terms of possible pop-out effects. Critically, all the faces were restricted to the same limited range of intensities (e.g., all medium intensity) within any specific trial (e.g., 3 angry faces and 5 happy faces, all from the respective medium intensity distributions).

The selected emotional intensity distribution, however, varied between trials. This design allowed us to isolate a potential effect of emotion ratio while keeping intensity constant.

Each set was split into three ranges or distributions of emotional intensity: Low (morphs 2%-32% intensity), Medium (morphs 34%-68% intensity), and High (morphs 70%-100% intensity). Faces drawn for each trial were restricted to one of these distributions (e.g., 2 faces selected from the Low intensity happy set and 6 faces selected from the Low intensity angry set). On each trial a random mean was generated that was within 6% of the median (e.g., between 10% and 22% for the Low condition), and then faces were randomly selected (e.g., between 2% and 22% for the Low condition) that would generate that mean from a uniform distribution. We chose this approach because the variability in individual faces eliminated the observer's ability to use a single face as the exact representative of the crowd, thus introducing additional uncertainty.

Faces were randomly distributed within the bounds of an annulus around a centrally presented fixation point (i.e., all the faces were the same distance from fixation, in a donut-shape). The reason for this more uniform organization of faces was to ensure that any potential differences in the weighting of angry versus happy faces was not due to differences in visual acuity (e.g., if the 2 angry faces in a crowd of mostly happy faces appeared near the fixation point, potentially biasing the response toward angry). To

prevent observers from focusing on specific locations, we randomly and independently jittered the location of each face by 1-to-15 pixels both horizontally and vertically on each trial. Crowds were presented for 200-msec, a presentation time brief enough to prevent observers from making deliberate saccades to individual faces (Findlay & Walker, 1999) as well as to prevent the counting of faces to determine the ratio. Each face was followed by a 200-msec blank screen and then followed by a pattern mask (70 rectangular pieces derived from the preceding image, reshuffled into a new image) shown for 250-msec. This approach ensured that the emotional faces and scrambled masks matched in terms of low-level image characteristics, thus decreasing the visibility of the masked image (Enns & Oriet, 2007) and preventing residual visual processing (Rolls, Tovee, & Panzeri, 1999). Observers were required to categorize the crowd as “happy” or “angry” “on average” by pressing the right or left arrow key (counterbalanced across observers). Observers were free to respond as quickly or as slowly as they wished. Observers also received three breaks, which were evenly distributed across the task. Trials were randomized for each block and observer. Each observer viewed and evaluated 40 repetitions of each trial type for a total of 840 trials.

Analyses and Results. We used Signal Detection Theory (SDT) to analyze our data, as it allows for analysis of sensitivity and bias when evaluating whether a signal (in this case “anger”) is present. In this study, there were two possible responses (angry or happy) on trials in which the number of emotional faces in a crowd was imbalanced (e.g., 5 angry, 3 happy faces) resulting in four possible outcomes (see table below): hits (predominantly angry crowds identified as angry), misses (predominantly angry crowds identified as happy), false alarms (predominantly happy crowds identified as angry), and

correct rejections (predominantly happy crowds identified as happy). Based on these four calculations, we derived d' , which is an index of sensitivity, which in this case measures an observer's ability to correctly discriminate between angry and happy. We also calculated criterion, which measures an observer's bias to respond a certain way (i.e., "angry") independent of whether the signal was present or not. The majority of the crowds contained some level of signal (i.e., with the exception of the 0:8 crowd, even the predominantly happy crowds contained some angry faces). We thus categorized crowds of 8:0 (All), 6:2 (Most); and 5:3 (Some) angry to happy faces as "signal-present crowds" and 3:5, 2:6, and 0:8 as "signal-absent crowds" (the 4:4 crowd was excluded from this analysis).

Table 1: SDT analysis parameters for crowds of angry and happy faces.

Response	Signal (Angry)	Noise (Happy)
Yes: "Angry"	Hit	False Alarm
No: "Happy"	Miss	Correct Rejection

The figure and analyses below pertain to findings related to sensitivity (d').

Figure 1 shows two trends: (1) Observers were the most accurate at categorizing emotions when the faces in the crowds exhibited the same emotion and at a high intensity of expression (e.g., all the faces were happy or angry and the intensity of these facial expressions was high). (2) As the ratio of the faces within the crowd became more balanced, observers experienced more difficulty in discriminating the average emotion of

the crowd. Both of these observations were confirmed via the statistical analyses reported below.

Figure 1. Sensitivity (d') for classification of crowds. Each panel depicts performance across changes in the ratio of angry to happy faces, separately for each intensity of expression. Error bars in each panel represent 95% confidence intervals.

We ran a repeated-measures ANOVA using d' as our dependent variable with factors of Ratio [All (8:0/0:8), Most (6:2/2:6), and Some (5:3/3:5)] and Intensity (Low, Medium, and High). There was a significant main effect of Ratio ($F[2, 92] = 132.87, p < .001, \eta_p^2 = 0.74$) and a main effect of Intensity ($F[2, 92] = 122.73, p < .001, \eta_p^2 = 0.73$). Additionally, there was a significant interaction between Ratio and Intensity ($F[4, 184] = 54.38, p < .001, \eta_p^2 = 0.54$). Specific comparisons then confirmed that in the All condition, observers were less sensitive to the Low intensity faces ($M = 0.65, SD = 0.48$) relative to Medium ($M = 2.05, SD = 1.21; t[46] = 9.79, p < .001, d = 1.65$) and High (M

= 2.60, SD = 1.36; $t[46] = 11.63$, $p < .001$, $d = 2.12$) intensity faces. Observers were also more sensitive to the High intensity faces relative to the Medium intensity faces ($t[46] = 7.24$, $p < .001$, $d = 0.43$). We observed the same pattern in the Most condition: observers were less sensitive to the Low intensity faces ($M = 0.32$, $SD = 0.35$) relative to Medium ($M = 0.81$, $SD = 0.46$; $t[46] = 7.91$, $p < .001$, $d = 1.21$) and High ($M = 1.02$, $SD = 0.51$; $t[46] = 10.49$, $p < .001$, $d = 1.65$) intensity faces. Observers were also more sensitive to the High intensity faces relative to the Medium intensity faces ($t[46] = 4.15$, $p < .001$, $d = 0.45$). However, we observed a slightly different pattern in the Some condition: observers were less sensitive to the Low intensity faces ($M = 0.16$, $SD = 0.41$) relative to Medium ($M = 0.37$, $SD = 0.35$; $t[46] = 2.98$, $p = .005$, $d = 0.54$) and High ($M = 0.52$, $SD = 0.38$; $t[46] = 4.49$, $p < .001$, $d = 0.90$) intensity faces. Observers did not significantly differ in sensitivity to the High intensity faces relative to the Medium intensity faces ($t[46] = 2.42$, $p = .019$, $d = 0.42$).

The figure and analyses below pertain to findings related to bias (criterion). Criterion values can be negative or positive, with negative values reflecting a liberal bias to report anger. Figure 2 below shows three trends: (1) Observers were most likely to misclassify crowds as angry when the faces in the crowds exhibited low intensities of facial expressions. (2) As the intensity of facial expressions increased, the bias to misclassify these crowds as angry decreased. (3) Ratio of angry to happy faces did not appear to impact anger bias. These observations were confirmed via the statistical analyses reported below.

Figure 2. Bias (criterion) for classification of crowds. Each panel depicts performance across changes in the ratio of angry to happy faces, separately for each intensity of expression. Error bars in each panel represent 95% confidence intervals.

A repeated-measures ANOVA with criterion as the dependent variable revealed significant main effects of Intensity, $F(2, 92) = 17.43$, $p < .001$, $\eta_p^2 = 0.28$, but no main effect of Ratio, $F(2, 92) = 0.31$, $p = .74$, $\eta_p^2 = 0.01$, and no interaction between Ratio and Intensity $F(4, 184) = 1.23$, $p = .30$, $\eta_p^2 = 0.03$. Observers showed more anger bias in the Low condition ($M = -0.37$, $SD = 0.57$) relative to the Medium ($M = -0.12$, $SD = 0.28$; $t[140] = 7.19$, $p < .001$, $d = 0.60$) and High ($M = -0.07$, $SD = 0.22$; $t[140] = 6.71$, $p < .001$, $d = 0.77$) conditions, which also differed from one another ($t[140] = 2.29$, $p = .02$, $d = 0.19$).

In order to evaluate classification for crowds with an angry and happy majority of faces, we also plotted accuracy as proportion of correct responses separately for each type

of crowd (Fig. 3). The purpose of assessing data this way was to double check that observers were in fact able to recognize happy and angry faces. We were especially interested in comparing the homogenous crowds because these crowds were the only ones without any contradictory information (all faces were angry or happy). Planned comparisons between conditions where the Number of Angry Faces were 8 and 0 revealed that in the High intensity condition, there was no difference in accuracy between an all-happy crowd ($M = 0.85$, $SD = 0.18$) and an all-angry crowd ($M = 0.88$, $SD = 0.16$; $\alpha_{corrected} = .017$, $t[46] = 2.21$, $p = .032$, $d = 0.19$). Thus, if d' is low in other conditions where there is contradictory information but at a high intensity of expression (e.g., multiple high intensity happy and angry faces within the same crowd), this would be unlikely due to an inability to perceive happy or angry expressions. However, there was a significant difference in both the Medium intensity condition (all-happy crowd: $M = 0.77$, $SD = 0.18$; all-angry crowd: $M = 0.84$, $SD = 0.16$; $t[46] = 2.92$, $p = .005$, $d = 0.38$) and the Low intensity condition (all-happy crowd: $M = 0.48$, $SD = 0.20$; all-angry crowd: $M = 0.71$, $SD = 0.21$; $t[46] = 4.17$, $p < .001$, $d = 1.15$). That is, unanimously angry crowds were recognized better than unanimously happy crowds when their expressions were medium or low intensity. These results inform our results of bias at lower intensities of expression; a higher proportion of correct responses for angry crowds and a lower proportion of correct responses for happy crowds reflect a liberal use of the “anger” response category and a tendency to misclassify happy crowds as angry crowds.

