A BAYESIAN CUE INTEGRATION APPROACH TO RACIAL BIAS IN PAIN ASSESSMENT AND TREATMENT

by

Azaadeh Goharzad

A thesis submitted to the Faculty of the University of Delaware in partial fulfillment of the requirements for the degree of Master of Arts in Psychology

Spring 2023

A BAYESIAN CUE INTEGRATION APPROACH

TO RACIAL BIAS IN PAIN ASSESSMENT AND TREATMENT

by

Azaadeh Goharzad

Approved: _____

Peter Mende-Siedlecki, Ph.D
Professor in charge of thesis on behalf of the Advisory Committee

Approved: _____

Tania Roth, Ph.D
Chair of the Department of Psychology and Brain Sciences

Approved: _____

John A. Pelesko PhD
Dean of the College of Arts and Sciences

Approved: _____

Louis F. Rossi, Ph.D.
Vice Provost for Graduate and Professional Education and
Dean of the Graduate College

TABLE OF CONTENTS

ABSTRACT . . . . . **v**

Chapter

1 INTRODUCTION **1**

1.1 Bayesian Cue Integration in Social Cognition 3
1.2 Cue Integration in Pain Care 6
1.3 Proposal 8

2 METHODS **9**

2.1 Participants 9

2.1.0.1 Exclusion Criteria 9

2.2 Stimuli 10
2.3 Procedure 11

2.3.1 Baseline Pain Treatment Task 11
2.3.2 Pain Assessment Task 12
2.3.3 Cue Combination Treatment Recommendations Task. 14

2.4 Analytic Plan 14

2.4.1 Pain Assessment Models 14
2.4.2 Pain Treatment Models 15
2.4.3 Model Fit Assessment 16

3 RESULTS **17**

3.1 Pain Assessment Models 17
3.2 Pain Treatment Models 20
3.3 Subject Level Analysis 21

4 DISCUSSION **23**

4.1 Limitations and Future Directions 25
4.2 Conclusion 27

REFERENCES 28

Appendix

A IRB/HUMAN SUBJECTS APPROVAL LETTER, 2019 37
B IRB/HUMAN SUBJECTS APPROVAL LETTER, 2020 38

ABSTRACT

Patients' reports of subjective pain experience are at least nominally a primary diagnostic cue in assessment and treatment of pain. Despite this, there is low concordance between provider assessments and self-reported pain ratings, such that patient pain is regularly underestimated and undertreated. Such discrepancies in care are particularly stark for Black patients, who receive less adequate pain care compared to White patients. While attending to facial expressions of pain marginally improves concordance in patient-provider pain ratings, it is not clear that this would improve concordance for Black patients given previous work demonstrating blunted recognition of pain on Black faces. Moreover, it is unclear how self-reported pain information is integrated with facial expressions of pain, and whether this integration is similarly biased as a function of patient race. In the present paper we construct three models of pain assessment for both Black and White targets to examine how individuals use facial expression and self-reported pain cues in making holistic judgements of pain intensity as well as subsequent treatment decisions. Overall, we find that the Bayesian Cue Integration model (compared to Face Dominant and Self-Report Dominant models) best predict participant assessments of pain as well as treatment outcomes for both Black and White targets, suggesting that both facial expression and self-report are integrated in pain assessment.

Word Count: 214

Keywords: Social cognition, pain, race, cue integration

Chapter 1

INTRODUCTION

The pain of Black Americans is consistently underdiagnosed and undertreated (Anderson et al., 2009) across a variety of medical contexts (Bonham, 2001; Burgess et al., 2008; Haider et al., 2013; Pletcher et al., 2008; van Ryn, 2002) and intersecting patient identities (Badreldin et al., 2019; Goyal et al., 2015; Lawrence et al., 2020; Weisse et al., 2001). These present-day disparities stand on a foundation of medical neglect, experimentation, and abuse of Black people that historically are at the core of American medical care (Gamble, 1993; Washington, 2006). Current healthcare structures and practices are built on these histories (Bailey et al., 2021; Trawalter et al., 2020) and are modelled after the needs of White patients (Feagin Bennefield, 2014). Medical racism and inequities manifest at myriad macro- and meso-levels of the American healthcare system, including access and utilization (Crocker et al., 2009; Waidmann Rajan, 2000), residential segregation (Acevedo-Garcia et al., 2003; Gaskin et al., 2009; Williams and Collins, 2001), and medical school training (Nieblas-Bedolla et al., 2020; Onyeador et al., 2020; van Ryn et al., 2015).

In concert with broader systemic forces, the devaluation of Black patients' experiences is also reflected at the level of individual patient-provider interactions (Bailey et al., 2021; Bailey et al., 2017; Mateo Williams, 2021). Previous work examining racial biases in providers has primarily focused on few key factors, including concordance of patient/provider identities (Cooper et al., 2003; Hagiwara et al., 2017; Schnittker Liang, 2006; Strumpf, 2011), implicit racial biases (Chapman et al., 2013; Moskowitz et al., 2012; Sabin et al., 2008), and endorsement of stereotypic beliefs about pain tolerance (Hoffman et al., 2016, Wandner et al., 2012). However, more recently, perceptual disparities in pain recognition have been shown to reliably contribute to biases

in care, such that when individuals perceive pain less readily on Black faces (compared to White counterparts) they, in turn, treat them with less pain reliever (Lin et al., under review; Drain et al., in prep; Goharzad et al., in prep; Mende-Siedlecki et al., 2019; Mende-Siedlecki et al., 2021; Mende-Siedlecki et al., 2022). Of course, in real clinical encounters, pain is assessed using a multitude of factors, including both observed facial and bodily expressions of pain, as well as the patient's own self-reported experience of pain (Tait et al., 2009). Indeed, although self-reported pain is often considered the "gold standard" in pain assessment and treatment (Wells et al., 2008; Pasero McCaffery, 2010) and the assessment of self-reported pain is standard clinical practice (Karcioglu, Topacoglu, Dikme, Dikme, 2018), it is not clear that this should minimize racial disparities in pain care. Race-based perceptual disparities are evident even when self-reported pain descriptions preempt judgments of pained faces. In other words, self-reported pain does not appear to exacerbate or attenuate racial bias in thresholds for perceiving pain on Black versus White faces (Goharzad et al., in prep). That said, self-reported pain influences overall thresholds for perceiving pain independent of target race: participants' thresholds for perceiving pain are lower across all faces when self-reported pain was higher. This may suggest that self-reported pain exerts a demonstrable contextual effect on the visual perception of pain, and that it does so in a manner that maintains differences in racial bias (i.e., it does not produce a corrective impact on disruptions in pain recognition on Black faces by White observers).

This prior work assessed whether contextual information such as self-report has a top-down impact on recognition of pain on faces, such that one's perceptual measurement is tuned by additional contextual information. However, it remains unclear whether and how these sources of pain information are integrated to make an overall assessment of pain. Moreover, it is not clear how the evaluation of self-reported pain as a singular cue (e.g., when considered separately from painful facial expressions) may be biased differentially for Black versus White targets, impacting its potential corrective or exacerbating effects on pain assessment. In the present paper, I draw on a vast literature within perception (Körding et al., 2007; Knill Richards 1996; Griffiths et

al., 2010; Tabor et al., 2017) and more recently, affective inference (Saxe Houlihan 2017; Anzellotti et al., 2021; Ong et al., 2015; Zaki, 2013), that proposes Bayesian cue integration models may effectively formalize how we navigate complex and often conflicting social information to arrive at a unitary judgment. In this paper, I extend this framework to pain inference and treatment within the context of racial disparities in pain care.

1.1 Bayesian Cue Integration in Social Cognition

Imagine a situation in which a friend falls off their bike and severely scrapes up their knee. When asked if they are in pain, your friend responds through tensely pursed lips and furrowed brows that their knee is fine and barely hurts at all. You may decide that their knee is indeed fine and barely hurts at all, directly reflecting what they told you. Alternatively, you might think that their pained face, the context of the injury, and the visible condition of their knee belie their insistence that they are not hurt. How do you decide how much pain they are in and in turn, if they need medical attention? While it may not be your responsibility to make treatment decisions for your friend, such inferences about pain and subsequent treatment are made by medical providers every day. Given the variability in pain interpretation both in facial expression and self-report (Prkachin et al., 2007; Kappesser et al., 2006; Deyo et al., 2004), it is unclear how providers incorporate these sources of information to make coherent decisions about a patient's health state and course of treatment. When making treatment decisions, providers may weigh a variety of details, including patients' health history and insurance coverage, the effectiveness of treatment options, and other diagnostic information. However, unlike other forms of treatment, which aim to cure or minimize disease progress, pain treatment serves almost entirely to minimize patient suffering. As such, the patient's inferred internal pain state or extent of suffering is the primary consideration for course of treatment, and self-reported pain is at least nominally a primary factor in pain assessment (Pasero McCaffery 2010; Wells et al., 2008; Dowell et al., 2016; Karcioglu et al., 2018). Despite this, many sources of

information are incorporated into pain assessment beyond self-reported pain—including social factors (Bonham, 2001; De Ruddere et al., 2011; Werner Malterud, 2003), as well as bodily and facial expressions of pain (Fordyce et al., 1984; Deyo et al., 2004; Ruben et al., 2018). Both self-reported pain and facial expressions of pain may be conceptualized as distinct cues that are integrated to make inferences about the patient's internal pain state (Kappesser et al., 2006). In both cases, these cues or sources of information are noisy and incomplete, such that inference about another's pain becomes a process of probabilistic estimation of the most likely internal state of the patient. We may formalize such estimation processes through Bayesian models of cognition and perception, wherein the observer's experience of a stimulus is conceived as a probabilistic estimation based on noisy sensory information as well as current and prior beliefs about the inferred signal or situation (Clark, 2013; Friston, 2010; Hohwy, 2013). What's more, these sources of information may come from multiple processing streams, as in our pain inference example, where facial expressions of pain and self-reported pain experience both serve as cues to the patient's true internal pain state.

Bayesian cue integration frameworks have been used to successfully formalize the probabilistic combination of multimodal sensory signals across a variety of domains and real-world applications (Acerbi et al., 2018; Beierholm et al., 2009; Dobs et al., 2017; Hillis et al., 2004; Knill Richards, 1996; Körding et al., 2004, 2007; Oruç et al., 2003; Rock Victor, 1964; Shams et al., 2000; Welch Warren, 1980). For example, say you are on a hike in an area known to have a population of wild mountain lions. As you're walking you hear some rustling and see some bushes move along the trail- to be safe, you want to walk or run away from what might be a mountain lion. To determine where an observer perceives the mountain lion is hiding, we may derive the joint-cue posterior of the position, where the posterior is proportional to the independent conditional probabilities of the visual (moving bushes) and auditory cues (rustling), normalized by the prior probability of the position (Vilares Kording 2011):

P(position — vision,audition) (P(position — vision) *P (position — audition))/(P

(position))

Importantly, the reliability of each cue is considered independently of the other. Previous models of multisensory perception took a "winner takes all approach" suggesting that perceivers consider only the most reliable cue, ultimately yielding uncertainty that is higher than that of any one cue alone. While in the Bayesian Cue Integration model, the posterior probability of position integrates the reliability of each cue, with the probability that the individual cues are believed to share a common cause. As such, Bayesian Cue-Integration models allow us to estimate not only the perceiver's estimation of the mountain lion position, or some other stimulus value, but the optimal model by which the estimation is generated (Gershman Beck, 2017).

More recently, Bayesian cue integration models have been used to formalize the probabilistic integration of multiple complex social cues toward the inference of emotions and intentions of others (Anzellotti et al., 2021; Baker et al., 2009; Jospe et al., 2020; Ong et al., 2015; Saxe Houlihan, 2017; Wu et al., 2018; Zaki, 2013). For example, when observing the outcome of a social agent playing a standard gambling task, observers made inferences about the agent's emotional response (e.g., happy, angry, or sad) to the gambling outcome (i.e., the value of wins and losses) in an optimally Bayesian fashion (Ong et al., 2015). Observers inferred the probability of an emotion (e.g., P(happy)), given the joint or integrated probability (e.g., P(happy — smiling, win) of the agents facial expression (e.g., P(happy — smiling) and the gambling outcome (e.g., P(happy — win). While paradigms optimized for the assessment of a singular modality of social inference offer powerful examinations of these processes in isolation (e.g., face evaluation, mentalizing, emotion sharing), real world social inferences happen in the context of multiple, sometimes conflicting sources of information most accurately understood in the context of the other cues (Zaki, 2013). For example, while seeing someone crying may typically suggest that they are sad, seeing someone crying at their child's wedding may suggest that they are genuinely happy, given the probability that they are likely to be happy at their child's wedding. Indeed, integrating multiple modes of affective inference improves empathic accuracy (Jospe et al.,

2020), compared to reliance on a single channel of information.

1.2 Cue Integration in Pain Care

Returning to our friend who has fallen off their bike, if they texted you to tell you that what happened and said that they were fine, you may be inclined to believe them given this particular source of information. Alternatively, you may see a picture of their knee and determine that the injury does not look so bad, so their pain state may be correspondingly low. However, if they share this information in a context where you can see their pained expression (e.g., face-to-face), you may be inclined to believe that they are genuinely in some pain, despite their mild self-report. Alternatively, you may be more likely to dismiss their claim all together if you know them to be a particularly prideful individual. In this way, your prior knowledge or beliefs about your friend's personality may influence your beliefs about the causal relationship between these two sources of information (their self-report and their facial expression) and inform your overall assessment of their internal pain state.

While we may not expect such lay theories of personality, mental state inference, or emotion to influence medical decision making, given the availability of prescriptive guidelines for pain treatment (Dowell et al., 2016), providers rely heavily on factors outside of formal medical guidelines (Prkachin et al., 2007; Tait et al., 2009). Providers are more likely to rely on internal inferences about patients' pain experience when pain etiology is ambiguous or when "objective" diagnostic evidence is lacking (e.g., an MRI that could support or validate the existence of back pain). In such cases, providers tend to incorporate other information, such as observed pain behaviors (i.e., facial expressions or motor movements), self-reported pain, and patient social characteristics (e.g., race, gender, or age) into their assessment. For example, nurses are more likely to rely on their own internal assessment of a patient's pain experience based on the patient's facial expression than the patient's self-report (McCaffery et al., 2000). Consequently, providers' inferences are often inconsistent with the patient's own assessment of their pain. Indeed, a considerable clinical literature demonstrates that providers regularly

discount patient pain (De Ruddere et al., 2013; Prkachin et al., 2007; Ruben et al., 2018; Tait et al., 2009).

Not only are patient-provider divergences in pain assessment well-documented, but these gaps are further exacerbated by patient social characteristics such as age (Tait Chibnall, 2002), race (Tait Chibnall, 2014), gender (Hoffmann Tarzian, 2001), and the extent to which providers "like" the patient (De Ruddere et al., 2011). Moreover, inconsistency between patient reports and provider assessments of pain are often higher for high severity self-reported pain, suggesting that providers are often credulous of patient pain reports when severity is high, and particularly so if they believe "objective" medical evidence is unavailable, ambiguous, or inconsistent with severity of patient's reported pain (Chibnall et al., 1997; De Ruddere et al., 2014; Lieberman et al., 1996; Zalon, 1993). Notably, the effects of medical evidence ambiguity appear to primarily affect assessment of pain for White patients in experimental settings, such that while White patient pain is assessed more conservatively when medical evidence is ambiguous, assessments of Black patients' pain are unmoved (Hirsh et al., 2015). This may suggest that for providers treating Black patients in pain, neither facial expressions nor self-reported pain are considered reliable cues to pain assessment, whereas for White patients, information about pain experience is weighted in overall assessments, such that conflict between multiple cues to pain impacts assessments of overall pain in White patients more than in Black patients.

Despite emphasis on self-reported pain as a primary diagnostic cue for pain assessment (Max et al., 1995; Mularski et al., 2006; Phillips, 2000; Wells et al., 2008), evidence suggests that evaluations of self-reported pain do not provide improvements in accurate assessment of patient pain (Mularski et al., 2006; Ruben et al., 2018; van Dijk et al., 2012). In fact, reliance on facial expressions appears to provide marginally more accurate assessments of patient pain, and the integration of multiple cues (self-reported pain assessed in multiple ways, facial expressions, etc.) provide most accurate assessments (Ruben et al., 2015, 2018). However, the relative efficacy of these different cues has been assessed primarily in White patients, or without explicit analysis of

patient race/ethnicity. Given that pain is observed less readily on Black faces (Lin et al., under review), it is not clear that reliance on facial cues should provide improvements in pain assessment for Black patients as it may for White. What's more, the role of self-reported pain for Black patients is unclear. Self-reported pain experience or contextual information about pain may not be weighed at all in pain assessment (Hirsh et al., 2015), it may be weighed such that it does not influence the magnitude of racial bias in the perception of painful expressions (Mende-Siedlecki et al., 2021), or it may be weighed differentially based on varying expectations about pain and race (Anastas et al., 2020).

1.3 Proposal

The present paper quantifies the integration of multimodal sources of congruent and incongruent information (self-report and facial expression) about patient pain to make holistic inferences about an individual's internal pain state based on the race of the patient. We propose that a Bayesian cue integration model, wherein each cue or source of information is weighted proportionally to the perceived reliability of either cue, will most tightly track empirical pain judgements when both cues are presented simultaneously. We additionally consider two other cue segregation models where either only facial expressions or only self-reported pain experience is incorporated into overall pain assessments We construct and test each of these three models for Black and White targets separately to assess whether observers rely on different models for pain assessment across race. While we predict that the Bayesian cue integration model will best fit participant responses for both Black and White targets, we expect the relative weight of either self-report or face cues to vary between the two models.

Chapter 2

METHODS

2.1 Participants

We recruited 229 White/Caucasian (Mage = 33.88, SDage = 10.92; 122 men, 96 women, 11 non-binary) participants living in the United States [1] from Prolific. Fifty-eight participants were excluded for failing attention and/or testing environment checks, yielding a final sample of 171 participants.

2.1.0.1 Exclusion Criteria

Sample size, exclusion criteria, procedures, stimuli, and hypotheses were pre-registered (https://osf.io/zb4rm). Given recent spikes in fraudulent participation in online recruitment platforms, involving non-US participants using virtual private servers (Kennedy et al., 2020), we screened out any individuals using VPS/VPNs prior to data collection, using an established procedure (Winter, Burleigh, Kennedy, and Clifford, 2019). Pre-registered exclusion criteria included 1) non-differentiation of responses, where non-differentiation was defined as a response between 299 and 301 on 90% of treatment trials or greater. This value was chosen as the slider marker on treatment tasks reset to the center of the scale (300 mg) on each trial.

In addition, we excluded participants based on a series of post-task questions regarding 1) task interruption, 2) failure to comply with presentation requirements,

[1] We did not restrict based on participants' duration of residence in the United States, which may have possibly increased between-subject variability given differences in a) self-reported adoption of racial stereotypes and b) interracial contact. Indeed, meta-analysis across our previous work on racial bias in pain perception and treatment suggests that while both factors are positively associated with these biases, those these effects are rather small.

and 3) failure to comply with attention requirements. Participants were asked if they completed the task alone and excluded from analyses if they responded with: "I was interrupted several times while completing the main task." When asked about their lighting and brightness, participants who answered, "I completed the main task in a dimly-lit room but left my brightness level below 100%," or "I completed the main task either outdoors or in full lighting, and left my brightness level below 100%" were also excluded. Finally, participants were excluded from the analyses if when asked about their attention, answered "I was actively watching or listening to something else while completing the main task," "I was working on other things while completing the main task," or "I had music or other audio on in the background while completing the main task". (No additional attention checks were embedded during the experiment itself.)