Figure 3. Proportion of correct responses for classification of crowds. Each panel depicts performance across changes in the number of angry faces in the crowd, separately for each intensity of expression. Error bars in each panel represent 95% confidence intervals.

We calculated Positive and Negative Affect for each observer based on their responses to the PANAS questionnaire, which they completed before starting the task (Pre-) and after completing the task (Post-). Comparing responses on the Pre- and Post-PANAS questionnaires, there was a reduction in both the Positive (Pre: $M = 28.53$, $SD = 7.20$; Post: $M = 22.89$, $SD = 8.12$; $t[46] = 4.58$, $p < .001$, $d = 0.73$) and Negative affect scales (Pre: $M = 16.21$, $SD = 5.65$; Post: $M = 14.32$, $SD = 4.22$; $t[46] = 3.05$, $p = .004$, $d = 0.38$). These patterns reflect an overall reduction in affect, possibly due to boredom experienced over the course of a lengthy and repetitive computer-based task. To examine if there was a relationship between observers' mood and bias to report faces as angry, we compared their affect scores as measured by the PANAS with their overall criterion

values. There were no significant correlations between observers' criterion and the PANAS scales [all $r_s[47] < 0.19$, all $p_s > .20$].

Experiment 2: How does the anger bias interact with differing intensities of emotions within emotionally heterogeneous crowds?

In Experiment 2, we evaluated to what extent heterogeneity in emotion and intensity of expression impacts anger bias. Specifically, the ratio of angry to happy faces was locked at 4:4, but the intensities of the expressions within the crowd varied (e.g., four faces expressing moderate anger and four faces expressing mild or intense happiness). We predicted that anger bias would be stronger when the two emotions were closer in intensity (e.g., half of the crowd expressing 50% intensity and the other half of the crowd expressing 60% intensity) than when they were further apart in intensity (e.g., half of the crowd expressing 50% intensity and the other half of the crowd expressing 100% intensity). We also predicted that bias would be stronger when the two emotions were in the lower intensity range (low-medium) and weaker in the higher intensity range (medium-high). By conducting this study, we hoped to gain additional insight into whether observers engage anger bias in the context of perceptual uncertainty (i.e., low intensity facial expressions) or if they also engage it when perceptual information is clear, but contradictory (i.e., a crowd of facial expressions that are all higher intensity, but expressing different emotions).

Method

Observers. Forty-seven observers (34 identified as female, 12 identified as male, and one identified as other) were recruited from the Psychology Department's participant

pool at the University of Denver. All participants provided informed consent and received course credit for their participation. Each observer had normal or corrected-to-normal visual acuity and completed the experiment in a dimly lit room.

Stimuli and Procedure. The stimuli and procedures were identical to those in Experiment 1 with the exception that the emotional intensity of happy versus angry faces varied within each crowd while the ratio did not vary—the opposite of Experiment 1. All crowds were restricted to a balanced ratio of 4:4 angry to happy faces across all trials. The intensity of the angry face set varied from the intensity of the happy face set within each crowd. There were two crowd types. The first crowd type featured four happy faces that always had a constant emotion intensity of 50%. The intensity of the angry expressions in these crowds varied from 2%, 10%, 20%, 40%, 60%, 80%, 90%, to 100% (although intensity was kept consistent within a trial). The second crowd type featured four angry faces that always had a constant emotion intensity of 50%. The intensity of the happy expressions in these crowds varied from 2%, 10%, 20%, 40%, 60%, 80%, 90%, to 100%. Table 2 summarizes the combinations of crowds that were included in this experiment. Observers were required to categorize each crowd as “happy” or “angry” “on average” by pressing the right or left arrow key (counterbalanced across observers). Each observer viewed and evaluated 60 repetitions of each trial type for a total of 960 trials.

Table 2. All crowd combinations of emotional intensities included in Experiment 2.

			More Intensely Happy Crowd				More Intensely Angry Crowd			
Crowd Type 1	Constant Emotion	Happy Intensity	50%	50%	50%	50%	50%	50%	50%	50%
	Varied Emotion	Angry Intensity	2%	10%	20%	40%	60%	80%	90%	100%
Crowd Type 2	Constant Emotion	Angry Intensity	50%	50%	50%	50%	50%	50%	50%	50%
	Varied Emotion	Happy Intensity	100%	90%	80%	60%	40%	20%	10%	2%

Analyses and Results. We used SDT analyses to evaluate sensitivity and bias in this experiment as well. As in Experiment 1, we calculated both d' as a measure of sensitivity and criterion as a measure of bias. In our previous experiments we treated crowds with a majority of angry expressions as the signal and crowds with a majority of happy expressions as the noise in order to calculate hits, false alarms, correct rejections, and misses. In the current experiment, however, each crowd contained an equal number of faces with happy and angry expressions. We thus used the relative intensity of the happy and angry faces to classify trials into these categories. For example, a trial in which the angry faces were more intense than the happy faces would be classified as a “signal present” trial, and a judgment of that crowd as “angry” would be recorded as a “hit.” Signal-present crowds included: (1) trials where the happy expressions were set at 50% intensity and the intensity of the angry expressions varied between 60%-100%, and (2) trials where the angry expressions were set at 50% intensity and the intensity of the happy expressions varied between 2%-40%. We wanted to determine sensitivity and bias

as a function of the intensity of the variable emotion. In other words, how biased were observers when one category of emotional expressions (happy or angry) was seen at 50% intensity and the other set of expressions (angry or happy) was seen at 2% intensity? Thus, each crowd type (with the constant emotion set at 50%) served as each other's counterpart when deriving hits, false alarms, misses, and correct rejections (e.g., identifying a crowd with 50% angry and 2% happy faces as "angry" would be a hit whereas identifying a crowd with 50% happy and 2% angry faces as "angry" would be a "false alarm"). Table 3 illustrates an example of how hits, misses, false alarms, and correct rejections were derived for this example of a crowd:

Table 3: An example of how d' and criterion parameters were derived from each crowd type.

Response:	Constant Emotion: Angry 50% Varied Emotion: Happy 2%	Constant Emotion: Happy 50% Varied Emotion: Angry 2%
"Angry"	Hit	False Alarm
"Happy"	Miss	Correct Rejection

Figure 4 below shows that observers were able to better discriminate crowds in which the four faces from the varied-intensity set displayed low levels of expression intensity, and sensitivity decreased as intensity of the varied emotion faces increased. This observation was confirmed with statistical analyses. A repeated-measures ANOVA with d' as the dependent variable revealed the main effect of varied emotion, ($F[7, 322]=63.495, p<.001, \eta_p^2=0.54$). Comparing neighboring emotional intensity points to one another revealed that there were statistically significant differences in sensitivity

between 2% ($M = 0.97$, $SD = 0.48$) and 10% ($M = 0.81$, $SD = 0.44$; $t[46] = 3.01$, $p = .004$, $d = 0.34$); 10% and 20% ($M = 0.53$, $SD = 0.32$; $t[46] = 5.36$, $p < .001$, $d = 0.73$); and 20% and 40% ($M = 0.12$, $SD = 0.24$; $t[46] = 6.24$, $p < .001$, $d = 1.43$), even after correcting for multiple comparisons ($\alpha_{\text{corrected}} = .007$). No other neighboring intensity values significantly differed from one another (all $ts < 1.38$, $ps > 0.17$, $ds < 0.31$).

Figure 4 below shows that observers displayed more anger bias toward crowds in which the four faces from the varied intensity set displayed low levels of expression intensity, and bias decreased as intensity of the varied emotion faces increased. This observation was confirmed with statistical analyses. A repeated-measures ANOVA with criterion as our dependent variable revealed the main effect of varied emotion, ($F[7, 322] = 35.83$, $p < .001$, $\eta_p^2 = 0.44$). Comparing neighboring emotional intensity points to one another revealed that there were statistically significant differences in criterion between 10% ($M = -0.38$, $SD = 0.27$) and 20% ($M = -0.28$, $SD = 0.22$; $t[46] = 4.29$, $p < .001$, $d = 0.38$); and 20% and 40% ($M = -0.13$, $SD = 0.20$; $t[46] = 4.97$, $p < .001$, $d = 0.75$), even after correcting for multiple comparisons ($\alpha_{\text{corrected}} = .007$). No other neighboring intensity values significantly differed from one another (all $ts < 2.18$, $ps > 0.03$, $ds < 0.27$).

Figure 4. Sensitivity (d') and bias (criterion) for classification of crowds. Each panel depicts performance across changes in the varied emotional category's intensity. Error bars in each panel represent 95% confidence intervals.