2.2 Stimuli

We selected three digitally rendered pain expressions from the Delaware Pain Database (DPD; Mende-Siedlecki, Qu-Lee, Goharzad, & Drain, 2020) and created Black and White versions of 38 randomly generated male target heads in FaceGen. [2] Here we use "expressions" to refer to individual combinations of changes in facial action units associated with pain (e.g., Kunz, Meixner, & Lautenbacher, 2019; Prkachin & Solomon, 2008; Williams, 2002), each of which has robustly been rated as resembling pain (see below). Rather than use one "canonical" expression of pain, we typically use multiple expressions in our work to account for the variability across pain displays (e.g., eyes open versus eyes closed, teeth gritted and bared versus mouth clenched,

[2] Target race (Black versus White) and gender (men or male presenting) here are defined in terms of slider values in FaceGen Modeller Pro (v3.18). Black target faces were set to a value of 1.5 on the 'African' labeled slider and-1.5 on the 'European', 'East Asian', and 'South Asian' labeled sliders. White target faces were set to a value of 1.5 on the 'European' slider and -1.5 on the 'African', 'East Asian', and 'South Asian' sliders. All faces were set to a value of -1 on the gender slider, which ranges from masculine (negative values) to feminine (positive values). Race categorizations and evaluations of masculinity/femininity in prior work bear out the effectiveness of these manipulations (e.g., Drain et al., in prep; Goharzad et al., in prep).

etc.) and to enhance the generalizability of our findings (Mende-Siedlecki et al., 2019; Mende-Siedlecki et al., 2021; Mende-Siedlecki et al., 2022).

We then created a set of stimuli by combining each version of each target face with each pain expression and rendering each combination at various intensities using FaceGen Modeller Pro. The resulting images were vignetted to remove the appearance of baldness (as in Freeman, Stolier, Ingbretsen, & Hehman, 2014). Based on previous norming of the DPD (N = 81; see Study 2 in Mende-Siedlecki et al., 2020), the three pain expressions selected have been rated as looking more like pain on average than any other emotion (MPain = 5.25 out of 7; all comparison emotion Ms ¡ 3.26, all comparison ps ¡ .0001). Using these targets and expressions, we created 96 individual stimuli which varied in race (48 Black, 48 White) and pain intensity, such that no individual target was repeated in any trial across all pain judgement and treatment blocks. Each target varied in pain intensity (0% [Neutral], 40%, 50% [Ambiguous], 80%, 90% [High]). Target/expression pairs varied between task versions, such that all participants saw a given target and expression pair but did not see both Black and White versions of said target. Varying these combinations across participants was done simply to minimize suspicion that could have arisen if a given participant saw similar-looking targets varying only in race and did not constitute a manipulation of interest. Participants saw neutral versions of all stimuli, followed by Fourier phase scrambled images of each individual face stimuli as visual masks for the respective neutral face (created using the Matlab function imscramble; Hebart, 2009).

2.3 Procedure

2.3.1 Baseline Pain Treatment Task

Participants first completed a treatment task to assess their baseline treatment recommendations for a given self-reported amount of pain, independent of any specific patient (e.g., target face) experiencing it or facial display of pain (e.g., expression intensity). Participants were asked how much pain reliever they believe is appropriate for a given amount of pain ranging from 1 to 10. While factors like age and weight may

influence such prescriptions in treatment contexts, our goal here was to understand participants' initial reference points for prescription at each step of this scale.Participants were oriented to a common 0-to-10 Numeric Rating Scale (NRS; Bijur, Lattimer, & Gallagher, 2002), where 0 was described as representing "no pain at all" and 10 was described as representing "the worst imaginable pain" before beginning the task.

Each level of pain from 1 to 10 (inclusive) appeared twice across the course of the baseline treatment task. Participants were given no additional information about the nature of the pain in this task (i.e., type of pain, individual experiencing the pain, facial expression, etc.). They were instructed to make a recommendation based purely on what they think is appropriate for each level of pain. In both the baseline and cued treatment recommendations task, participants were asked to recommend treatment on a scale from 0-600mg of a generic pain reliever similar to Advil or Tylenol. Stimulus information pertaining to pain level was available on screen throughout participants response, and response times were not limited. Each trial concluded only after a response was registered on the treatment slider. Each trial was followed by an inter-trial interval (ITI) of 1500ms.

2.3.2 Pain Assessment Task

Participants then completed a pain judgement task in which they saw 64 targets (32 Black men, 32 White men). On each trial, participants first saw the target face with a neutral expression and a randomly generated patient ID number for 1500ms, followed by a phase scrambled mask of the neutral face for 1000ms. Following the mask, participants were asked to judge how much pain they thought the patient was in based on either single or combined cues. On single cue (self-report only or face only) trials, participants received either the patients' self-reported pain judgement alone (e.g., "This person said their pain was around a 5."; where pain severity was either ambiguous (4 or 5) or high (8 or 9)), or the patient's facial expression of pain alone (either ambiguous [40-50% intensity morphs], or high [80-90% intensity morphs] in severity). Single cue conditions consisted of 16 trials each, where 8 trials presented

Black targets and 8 presented White targets. On joint cue trials, participants saw both self-reported pain and facial expressions presented together on each trials. Trial order randomized within subjects (32 Joint Cue trials, 16 Face Only trials, and 16 Self-Report Only trials). Within Joint Cue trials, cues were either completely congruent (such that facial expression and self-reported pain were at the same level [e.g., a 40% pain expression combined with a self-report of 4 out of 10 pain]) or incongruent (such that ambiguous expressions were paired with high severity expressions and vice versa [e.g., a 40% pain expression combined with a self-report of 8 out of 10 pain]). The disparity in pain cues always amounted to a difference of four units. The Joint Cue condition consisted of 8 congruent (High, Low pain) and 8 incongruent trials (High SR pain, Low Expression pain; Low SR pain; High Expression pain) for each target race (i.e., 16 incongruent and 16 congruent trials in total).

As previously described, targets were never repeated for a given subject throughout the experiment (e.g., if a participant saw a particular target stimulus in the self-report only condition, they would not see that target in any other pain rating or treatment condition). Target race was counterbalanced across two versions of the task such that target heads which were Black in version 1 were White in version 2 for all conditions.

Participants were instructed to rate how much pain they believed the target was in, again orienting them to a typical 0-to-10 NRS pain scale (0 = "no pain at all," 10 = "the worst imaginable pain"). Participants made their ratings using an unmarked slider where the left side of the scale was anchored as "less painful" and the right as "more painful". This level of scale abstraction was necessary to avoid participants defaulting to the exact or proportional self-report cue values. As in the baseline treatment task, stimulus information pertaining to pain level (i.e., facial expression of pain, self-reported pain, or both) was available on screen throughout participants response, and response times were not limited. Each trial concluded only after a response was registered on the treatment slider. Each trial was followed by an inter-trial interval (ITI) of 1500ms.

2.3.3 Cue Combination Treatment Recommendations Task.

The cue-combination treatment recommendation task followed the Pain Assessment Task. Participants were given a short 3500ms break between tasks. The structure of the cue combination treatment trials mimicked the joint cue trials of pain rating task completely for 32 total trials (16 Black, 16 White). The only divergence from the pain rating task was that participants were asked to give treatment recommendations using the 0-600 mg scale used in the baseline treatment task as opposed to explicitly assessing how much pain the target was in. Again, participants saw an entirely new set of targets on treatment trials, as no target was repeated across the experiment.

2.4 Analytic Plan

2.4.1 Pain Assessment Models

We constructed three theoretical observer models (one Bayesian Cue Integration model, and two single cue models: Face Dominant model and Self-Report Dominant model). Single cue models approximate either Face dominant or Self-Report dominant pain assessment styles in line with 'Winner Takes All' models of multisensory decision (Burr, 2004), such that the model associated with the more reliable cue should predict pain assessments on Joint Cue trials regardless of the value of the other cue or congruency between the cues. We were underpowered to construct a full integration model for each subject individually instead, all models were constructed by collapsing equivalent trials across participants (Ong et al., 2015). We constructed each pain assessment model separately for Black and White targets in order to assess differences in pain assessment style as a function of target race, yielding six models in total. For all models, we assumed that an ideal observer would choose the pain rating value closest to their internal estimate of how much pain they assessed the target to be in (e.g., that when participants responded with a pain assessment of 6, that this value reflected an internal estimate with a gaussian distribution around 6) (Ma, 2019).

To test Face Dominant and Self-Report Dominant pain assessment styles, we constructed empirical distributions for P(Pain rating — Face Intensity) and P(Pain

rating — Self-Report) for each condition (i.e., for each possible value of Face or Self-Report) from raw density smoothed ratings on the Face only and Self-report only trials. We performed density estimation using SciPy's Gaussian kernel density estimation function with default settings (Virtanen et al., 2020). . As such, the models reflect not the objective value of the stimulus as set by the experiment, but rather the optimal response on a Joint Cue trial, based on ratings from the single cue trials for a given pain intensity. For example, if a given joint cue trial consists of a Face objectively set to pain intensity of 8, and a self-reported pain value of 4, the Face Dominant model estimates the probability of a pain rating (1-10), based on the probability of that pain rating given that it were a Face only trial. Using the single cue (Face Dominant, and Self-Report Dominant) models we constructed an empirical density smoothed probability distribution reflecting the full Bayesian cue integration model for the conditional probabilities of each combination of Face and Self-Report (see Ong et al., 2015).

Using the full Bayesian Cue Integration model, we then calculated expected pain ratings (i.e., a value from 0 to 10) for each Face and Self-Report combination, for Black and White targets. Finally, we used the same Gaussian kernel density estimation procedure to derive the true distribution on Joint Cue trials for each condition collapsed across participants; it is this distribution to which we can compared our model predictions (Körding et al., 2007; Ma, 2019).

2.4.2 Pain Treatment Models

For pain treatment, we first modelled the empirical distribution for baseline treatment recommendations for each pain value (1-10) using the Gaussian density procedure described previously. Again, all models were collapsed across participants and each model was calculated separately for Black and White targets. We then used the optimal response estimates from each pain judgement model (Face Only, Self-Report Only, and Bayesian Cue Integration), conditioned on the probability distribution for baseline treatment recommendations, to obtain treatment estimates for each model,

for each condition (i.e., the expected treatment recommendation under a particular joint cue condition, given the pain estimate from a specific model, e.g., P(Treatment Amount — Face Only estimate)). Just as with the pain judgment models, we derived the empirical response distribution for treatment recommendations to which we compared model estimates.

2.4.3 Model Fit Assessment

We calculated correlations and the root mean squared error (RMSE) to assess model fit against the empirical distribution on Joint Cue trials for both judgement and treatment tasks respectively. All reported estimates and 95% confidence intervals were obtained from bootstrapped calculations with 5000 replicates. Finally, we assess cue reliability using Shannon's Information theoretic Entropy (Shannon, 1948), H, to expand on potential mechanisms underlying each model.

Chapter 3

RESULTS

3.1 Pain Assessment Models

Comparing overall RMSE for each model[1], collapsed across cue conflict conditions (both congruent and incongruent), the *Cue-Integration* model had the lowest RMSE of all three pain assessment models for both Black ($RMSE_{BlackTargets}$.442 [.30, .58]), and White targets ($RMSE_{WhiteTargets}$ =.347 [.26, .43]), outperforming both *Self-Report Dominant* $RMSE_{BlackTargets}$ =1.05 [.85, 1.24]; $t(31)$ = -4.89, p <.001; $RMSE_{WhiteTargets}$ =.959 [.78, 1.12]), $t(31)$ = -6.14, p <.001)[2] , and Face-Dominant models $RMSE_{BlackTargets}$ =.967 [.79, 1.12], $t(31)$ = -4.61, p <.001; $RMSE_{WhiteTargets}$= 1.10 [.87, 1.33], $t(31)$ = -6.21, p <.001). In kind, the *Cue-Integration* model was most strongly correlated with the empirical data ($r_{BlackTargets}$ = .941 [.86, .98]; $r_{WhiteTargets}$ =.962 [.92, .98]) compared to both, *Self-Report Dominant* ($r_{BlackTargets}$ =.671 [.47, .81], $t(31)$ = 3.25, p= .003; $r_{WhiteTargets}$ =. 732 [.56, .84], $t(31)$ = 3.51, p <.001) and Face-Dominant models ($r_{BlackTargets}$ =.70 [.55, .81], $t(31)$ = 3.48, p= .002; $r_{WhiteTargets}$ =.619 [.41, .78]; $t(31)$ = 3.10, p= .004).

We performed additional exploratory analyses, subsetting trials by conflict condition. When the self-report and facial pain cues were congruent the Bayesian *Cue-Integration* model again had the lowest RMSE, and highest correlation with empirical

[1] For all RMSE and r values related to model fit for both pain assessment and treatment, we report bootstrapped parameter estimates and 95% confidence intervals that provide non-parametric assessment of statistical significance. However, we also provide t-statistics, for ease of comparison between models, noting that these tests may not meet parametric assumptions (see Ong et al., 2015).

[2] Because analyses are collapsed across participants, degrees of freedom reflect trials unless otherwise noted

responses among all three pain assessment models for both Black ($RMSE_{BlackTargetsCongruent}$ = .276 [.18, .37]; $r_{BlackTargetsCongruent}$ = .989 [.98, 1.00]), and White targets ($RMSE_{WhiteTargetsCongruent}$ = .271 [.20, .34]; $r_{WhiteTargetsCongruent}$ = .988 [.98, 1.00]). The Bayesian *Cue-Integration* model fit was significantly better than the Face-Dominant model within both Black ($RMSE_{BlackTargetsCongruent}$ = .629 [.45, .82], $t(15)$ = -3.47, p= .003; $r_{BlackTargetsCongruent}$ = .928 [.86, .98], $t(15)$ =2.28, p= .038) and White targets ($RMSE_{WhiteTargetsCongruent}$ = .664 [.46, .86], $t(15)$ = -3.64, p=.002; $r_{WhiteTargetsCongruent}$ = .926 [.85, .98], $t(15)$ = 2.27, p=.038). However, there was no significant difference in RMSEs between the *Cue-Integration* and *Self-Report Dominant* models within both Black $RMSE_{BlackTargets}$Congruent = .33 [.23, .43]; $t(15)$ = -.817, p= .427; $r_{BlackTargetsCongruent}$ = .989 [.98, 1.00], $t(15)$ = -0.012, p= .990) and White targets($RMSE_{WhiteTargetsCongruent}$ = .481 [.31, .66]), $t(15)$ = -2.25, p= .040; $r_{WhiteTargetsCongruent}$ = .974 [.95, 1.00], $t(15)$ = 1.80, p=.09).

When trial types were subset to incongruent conditions (i.e., facial expression pain intensity was high while self-reported pain was ambiguous and vice versa), again the *Cue-Integration* model had the lowest RMSE of all three models for both Black $RMSE_{BlackTargetsIncongruent}$ = .561 [.37, .75]) and White targets ($RMSE_{WhiteTargetsIncongruent}$ = .410 [.28, .53]) and differences in RMSE were significant for both *Self-Report Dominant* ($RMSE_{BlackTargetsIncongruent}$ = 1.45 [1.30, 1.60], $t(15)$ = -7.08, p <.001; $RMSE_{WhiteTargetsIncongruent}$ = 1.26 [1.09, 1.42], $t(15)$ = -7.82, p <.001) and *Face Dominant* models ($RMSE_{BlackTargetsIncongruent}$ = 1.21 [1.01, 1.40], $t(15)$ = -4.59, p <.001; $RMSE_{WhiteTargetsIncongruent}$ = 1.40 [1.13, 1.70], $t(15)$ = -6.67, p <.001) across both Black and White targets. Similarly, *Cue-Integration* model correlations were highest within Black ($r_{BlackTargetsIncongruent}$ = .358 [- .12, .78]) and White targets ($r_{WhiteTargetsIncongruent}$ = .707 [.06, .92]), but significantly better than only the *Face Dominant* Model within White targets ($r_{WhiteTargetsIncongruent}$ = -.001 [-.40, .46], $t(15)$ = 2.35, p= .033).

We further decomposed the incongruent cue condition into subtypes (i.e., self-reported pain high + face pain intensity ambiguous, and self-reported pain ambiguous

+ face pain intensity high). Across race, when self-reported pain intensity was ambiguous (level 4 or 5), and was paired with a high pain facial expression (level 8 or 9),the *Cue-Integration* model ($r_{BlackTargets}$ =.521 [-.37, .95]; $r_{WhiteTargets}$ =.576 [-.88, .93]) was not significantly different from the *Self-Report Dominant* model ($r_{BlackTargets}$ =.804 [.23, .93], $t(7)$ = -1.73, p= .128; $r_{WhiteTargets}$ =.850 [.53, .99], $t(7)$ = -1.52, p= .172) or *Face Dominant* Model ($r_{BlackTargets}$ =.382 [-.44, .86], $t(7)$ = .245, p= .813; $r_{WhiteTargets}$ =.458 [-.57, .94], $t(7)$ = .254, p= .801).

Although underpowered, we note for future research a divergence in pattern between Black and White targets in the inverse condition (self-reported pain high, and face pain intensity ambiguous), such that for White targets, both the *Cue-Integration* model ($r_{WhiteTargets}$ =.904 [.49, .99]), and *Self-Report Dominant* model ($r_{WhiteTargets}$ =.855 [.27, .99), performed well in predicting the empirical data, though neither was significantly better than the other ($t(7)$ = .440 p= .673). While the *Face Dominant* model was not significantly correlated with the empirical data ($r_{WhiteTargets}$ =.200 [-.90, .85], p= .634. However, for Black targets within the same condition, model performance was not significant across any of the three models ($r_{CueIntegration}$ = .594 [.06, .95], p= .158; $r_{Self-ReportDominant}$ = .603[.01, .94], p= .114; $r_{FaceDominant}$ = .550 [.19, .93], p= .120).