Because d' and criterion are derived from using data from both crowd types (angry majority and happy majority), it can be difficult to see how each type of crowd (with the constant emotion of angry versus happy) was evaluated by observers. We thus plotted and analyzed accuracy as the proportion of correct responses separately for each constant emotion (Fig. 5). A repeated-measures ANOVA with accuracy as the dependent variable and factors of Constant Emotion (angry, happy) and Intensity (2%, 10%, 20%, 40%, 60%, 80%, 90%, 100%) revealed significant main effects of both Constant Emotion, ($F[1, 46]=73.91, p<.001, \eta_p^2=0.62$) and Intensity ($F[7, 322]=51.86, p<.001, \eta_p^2=0.53$), as well as a significant interaction between Constant Emotion and Intensity ($F[7, 322]=44.44, p<.001, \eta_p^2=0.49$). Planned comparisons between corresponding varied emotional intensity values across Angry and Happy conditions revealed that there were statistically significant differences the two conditions at 2% (Angry: $M = 0.78, SD = 0.12$; Happy: $M = 0.54, SD = 0.11$; $t[46] = 8.69, p < .001, d = 2.08$); 10% (Angry: $M =$

0.76, SD = 0.11; Happy: M = 0.51, SD = 0.10; $t[46] = 10.22$, $p < .001$, $d = 2.41$); 20% (Angry: M = 0.70, SD = 0.10; Happy: M = 0.49, SD = 0.08; $t[46] = 9.16$, $p < .001$, $d = 2.22$); 40% (Angry: M = 0.57, SD = 0.09; Happy: M = 0.48, SD = 0.08; $t[46] = 4.40$, $p < .001$, $d = 1.09$); and 90% (Angry: M = 0.49, SD = 0.07; Happy: M = 0.57, SD = 0.07; $t[46] = 4.74$, $p < .001$, $d = 1.01$), even after correcting for multiple comparisons ($\alpha_{\text{corrected}} = .006$). No other comparisons were statistically significant (all $t_s \leq 2.63$, $p_s > .01$, $d_s < 0.67$).

Figure 5. Proportion of correct responses for classification of crowds. Each panel depicts performance across changes in varied emotion intensity, separately for face pairs in which the constant emotion was either angry or happy. Error bars in each panel represent 95% confidence intervals.

We also compared accuracy for each Constant Emotion at each level of Intensity against chance performance (i.e., 50% accuracy) using a series of one-sample t-tests.

When constant emotion was set to angry, the following intensity levels had accuracies rates statistically significantly above chance performance: 2%, 10%, 20%, 40% ($\alpha_{\text{corrected}}$

= .006, all $t_s \geq 5.30$, $p_s < .001$, $d_s \geq 0.77$). No other levels significantly differed from chance ($\alpha_{\text{corrected}} = .006$, all $t_s \leq 1.77$, $p_s \geq .083$, $d_s \leq 0.26$). When constant emotion was set to happy, the following intensity levels had accuracies rates statistically significantly above chance performance: 60%, 80%, 90%, and 100% ($\alpha_{\text{corrected}} = .006$, all $t_s \geq 3.93$, $p_s < .001$, $d_s \geq 0.57$). No other levels significantly differed from chance ($\alpha_{\text{corrected}} = .006$, all $t_s \leq 2.54$, $p_s \geq .015$, $d_s \leq 0.37$).

We calculated Positive and Negative Affect for each observer based on their responses to the PANAS questionnaire, which they completed before starting the task (Pre-) and after completing the task (Post-). Comparing responses on the Pre- and Post-PANAS questionnaires, there was a reduction in the Positive affect scale (Pre: $M = 26.17$, $SD = 6.45$; Post: $M = 21.11$, $SD = 5.78$; $t[46] = 5.14$, $p < .001$, $d = 0.63$). There was no statistically significant difference between the Pre- and Post- Negative affect scale (Pre: $M = 15.51$, $SD = 6.21$; Post: $M = 14.74$, $SD = 5.25$; $t[46] = 1.02$, $p = .31$, $d = 0.13$). To check if there was a relationship between observers' mood and bias to report faces as angry, we compared their affect scores as measured by the PANAS with their overall criterion. There were no significant correlations between observers' criterion and the PANAS scales [all $r_s[47] < 0.24$, all $p_s > .11$].

Discussion

In Experiment 1, observers navigated crowds of varying ratios of angry to happy faces. Observers were most accurate in classifying crowds of faces expressing the same emotion (e.g., all eight faces displaying anger or happiness) and were the least accurate when categorizing crowds with more balanced ratios (e.g., a crowd of three angry faces

and five happy faces). Observers' ability to discriminate between the two emotions also depended on the intensity of expression; they categorized the emotions of crowds with high intensity of facial expression better than crowds with low intensity of expression. Anger bias emerged most strongly for judgments of crowds with low intensity of expression. However, the ratio of angry to happy faces in a crowd did not impact the extent of this bias. In Experiment 2, observers navigated crowds in which the number of happy and angry faces were always matched, but where one category's expressive intensity was greater than that of the other. In general, observers were most sensitive and most biased when discriminating between crowds with low and medium intensity of expression. More specifically, observers were able to accurately classify crowds with medium intensity of anger and low intensity of happiness, but they struggled with crowds with medium intensity of happiness and low intensity of anger. Although observers showed difficulty discriminating between mixed emotion crowds with medium and high intensities, they did not engage anger bias.

Our main finding, that anger bias emerges in the context of evaluating crowds with low intensities of expressions, is consistent with our hypotheses and previous work (Mihalache, Lamer, Allen, Maher, & Sweeny, in revision). However, our hypothesis that observers would further engage anger bias in response to more balanced ratios of angry to happy faces (e.g., 5:3 compared with 6:2) was not supported. There are several possible explanations: First, when judging crowds with low intensities of expression, observers may be unable to reliably extract the ratio of angry to happy faces. To be able to obtain a count of facial expressions, one needs to be able to identify the expressions first, and observers showed difficulty categorizing even homogenous crowds when their faces

displayed emotion with weak intensity. Second, bias seems to be engaged in the context of perceptual uncertainty, but not necessarily in the context of contradictory information. In other words, observers may have been confident in their perception of both “angry” and “happy” faces within a crowd, but not in their ability to extract an accurate ratio of each. Thus, sensitivity was low in response to more balanced ratios while bias also remained low as faces displayed high intensity facial expressions. Not only are perceivers more likely to make errors under these conditions and in turn reveal their biases, loss of perceptual sensitivity may actually cause latent biases to emerge (Lynn & Feldman Barrett, 2014).

A similar pattern was found in Experiment 2 with bias emerging in response to lower intensity of expression while remaining low in the context of clear, but contradictory information. Specifically, observers had a tendency to misclassify crowds with mildly angry and moderately happy faces as more angry than happy. These findings are consistent with previous research that found that perceivers may accumulate information about low-intensity facial expressions slowly, which in turn allows for random noise to push toward a false alarm (Mihalache, Lamer, Allen, Maher, & Sweeny, in revision). Similar to Experiment 1, observers made many errors in response to mixed-emotion crowds with higher intensity faces, and yet, these errors were not systematically. It is not surprising that observers had difficulty discriminating between medium and high intensity expressions as visual processes tend to become less sensitive farther away from category boundaries (Hamad, 1987; Sweeny, Haroz, & Whitney, 2012), and because facial expressions are categorically perceived (Fugate, 2013). Observers, therefore, are likely encoding these faces as distinctly “angry” and “happy” and are not as sensitive to

their differences in intensity. In other words, they may be confident in their perception of these multiple facial expressions and yet mostly guessing when required to categorize emotionally heterogeneous crowds as either simply “angry” or “happy.” Thus, like in Experiment 1 where observers exhibited low sensitivity and low bias when judging imbalanced but high intensity crowds, observers in Experiment 2 also exhibited low sensitivity and low bias when judging crowds with high intensity and contradictory information. In short, when perceivers are uncertain about their perceptual information, bias emerges. However, in the context of certain but potentially contradictory information, perceivers do not engage systematic bias despite making many errors.

It is important to note that observers in these studies did not interact with the crowds and, therefore, there were no significant consequences from their errors. It is possible that in the context of real-world situations, even if observers are certain of perceptual information, they may nonetheless be hesitant to engage with a crowd with multiple angry members. Additionally, although we did not find a significant relationship between a perceiver’s bias and their affective state at the time of the experiment (as measured by the PANAS), it is possible that heightened affect, whether due to an encounter with an actual crowd or through manipulation in the lab, may exacerbate anger bias. For example, individuals with anxiety tend to interpret ambiguous information as hostile (Bell et al., 2011), are quick to push away images of angry crowds (Lange et al., 2008), and attribute higher emotional cost to interacting with groups than individuals (Douilliez et al., 2012). They also tend to interpret ambiguous information as being hostile (Bell et al., 2011; Yoon & Zinbarg, 2008), and are faster to push away images of crowds when the number of angry members increases (Lange et al., 2008). Aggressive

individuals are also biased to report ambiguously emotional faces as angry (Schönenberg & Jusyte, 2014) and infer hostile intent in behaviors of others (Dodge, 2006; Knyazev et al., 2008). Thus, further studies should examine how state or trait anxiety as well as other characteristics like aggression impact anger bias.

Additionally, the faces in our crowds consisted only of White men, which allowed us to compare anger bias in the context of mixed-emotion crowds to our own previous research on emotionally homogenous crowds (Mihalache, Lamer, Allen, Maher, & Sweeny, in revision), which featured the same stimulus set. However, it is possible that anger bias may diminish or reverse for judgments of crowds of facial expressions of women. Gender is differentially associated with emotion and affordance of threat (Becker et al., 2007; Neel et al., 2012). Masculine faces tend to be associated with angry expressions whereas feminine faces are more likely to be associated with happy or fearful expressions (Hess et al., 2004). Indeed, our own preliminary research suggests that perceivers engage anger bias in the context of low intensity female crowds, but they endorse a happy bias in the context of high intensity female crowds. It is thus important to evaluate how gender, both of the crowd and of the perceiver, may impact anger bias.