Finally, we assessed cue reliability by calculating Shannon's information theoretic entropy (Shannon, 1948) for each unique target within a given cue type (Self-Report and Face)[3]. Overall (collapsed across pain cue intensity), we did not observe any significant difference in mean entropy of a given single cue distribution for either Black ($M_{EntropySR}$ = 1.10 bits, $SD_{EntropySR}$= .56 bits; $M_{EntropyFace}$ = 1.38 bits, $SD_{EntropyFace}$ = .83 bits, $t(15)$ = 1.02, p= .324) or White ($M_{EntropySR}$ = 1.17 bits, $SD_{EntropySR}$ = .62 bits; $M_{EntropyFace}$ = 1.21 bits, $SD_{EntropyFace}$ = 1.00 bits, $t(15)$ = .112, p= .912) targets, such that neither cue was invariably more reliable than the other. This is perhaps unsurprising given that, discordance between patient and provider pain assessments is

[3] Unless otherwise noted, analyses are collapsed across subjects.

larger when pain severity is high (De Ruddere et al., 2014), suggesting that the reliability of the cue itself may vary as a function of pain severity rather than the form of the cue itself. Given this and previous work demonstrating a valence dominance effect (as opposed to cue dominance; Ong et al., 2015) in affective decision making, we calculated entropy for each cue distribution as a function of pain severity. When broken down by pain severity, we observed equivalent entropy for ambiguous (compared to high) self-reported pain intensity for both Black ($M_{EntropySRAmbi}$ = 1.41 bits, $SD_{EntropySRAmbi}$ = .53 bits; $M_{EntropySRHigh}$ = .79 bits, $SD_{EntropySRHigh}$ = .39 bits, $t(7) = 1.02$, p= .071) and White targets ($M_{EntropySRAmbi}$ = 1.29 bits, $SD_{EntropySRAmbi}$ =.70 bits; $M_{EntropySRHigh}$ = 1.05 bits, $SD_{EntropySRHigh}$ = .49 bits, $t(7) = 1.02$, p= .359). In a similar pattern, we observed significantly higher mean entropy for ambiguous (compared to high) face pain intensity for both Black ($M_{EntropyFaceAmbi}$ = 2.02 bits, $SD_{EntropyFaceAmbi}$ = .66 bits; $M_{EntropyFaceHigh}$ = .74 bits, $SD_{EntropySRHigh}$ = .37 bits, $t(7) = 4.18$, p= .004) and White targets ($M_{EntropyFaceAmbi}$ = 2.10 bits, $SD_{EntropyFaceAmbi}$ = .54 bits; $M_{EntropyFaceHigh}$ = .31 bits, $SD_{EntropyFaceHigh}$ = .33 bits, $t(7) = 6.91$, p <.001).

3.2 Pain Treatment Models

In comparing overall (collapsed across conditions) RMSE for each model, we found that the *Cue-Integration* model performed better than the Face-Dominant model for both Black($RMSE_{BlackTargets}$= 48.31 [39.86, 56.49]; $t(31)$ = -3.97, p <.001) and White targets ($RMSE_{WhiteTargets}$ = 47.50 [37.37, 58.46]; $t(31)$ = -3.84, p<.001), but not compared to the *Self-Report Dominant* model for either Black or White targets ($RMSE_{BlackTargets}$ = 48.31 [39.86, 56.49]; $t(31)$ = -.213, p= .481; ($RMSE_{WhiteTargets}$ = 47.50 [37.37, 58.46]), $t(31)$ = -.61, p=.546). Similarly, there was no significant difference between the *Cue-Integration* model ($r_{BlackTargets}$ = .913 [.81, .96]; $r_{WhiteTargets}$ =.914 [.83, .96]) and *Self-Report Dominant* ($r_{BlackTargets}$ =.840 [.75, .90], $t(31)$ = 1.79, p= .083; $r_{WhiteTargets}$ =. 833 [.72, .91], $t(31)$ = 1.70, p= .098) in predicting treatment decisions overall. However, we found that the *Cue-Integration* model again performed better than the Face-Dominant model for both Black ($r_{BlackTargets}$ =.476 [.17, .71],

$t(31) = 3.33$, $p = .002$) and White Targets ($r_{WhiteTargets}$ =.448 [.13, .69]; $t(31) = 3.54$, $p = .001$).

We performed additional exploratory analyses subsetting trials into congruent and incongruent conflict conditions as in the pain assessment analysis. When the self-report and face pain intensity were congruent, the *Cue-Integration* model was again more strongly correlated with empirical treatment outcomes ($r_{BlackTargets} = .986$ [.97, .99]; $r_{WhiteTargets}$ =.990 [.98, .99]) compared to the Face-Dominant model ($r_{BlackTargets}$ =.941 [.88, .97], $t(15) = 2.71$, $p = .016$; $r_{WhiteTargets}$ =.940 [.87, .98], $t(15) = 2.42$, $p = .030$). While the comparison against the *Self-Report Dominant* models were not statistically significant ($r_{BlackTargets}$ =.980 [.96, .99], $t(15) = 1.08$, $p = .296$; $r_{WhiteTargets}$ =. 972 [.93, .99], $t(15) = 1.95$, $p = .070$).

When trial types were subset to incongruent conditions, the *Self-Report Dominant* model ($r_{BlackTargetsIncongruent} = .805$ [.57, .91]; $r_{WhiteTargetsIncongruent} = .749$[.44, .89]) performed similarly to the *Cue-Integration* model ($r_{BlackTargetsIncongruent} = .375$ [-.11, .79], $t(15) = -1.96$, $p = .069$; $r_{WhiteTargetsIncongruent} = .385$ [-.04, .66], $t(15) = -1.70$, $p = .110$) for both Black and White targets. Both models performed significantly better than the *Face Dominant* model across race ($r_{BlackTargetsIncongruent} = -.653$ [- .86, -.34], $t(15) = 3.46$, $p = .003$; $r_{WhiteTargetsIncongruent} = -.651$ [-.85, -.35], $t(15) = 3.89$, $p = .001$).

3.3 Subject Level Analysis

Finally, because the *Cue-Integration* and Single Cue models could not be constructed for each subject with sufficient power (and trials were instead collapsed across subjects), we analyzed average pain assessment and treatment for Black and White targets within subjects. Overall, we did not observe typical disparities in pain assessment. In fact, counter to our predictions and other work within the literature, subjects assessed the pain of Black targets as higher than White targets on average within the Self-Report Only condition ($t(170) = 2.06$, $p = .041$; $M_{SRBlackTargets}$

=6.60, $SD_{SRBlackTargets}$ =1.04; $M_{SRWhiteTargets}$ =6.51, $SD_{SRWhiteTargets}$ =1.07). Similarly, we did not observe any bias in pain treatment recommendations ($t(170) = 1.38$, $p = .170$; $M_{TreatBlackTargets} = 391.28$, $SD_{TreatBlackTargets} = 66.62$, $M_{TreatWhiteTargets} = 388.10$, $SD_{TreatWhiteTargets} = 64.76$).

Chapter 4

DISCUSSION

Previous work has demonstrated that perceivers display a tendency to see pain less readily on Black versus White faces, facilitating racial bias in treatment (Lin et al., under review). However, given that self-report and other contextual information are commonly used to make treatment decisions, it is unclear how self-reported pain information might be integrated with perception of pain on faces. While recent work has shown that self-reported pain reliably shapes thresholds for perceiving pain on faces, race disparities in pain perception do not appear to be affected (Mende-Siedlecki et al., 2021). Moreover, while this previous work demonstrates that self-reported pain may shape the visual perception of pain, it is not clear how biases in face perception and self-report are integrated into a coherent assessment of the target's inferred pain state. The present study proposed and tested a Bayesian cue integration model for pain inference where observers integrated self-report and face perception cues concomitant with each cue's perceived reliability. This model was compared to two cue segregation models of pain inference where observers used only facial expression or self-report cues to make pain inferences even when both cues are present. Overall, we demonstrate that for both Black and White targets, the Bayesian cue integration model best predicts observer responses overall. This is consistent with previous work demonstrating that Bayesian cue integration models are well suited to model the generative process underlying mental state inferences of emotion (Anzellotti et al., 2021; Ong et al., 2015; Saxe Houlihan, 2017), particularly when multiple channels of information are available. Contextual information has been shown to reliably shape emotion inferences across a variety of emotion cues and contextual paradigms (Aviezer et al., 2012; Barrett Kensinger, 2010; Carroll Russell, 1996; Doyle et al., 2021); however,

the processes underlying these contextual shifts have been opaque or challenging to generalize to other cognitive and neural processes. Bayesian cue integration models afford the ability to formalize statistical models of these inferences as well as the uncertainty underlying them. This is of particular interest for emotion inferences such as pain, where uncertainty may vary as a function of both the cue type (various medical and diagnostic information) and social factors (such as patient identity, injury context, and provider beliefs). While cue ambiguity is typically operationalized as an attribute of the cue itself, Bayesian cue integration models allow us to model the ambiguity the observer associates with a particular cue. For example, while high self-reported pain could be considered objectively more reliable based on its signal intensity, beliefs that a patient may be lying, exaggerating, or less sensitive to pain may impact cue reliability, particularly in the context of conflicting cue information. Indeed, when trials were subset by conflict conditions, such that self-reported pain was high and facial expression intensity was ambiguous, the Self-Report Dominant model performed nearly as well as the Cue Integration model for White targets, while the Face Dominant model trailed significantly- perhaps due to relatively high entropy for ambiguous White faces. However, within the same condition, when self-reported pain was high and facial expression was ambiguous, none of the three models significantly predicted empirical responses. While we are underpowered to assess this possibility, this may suggest, as other work has demonstrated (Hirsh et al., 2015), that White patients' word is meaningfully incorporated into the providers' pain assessment such that even when it is ambiguous, it is weighted with fidelity to their report. Conversely, Black patients' self-reported pain may be weighted only or mostly when lower or ambiguous. When self-reported pain was high, neither single cue or Cue Integration model reliability correlated with empirical responses, suggesting perhaps a general discounting of Black patient's word when it comes to their own pain management (O'Connor, 2021; Ellis, 2021). While we did not observe significant differences in pain assessment as a function of target race, this may have been due in part to the abstractness of the scale used, or because we were underpowered to detect nuance variation in how each model operated within

conflicting cue conditions. Notably, though we did not observe significant differences in treatment recommendation as a function of race, treatment predictions produced by both the Bayesian Cue Integration model and Self-Report Dominant were significantly correlated with empirical responses, demonstrating not only that observers reliably integrate information from self-report and face cues to pain, but that these models generalize to treatment decisions.