Lastly, it is unclear whether anger bias would persist, especially in the context of mixed-emotion crowds, if observers had more time to evaluate them or render their judgments. For example, delaying judgments and using reappraisal strategies can diminish bias (Neta & Tong, 2016). Additionally, when perceivers have the opportunity to scan the crowds serially, they may attend to some facial expressions more than others. For example, some studies have found that perceivers tend to focus more on happy faces (Bucher & Voss, 2019) in a crowd while others found that observers may have difficulty

disengaging from angry faces (Goldenberg, Weisz, Sweeny, Cikara, & Gross, in press). It is thus unclear how attentional processes, both automatic and conscious, may impact anger bias.

In summary, we demonstrated that anger bias emerges when evaluating crowds with low intensities of expression. Although observers had difficulty navigating crowds with emotionally contradicting information, they did not endorse bias if these crowds displayed high intensity facial expressions. In other words, anger bias seems to emerge in the context of perceptual uncertainty, where perceivers are unsure of the content in the faces, but it does not necessarily emerge in the context of any decision uncertainty, such as when a crowd contains clear but contradictory information.

CHAPTER 4: Summary and Limitations

This series of studies demonstrated that people are biased to classify facial expression as angry, particularly in the context of perceptual uncertainty (in terms of subtlety or visibility of facial expressions). When uncertainty was high, observers showed a tendency to over-endorse anger for single faces and especially for crowds. This anger bias was non-linearly related to crowd size, peaking when sets had approximately four constituents. Bias to report anger occurred regardless of whether angry faces were discriminated against positive (happy) or negative (fearful) expressions, indicating that it did not simply reflect a generic negativity bias. Anger bias also occurred with equal strength for crowds made up of identical or unique identities. When perceptual uncertainty was high, anger bias presented itself independent of gender, although it was strongest for evaluations of male faces.

When observers evaluated crowds with both angry and happy constituents, they once again endorsed anger bias in response to crowds with low intensity of facial expression. This bias did not depend on the ratio of angry to happy faces, possibly due to difficulty extracting this information at low intensities of facial expression. Importantly, anger bias emerged when judgments were difficult in the context of unclear or ambiguous perceptual information, but not in the context of visibly clear yet contradictory information. Specifically, perceivers had a tendency to judge crowds as angry when those crowds consisted of happy and angry constituents expressing these emotions with low

intensity. However, despite having difficulty classifying crowds that were closely balanced in the number of angry and happy faces at higher intensities of expression, perceivers did not produce systematic biases toward any particular emotion for their judgments of these crowds. In summary, simply seeing multiple faces or contradictory information about emotion in a crowd did not invariably trigger endorsements of anger. Instead anger bias for crowds seemed to emerge when perceptual sensitivity was low as the result of diagnostic visual information being scarce. Not only are perceivers more likely to make errors under these conditions and in turn reveal their biases, but loss of sensitivity may actually cause latent biases to emerge (Lynn & Feldman Barrett, 2014).

Limitations and future directions

It is important to evaluate the extent to which anger bias would persist outside the lab. On the one hand, real-world events are often fraught with fleeting, chaotic, and unpredictable information, requiring quick action. Being faced with a physical crowd also poses more consequences and potential for harm in a way that judging faces on a screen does not. On the other hand, in the real world, perceivers may also have more time to evaluate the crowd as well as engage other processes. For example, delaying judgments and using reappraisal strategies can diminish bias (Neta & Tong, 2016). Additionally, when observers have the time to serially search the crowds, they may be susceptible to the Angry Superiority Effect (ASE; Hansen & Hansen, 1988; Horstmann & Bauland, 2006; Öhman, Lundqvist, & Esteves, 2001) or the Happy Superiority Effect (HSE; Hunt et al., 2007; Purcell, Stewart & Skov, 1996). Perceivers have been shown as more likely to attend to happy faces, focus on them for longer durations, and evaluate them more accurately (Bucher & Voss, 2019). However, they have also been shown to linger more

on the angry faces in the crowd (Heinrichs and Hofmann, 2001; Goldenberg et al., in press). Additionally, observers in my studies were limited to a binary decision. This was a useful design as it allowed me to conduct signal detection analyses on my data, but it is also true that decisions about behavior are often made based on binary, categorical judgments. However, it is unclear whether perceivers would have engaged in anger bias if they had not been forced into categorical decisions. Research on labeling of facial expressions suggests that when perceivers are allowed to assign labels freely, they often provide multi-dimensional evaluations, attributing a mixture of facial expressions to the same face (Hall & Matsumoto, 2004; Riediger et al., 2011). In short, if perceivers had been given more time and freedom to perceive and evaluate crowds, they may have relied on additional processes that could have alleviated or exacerbated anger bias.

My stimuli were also limited. This choice was intentional. Using a restricted and highly controlled stimulus set provided greater control to test anger bias under strict conditions with high internal validity. This sort of approach has merit for initial investigations of new phenomena, even though it may sacrifice generalizability in the short run (Mook, 1983). Although I included some heterogeneity within crowds in terms of identity, gender, and emotion in later experiments, these variables need to be evaluated more thoroughly. For example, Young and colleagues (2013) found a positivity bias in non-anxious individuals who evaluated crowds containing angry and happy faces as well as male and female faces with full intensity facial expressions. To understand the potential impact of gender on anger bias, it would be important to evaluate more heterogeneous crowds in terms of gender as well as measure gender attitudes in perceivers. The majority of the crowds in my experiments also consisted of male identities while the

majority of my observers were female. Thus, it is also important to examine the extent to which anger bias would be diminished or amplified in the context of in-group and out-group judgments, including gender and race, as well as context dependent in- and out-groups (e.g., fans of opposing sport teams or protesters versus police; Ackerman et al., 2006, other?).

When evaluating the relationship between the perceiver and the crowd, it may also be important to understand the extent to which perceivers view the crowd as a cohesive group with shared goals and values. The elaborated social identity model (ESIM) describes crowd behavior in the context of intergroup relations—not only does the crowd influence its internal members, but the way an outsider views and acts toward the crowd may also impact that crowd's behavior (Drury & Reicher, 2000; Reicher, 1996; Stott & Reicher, 1998). For example, a few researchers have suggested police officers may view crowds, such as sports fans and demonstration groups, as heterogeneous and yet dichotomous, with the majority of constituents being representative of the general public with peaceful intentions and with the minority of the constituents as possessing intent to incite violence (Prati & Pietrantonio, 2009; Stott & Reicher, 1998). They may also view the majority as easily susceptible to be manipulated into violence, and, therefore, the entire crowd is seen as potentially dangerous. In other words, a crowd is perceived as heterogeneous and yet treated as homogenous. According to these studies, police treatment can then influence the crowd members to shift toward collective identity and resist the police. This behavior, in turn, can reinforce the police officers' initial perception that crowds are pathological, and that peaceful individuals can

easily lose themselves and become “irrational and violent” in a crowd. Thus, additional research is necessary to understand the dynamic between the perceiver and the crowd.

Although I gathered some information about the observers’ positive and negative affect at the time of testing, I did not manipulate perceivers’ moods or gather any additional information about their traits. For example, people with social anxiety tend to attribute a higher emotional cost to interacting with groups than individuals (Douilliez et al., 2012). They also tend to interpret ambiguous information as being hostile (Bell et al., 2011; Yoon & Zinbarg, 2008), and are faster to push away images of crowds when the number of angry members increases (Lange et al., 2008). Aggressive individuals are also biased to report ambiguously emotional faces as angry (Schönenberg & Jusyte, 2014) and infer hostile intent in behaviors of others (Dodge, 2006; Knyazev, 2008). Further studies should examine state and trait anxiety as well as other characteristics like aggression or risk-averse attitudes to better understand which individuals may be particularly susceptible to anger bias.

Finally, signal detection theory, conceptualized as a model of optimal decision making, involves three parameters that impact uncertainty and risk (Lynn & Barrett, 2014). The first is the similarity between the targets and foils and underlies perceptual sensitivity. I mainly manipulated this parameter through varying intensity of facial expressions. However, I did not assess the impact of the second parameter—base rates—as both angry and happy crowds appeared in equal frequency. Although people may not have access to many base rates in the real world, it is possible that perceivers calibrate their biases depending on assumptions of how frequently angry or happy expressions are seen. For example, the assumed probability of encountering an angry or violent group at a

rival sports event may differ from a classroom, and thus in turn may impact perceivers' bias. Lastly, future studies should further examine the third parameter—payoff—which refers to the relative costs of misses and false alarms. Although I suggested that miscategorizing a crowd is more costly than an individual, I did not directly assess this. Future studies should manipulate costs associated with different types of errors and how this may impact bias in response to crowds. Additionally, it may be important to consider anger bias in contexts in which asymmetry of costs may be inverted. For example, in some circumstances it may be more adaptive to recalibrate anger bias. If one is starving on an island or fleeing from an attacker at night and then stumbles upon a crowd, it may be more adaptive to assume that the crowd is friendly and may help, as avoiding it may pose greater consequences. In short, additional research is required to understand how different parameters may exaggerate, diminish, or even reverse anger bias.

I demonstrated a bias to report anger on faces when visual information is scarce or ambiguous, and I showed that this bias is amplified when perceivers evaluate crowds relative to individuals. These findings complement the many examples of biased processing of threat when it is more conspicuous, such as in the context of prioritized attention (Dominguez-Borrás & Vuilleumier, 2013), access to visual awareness (Capitão et al., 2014), or during visual search (Gilbert et al., 2011; Huang et al., 2011). More generally, this work illustrates the value of considering bias in addition to sensitivity, especially for understanding how people see and understand information at the crowd level.