4.1 Limitations and Future Directions

While this work demonstrates the viability of a Bayesian Cue Integration model in the assessment of pain and pain treatment, we note several key limitations that should be addressed in future work. First, it is important to note that contrary previous work, we do not observe biases in either pain assessment or treatment as a function of race. This may be due, in part, to relatively explicit measures of pain assessment and treatment compared to previous studies. Indeed, biases in pain perception appear to be relatively automatic (Mende-Siedlecki et al., 2021), appearing as early as 33ms, and becoming less robust with unconstrained viewing times. Moreover, biases in treatment, which depend on a more explicit measure compared to bias in perception, are less pronounced across studies (for meta-analysis see: Lin et al., under review), at least within the context of lab-based work. As such, future work should assess the fit of these models where racial biases are more pronounced. Second, while it is not uncommon in such paradigms to construct the Bayesian model collapsed across participants (Körding et al., 2007; Ong et al., 2015), doing so likely minimized potentially informative variability between and within subjects, such that subtle biases in pain assessment or treatment recommendations by race may have been lost. Future work should consider a smaller participant pool but incorporate more trials (with greater variability in terms of both self-report and expression intensity), where each model could be constructed for each participant individually and sufficiently powered within each condition. Notably, while analyses of the conflict conditions were exploratory, they were underpowered, and any inferences drawn from these analyses should be weighted accordingly. Additionally,

the range of stimuli for both faces and self-report were limited to ambiguous (4 and 5 out of 10) and high pain (8 and 9 out of 10) ranges where the edge cases of those categories (6 and 7 out of 10) were left out. This created a potentially artificial bimodal distribution in these data that may have contributed to poor model fit across models. Moreover, both the Face Only and Self-Report Only models were extreme causal models, where the observer used only the single cue (despite the presence of the other cue) and the model assumed complete causal dissociation between the self-report and face cues. This may be better modeled by a softer or more realistic single cue model where the probability that these cues have a common cause is generated with more space for variability. Future work could also explicitly ask subjects about whether they believe the two cues to share a common cause to better assess beliefs about the association between these pain cues that may diverge based on race-based stereotypes and attitudes.

We also note that because participants were only able to respond with a rating within a single emotion modality, the context of the experiment itself is one where pain and only pain is expected. Greater variation in response options (i.e., distributional ratings across different emotions (Anzellotti et al., 2021), as well as greater variation in pain context (i.e., how and why the injury occurred) may yield a more robust model. This may be particularly important for elucidating disparities between race, such that racially stereotypic injury contexts may shift the reliability of either cue.

Finally, this study was conducted with only Black and White men or male presenting faces and White identified observers. Previous work has demonstrated that perceptual bias is stronger within men/male targets than women/female targets, but that Black women are nevertheless given less treatment than any other category (Goharzad et al., in prep), suggesting that the generative model for treatment may vary fundamentally for the treatment of Black women. Indeed, the self-reported pain of Black women is consistently discounted or ignored in treatment decisions (Badreldin et al., 2019), and with devastating consequence (McKinnish et al., 2021; Cottom, 2018; Lu Halfon, 2003). With respect to White (particularly women) patients, there has

been considerable attention given to the potential value of race and gender concordant patient-provider interactions (Cooper et al., 2003; Hagiwara et al., 2017; Penner et al., 2013; Schnittker Liang, 2006; Strumpf, 2011). Formalizing how pain inference models vary as a function of patient-provider concordance may be valuable in further understanding sources of racial bias in pain assessment and treatment.

4.2 Conclusion

Overall, the present work demonstrates that observers reliably integrate self-report and facial expression cues to pain inference, and that reliance on one cue or another may vary as a function of cue severity and race. By extension, this work demonstrates that the Bayesian integration model may be flexibly integrated over treatment space to make treatment recommendations. While this work did not observe the disparities in pain assessment and treatment that are typical of these paradigms, it presents a framework for which such disparities may be examined in the future.

REFERENCES

Acerbi, L., Dokka, K., Angelaki, D. E., & Ma, W. J. (2018). Bayesian comparison of explicit and implicit causal inference strategies in multisensory heading perception. PLoS Computational Biology, 14(7), e1006110.

Acevedo-Garcia, D., Lochner, K. A., Osypuk, T. L., & Subramanian, S. V. (2003). Future directions in residential segregation and health research: a multilevel approach. American Journal of Public Health, 93(2), 215–221.

Anderson, K. O., Green, C. R., & Payne, R. (2009). Racial and ethnic disparities in pain: causes and consequences of unequal care. The Journal of Pain: Official Journal of the American Pain Society, 10(12), 1187–1204.

Anzellotti, S., Houlihan, S. D., Liburd, S., & Saxe, R. (2021). Leveraging facial expressions and contextual information to investigate opaque representations of emotions. Emotion , 21(1), 96–107.

Aviezer, H., Trope, Y., & Todorov, A. (2012). Body cues, not facial expressions, discriminate between intense positive and negative emotions. Science, 338(6111), 1225–1229.

Badreldin, N., Grobman, W. A., & Yee, L. M. (2019). Racial Disparities in Postpartum Pain Management. Obstetrics and Gynecology, 134(6), 1147–1153.

Bailey, Z. D., Feldman, J. M., & Bassett, M. T. (2021). How structural racism works—Racist policies as a root cause of US racial health inequities.

Bailey, Z. D., Krieger, N., Agénor, M., Graves, J., Linos, N., & Bassett, M. T. (2017). Structural racism and health inequities in the USA: evidence and interventions. The Lancet, 389(10077), 1453–1463.

Baker, C. L., Saxe, R., & Tenenbaum, J. B. (2009). Action understanding as inverse planning. Cognition, 113(3), 329–349.

Barrett, L. F., & Kensinger, E. A. (2010). Context is routinely encoded during emotion perception. Psychological Science, 21(4), 595–599.

Beierholm, U. R., Quartz, S. R., & Shams, L. (2009). Bayesian priors are encoded independently from likelihoods in human multisensory perception. Journal of Vision, 9(5), 23.1–9.

Bonham, V. L. (2001). Race, ethnicity, and pain treatment: striving to understand the causes and solutions to the disparities in pain treatment. The Journal of Law, Medicine & Ethics: A Journal of the American Society of Law, Medicine & Ethics, 29(1), 52–68.

Burgess, D. J., Crowley-Matoka, M., Phelan, S., Dovidio, J. F., Kerns, R., Roth, C., Saha, S., & van Ryn, M. (2008). Patient race and physicians' decisions to prescribe opioids for chronic low back pain. Social Science & Medicine, 67(11), 1852–1860.

Carroll, J. M., & Russell, J. A. (1996). Do facial expressions signal specific emotions? Judging emotion from the face in context. Journal of Personality and Social Psychology, 70(2), 205–218.

Chapman, E. N., Kaatz, A., & Carnes, M. (2013). Physicians and implicit bias: how doctors may unwittingly perpetuate health care disparities. Journal of General Internal Medicine, 28(11), 1504–1510.

Chibnall, J. T., Tait, R. C., & Ross, L. R. (1997). The effects of medical evidence and pain intensity on medical student judgments of chronic pain patients. Journal of Behavioral Medicine, 20(3), 257–271.

Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. The Behavioral and Brain Sciences, 36(3), 181–204.

Cooper, L. A., Roter, D. L., Johnson, R. L., Ford, D. E., Steinwachs, D. M., & Powe, N. R. (2003). Patient-centered communication, ratings of care, and concordance of patient and physician race. Annals of Internal Medicine, 139(11), 907–915.

Crocker, D., Brown, C., Moolenaar, R., Moorman, J., Bailey, C., Mannino, D., & Holguin, F. (2009). Racial and ethnic disparities in asthma medication usage and health-care utilization: data from the National Asthma Survey. Chest, 136(4), 1063–1071.

De Ruddere, L., Goubert, L., Prkachin, K. M., Louis Stevens, M. A., Van Ryckeghem, D. M. L., & Crombez, G. (2011). When you dislike patients, pain is taken less seriously. Pain, 152(10), 2342–2347.

De Ruddere, L., Goubert, L., Stevens, M. A. L., Deveugele, M., Craig, K. D., & Crombez, G. (2014). Health care professionals' reactions to patient pain: impact of knowledge about medical evidence and psychosocial influences. The Journal of Pain: Official Journal of the American Pain Society, 15(3), 262–270.

De Ruddere, L., Goubert, L., Stevens, M., de C Williams, A. C., & Crombez, G. (2013). Discounting pain in the absence of medical evidence is explained by negative evaluation of the patient. Pain, 154(5), 669–676.

Dobs, K., Ma, W. J., & Reddy, L. (2017). Near-optimal integration of facial form and motion. Scientific Reports, 7(1), 11002.

Dowell, D., Haegerich, T. M., & Chou, R. (2016). CDC Guideline for Prescribing Opioids for Chronic Pain–United States, 2016. JAMA: The Journal of the American Medical Association, 315(15), 1624–1645.

Doyle, C. M., Gendron, M., & Lindquist, K. A. (2021). Language Is a Unique Context for Emotion Perception. Affective Science, 2(2), 171–177.

Drain, A., Goharzad, A., Qu-Lee, J., Lin, J., Mende-Siedlecki, P. (in prep). Skin deep: How racially prototypic features exacerbate bias in pain perception.

Feagin, J., & Bennefield, Z. (2014). Systemic racism and U.S. health care. Social Science & Medicine, 103, 7–14.

Freeman, J. B., Stolier, R. M., Ingbretsen, Z. A., & Hehman, E. A. (2014). Amygdala responsivity to high-level social information from unseen faces. Journal of Neuroscience, 34(32), 10573-10581. Friston, K. (2010). The free-energy principle: a unified brain theory? Nature Reviews. Neuroscience, 11(2), 127–138.