References

- Ackerman, J. M., Shapiro, J. R., Neuberg, S. L., Kenrick, D. T., Becker, D. V.,
Griskevicius, V., Maner, J. K., & 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(10), 836–840.
<http://pss.sagepub.com/content/17/10/836.short>
- Adams, R. B., Ambady, N., Macrae, C. N., & Kleck, R. E. (2006). Emotional expressions
forecast approach-avoidance behavior. *Motivation and Emotion*, 30, 179–188.
- Allport, F. H. (1924). The group fallacy in relation to social science. *American Journal of
Sociology*, 29(6), 688-706.
- Alvarez, G. A. (2011). Representing multiple objects as an ensemble enhances visual
cognition. *Trends in Cognitive Sciences*, 15(3), 122–131.
<https://doi.org/10.1016/j.tics.2011.01.003>
- Alvarez, G. A., & Oliva, A. (2008). The representation of simple ensemble visual
features outside the focus of attention. *Psychological Science*, 19(4), 392–398.
<https://doi.org/10.1111/j.1467-9280.2008.02098.x>
- Ariely, D. (2001). Seeing sets: Representation by statistical properties. *Psychological
Science*, 12, 157-162.
- Ashby, F. G., & Townsend, J. T. (1986). Varieties of perceptual independence.
Psychological Review, 93, 154–179.
- Awh, E., Barton, B., & Vogel, E. K. (2007). Visual working memory represents a fixed
number of items regardless of complexity. *Psychological Science*, 18(7), 622–628.
- Bar, M., Neta, M., & Linz, H. (2006). Very first impressions. *Emotion*, 6(2), 269.

- Bar-Haim, Y., Lamy, D., Pergamin, L., Bakermans-Kranenburg, & van IJzendoorn, M. H. (2007). Threat-related attentional bias in anxious and non-anxious individuals: A meta-analytic study. *Psychological Bulletin*, 133, 1–24.
- Barrett, L. F. (2006). Are emotions natural kinds? *Perspectives on Psychological Science*, 1(1), 28–58. doi:10.1111/j.1745-6916.2006.00003.x
- Bauer, B. (2009). Does Stevens’s power law for brightness extend to perceptual brightness averaging?. *The Psychological Record*, 59(2), 171-185.
- Baumeister, R. F., Ainsworth, S. E., & Vohs, K. D. (2016). Are groups more or less than the sum of their members? The moderating role of individual identification. *Behavioral and Brain Sciences*, 39.
- Baumeister, R. F., Bratslavsky, E., Finkenauer, C., & Vohs, K. D. (2001). Bad is stronger than good. *Review of general psychology*, 5(4), 323-370.
- Bayle, D. J., Schoendorff, B., Henaff, M. A., & Krolak- Salmon, P. (2011). Emotional facial expression detection in the peripheral visual field. *PLoS one*, 6(6), e21584. doi:10.1371/journal.pone.0021584.g003
- Becker, D. V., Kenrick, D. T., Neuberg, S. L., Blackwell, K. C., & Smith, D. M. (2007). The confounded nature of angry men and happy women. *Journal of Personality and Social Psychology*, 92(2), 179–190. <https://doi.org/10.1037/0022-3514.92.2.179>
- Becker, D. V., & Srinivasan, N. (2014). The vividness of the happy face. *Current Directions in Psychological Science*, 23, 189 –194. <http://dx.doi.org/10.1177/0963721414533702>
- Bell, C., Bourke, C., Colhoun, H., Carter, F., Frampton, C., & Porter, R. (2011). The

- misclassification of facial expressions in generalised social phobia. *Journal of Anxiety Disorders*, 25(2), 278–283. <https://doi.org/10.1016/j.janxdis.2010.10.001>
- Biernat, M., & Manis, M. (1994). Shifting standards and stereotype-based judgments. *Journal of Personality and Social Psychology*, 66, 5–20.
- Blake, R., Martens, W., Garrett, a, & Westendorf, D. (1980). Estimating probability summation for binocular reaction time data. *Perception & Psychophysics*, 27(4), 375–378. <https://doi.org/10.3758/BF03206129>
- Bombardi, D., Schmid, P. C., Schmid-Mast, M., Birri, S., Mast, F. W., & Lobmaier, J. S. (2013). Emotion recognition: The role of featural and configural face information. *The Quarterly Journal of Experimental Psychology*, 66, 2426–2442.
- Bucher, A., & Voss, A. (2019). Judging the mood of the crowd: Attention is focused on happy faces. *Emotion*, 19(6), 1044.
- Butler, T., Pan, H., Tuescher, O., Engelen, A., Goldstein, M., Epstein, J., Weisholtz, D., Root, J.C., Protopopescu, X., Cunningham-Bussel, A.C., Chang, L., Xie, X.H., Chen, Q., Phelps, E.A., LeDoux, J.E., Stern, E., Silbersweig, D.A., 2007. Human fear-related motor neurocircuitry. *Neuroscience*, 150, 1–7.
- Cacioppo, J. T., & Berntson, G. G. (1994). Relationship between attitudes and evaluative space: A critical review, with emphasis on the separability of positive and negative substrates. *Psychological bulletin*, 115(3), 401.
- Calvo, M. G., & Lundqvist, D. (2008). Facial expressions of emotion (KDEF): Identification under different display-duration conditions. *Behavior Research Methods*, 40(1), 109–115. doi:10.3758/BRM.40.1.109
- Calvo, M. G., & Nummenmaa, L. (2016). Perceptual and affective mechanisms in facial

- expression recognition: An integrative review. *Cognition and Emotion*, 30(6), 1081–1106. <https://doi.org/10.1080/02699931.2015.1049124>
- Carré, J. M., McCormick, C. M., & Mondloch, C. J. (2009). Facial structure is a reliable cue of aggressive behavior. *Psychological Science*, 20, 1194–1198. <http://dx.doi.org/10.1111/j.1467-9280.2009.02423.x>
- Capitão, L. P., Underdown, S. J. V., Vile, S., Yang, E., Harmer, C. J., & Murphy, S. E. (2014). Anxiety increases breakthrough of threat stimuli in continuous flash suppression. *Emotion*, 14(6), 1027–1036. <https://doi.org/10.1037/a0037801>
- Chong, S. C., & Treisman, A. (2005). Statistical processing: computing the average size in perceptual groups. *Vision Research*, 45(7), 891–900. <https://doi.org/10.1016/j.visres.2004.10.004>
- Cisler, J. M., & Koster, E. H. (2010). Mechanisms of attentional biases towards threat in anxiety disorders: An integrative review. *Clinical psychology review*, 30(2), 203–216
- Cosmides, L., Tooby, J., Fiddick, L., & Bryant, G. A. (2005). Detecting cheaters. *Trends in cognitive sciences*, 9, 505–506.
- Davis, C. F., Neta, M., Kim, M. J., Moran, J. M., & Whalen, P. J. (2016). Interpreting ambiguous social cues in unpredictable contexts. *Social Cognitive and Affective Neuroscience*, 5(1), 775–782.
- Dodge, K. A. (2006). Translational science in action: Hostile attributional style and the development of aggressive behavior problems. *Development and Psychopathology*, 18(3), 791–814.
- Dominguez-Borras, J., & Vuilleumier, P. (2013). Affective biases in attention and

- perception. In *Handbook of human affective neuroscience* (pp. 331–356).
Cambridge University Press.
- Douilliez, C., Yzerbyt, V., Gilboa-Schechtman, E., & Philippot, P. (2012). Social anxiety biases the evaluation of facial displays: Evidence from single face and multi-facial stimuli. *Cognition & Emotion*, 26(6), 1107–1115.
<https://doi.org/10.1080/02699931.2011.632494>
- Druckman & R. A. Bjork (Eds.), *Learning, remembering, believing: Enhancing human performance* (pp. 251–276). Washington, DC: National Academy Press.
- Drury, J., & Reicher, S. D. (2000). Collective action and psychological change: The emergence of new social identities. *British Journal of Social Psychology*, 39, 579–604.
- Ekman, P. (1992). Are there basic emotions? *Psychological Review*, 99, 550–553.
doi:10.1037/0033-295X.99.3.550
- Ekman, P., & Cordaro, D. (2011). What is meant by calling emotions basic. *Emotion Review*, 3, 364–370. doi:10.1177/1754073911410740
- Ekman, P., Friesen, W. V., & Hager, J. C. (2002). *Facial action coding system. Investigator's guide*. Salt Lake City, UT: Human Face.
- Elfenbein, H. A., & Ambady, N. (2002). On the universality and cultural specificity of emotion recognition: A meta-analysis. *Psychological Bulletin*, 128, 203–235.
doi:10.1037/0033-2909.128.2.203
- Elias, E., Dyer, M., & Sweeny, T. D. (2017). Ensemble perception of dynamic emotional groups. *Psychological Science*, 28(2), 193–203.
- Enns, J., & Oriet, C. (2007). Visual similarity in masking and priming: The critical role

- of task relevance. *Advances in Cognitive Psychology*, 3(1–2), 211–226.
- <https://doi.org/10.2478/v10053-008-0026-z>
- Fabes, R. A., & Martin, C. L. (1991). Gender and age stereotypes of emotionality. *Personality and Social Psychology Bulletin*, 17, 532–540.
- Fiddick, L., 2011. There is more than the amygdala: potential threat assessment in the cingulate cortex. *Neurosci. Biobehav. Rev.* 35, 1007–1018.
- Finkenauer, C., & Rimé, B. (1998). Socially shared emotional experiences vs. emotional experiences kept secret: Differential characteristics and consequences. *Journal of Social and Clinical Psychology*, 17(3), 295–318.
- Fischer, A. H. (1993). Sex differences in emotionality: Fact or stereotype? *Feminism & Psychology*, 3, 303–318.
- Fiske, S. T. (1980). Attention and weight in person perception: The impact of negative and extreme behavior. *Journal of Experimental Research in Personality*, 22, 889–906.
- Fowler, J., & Johnson, D. (2011). The evolution of overconfidence. *Nature*, 477(7364), 317–320. <https://doi.org/10.1038/nature10384>
- Fugate, J. M. (2013). Categorical perception for emotional faces. *Emotion Review*, 5(1), 84–89.
- Gallup, A. C., Hale, J. J., Sumpter, D. J., Garnier, S., Kacelnik, A., Krebs, J. R., & Couzin, I. D. (2012). Visual attention and the acquisition of information in human crowds. *Proceedings of the National Academy of Sciences*, 109(19), 7245–7250.
- Garcia, J., Hankins, W. G., Rusiniak, K. W., & Bitterman, M. E. (1976). Flavor aversion studies. *Science*, 192(4236), 265–267.

- Garver-Apgar, C. E., Gangestad, S. W., & Simpson, J. A. (2007). Women's perceptions of men's sexual coerciveness change across the menstrual cycle. *Acta Psychologica Sinica*, 39(3), 536-540.
- Gibson, J. (1979). The theory of affordances. In *The Ecological Approach to Visual Perception* (pp. 127–143). Boston: Houghton Mifflin.
- Gilbert, D. T., Pinel, E. C., Wilson, T. D., Blumberg, S. J., & Wheatley, T. P. (1998). Immune neglect: a source of durability bias in affective forecasting. *Journal of personality and social psychology*, 75(3), 617.
- Gilboa-Schechtman, E., Presburger, G., Marom, S., & Hermesh, H. (2005). The effects of social anxiety and depression on the evaluation of facial crowds. *Behaviour research and therapy*, 43(4), 467-474.
- Gilbert, T., Martin, R., & Coulson, M. (2011). Attentional biases using the body in the crowd task: Are angry body postures detected more rapidly? *Cognition and Emotion*, 25(4), 700–708. <https://doi.org/10.1080/02699931.2010.495881>
- Goren, D., & Wilson, H. R. (2006). Quantifying facial expression recognition across viewing conditions. *Vision Research*, 46, 1253–1262. doi:10.1016/j.visres.2005.10.028
- Hall, J. A., & Matsumoto, D. (2004). Gender differences in judgments of multiple emotions from facial expressions. *Emotion*, 4(2), 201–206. <https://doi.org/10.1037/1528-3542.4.2.201>
- Haberman, J., Lee, P., & Whitney, D. (2015). Mixed emotions: Sensitivity to facial variance in a crowd of faces. *Journal of Vision*, 15(4), 16-16.
- Haberman, J., & Whitney, D. (2007). Rapid extraction of mean emotion and gender from

- sets of faces. *Current Biology*, 17(17), R751-3.
<https://doi.org/10.1016/j.cub.2007.06.039>
- Haberman, J., & Whitney, D. (2010). The visual system discounts emotional deviants when extracting average expression. *Attention, Perception, & Psychophysics*, 72(7), 1825-1838.
- Hamilton, D. L., & Zanna, M. P. (1972). Differential weighting of favorable and unfavorable attributes in impressions of personality.
- Hampton, C., Purcell, D. G., Bersine, L., Hansen, C. H., & Hansen, R. D. (1989). Probing “pop-out”: Another look at the face-in-the-crowd effect. *Bulletin of the Psychonomic Society*, 27(6), 563-566.
- Hansen, C. H., & Hansen, R. D. (1988). Finding the face in the crowd: An anger superiority effect. *Journal of Personality and Social Psychology*, 54, 917–924.
doi:10.1037/0022-3514.54.6.917
- Harnad, S. (1987). Psychophysical and cognitive aspects of categorical perception: A critical overview. In *Categorical perception: The groundwork of cognition* (pp. 1-52). Cambridge University Press.
- Hess, U., Adams, R. B., Grammer, K., & Kleck, R. E. (2009). Face gender and emotion expression: Are angry women more like men?. *Journal of Vision*, 9(12), 19-19.
- Hess, U., Blairy, S., & Kleck, R. E. (2000). The influence of facial emotion displays, gender, and ethnicity on judgments of dominance and affiliation. *Journal of Nonverbal behavior*, 24(4), 265-283.
- Haselton, M. G., Bryant, G. A., Wilke, A., Frederick, D. A., Galperin, A., Frankenhui, W. E., & Moore, T. (2009). Adaptive rationality: An evolutionary perspective on

- cognitive bias. *Social Cognition*, 27, 733–763.
- Hatfield, E., Cacioppo, J. T., & Rapson, R. (1993). Emotional contagion. *Current Directions in Psychological Science*, 2, 96–99.
- Hess, U., Adams, R. B., & Kleck, R. E. (2004). Facial Appearance , Gender , and Emotion Expression. 4(4), 378–388. <https://doi.org/10.1037/1528-3542.4.4.378>
- Hess, U., Blairy, S., & Kleck, R. E. (2000). The influence of facial emotion displays, gender, and ethnicity on judgments of dominance and affiliation. *Journal of Nonverbal behavior*, 24(4), 265-283.
- Hogg, M. A., & Tindale, S. (Eds.). (2008). *Blackwell handbook of social psychology: Group processes*. John Wiley & Sons.
- Holbrook, C., Galperin, A., Fessler, D. M. T., Johnson, K. L., Bryant, G. a, & Haselton, M. G. (2014). If looks could kill: Anger attributions are intensified by affordances for doing harm. *Emotion*, 14(3), 455–461. <https://doi.org/10.1037/a0035826>
- Horstmann, G., & Bauland, A. (2006). Search asymmetries with real faces: Testing the anger–superiority effect. *Emotion*, 6, 193–207. doi:10.1037/ 1528-3542.6.2.193
- Huang, S. L., Chang, Y. C., & Chen, Y. J. (2011). Task-irrelevant angry faces capture attention in visual search while modulated by resources. *Emotion*, 11(3), 544–552. <https://doi.org/10.1037/a0022763>
- Hunt, A. R., Cooper, R. M., Hungr, C., & Kingstone, A. (2007). The effect of emotional faces on eye movements and attention. *Visual Cognition*, 15, 513–531. doi:10.1080/13506280600843346
- Ito, T. A., Larsen, J. T., Smith, N. K., & Cacioppo, J. T. (1998). Negative information weighs more heavily on the brain: the negativity bias in evaluative

- categorizations. *Journal of personality and social psychology*, 75(4), 887.
- Izard, C. E. (1994). Innate and universal facial expressions: Evidence from developmental and cross-cultural research. *Psychological Bulletin*, 115(2), 288-299. <https://doi.org/10.1037/0033-2909.115.2.288>
- Jackson, R. E., & Cormack, L. K. (2007). Evolved navigation theory and the descent illusion. *Perception & Psychophysics*, 69(3), 353-362.
- Johansen, J.P., Hamanaka, H., Monfils, M.H., Behnia, R., Deisseroth, K., Blair, H.T., LeDoux, J.E., 2010. Optical activation of lateral amygdala pyramidal cells instructs associative fear learning. *Proceedings of the National Academies of Science*, 107, 12692–12697.
- Johnson, D. D. P., Blumstein, D. T., Fowler, J. H., & Haselton, M. G. (2013). The evolution of error: Error management, cognitive constraints, and adaptive decision-making biases. *Trends in Ecology & Evolution*, 28(8), 474–481.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decisions under risk. *Econometrica*, 47, 263-291.
- Kahneman, D., & Tversky, A. (1984). Choices, values, and frames. *American Psychologist*, 39(4), 341–350.
- Kanaya, S., Hayashi, M. J., & Whitney, D. (2018). Exaggerated groups: Amplification in ensemble coding of temporal and spatial features. *Proceedings of the Royal Society B: Biological Sciences*, 285(1879), 1–9. <https://doi.org/10.1098/rspb.2017.2770>
- Kayyal, M. H., & Russell, J. A. (2013). Americans and Palestinians judge spontaneous facial expressions of emotion. *Emotion*, 13, 891–904. doi:10.1037/a00 33244
- Knutson, B. (1996). Facial expressions of emotion influence interpersonal trait

- inferences. *Journal of Nonverbal Behavior*, 20, 165-182.
- Knyazev, G. G., Bocharov, A. V., Slobodskaya, H. R., & Ryabichenko, T. I. (2008). Personality-linked biases in perception of emotional facial expressions. *Personality and individual differences*, 44(5), 1093-1104.
- Krull, D. S., & Dill, J. C. (1998). Do smiles elicit more inferences than do frowns? The effect of emotional valence on the production of spontaneous inferences. *Personality and Social Psychology Bulletin*, 24, 289-300.
- Kurzban, R., Leary, M.R., 2001. Evolutionary origins of stigmatization: the functions of social exclusion. *Psychological Bulletin*, 127, 187–208.
- Lange, W. G., Heuer, K., Langner, O., Keijsers, G. P., Becker, E. S., & Rinck, M. (2011). Face value: Eye movements and the evaluation of facial crowds in social anxiety. *Journal of Behavior Therapy and Experimental Psychiatry*, 42(3), 355-363.
- Lange, W. G., Keijsers, G., Becker, E. S., & Rinck, M. (2008). Social anxiety and evaluation of social crowds: Explicit and implicit measures. *Behaviour Research and Therapy*, 46(8), 932-943.
- Lazerus, T., Ingbreetsen, Z. A., Stulier, R. M., Freeman, J. B., & Cikara, M. (2016). Positivity bias in judging ingroup members' emotional expressions. *Emotion*, 16(8), 1117–1125. <https://doi.org/10.1037/emo0000227>
- Le Bon, G. (1960). *Psychology of crowds*. (R. K. Merton, Trans.). New York: Viking Press. (Original work published 1896). McDougall, W. (1920). *The group mind: A sketch of the principles of collective psychology with some attempt to apply them to the interpretation of national life and character*. Cambridge, UK: Cambridge University Press.