Gamble, V. N. (1993). A legacy of distrust: African Americans and medical research. American Journal of Preventive Medicine, 9(6 Suppl), 35–38. https://www.ncbi.nlm.nih.gov/p
Gaskin, D. J., Price, A., Brandon, D. T., & Laveist, T. A. (2009).

Segregation and disparities in health services use. Medical Care Research and Review: MCRR, 66(5), 578–589.

Gershman, S. J., & Beck, J. M. (n.d.). Complex Probabilistic Inference: From Cognition to Neural Computation.

Goharzad, A., Drain, A., Qu-Lee, J., Lin, J., Mende-Siedlecki, P., (in prep). Invisible identities, invisible pain: How race and gender prototypicality inform pain perception & treatment. =

Goharzad, A., Drain, A., Tuerxuntuoheti, A., Reyes, P., & Mende-Siedlecki, P. (in prep). Does self-reported pain experience reduce racial bias in pain perception?

Goyal, M. K., Kuppermann, N., Cleary, S. D., Teach, S. J., & Chamberlain, J. M. (2015). Racial Disparities in Pain Management of Children With Appendicitis in Emergency Departments. JAMA Pediatrics, 169(11), 996–1002.

Hagiwara, N., Slatcher, R. B., Eggly, S., & Penner, L. A. (2017). Physician Racial Bias and Word Use during Racially Discordant Medical Interactions. Health Communication, 32(4), 401–408.

Haider, A. H., Weygandt, P. L., Bentley, J. M., Monn, M. F., Rehman, K. A., Zarzaur, B. L., Crandall, M. L., Cornwell, E. E., & Cooper, L. A. (2013). Disparities in trauma care and outcomes in the United States: a systematic review and meta-analysis. The Journal of Trauma and Acute Care Surgery, 74(5), 1195–1205.

Hebart, M. (2009). IMScramble. Retrieved from http://martin-hebart.de/code/imscramble.m.

Hillis, J. M., Watt, S. J., Landy, M. S., & Banks, M. S. (2004). Slant from texture and disparity cues: optimal cue combination. Journal of Vision, 4(12), 967–992.

Hirsh, A. T., Hollingshead, N. A., Ashburn-Nardo, L., & Kroenke, K. (2015). The interaction of patient race, provider bias, and clinical ambiguity on pain management decisions. The Journal of Pain: Official Journal of the American Pain Society, 16(6), 558–568.

Hoffman, K. M., Trawalter, S., Axt, J. R., & Oliver, M. N. (2016). Racial bias in pain assessment and treatment recommendations, and false beliefs about biological differences between blacks and whites. Proceedings of the National Academy of Sciences of the United States of America, 113(16), 4296–4301.

Hoffmann, D. E., & Tarzian, A. J. (2001). The girl who cried pain: a bias against women in the treatment of pain. The Journal of Law, Medicine & Ethics: A Journal of the American Society of Law, Medicine & Ethics, 29(1), 13–27.

Hohwy, J. (2013). The Predictive Mind. Oxford University Press. https://play.google.com/store/book

Jospe, K., Genzer, S., Klein Selle, N., Ong, D., Zaki, J., & Perry, A. (2020). The contribution of linguistic and visual cues to physiological synchrony and empathic accuracy. Cortex; a Journal Devoted to the Study of the Nervous System and Behavior, 132, 296–308.

Kappesser, J., Williams, A. C. de C., & Prkachin, K. M. (2006). Testing two accounts of pain underestimation. Pain, 124(1-2), 109–116.

Karcioglu, O., Topacoglu, H., Dikme, O., & Dikme, O. (2018). A systematic review of the pain scales in adults: which to use?. The American journal of emergency medicine, 36(4), 707-714.

Knill, D. C., & Richards, W. (1996). Perception as Bayesian Inference. Cambridge University Press. https://play.google.com/store/books/details?id=$_cTLCgAAQBAJ$

Körding, K. P., Beierholm, U., Ma, W. J., Quartz, S., Tenenbaum, J. B., & Shams, L. (2007). Causal inference in multisensory perception. PloS One, 2(9), e943.

Körding, K. P., Ku, S.-P., & Wolpert, D. M. (2004). Bayesian integration in force estimation. Journal of Neurophysiology, 92(5), 3161–3165.

Kunz, M., Meixner, D., & Lautenbacher, S. (2019). Facial muscle movements encoding pain—a systematic review. Pain, 160(3), 535-549.

Lawrence, A. E., Deans, K. J., Chisolm, D. J., Wrona, S. K., Minneci, P. C., & Cooper, J. N. (2020). Racial Disparities in Receipt of Postoperative Opioids After Pediatric Cholecystectomy. The Journal of Surgical Research, 245, 309–314.

Lieberman, J. R., Dorey, F., Shekelle, P., Schumacher, L., Thomas, B. J., Kilgus, D. J., & Finerman, G. A. (1996). Differences between Patients and Physicians Evaluations of Outcome after Total Hip Arthroplasty*. In The Journal of Bone & Joint Surgery (Vol. 78, Issue 6, pp. 835–838).

Lin, J., Drain, A., Goharzad, A., & Mende-Siedlecki, P. (under review). What factors fuel racial bias in pain perception and treatment?: A meta-analysis across 40 experimental studies. Under review at Social and Personality Psychology Compass.

Mateo, C. M., & Williams, D. R. (2021). Racism: a fundamental driver of racial disparities in health-care quality. Nature Reviews. Disease Primers, 7(1), 20.

Ma, W. J. (2019). Bayesian Decision Models: A Primer. Neuron, 104(1), 164–175.

Max, M. B., Donovan, M., Miaskowski, C. A., Ward, S. E., Gordon, D., Bookbinder, M., Cleeland, C. S., Coyle, N., Kiss, M., Thaler, H. T., Janjan, N., Anderson, M. D., Weinstein, S., Thomas Edwards, W., & American Pain Society Quality of Care Committee. (1995).

Quality Improvement Guidelines for the Treatment of Acute Pain and Cancer Pain. JAMA: The Journal of the American Medical Association, 274(23), 1874–1880.

McCaffery, M., Ferrell, B. R., & Pasero, C. (2000). Nurses' personal opinions about patients' pain and their effect on recorded assessments and titration of opioid doses. Pain Management Nursing: Official Journal of the American Society of Pain Management Nurses, 1(3), 79–87.

Mende-Siedlecki, P., Goharzad, A., Tuerxuntuoheti, A., Reyes, P. G. M., Lin, J., & Drain, A. (2022). Assessing the speed and spontaneity of racial bias in pain perception. Journal of Experimental Social Psychology, 101, 104315.

Mende-Siedlecki, P., Lin, J., Ferron, S., Gibbons, C., Drain, A., & Goharzad, A. (2021). Seeing no pain: Assessing the generalizability of racial bias in pain perception. Emotion .

Mende-Siedlecki, P., Qu-Lee, J., Backer, R., & Van Bavel, J. J. (2019). Perceptual contributions to racial bias in pain recognition. Journal of Experimental Psychology. General, 148(5), 863–889.

Moskowitz, G. B., Stone, J., & Childs, A. (2012). Implicit stereotyping and medical decisions: unconscious stereotype activation in practitioners' thoughts about African Americans. American Journal of Public Health, 102(5), 996–1001. https://doi.org/10.2

Mularski, R. A., White-Chu, F., Overbay, D., Miller, L., Asch, S. M., & Ganzini, L. (2006). Measuring pain as the 5th vital sign does not improve quality of pain management. Journal of General Internal Medicine, 21(6), 607–612.

Nieblas-Bedolla, E., Christophers, B., Nkinsi, N. T., Schumann, P. D., & Stein, E. (2020). Changing How Race Is Portrayed in Medical Education: Recommendations From Medical Students. Academic Medicine: Journal of the Association of American Medical Colleges, 95(12), 1802–1806.

Ong, D. C., Zaki, J., & Goodman, N. D. (2015). Affective cognition: Exploring lay theories of emotion. Cognition, 143, 141–162.

Onyeador, I. N., Wittlin, N. M., Burke, S. E., Dovidio, J. F., Perry, S. P., Hardeman, R. R., Dyrbye, L. N., Herrin, J., Phelan, S. M., & van Ryn, M. (2020). The Value of Interracial Contact for Reducing Anti-Black Bias Among Non-Black Physicians: A Cognitive Habits and Growth Evaluation (CHANGE) Study Report. Psychological Science, 31(1), 18–30.

Oruç, I., Maloney, L. T., & Landy, M. S. (2003). Weighted linear cue combination with possibly correlated error. Vision Research, 43(23), 2451–2468.

Penner, L. A., Hagiwara, N., Eggly, S., Gaertner, S. L., Albrecht, T. L., & Dovidio, J. F. (2013). Racial Healthcare Disparities: A Social Psychological Analysis. European Review of Social Psychology, 24(1), 70–122.

Phillips, D. M. (2000). JCAHO Pain Management Standards Are Unveiled. JAMA: The Journal of the American Medical Association, 284(4), 428–429.

Pletcher, M. J., Kertesz, S. G., Kohn, M. A., & Gonzales, R. (2008). Trends in opioid prescribing by race/ethnicity for patients seeking care in US emergency departments. JAMA: The Journal of the American Medical Association, 299(1), 70–78.

Prkachin, K. M., Solomon, P. E., & Ross, J. (2007). Underestimation of pain by healthcare providers: towards a model of the process of inferring pain in others. The Canadian Journal of Nursing Research = Revue Canadienne de Recherche En Sciences Infirmieres, 39(2), 88–106.

Prkachin, K. M., & Solomon, P. E. (2008). The structure, reliability and validity of pain expression: Evidence from patients with shoulder pain. Pain, 139(2), 267-274.

Rock, I., & Victor, J. (1964). Vision and touch: An experimentally created conflict between the two senses. Science, 143(3606), 594–596.

Ruben, M. A., Blanch-Hartigan, D., & Shipherd, J. C. (2018). To Know Another's Pain: A Meta-analysis of Caregivers' and Healthcare Providers' Pain Assessment Accuracy. Annals of Behavioral Medicine: A Publication of the Society of Behavioral Medicine, 52(8), 662–685.

Ruben, M. A., van Osch, M., & Blanch-Hartigan, D. (2015). Healthcare providers' accuracy in assessing patients' pain: A systematic review. Patient Education and Counseling, 98(10), 1197–1206.

Sabin, J. A., Rivara, F. P., & Greenwald, A. G. (2008). Physician implicit attitudes and stereotypes about race and quality of medical care. Medical Care, 46(7), 678–685.

Saxe, R., & Houlihan, S. D. (2017). Formalizing emotion concepts within a Bayesian model of theory of mind. Current Opinion in Psychology, 17, 15–21.

Schnittker, J., & Liang, K. (2006). The promise and limits of racial/ethnic concordance in physician-patient interaction. Journal of Health Politics, Policy and Law, 31(4), 811–838.

Shams, L., Kamitani, Y., & Shimojo, S. (2000). Illusions. What you see is what you hear. Nature, 408(6814), 788.

Strumpf, E. C. (2011). Racial/ethnic disparities in primary care: the role of physician-patient concordance. Medical Care, 49(5), 496–503.

Tait, R. C., & Chibnall, J. T. (2002). Pain in older subacute care patients: associations with clinical status and treatment. Pain Medicine , 3(3), 231–239.

Tait, R. C., & Chibnall, J. T. (2014). Racial/ethnic disparities in the assessment and treatment of pain: Psychosocial perspectives. The American Psychologist, 69(2), 131–141.

Tait, R. C., Chibnall, J. T., & Kalauokalani, D. (2009). Provider judgments of patients in pain: seeking symptom certainty. Pain Medicine , 10(1), 11–34.

Trawalter, S., Bart-Plange, D.-J., & Hoffman, K. M. (2020). A socioecological psychology of racism: making structures and history more visible. Current Opinion in Psychology, 32, 47–51.

van Dijk, J. F. M., van Wijck, A. J. M., Kappen, T. H., Peelen, L. M., Kalkman, C. J., & Schuurmans, M. J. (2012). Postoperative pain assessment based on numeric ratings is not the same for patients and professionals: a cross-sectional study. International Journal of Nursing Studies, 49(1), 65–71.

van Ryn, M. (2002). Research on the provider contribution to race/ethnicity disparities in medical care. Medical Care, 40(1 Suppl), I140–I151.

van Ryn, M., Hardeman, R., Phelan, S. M., Burgess, D. J., Dovidio, J. F., Herrin, J., Burke, S. E., Nelson, D. B., Perry, S., Yeazel, M., & Przedworski, J. M. (2015). Medical School Experiences Associated with Change in Implicit Racial Bias Among 3547 Students: A Medical Student CHANGES Study Report. Journal of General Internal Medicine, 30(12), 1748–1756.

Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., ... SciPy 1.0 Contributors. (2020). SciPy 1.0: fundamental algorithms for scientific computing in Python. Nature Methods, 17(3), 261–272.

Waidmann, T. A., & Rajan, S. (2000). Race and ethnic disparities in health care access and utilization: an examination of state variation. Medical Care Research and Review: MCRR, 57 Suppl 1, 55–84.

Wandner, L. D., Scipio, C. D., Hirsh, A. T., Torres, C. A., & Robinson, M. E. (2012). The perception of pain in others: how gender, race, and age influence pain expectations. The Journal of Pain: Official Journal of the American Pain Society, 13(3), 220–227.

Washington, H. A. (2006). Medical Apartheid: The Dark History of Medical Experimentation on Black Americans from Colonial Times to the Present. Doubleday. https://play.google.com/store/books/details?id=q0bMejttqgEC

Weisse, C. S., Sorum, P. C., Sanders, K. N., & Syat, B. L. (2001). Do gender and race affect decisions about pain management? Journal of General Internal Medicine, 16(4), 211–217.

Welch, R. B., & Warren, D. H. (1980). Immediate perceptual response to intersensory discrepancy. Psychological Bulletin, 88(3), 638–667.

Wells, N., Pasero, C., & McCaffery, M. (2008). Improving the Quality of Care Through Pain Assessment and Management. Agency for Healthcare Research and Quality (US).

Williams, A. C. D. C. (2002). Facial expression of pain: an evolutionary account. Behavioral and brain sciences, 25(4), 439-455.

Williams, D. R., & Collins, C. (2001). Racial residential segregation: a fundamental cause of racial disparities in health. Public Health Reports, 116(5), 404–416.

Wu, Y., Baker, C. L., Tenenbaum, J. B., & Schulz, L. E. (2018). Rational Inference of Beliefs and Desires From Emotional Expressions. Cognitive Science, 42(3), 850–884.

Zaki, J. (2013). Cue Integration: A Common Framework for Social Cognition and Physical Perception. Perspectives on Psychological Science: A Journal of the Association for Psychological Science, 8(3), 296–312.

Zalon, M. L. (1993). Nurses' assessment of postoperative patients' pain. Pain, 54(3), 329–334.

Appendix A

IRB/HUMAN SUBJECTS APPROVAL LETTER, 2019

Institutional Review Board
210H Hullihen Hall
Newark, DE 19716
Phone: 302-831-2137
Fax: 302-831-2828

DATE: October 25, 2019

TO: Peter Mende-Siedlecki, PhD

FROM: University of Delaware IRB

STUDY TITLE: [958105-11] Leveraging perceptual and psychological mechanisms to understand racial bias in pain care [Social and Emotional Evaluation of Faces]

SUBMISSION TYPE: Continuing Review/Progress Report

ACTION: APPROVED

APPROVAL DATE: October 25, 2019

EXPIRATION DATE: October 20, 2020

REVIEW TYPE: Expedited Review

REVIEW CATEGORY: Expedited review category # *(4,7)*

Thank you for your Continuing Review/Progress Report submission to the University of Delaware Institutional Review Board (UD IRB). The UD IRB has reviewed and APPROVED the proposed research and submitted documents via Expedited Review in compliance with the pertinent federal regulations.

As the Principal Investigator for this study, you are responsible for and agree that:

- All research must be conducted in accordance with the protocol and all other study forms as approved in this submission. Any revisions to the approved study procedures or documents must be reviewed and approved by the IRB prior to their implementation. Please use the UD amendment form to request the review of any changes to approved study procedures or documents.
- Informed consent is a process that must allow prospective participants sufficient opportunity to discuss and consider whether to participate. IRB-approved and stamped consent documents must be used when enrolling participants and a written copy shall be given to the person signing the informed consent form.
- Unanticipated problems, serious adverse events involving risk to participants, and all non-compliance issues must be reported to this office in a timely fashion according with the UD requirements for reportable events. All sponsor reporting requirements must also be followed.

Oversight of this study by the UD IRB REQUIRES the submission of a CONTINUING REVIEW seeking the renewal of this IRB approval, which will expire on October 20, 2020. A continuing review/progress report form and up-to-date copies of the protocol form and all other approved study materials must be submitted to the UD IRB at least 45 days prior to the expiration date to allow for the required IRB review of that report.

If you have any questions, please contact the UD IRB Office at (302) 831-2137 or via email at hsrb-research@udel.edu. Please include the study title and reference number in all correspondence with this office.

Appendix B

IRB/HUMAN SUBJECTS APPROVAL LETTER, 2020

Institutional Review Board
210H Hullihen Hall
Newark, DE 19716
Phone: 302-831-2137
Fax: 302-831-2828

DATE:	November 11, 2020
TO:	Peter Mende-Siedlecki, PhD
FROM:	University of Delaware IRB
STUDY TITLE:	[958105-13] Leveraging perceptual and psychological mechanisms to understand racial bias in pain care [Social and Emotional Evaluation of Faces]
SUBMISSION TYPE:	Continuing Review/Progress Report
ACTION:	APPROVED
APPROVAL DATE:	November 11, 2020
EXPIRATION DATE:	October 20, 2021
REVIEW TYPE:	Expedited Review
REVIEW CATEGORY:	Expedited review category # *(4,7)*

Thank you for your Continuing Review/Progress Report submission to the University of Delaware Institutional Review Board (UD IRB). The UD IRB has reviewed and APPROVED the proposed research and submitted documents via Expedited Review in compliance with the pertinent federal regulations.

As the Principal Investigator for this study, you are responsible for and agree that:

- All research must be conducted in accordance with the protocol and all other study forms as approved in this submission. Any revisions to the approved study procedures or documents must be reviewed and approved by the IRB prior to their implementation. Please use the UD amendment form to request the review of any changes to approved study procedures or documents.
- Informed consent is a process that must allow prospective participants sufficient opportunity to discuss and consider whether to participate. IRB-approved and stamped consent documents must be used when enrolling participants and a written copy shall be given to the person signing the informed consent form.
- Unanticipated problems, serious adverse events involving risk to participants, and all non-compliance issues must be reported to this office in a timely fashion according with the UD requirements for reportable events. All sponsor reporting requirements must also be followed.

Oversight of this study by the UD IRB REQUIRES the submission of a CONTINUING REVIEW seeking the renewal of this IRB approval, which will expire on October 20, 2021. A continuing review/progress report form and up-to-date copies of the protocol form and all other approved study materials must be submitted to the UD IRB at least 45 days prior to the expiration date to allow for the required IRB review of that report.

If you have any questions, please contact the UD IRB Office at (302) 831-2137 or via email at hsrb-research@udel.edu. Please include the study title and reference number in all correspondence with this office.