- Le Bon, G. (1896). *The crowd: A study of the popular mind*. London: Ernest Benn.
- LeDoux, J.E., 2000. Emotion circuits in the brain. *Annual Review of Neuroscience*, 23, 155–184.
- Leith, K. P., & Baumeister, R. F. (1996). Why do bad moods increase self-defeating behavior? Emotion, risk taking, and self-regulation. *Journal of Personality and Social Psychology*, 71(6), 1250–1267. <https://doi.org/10.1037/0022-3514.71.6.1250>
- Levine, M. (1959). A model of hypothesis behavior in discrimination learning set. *Psychological Review*, 66(6), 353–366. <https://doi.org/10.1037/h0044050>
- Lindzey, G. (1954). *Handbook of social psychology*. Cambridge, MA: Addison-Wesley.
- Lira, J. Y. K., Sang, Y., & Chong, C. (2013). Accurate but pathological: Social anxiety and ensemble coding of emotion. *Cognitive Therapy and Research*, 37, 572–578. <https://doi.org/10.1007/s10608-012-9500-5>
- Lynn, S. K., & Feldman Barrett, L. (2014). “Utilizing” Signal Detection Theory. *Psychological Science*, 25(9), 1663–1673. <https://doi.org/10.1177/0956797614541991>
- Maoz, K., Eldar, S., Stoddard, J., Pine, D. S., Leibenluft, E., & Bar-Haim, Y. (2016). Angry-happy interpretations of ambiguous faces in social anxiety disorder. *Psychiatry Research*, 241, 122-127.
- Mareschal, I., Morgan, M. J., & Solomon, J. A. (2010). Attentional modulation of crowding. *Vision Research*, 50(8), 805-809.
- Matsumoto, D., & Hwang, H. C. (2014). Judgments of subtle facial expressions of emotion. *Emotion*, 14, 349–357. doi:10.1037/a0035237
- McIntosh, D. N., Druckman, D., & Zajonc, R. B. (1994). Socially induced

- affect. Learning, remembering, believing: Enhancing human performance, 251-276.
- Milders, M., Sahraie, A., & Logan, S. (2008). Minimum presentation time for masked facial expression discrimination. *Cognition and Emotion*, 22(1), 63–82.
doi:10.1080/02699930701273849
- Milgram, S., Bickman, L., & Berkowitz, L. (1969). Note on the drawing power of crowds of different size. *Journal of Personality and Social Psychology*, 13(2), 79–82.
- Mook, D. G. (1983). In defense of external invalidity. *American Psychologist*, 38(4), 379–387. <https://doi.org/10.1037/0003-066X.38.4.379>
- Morgan, M., Dillenburg, B., Raphael, S., & Solomon, J. A. (2012). Observers can voluntarily shift their psychometric functions without losing sensitivity. *Attention, Perception, & Psychophysics*, 74(1), 185–193.
- Motley, M. T., & Camden, C. T. (1988). Facial expression of emotion: A comparison of posed expressions versus spontaneous expressions in an interpersonal communication setting. *Western Journal of Speech Communication*, 52(1), 1–22.
doi:10.1080/ 10570318809389622
- Naab, P. J., & Russell, J. A. (2007). Judgments of emotion from spontaneous facial expressions of New Guineans. *Emotion*, 7, 736–744. doi:10.1037/1528-3542.7.4.736
- Neel, R., Becker, D. V., Neuberg, S. L., & Kenrick, D. T. (2012). Who expressed what emotion? Men grab anger, women grab happiness. *Journal of Experimental Social Psychology*, 48, 583–586. <https://doi.org/10.1016/j.jesp.2011.11.009>
- Nesse, R. M. (2005). Natural selection and the regulation of defenses: A signal detection analysis of the smoke detector principle. *Evolution & Human Behavior*, 26(1), 88–

105.

- Neta, M., & Tong, T. (2016). Don't like what you see? Give it time: Longer reaction times associated with increased positive affect. *Emotion*, 16(5), 730–739.
- Neuberg, S. L., Kenrick, D. T., & Schaller, M. (2011). Human threat management systems: Self-protection and disease avoidance. *Neuroscience & Biobehavioral Reviews*, 35(4), 1042-1051.
- Neuberg, S. L., & Schaller, M. (2016). An evolutionary threat-management approach to prejudices. *Current Opinion in Psychology*, 7, 1-5.
- Neuhoff, J. G. (1998). Perceptual bias for rising tones. *Nature*, 395(September), 1991–1992.
- Öhman, A., Lundqvist, D., & Esteves, F. (2001). The face in the crowd revisited: A threat advantage with schematic stimuli. *Journal of Personality and Social Psychology*, 80, 381–396. doi:10.1037/0022-3514.80.3.381
- Palermo, R., & Coltheart, M. (2004). Photographs of facial expression: Accuracy, response times, and ratings of intensity. *Behavior Research Methods, Instruments, and Computers*, 36, 634–638. doi:10.3758/BF03206544
- Panksepp, J., & Watt, D. (2011). What is basic about basic emotions? Lasting lessons from affective neuroscience. *Emotion Review*, 3, 387–396. doi:10.1177/1754073911410741
- Peeters, G., & Czapinski, J. (1990). Positive-negative asymmetry in evaluations: The distinction between affective and informational negativity effects. *European review of social psychology*, 1(1), 33-60.
- Phillips, L. T., Weisbuch, M., & Ambady, N. (2014). People perception: Social vision of

- groups and consequences for organizing and interacting. *Research in Organizational Behavior*, 34, 101-127.
- Prati, G., & Pietrantonio, L. (2009). Elaborating the police perspective: The role of perceptions and experience in the explanation of crowd conflict. *European journal of social psychology*, 39(6), 991-1001.
- Pratto, F., & John, O. P. (1991). Automatic vigilance: The attention-grabbing power of negative social information. *Journal of Personality and Social Psychology*, 61, 380-391.
- Proffitt, D. R., "Embodied perception and the economy of action." *Perspectives on psychological science* 1, no. 2 (2006): 110-122.
- Purcell, D. G., Stewart, A. L., & Skov, R. B. (1996). It takes a confounded face to pop out of a crowd. *Perception*, 25, 1091–1108. doi:10.1068/ p251091
- Ratcliff, R., & Childers, R. (2015). Individual differences and fitting methods for the two-choice diffusion model of decision making. *Decision*, 2(4), 237–239.
- Recio, G., Schacht, A., & Sommer, W. (2013). Classification of dynamic facial expressions of emotion presented briefly. *Cognition and Emotion*, 27, 1486–1494. doi:10.1080/02699931.2013.794128
- Reicher, S. (1996). “The battle of Westminster”: Developing the social identity model of crowd behaviour in order to explain the initiation and development of collective conflict. *European Journal of Social Psychology*, 26, 115–134.
- Resnick, L., Levine, J., & Teasley, S. D. (Eds.). (1991). *Perspectives on socially shared cognitions*. Washington, DC: American Psychological Association.
- Riediger, M., Voelkle, M. C., Ebner, N. C., & Lindenberger, U. (2011). Beyond “Happy,

- angry, or sad?": Age-of-poser and age-of-rater effects on multi-dimensional emotion perception. *Cognition and Emotion*, 25(6), 968–982.
- <https://doi.org/10.1080/02699931.2010.540812>
- Rolls, E. T., Tovee, M. J., & Panzeri, S. (1999). The neurophysiology of backward masking: information analysis. *Journal Cognitive Neuroscience*, 11(3), 300–311.
- Rozin, P., & Royzman, E. B. (2001). Negativity bias, negativity dominance, and contagion. *Personality and social psychology review*, 5(4), 296-320.
- Russell, J. A. (1994). Is there universal recognition of emotion from facial expression? A review of the cross-cultural studies. *Psychological Bulletin*, 115(1),
- Said, C. P., Sebe, N., & Todorov, A. (2009). Structural resemblance to emotional expressions predicts evaluation of emotionally neutral faces. *Emotion*, 9, 260–264.
- <http://dx.doi.org/10.1037/a0014681>
- Schönenberg, M., & Jusyte, A. (2014). Investigation of the hostile attribution bias toward ambiguous facial cues in antisocial violent offenders. *European archives of psychiatry and clinical neuroscience*, 264(1), 61-69.
- Schwarz, N. (2001). Feelings as information: Implications for affective influences on information processing. *Theories of mood and cognition: A user's guidebook*, 159-176.
- Senior, C., Phillips, M. L., Barnes, J., & David, A. S. (1999). An investigation into the perception of dominance from schematic faces: A study using the World- Wide Web. *Behavior Research Methods, Instruments, & Computers*, 31, 341–346.
- Shasteen, J. R., Sasson, N. J., & Pinkham, A. E. (2015). A detection advantage for facial threat in the absence of anger. *Emotion*, 15(6), 837–845.

- Sidanius, J., & Pratto, F. (1999). *Social dominance: An intergroup theory of social hierarchy and oppression*. New York: Cambridge University Press.
- Sinclair, R. C., & Mark, M. M. (1992). The influence of mood state on judgment and action: Effects on persuasion, categorization, social justice, person perception, and judgmental accuracy.
- Stillman, T. F., Maner, J. K., & Baumeister, R. F. (2010). A thin slice of violence: Distinguishing violent from nonviolent sex offenders at a glance. *Evolution and Human Behavior*, 31, 298–303. <http://dx.doi.org/10.1016/j.evolhumbehav.2009.12.001>
- Stott, C., & Reicher, S. D. (1998). Crowd action as intergroup process: Introducing the police perspective. *European Journal of Social Psychology*, 28, 509–529.
- Sweeny, T. D., Grabowecky, M., Kim, Y. J., & Suzuki, S. (2011). Internal curvature signal and noise in low-and high-level vision. *Journal of Neurophysiology*, 105(3), 1236-1257.
- Sweeny, Haroz, & Whitney, 2012, Reference repulsion in the categorical perception of biological motion. *Vision Research*, 64, 26-34.
- Sweeny, T. D., Haroz, S., & Whitney, D. (2013). Perceiving group behavior: Sensitive ensemble coding mechanisms for biological motion of human crowds. *Journal of Experimental Psychology: Human Perception and Performance*, 39(2), 329–337.
- Sweeny, T. D., & Whitney, D. (2014). Perceiving crowd attention: Ensemble perception of a crowd's gaze. *Psychological Science*, 25(10), 1903–1914. <https://doi.org/10.1177/0956797614544510>
- Taylor, S. E. (1991). Asymmetrical effects of positive and negative events: the

- mobilization-minimization hypothesis. *Psychological bulletin*, 110(1), 67.
- Thompson, L., & Fine, G. A. (1999). Socially shared cognition, affect, and behavior: A review and integration. *Personality and Social Psychology Review*, 3, 278–302.
- Tooby, J., & Cosmides, L. (2005). Conceptual foundations of evolutionary psychology. *The handbook of evolutionary psychology*, 5-67.
- Tottenham, N., Phuong, J., Flannery, J., Gabard-Durnam, L., & Goff, B. (2013). A negativity bias for ambiguous facial-expression valence during childhood: Converging evidence from behavior and facial corrugator muscle responses. *Emotion*, 13(1), 92.
- Tottenham, N., Tanaka, J. W., Leon, A. C., McCarry, T., Nurse, M., Hare, T. A., ... & Nelson, C. (2009). The NimStim set of facial expressions: judgments from untrained research participants. *Psychiatry research*, 168(3), 242-249.
- Triplett, N. (1898). The dynamogenic factors in pace-making and competition. *American Journal of Psychology*, 9, 507–533.
- Van Vugt, M., & Kameda, T. (2012). Evolution and groups. *Group processes*, 297-332.
- Weinstein, N.D. (1980). Unrealistic optimism about future life events. *Journal of Personality and Social Psychology*, 39, 806-820.
- Whitney, D., Haberman, J., & Sweeny, T. D. (2014). From Textures to Crowds: Multiple Levels of Summary Statistical Perception. 695–710.
- Wiltermuth, S. S., & Heath, C. (2009). Synchrony and cooperation. *Psychological Science*, 20, 1–5.
- Witt, J. K., Taylor, J. E. T., Sugovic, M., & Wixted, J. T. (2015). Signal detection measures cannot distinguish perceptual biases from response biases. *Perception*, 44,

289–300. <https://doi.org/10.1068/p7908>

- Woody, E. Z., & Szechtman, H. (2011). Adaptation to potential threat: the evolution, neurobiology, and psychopathology of the security motivation system. *Neuroscience & Biobehavioral Reviews*, 35(4), 1019-1033.
- Yamanashi Leib, A., Landau, A. N., Baek, Y., Chong, S. C., & Robertson, L. (2012). Extracting the mean size across the visual field in patients with mild, chronic unilateral neglect. *Frontiers in Human Neuroscience*, 6(October), 1–11.
<https://doi.org/10.3389/fnhum.2012.00267>
- Yang, J. W., Yoon, K. L., Chong, S. C., & Oh, K. J. (2013). Accurate but pathological: Social anxiety and ensemble coding of emotion. *Cognitive Therapy and Research*, 37, 572–578. <https://doi.org/10.1007/s10608-012-9500-5>
- Yzerbyt, V. Y., & Leyens, J. P. (1991). Requesting information to form an impression: The influence of valence and confirmatory status. *Journal of Experimental Social Psychology*, 27(4), 337-356.
- Yik, M. S. M., & Russell, J. A. (1999). Interpretation of faces: A cross-cultural study of a prediction from Fridlund's theory. *Cognition and Emotion*, 13, 93-104.
- Yoon, K. L., & Zinbarg, R. E. (2008). Interpreting Neutral Faces as Threatening Is a Default Mode for Socially Anxious Individuals. *Journal of Abnormal Psychology*, 117(3), 680–685. <https://doi.org/10.1037/0021-843X.117.3.680>

Appendices

Appendix A: Investigation 1

Below are plots displaying Hit and False Alarm Rates for every observer in each condition in Experiments 1-6. Each circle represents a Hit and False Alarm Rate combination from a unique observer. Circles to the left of the dashed diagonal line represent combinations that yielded positive d' values. Circles closest to the upper left corner reflect the highest d' values. Circles to the right of the solid diagonal line represent combinations that yielded negative criterion values (anger bias). Circles closest to the upper right corner reflect the most liberal use of the anger category. These scatterplots illustrate how any individual observer can produce high perceptual sensitivity and liberal bias, at the same time.

EXPERIMENT 1:

EXPERIMENT 2:

EXPERIMENT 3:

EXPERIMENT 4:

EXPERIMENT 5:

EXPERIMENT 6:

Appendix B: Investigation 2

Experiment 1: Additional Analyses

The following are supplemental analyses for Experiment 1 from Investigation 2, which allowed me to analyze the equally balanced heterogeneous crowds of 4:4 angry to happy faces. This was not possible with a signal detection theory (SDT) approach, because SDT measures observers' responses relative to an objective "correct response," and a crowd of 4:4 angry to happy faces does not have a correct response to the question, "Is this crowd more angry or more happy?" I calculated point of subjective equality (PSE) and just-noticeable difference (JND) measures (described below). For each ratio of happy to angry faces in a crowd, I calculated the proportion of trials in which each observer indicated a crowd as "angry" on average. I predicted that the likelihood of categorizing crowds as "angry" would increase with ratio of angry to happy faces. However, perceptual judgments are typically nonlinear due to the sensitivity required to make discriminations around category boundaries. In other words, humans tend to be more sensitive to changes around the neutral point and within slightly angry or happy faces (around the categorical boundaries) and tend to perceive very angry and extremely angry faces as just "angry" (less sensitive to changes far away from the boundary). Thus, a logistic function, $y = 1/(1+10^{-(ax+b)})$, with a reflecting the slope of the line (sensitivity) and b reflecting the shift of the line (bias) was likely to be the best fit for observed data.

These logistic fits allowed me to derive two important metrics—the point of subjective equality (PSE) and just-noticeable difference (JND). The PSE reflects the

physical intensity of the stimulus (in this case the number of angry faces) at which an observer classifies the crowd as angry 50% of the time. The JND refers to the bare-minimum amount of change that an observer can reliably detect, measured by the amount of change in the stimulus that leads to a 25% change in the observer's response (the lower the JND the more sensitive the observer). The JND is calculated by measuring the amount of stimulus change required for the observer to shift from responding "angry" 25% of the time to 75% of the time, divided by 2.

I conducted one sample t-tests comparing the average PSE to four faces. If an observer is unbiased, the PSE should be four (i.e., in a crowd of 4:4 angry to happy faces, unbiased observers should be responding "angry" 50% of the time). I calculated PSEs and JNDs for each observer and ran statistical tests over those distributions. In the high intensity condition, the average PSE ($M = 3.61$; $SD = 1.10$) was significantly lower than four, $t(46) = 2.42$, $p = .019$, $d = 0.35$). In the medium intensity condition, the average PSE ($M = 4.75$; $SD = 8.66$) was not significantly different than four, $t(46) = 0.59$, $p = .558$, $d = 0.086$. In the low intensity condition, the average PSE ($M = 1.64$; $SD = 17.16$) was also not significantly different than four, $t(46) = 0.94$, $p = .351$, $d = 0.14$). Additionally, I ran a repeated measures ANOVA with a dependent variable of the PSE shift and the independent factor of emotional intensity (with 3 levels of low, medium, high). There was no main effect of average PSE, $F[2, 90] = 0.68$, $p = .51$, $\eta_p^2 = 0.02$). I also ran a repeated measures ANOVA with a dependent variable JND and the independent factor of emotional intensity. There was no main effect of JND (High: $M = 2.74$, $SD = 2.26$; Medium: $M = 1.53$, $SD = 11.83$; Low: $M = 4.98$, $SD = 23.30$; $F[2, 90] = 0.58$, $p = .56$, $\eta_p^2 = 0.01$).

Below are plots displaying Hit and False Alarm Rates for every observer in each condition in Experiments 1-6. Each circle represents a Hit and False Alarm Rate combination from a unique observer. Circles to the left of the dashed diagonal line represent combinations that yielded positive d' values. Circles closest to the upper left corner reflect the highest d' values. Circles to the right of the solid diagonal line represent combinations that yielded negative criterion values (anger bias). Circles closest to the upper right corner reflect the most liberal use of the anger category. These scatterplots illustrate how any individual observer can produce high perceptual sensitivity and liberal bias, at the same time.

EXPERIMENT 1:

EXPERIMENT 2: