The Pennsylvania State University

The Graduate School

THE STRINGENT MORAL CIRCLE: SELF-OTHER DISCREPANCIES IN THE PERCEIVED EXPANSION OF MORAL CONCERN

A Dissertation in

Social Psychology

by

Eliana Hadjiandreou

Submitted in Partial Fulfillment

of the Requirements

for the Degree of

Doctor of Philosophy

May 2024

The dissertation of Eliana Hadjiandreou was reviewed and approved by the following:

Daryl C. Cameron
Associate Professor of Social Psychology and Senior Research Associate at
Rock Ethics Institute
Dissertation Advisor
Chair of Committee

Jes L. Matsick
Associate Professor of Social Psychology and Women's, Gender, & Sexuality Studies

Sean M. Laurent
Assistant Professor of Social Psychology

Rachel A. Smith
Professor of Communication Arts & Sciences

Charlie R. Crimston
Lecturer at Australian National University School of Medicine and Psychology
Special Member

Kristin A. Buss
Professor of Psychology and Human Development and Family Studies
Head of the Department of Psychology

Acknowledgments

I grew up in a pretty cynical household, where baseline assumptions were that people are ready to get you at every corner, that you have to fend for yourself, and that no one, who is given the chance to steal a full plate of food, won't jump at the chance to do so. This view of the world might not be atypical for a lower-middle class family, but it wasn't until I was in college, when my undergraduate thesis advisor, handed me a theory piece in American Psychologist by Dale Miller called "The Norm of Self Interest" when I really started thinking about how selfishness might be overrated.

I started grad school really wanting to explore more of this - are people seeing the world as selfish because society is telling them selfishness is the norm? I dug into the literature on perceptions of prosociality, and I started floating around different ideas on how to check if people see the world as more selfish than it actually is. In the first year of grad school, Daryl suggested looking at one's moral circle, as one of a few suggestions. We ran a study then (now Study 1a of my dissertation package), and while you can read the rest in the actual manuscript that follows, the results were promising enough for us to keep digging.

That is all to say that, while this is my dissertation, for which I have certainly worked hard on, it is by no means a product of just me. It took planted seeds harkening back to my time in college from my then advisor Dr. Seana Moran, it flourished with the supervision and guidance from Dr. Daryl Cameron, my PhD advisor, and it got enriched by the time I spent at Stanford amongst experts on the cynicism construct. As always, ideas work best when they are shared, and sharing ideas is my favorite part of doing research.

In addition to thanking Dr. Seana Moran for planting these seeds and to Dr. Daryl Cameron for helping them grow, I also thank my committee for helpful feedback on my proposal - former member Dr. Janet Swim, and current members Dr. Sean Laurent, Dr. Jes Matsick, Dr. Rachel Smith, and co-author and special committee member Dr. Charlie Crimston, who developed and validated the main measure I used in my dissertation. All of you have made me question my assumptions and made me realize how much I still don't really know, which is something to be grateful for.

I big thank you to the A.G. Leventis Foundation for funding the first three years of this work, and the Rock Ethics Institute for awarding me the Forrest S. Crawford Graduate Fellowship to further support me for a full year of dissertation work.

I'd also like to thank Dr. Bobby Moulder for statistical guidance, and Penelope Zakou who watched me struggle time and again with my time management skills and helped me stay afloat. This wouldn't be completed without you.

And of course, I couldn't end without acknowledging my partner, Andreas, who has kept me sane this past year, and my family and friends, who stuck by my side for all of my 6 years as a PhD student at Penn State, to see me reach this final stage of the doctoral program.

Abstract

Across contexts, people consistently underestimate others' prosociality. In this dissertation work, I attempt to theoretically link these findings to a broader underestimation: the moral landscapes of other people are imagined as more constricted compared to that of the self. In two pilot studies (Studies 1a & 1b) using the Moral Expansiveness Scale (MES; Crimston et al., 2016), we found that participants reported their moral circles to be more expansive compared to those they predicted for others. This was true for entities that are typically distant in the moral circle (that is: outgroup, stigmatized, animals of high sentience, animals of low sentience, plants, environment), but not entities that are typically placed proximal to the center (that is: family & friends, ingroup, revered). We aimed to assess whether this is because of lack of access to information about these distant categories from the perspective of general others. In order to do this, three studies (Studies 2, 3, and 4) asked participants to fill out the Moral Expansiveness Scale from three perspectives – self, close others, average others. We predicted that participants' moral circles will be comparably inclusive to the moral circles they reported for their close others due to having informational access to their moral worlds, but that average others will be reported as being less inclusive than the self and close others due to greater reliance on stereotypes about distant categories, that reflect cynical bias. Because we expected that the comparable scores between self and close other are not only due to feeling more similar to close others and liking them more than average others, we expected to see the same results even when participants were encouraged to reflect on being dissimilar from their close others (Study 2), and when participants were encouraged to think of reasons they dislike their close others (Study 3). Finally, in Study 4, we expected that assigning participants to an anti-cynicism condition which exposes them to information highlighting human prosociality will eliminate this discrepancy, which we expected to persist in the pro-cynicism and control conditions. Results showed that the moral inclusivity discrepancy for distant entities persisted but was smaller in the case of close others, and that similarity and liking partially accounted for the overlap between self and close other scores. Study 4 showed unique reductions in the perceived moral expansiveness of average others for distant entities in the cynicism condition, although statistical results were mixed. Implications for the potential cynicism explanation and future directions are discussed.

Keywords: morality, prosocial behavior, perceptions, self-other differences, cynicism

Table of Contents

List of Figures

List of Tables

Chapter 1: Introduction

"It may appear altruistic when you donate $100 to your local public radio station, but in exchange you get a year of guilt- free listening (and, if you're lucky, a canvas tote bag). U.S. citizens are easily the world's leaders in per-capita charitable contributions, but the U.S. tax code is among the most generous in allowing deductions for those contributions."

~ Steven D. Levitt & Stephen J. Dubner, SuperFreakonomics
Unbelievable Stories About Apathy and Altruism, The New York Times, 2009

Although the above quote was written over 15 years ago, its cynical message about highlighting selfishness in the world is very much alive today. People assume that others are more self-interested than they actually are, and are ready to attribute selfishness to seemingly selfless behaviors even when they are confronted with the contrary (Critcher & Dunning, 2011). The umbrella term for this phenomenon is often referred to as "cynicism" (Neumann & Zaki, 2022), and research has demonstrated its existence across a variety of social contexts. This dissertation aims to zoom out of individual selfish or selfless behaviors, and trace cynicism on a broader level. Specifically, my goal is to demonstrate that people's understanding of others' moral landscape reflects such general cynicism, and that accessibility of information plays a key role in shaping whether those landscapes are cynically construed.

Underestimating Others' Prosociality

Beliefs that others will not step up to help as much as oneself would span a range of different contexts. In classic experiments in the late 1990s, Miller and Ratner (1998) asked an American sample about their willingness to donate blood with and without a financial $15 incentive, and then additionally asked them to estimate whether other people would donate under both of those conditions. They found that people themselves would largely donate blood under both conditions (73% with incentive, 63% without incentive), but when it came to their estimations for other people, the effect they assumed the incentive would have was much larger

(62% with incentive, 33% without incentive), indicating that others were believed to be more motivated by self-interest to enact prosocial behavior (note that these findings have been replicated in a recent pre-registered attempt; Brick et al., 2021).

About a decade later, a study by Kogut and Beyth-Marom (2008) asked Israeli participants how likely they were to respond to a request for help by a student at the participants' university who was ostensibly struggling to access university materials. In another study in the same paper (2008) they asked participants about their willingness to contribute to obtaining a new expensive medical drug needed to save a child in a life-threatening condition. In one last study (2008), they asked participants their willingness to contribute to a new recycling governmental program by separating their garbage into the relevant bins (as opposed to paying tax to the government to take care of the separation). In all 3 studies, there was an additional condition in which participants were asked about the willingness of an average student/citizen (depending on the study) to do all aforementioned actions. Conceptually replicating the findings from Ratner and Miller (1998) some ten years later, they found that in all studies, participants show large same-direction discrepancies between the self and average-other conditions, such that in all cases, participants underestimated how likely average others would be to help the student, the sick child, and the environment.

Yet more research continues to replicate this effect in the 2020s: Zhao and Epley (2022) asked participants to imagine either requesting help or being the recipients of a request for help across a variety of small daily behaviors including borrowing someone's cell phone, helping to carry boxes down the stairs, giving away one's seat on the subway, giving away change at a food truck, giving someone directions, and demonstrating how to use a library kiosk. Across all helping scenarios, those who imagined being the recipients of help consistently underestimated

the helper's (in this case, the participants assigned to the helper perspective) likelihood of saying yes, their willingness to provide that help, their anticipated positive mood following helping, and their level of reported inconvenience. Importantly, participants underestimated helpers' prosocial motivation - that is, how much the helpers reported wanting to see the requester out of their trouble and believing that their help was meaningful. The results replicated not only with hypothetical imagined scenarios, but also with real incidents that participants were asked to recall from their past at times when they received (vs. provided) help to others. Critically, participants underestimated others' willingness to help them not only when those helping were strangers but also when they were close friends, speaking to how potentially far-reaching these cynical misestimations are.

Helping contexts are not the only ones in which this overestimation of self-interest has been demonstrated in. Mason and colleagues (2018) asked participants to espouse one of three possible views of negotiation: 1) a game, in which trickery and deceit should be used at all costs in order to win, 2) a cost-benefit analysis, in which people should decide whether or not it is worth to deceive, and 3) a professional domain governed by principles in which deception should not take place. They also asked all participants to estimate what percentage of people espoused each of these three views of negotiation. The participants' experience with negotiation and economic contexts varied: two American and one Chinese sample of business students, a sample of Turkish company executives, a sample of American non-profit executives, and an American sample recruited on Amazon Mechanical Turk. The authors found that while the actual percentage of people self-categorizing under the deceitful definition of negotiations hovered mostly in the single digits (in 4 out of the 6 studies it was 6% of the participants or less), the percentage that participants predicted for others across all samples far exceeded the

reality (in the majority of studies predicted at around 30%), with participants from all three negotiation categorizations inaccurately inflating this cynical outlook.

This finding is not unique to negotiations but generalizes to broader humanitarian, environmental, and public health crises. The extent to which Americans think others support policies to combat climate change, a threat impacting all of humanity, is similarly underestimated. Sparkman and colleagues find that while 66-80% of Americans are in support, the estimate is a fraction of that, at 37–43% (Sparkman, Geiger, & Weber, 2022) - a cynical false reality. While cynicism has first been pinpointed in highly competitive capitalist-driven economies in Western societies, this phenomenon is not constrained to Western participants. In a Chinese sample, Chen and colleagues (2022) asked participants to estimate how frequently most people in China engage in ten pro-environmental behaviors (such as recycling, energy conservation, and environmental protection volunteering), as well as how much most people in China approve of engaging in those behaviors. Consistent with the research I have cited thus far, participants underestimated the extent to which others both engaged with and approved of those behaviors.

During the peak of the coronavirus pandemic in 2020, American college adults self-reported compliance with behaviors recommended by the Center for Disease Control such as wearing a face mask in public, washing hands, and avoiding large gatherings at rates overwhelmingly over 80%. Their estimations of their peer's compliance with those same behaviors were consistently grossly underestimated by around 20 to 30 percentage points (Graupensperger, Lee, & Larimer, 2021). In Switzerland, when people were asked whether they approved or disapproved of same-sex marriage and same-sex parenting, they falsely believed that more than half of the people in their district disapproved of those acts, when in reality, the

majority of the people were in agreement about their approval (Eisner, Turner-Zwinkels, & Spini, 2021).

Thus, across contexts, people underestimate prosociality in others. From estimations of direct helping behavior, to donations, to other people's honesty, to pro-environmental attitudes and action, to individual preventative behaviors aimed to protect community health, and views on minoritized groups. We could assume that these are all standalone, individual behavioral domains for which this underestimation has causes that are domain-specific. For example, the reason that people underestimate others' likelihood of recycling (Kogut & Beyth-Marom, 2008; Chen, Wan, & Yang, 2022) might be unique to recycling and independent of the reason that there is underestimation of others' willingness to donate blood without a financial incentive (Miller & Ratner, 1998; Brick et al., 2021). Or, the reason that people think that others will be dishonest in negotiations (Mason, Wiley, & Ames, 2018) is specific to negotiations and might be unrelated to reasons why people think that others disapprove of sexual minorities (Eisner, Turner-Zwinkels, & Spini, 2021). In other words, these findings might be related only by chance. There is, however, the possibility that people engage in this sort of underestimation because they systematically cynically underestimate prosocial concern in others, en large.

The Moral Circle, its Significance, and Correlates

Before one begins to understand how prosociality in others is anticipated, perceived, and interpreted in complex social contexts, one needs to take a step back and understand how people imagine the moral landscape of others. In other words, to understand how people predict others would behave in contexts of helping and cooperation such as the ones cited above, we need to first ask people the more basic question of what others care about. Whether or not people would reasonably expect others to step up to help a poor civilian or care for the

environment, for example, should largely depend on whether people think others see those targets as worthy of moral regard. The moral circle is a way to map such moral boundaries - where and to what extent concern for the world around us should be extended. More deeply, other's moral landscapes go above and beyond any individual prosocial behavior because of what is included in them, with this inclusion having implications about how that entity should be ethically treated.

Harkening back to the first introduction of the concept within psychology, applied ethicist Peter Singer (1981) was the first to write extensively on the ethics of the expanding moral circle. Singer famously made the normative argument that our moral circle is naturally constricted to kin, but it doesn't have to be. Instead, Singer argued that humans are capable of continually expanding their moral circle to include strangers, outgroups, and non-human beings capable of pain and pleasure. Since then, the concept of the moral circle has been brought to the surface to discuss and address phenomena such as moral exclusion (Opotow, 1990a, 1990b; Opotow, 2000), and to explore relevant predictors of moral expansion such as the experience of awe (Song, Klebl, & Bastian, 2023), as well as moderators of moral expansion such as exclusionary and inclusionary mindset framing (Laham, 2009), and political ideology (Waytz, Iyer, Young, Haidt, & Graham, 2019).

Of note, Crimston and colleagues (2016) developed a measure assessing people's placement of different entities within the moral boundaries, with entities placed closer to the self (and thus closer to the center of the circle) conveying higher moral value, and entities placed further away from the self (and thus further away from the moral center) deserving lower moral treatment. Crimston and colleagues (2016) found that moral expansiveness scores on this measure predict people's self-reported empathy, their identification with humanity, their

connectedness to nature and sense of social responsibility, as well as their willingness to engage in prosocial behavior, the extent to which they are concerned with humanitarian and environmental matters, their willingness to sacrifice their life to save others, and volunteering (2016). On the flip side, scores on the moral expansiveness scale are negatively predicted by people's fear of feeling and enacting compassion (Crimston, Blessing, Gilbert, & Kirby, 2022). In sum, people's moral expansiveness scores can tell us a lot about the extent to which they enact prosociality in the world.

These averages on moral expansiveness are helpful, because they give us an insight of people's baseline for moral inclusion - as one would expect, most people include in their inner moral circle their friends and family, and tend to place entities such as animals and plants further out (Crimston, Bain, Hornsey & Bastian, 2016). Even though there is variation in this finding based on cross-cultural factors (Kirkland et al., 2022), environmental valuation factors (e.g. tree-huggers vs. human-lovers, Rottman, Crimston, & Syropoulos, 2021), as well as the interaction between participant and circle entity (Jaeger & Wilks, 2023), there is still a valuable benefit in knowing, on average, where people see their own moral boundaries; this helps us understand where they might expect those of others. Yet, no work has explored these expectations. In order to understand whether people have reasonable and relatively accurate assumptions about the extent that the world is morally expansive, they need to report more than just their own moral circle, they also need to share how they think others morally expand. While recent work has looked at the perception of prosocial breadth in people, such as whether we like those who extend moral concern to strangers over close others (McManus, Kleiman-Weiner, & Young, 2020; Law, Campbell, & Gaesser, 2022; Everett, Faber, Savulescu, & Crockett, 2018), or to targets that seem to not deserve moral concern (Wang & Todd, 2021), no

work has studied one's own prediction of the moral landscape of others in relation to one's own, and, on the flip side, the ways one's own moral landscape might be thought of by others.

People consistently perceive differences between their self and others across different domains ranging from what people eat (Sproesser, Klusmann, Schupp, & Renner, 2017), to how likely people are to be persuaded (Douglas & Sutton, 2004), to how much wealth people own (Batista, Sussman, & Trueblood, 2023), and how much they engage in pro-environmental behavior (Leviston & Uren, 2020). Based on the aforementioned work finding self-other differences in a variety of prosociality contexts (blood donation - Miller & Ratner, 1998; Brick et al., 2021; everyday requests for help and environmental behavior - Kogut & Beyth-Marom, 2008; Chen, Wan, & Yang, 2022; negotiations - Mason, Wiley, & Ames, 2018; everyday requests for help - Zhao & Epley, 2022; climate policy support - Sparkman, Geiger, & Weber, 2022; COVID-19 collective health preventative behavior - Graupensperger, Lee, & Larimer, 2021; same-sex marriage approval - Eisner, Turner-Zwinkels, & Spini, 2021), we should expect that self-other differences in moral expansiveness might precede these findings, given that moral inclusion is the condition for much of the prosocial behavior and concern we see in the world (Crimston et al., 2016). That is, people might not only expect others to be less likely to help or cooperate in these contexts, they should expect that others are less morally inclusive than themselves in the first place.

The Current Research

There are three important questions that emerge: first, would we see self-other discrepancies in the perception of others' moral landscapes, and specifically in how people report their moral expansiveness? I propose to study self-other differences by having participants report their own moral expansiveness and anticipate that of others. Through this

method, people can showcase their level of concern by reporting who and what they (and others) would include in their circles of moral concern. Second, if potential differences existed, do people perceive themselves as more morally expansive than others? I predict that, in line with the literature I have cited thus far, people would underestimate others' moral circles relative to their own. Third, to the extent that people perceive themselves as morally expansive, the question still remains as to why, and whether this is a result of cynicism – the belief that people are primarily motivated by self-interest (Neumann & Zaki, 2022) – that fills in the gaps.

In Studies 1a and 1b, I address the first two questions – whether people report different moral circles for others compared to themselves and whether this difference underestimates others' moral expansion. In Studies 1a and 1b I find that people do indeed report lower moral expansiveness for others, compared to themselves. However, in both studies, this discrepancy appeared only for entities that are, on average (i.e. based on prior work, Crimston et al., 2016), placed further away from the self. My proposed dissertation studies that follow Studies 1a and 1b address the last question, namely, whether this selective discrepancy can be attributed to simulation gaps resulting from lack of access to information (Studies 2 & 3) that can be informed following an anti-cynicism manipulation (Study 4).

Specifically, following Studies 1a & 1b, Studies 2 and 3 will introduce a close-other moral expansiveness perspective to be compared and contrasted against the expansiveness of the self, expecting that the moral landscapes of close others will not display the same discrepancy even when those close others are thought of as dissimilar (Study 2), or when they are disliked (Study 3). Finally, Study 4 will retain the self, close other, general other structure to test whether reversing cynical assumptions has unique reducing effects of the discrepancy only on the moral landscapes of general others. We begin with Studies 1a and 1b to first understand whether

discrepancies exist. To the extent that others' moral worlds are imagined as constricted, this might provide a starting point to understand consistent underestimations relative to the self, across domains of prosociality.

Chapter 2: Studies 1a & 1b

We collected data from 2 separate online samples, one recruited in April of 2018 (Study 1a), and a replication in April of 2022 (Study 1b), asking participants to report their own moral circles, and those of others, as well as meta-perceptions about how others view the participant's own moral circles. The latter perspective was added in an exploratory fashion in an attempt to obtain a meta-perception of a potential self-other difference – that is, would any self-other differences also be reflected in how people think others view them? This third perspective can potentially triangulate the results in an informative way, because people's meta-perceptions of their immoral behavior as observed by general others tends to hover between accurate to overly positive (Lees, Young, & Waytz, 2022). Thus, if people perceive others as less morally expansive, and they also report that others view them as less morally expansive (contrary to existing work), this can give us more confidence about the existence of a potential perceived discrepancy across not just one, but two vantage points.

Study 1a

Participants

For Study 1a we recruited N=202 participants on Amazon Mechanical Turk (Paolacci, Chandler, & Ipeirotis, 2010). We excluded repeat participants and those who did not complete the Moral Expansiveness measure in full, resulting in a final sample of N=186. A sensitivity analysis for a repeated measures within factors ANOVA using G* Power software (Faul, Erdfelder, Lang, & Buchner, 2007) revealed that with a sample size of N=186, Study 1a could detect an effect as small as f = .118.

In Study 1a, the sample's mean age was M_{age} = 34.45 (SD=10.19) years, 55.4% of participants identified their gender as male, 43% as female, 1.1% as other, and 0.5% did not

provide an answer to the gender question. For their ethnicity, 72% of participants identified as White or Caucasian, 12.4% as Black or African American, 4.8% as Hispanic or Latino, 3.2% as Native American, 2.7% as Asian or Pacific Islander, 4.3% reported a combination of ethnicity backgrounds, and 0.5% did not provide an answer to the ethnicity question. The sample hovered around the mid-point on a 7-point conservativism scale (M=3.81, SD=1.93) ranging from Extremely Liberal (1) to Extremely Conservative (7) and appeared to be slightly to somewhat religious (M=2.56, SD=1.43) on a 5-point religiosity scale, ranging from Not at All Religious (1) to Extremely Religious (5).

Measures

Participants read about the meaning and significance of the moral circle and its different levels, and then were asked to complete the Moral Expansiveness Scale (hereafter referred to as "MES") as developed and validated by Crimston and colleagues (2016) (see Appendix A for full set of instructions and visuals the participants saw). All participants completed the measure three times in counterbalanced order, from their own perspective ("*How much do I care about each of these entities?*"), from the perspective of others ("*How much do others care about each of these entities?*"), and from the perspective of others on how they would expand ("*How much do others think I care about each of these entities?*").

The scale presented participants with 30 separate entities and asked them to place them on the level of the moral circle they believed those entities belonged in (inner circle/highest moral concern - scores 3 points, outer circle/moderate moral concern - scores 2 points, fringes/minimal moral concern - scores 1 point, outside of the moral boundary/no moral concern - scores 0 points). The highest possible score a participant could have on moral expansiveness if placing all entities in the inner most moral circle closer to the self would be 90

(scoring a 3 on all 30 possible entities). The entities themselves are conceptually grouped under the categories of family and friends, ingroups, outgroups, revered, stigmatized, villains, animals of high-sentience, animals of low-sentience, plants, and the environment (full list of entities available in the Appendix A, see Crimston et al., 2016 for selection and validation of materials).

Results

Moral Expansiveness Scale (MES) Total Comparisons. Means, standard deviations, and correlations between the three MES perspectives for Study 1a can be found in Table 1.

Table 1*:Reliabilities, Means, & Pearson's correlations between variables in Study 1a*

Variable	α	*M* (SD)	1.	2.	3.
1. MES self	.877	*M* =44.47 (12.56)	--		
2. MES others	.877	*M* =40.23 (12.28)	.60***	--	
3. MES others for self	.876	*M* =42.73 (12.43)	.71***	.64***	--

Note. * $p<.05$, ** $p<.01$, *** $p<.001$

A repeated measures ANOVA with a Greenhouse-Geisser[1] correction revealed a significant difference between the three MES perspectives, F (1.93, 357.23) = 15.66, $p<.001$, partial η^2 =.078. Post hoc tests using the Bonferroni correction revealed a significant difference (M_{diff} = 4.25, $p<.001$, 95% CI [2.27, 6.22]) between the scores for MES self (M=44.47, SE=.92) and MES others (M=40.23, SE=.90) such that people perceive others as less morally expansive compared to themselves. This difference (M_{diff} = 4.25) was the equivalent of a little over one entity being completely excluded from others' moral circle, compared to the self. There was also

[1] We used the Greenhouse-Geisser correction at all times, due to the assumption of sphericity being violated as is often the result of repeated measures designs. Where the sphericity assumption is not violated, the Greenhouse-Geisser correction conveys the same information as when sphericity is assumed. We report all corrections (Greenhouse-Geisser, Huynh-Feldt correction, Lower-Bound) only when they convey contradictory significance level information.

a smaller but significant difference (M_{diff} = 1.74, p= .038, 95% CI [.69, 3.41]) between the scores for MES self and MES others for self (M=42.73, SE=.91) such that people perceived others view of their own moral expansiveness as tracking this difference, namely that others view them as less expansive than they report to be.

MES Comparisons Across Entity Categories. In order to first check whether we replicate the general moral expansiveness entity order found in Crimston and colleagues' original validation work of the measure (2016), we averaged the entity categories across all three perspectives (self, others, others for self) (see Figure 1). Consistent with this prior work, participants in Study 1a perceived themselves and others, as well as how others see them as most inclusive towards friends and family (M=2.71, SD=0.51), followed by people in their ingroup (M=1.88, SD=0.49), revered entities (M=1.72, SD=0.51), stigmatized entities (M=1.55, SD=0.53), outgroups (M=1.35, SD=0.57), the environment (M=1.21, SD=0.60), animals of high perceived sentience (M=1.18, SD=0.63), animals of low perceived sentience (M=1.05, SD=0.68), plants (M=0.99, SD=0.68), and lastly, villains (M=0.52, SD=0.71).

Figure 1: *Means across the three MES perspectives for entity categories in Study 1a*

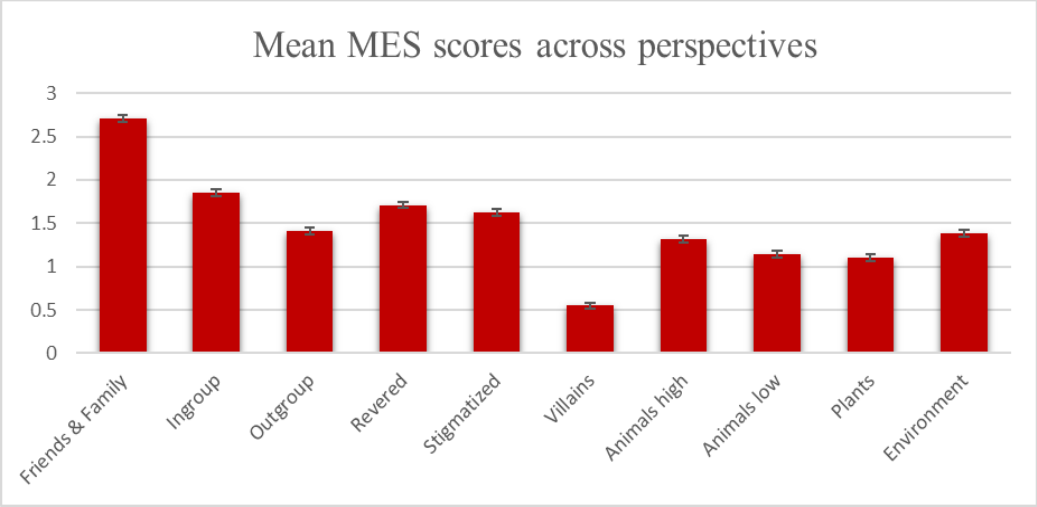

Note. *Error bars indicate standard error of the means.*

Close vs. Distant Comparison Across MES Self and MES Others. Given that we see differences across the three perspectives in the general moral expansiveness score, it is not yet clear whether these differences exist because people perceive themselves as more morally expansive than others across all entity categories, or whether there are specific categories that are driving the self-other discrepancy effect. We thus conducted separate self-other mean comparisons for all entity categories (see Appendix E in Supplementary materials for all entity comparisons). The overall pattern seemed to be that as people moved to the entities that are typically placed further from the center of the circle (particularly non-human entities), the self-other differences became stronger.

In order to avoid multiple comparisons between the three perspectives for all the entity categories, and to simplify the analysis of the above, we conducted a 2 (Entity: Close vs. Distant) x 3 (MES perspective: self, other, others for self) within-subjects repeated measures ANOVA. Within the Close entity variable, we included the categories of family and friends, ingroup, and revered entities, and within the Distant entity variable we included the categories of outgroups, stigmatized entities, animals of high sentience, animals of low sentience, plants, the environment, and villains. Beyond the fact that the distant grouping included the categories for which we saw significant self-other statistical differences (except the Villain category which suffered from floor effects, as is typical for this category, see Crimston et al., 2016), the decision to group the entities in this way was threefold.

First, there is a natural split between the entity categories by mean score, such that most of the entities in the Distant category (except the Stigmatized category), had inclusion scores below the midpoint (less than 1.50 out of the possible 3 points an entity could score). Second, and related to the first point, the split followed an ordinal logic, based on the sequence of

category scores (see Figure 1) - that is, in order from highest to lowest inclusion scores, the Close category entities are the first three, whereas the Distant category entities are the following and last seven. Lastly, we grouped these entities in this way because of theoretical reasons: friends and family, ingroups, and revered entities might have higher perceived overlap and similarity with the self than outgroups, stigmatized groups, non-human entities, and villains. The last reason also justifies why categories that wouldn't otherwise make it to the relevant group were placed there: even though the Stigmatized category and the Villain category did not fit all of the prior criteria, it made theoretical sense to include both in the Distant, as opposed to the Close category.

The entity x MES perspective interaction was significant, F (1.88, 347.80) = 17.42, p<.001, partial η^2 =.086, with Bonferroni corrections showing that the entities that are typically placed closer to the inner circles did not show the discrepancy between the participants' own expansiveness ($MES_{self_clo}$$SE$= 2.09, SE=0.03) and the predicted MES score for others ($MES_{others\_clo}$$SE$=2.11, SE=0.03) (M_{diff} = -0.02, p= 1.000, 95% CI [-0.09, 0.05]), or the predicted MES score for others for the self ($MES_{others\ for\ self\_close}$ =2.11, SE=0.03) (M_{diff} = -0.02, p= 1.000, 95% CI [-0.08, 0.04]), but the distant entity categories for self ($MES_{self\_distant}$= 1.22, SE=0.04) were significantly more included compared to the distant entity categories that were predicted for others ($MES_{others\_distant}$=1.01, SE=0.04) (M_{diff} = 0.21, p<.001, 95% CI [0.13, 0.29]), and what others would predict for the self ($MES_{others\ for\ self\_distant}$ =1.13, SE=0.04) (M_{diff} = 0.09, p= .004, 95% CI [0.02, 0.16]) (see Figure 2).

Study 1b

The main purpose of Study 1b was to pre-register and replicate the results of Study 1a. Our hypotheses for Study 1b were pre-registered on the Open Science Framework (access at

https://osf.io/kmqpd) and consisted of the following:

1) We expect that the average moral expansiveness score (MES) for the self will be higher than the one participants will report for others' MES.
2) We expect that the average moral expansiveness score (MES) for the self will be higher than the ones participants will report for others for self.
3) We expect a significant entity type (close vs. distant) x MES perspective (self, other, others for self) interaction, such that the entities that are typically placed closer to the inner circles (that is: family & friends, ingroup, revered) will not show the discrepancy between self MES and the other two MES perspectives, but the entity categories that are typically placed closer to the fringes (that is: outgroup, stigmatized, animals of high sentience, animals of low sentience, plants, environment, villains) will be higher for the self MES compared to other MES, as well as compared to the prediction of others for self MES.

Participants

For Study 1b we recruited *N*=299 participants on CloudResearch.com (Litman, Robinson & Abberbock, 2017). We again excluded repeat participants and those who did not complete the Moral Expansiveness measure in full, the sample for Study 1b after exclusions was *N*=286. A sensitivity analysis for a repeated measures within factors ANOVA using G* Power software (Faul, Erdfelder, Lang, & Buchner, 2007) revealed that with a sample size of *N*=286, Study 1b could detect an effect as small as f = .095.

In Study 1b, the sample's mean age was M_{age} = 39.15 (*SD*=11.03) years, 60.1% of participants identified their gender as male and 39.9% as female. For their ethnicity, 76.6% of participants identified as White or Caucasian, 7.7% as Asian or Pacific Islander, 7% as Black or

African American, 2.4% as Hispanic or Latino, 0.3% as Native American, and 5.8% reported a combination of ethnicity backgrounds. The sample leaned slightly liberal on a 7-point conservativism scale (M=3.29, SD=1.90) ranging from Extremely Liberal (1) to Extremely Conservative (7), as well as slightly religious (M=2.08, SD=1.42) on a 5-point religiosity scale, ranging from Not at All Religious (1) to Extremely Religious (5).

Measures

Participants saw the same Moral Expansiveness Scale instructions as with Study 1a, and identically filled out the measure three times in counterbalanced order, from their own perspective ("*How much do I care about each of these entities?*"), from the perspective of others ("*How much do others care about each of these entities?*"), and from the perspective of others on how they would expand ("*How much do others think I care about each of these entities?*"). (For additional exploratory measures and analyses included in Study 1b refer to Supplemental Analyses in Appendix E).

Results

Means, standard deviations, and correlations between the three MES perspectives for Study 1b can be found in Table 2.

Table 2: *Reliabilities, Means, & Pearson's correlations between variables in Study 1b*

Variable	α	*M (SD)*	1.	2.	3.
1. MES self	.920	*M* =43.92 (13.23)	--		
2. MES others	.921	*M* =39.82 (12.83)	.56***	--	
3. MES others for self	.917	*M* =42.51 (13.08)	.69***	.65***	--

Note. * p<.05, ** p<.01, *** p<.001

Moral Expansiveness Scale (MES) Total Comparisons. We conducted a repeated

measures ANOVA with a Greenhouse-Geisser correction which showed a significant difference between the three perspectives, F (1.92, 546.11) = 19.77, p<.001, partial η^2 =.065. Post-hoc tests with a Bonferroni correction revealed a significant difference between participants' reports for MES self (M=43.92, SE=.78) and MES others (M=39.82, SE=.76) in the same direction and magnitude as with Study 1a, such that people perceived more constricted moral circles for others compared to themselves by about 4 points (M_{diff} = 4.09, p<.001, 95% CI [2.35, 5.84]), which supported pre-registered hypothesis 1. Unlike Study 1a, the difference between MES self and MES others for self (M=42.51, SE=.77), was not significant, (M_{diff} = 1.41, p= .069, 95% CI [-.075, 2.89]), which did not support pre-registered hypothesis 2, although the effect was trending in the expected direction.

Close vs. Distant Comparison Across MES Self and MES Others. The order of moral inclusivity across the entity categories replicated in the same manner in Study 1b, so we moved directly into testing the comparison of the self and others MES categories across close and distant entity categories. The entity (Close vs. Distant) x MES perspective (self, others, others for self) interaction was significant, F (1.89, 537.61) = 37.08, p<.001, partial η^2 =.115, with close entities showing no discrepancy between the participants' own expansiveness ($MES_{self\_close=}$ 2.16, SE=0.03) and the other two MES perspectives of others ($MES_{others_close=}$2.21, SE=0.03) (M_{diff} = -0.05, p= .085, 95% CI [-0.10, 0.00]), and others for self ($MES_{others\ for\ self\_close}$ =2.20, SE=0.03) (M_{diff} = -0.03, p=.259, 95% CI [-0.08, 0.01]), while distant entities exhibited this difference again between participants' own MES score ($MES_{self\_distant}$= 1.16, SE=0.03) and both that of others ($MES_{others\_distant}$=0.95, SE=0.03) (M_{diff} = 0.21, p<.001, 95% CI [0.14, 0.29]) and the others for self perspective ($MES_{others\ for\ self\_distant}$ =1.08, SE=0.03) (M_{diff} = 0.08, p= .004, 95% CI [0.02, 0.14]) (see Figure 2). Thus, Study 1b replicated this finding and supported pre-registered

hypothesis 3.

Figure 2: *The entity x MES perspective interaction in Studies 1a & 1b*

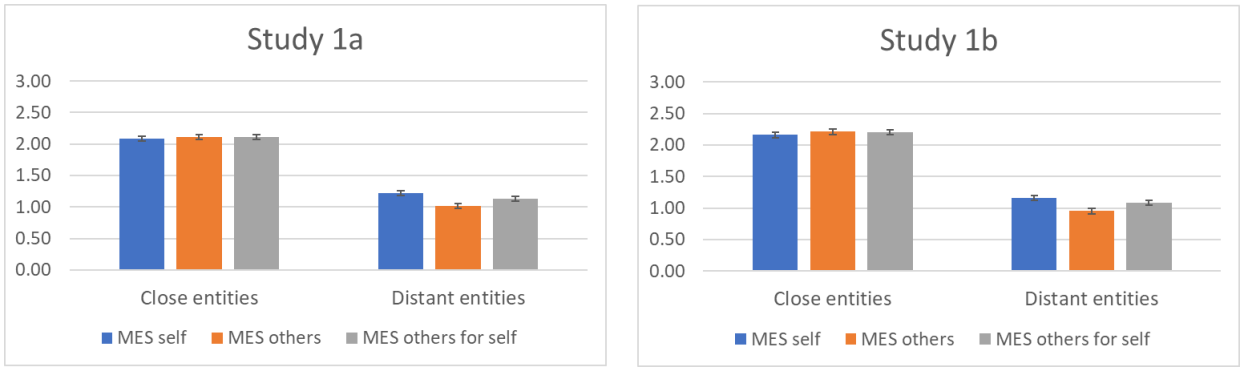

Note. *Plotted estimated marginal means. Error bars indicate standard error of the means.*

Heterogeneity in MES Discrepancy Scores

In order to understand whether the discrepancy we observe between the perceived MES scores for self and others is ubiquitous, or just driven by a few outliers in the Study 1a and 1b samples (see McManus, Young, & Sweetman, 2023 for discussing the importance of heterogeneity of responses beyond just averages), we created and took a closer look at the actual MES discrepancy score. We obtained this score (see Figure 3 for plotted discrepancy scores) by subtracting the MES score that participants reported for others, from the MES score they reported for themselves. A score of 0 means no difference between the two, a negative score indicates that participants, on average, overestimated others' expansiveness relative to their own, and a positive score means they underestimated others' expansiveness relative to their own.

Figure 3: *Discrepancy MES scores in Studies 1a (left) & 1b (right)*

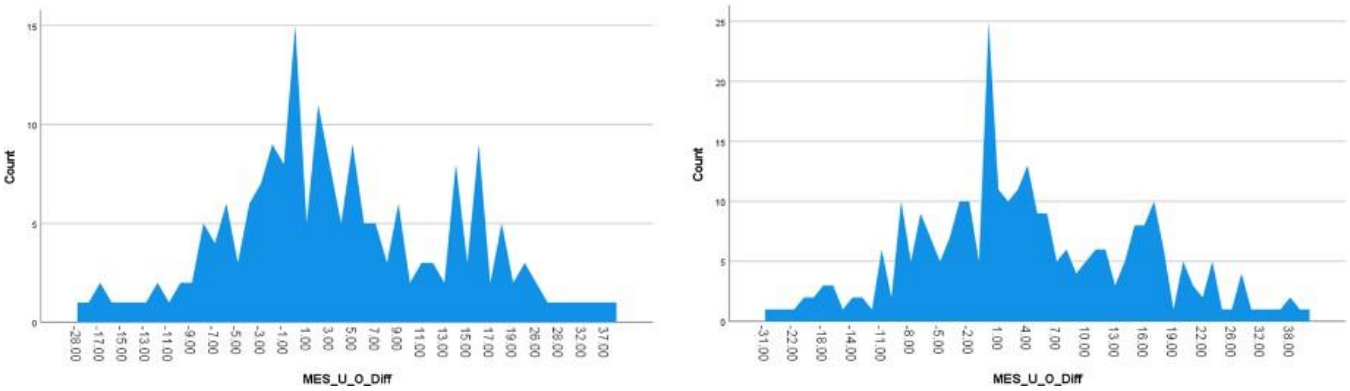

Note. *The X axis shows the amount of discrepancy and the Y axis the number of participants who displayed it. The spike on 0 indicates participants who, on average, matched their MES score with what they predicted for others. A positive skew is observed in both studies.*

In Study 1a, 41.9% of participants overestimated others' MES (difference score = −) or had no discrepancy (difference score = 0) between their self-reported MES and that of others, whereas 58.1% of participants underestimated others' MES compared to their own (difference score = +). A Chi square test comparing the two percentages against each other revealed a significant difference between the two, X^2= 4.84, p=.028, such that there were more people with an underestimation difference score, compared to those who reported a match and those who overestimated, combined. This pattern replicates in Study 1b, where 42.3% of participants similarly overestimated or reported the same MES for others as with their own, whereas 57.7% of participants underestimated other's MES compared to their own. A Chi square test comparing the two percentages against each other revealed a significant difference between the two, X^2= 6.77, p=.009.

Thus, while it appears true that there is heterogeneity in the sample in terms of the underestimation of others' MES compared to one's own, on average, people underestimated more than they overestimated or were accurate. This is further exacerbated when we break down the moral inclusion score to the specific Distant categories which showed the larger

discrepancies. To better exemplify, we have plotted the difference score between self and other for the category of friends and family as an example of a Proximal entity sub-category in the plots below (top two graphs in Figure 4). Virtually no one underestimates or overestimates others' moral expansiveness relative to the self in the friends and family category, as almost all scores can be seen hovering around 0 in both Studies 1a and 1b. This is in contrast to an example of a Distant entity sub-category, animals of low sentience, which presents a clear skew towards underestimation (bottom two graphs in Figure 4).

Figure 4: *Discrepancy MES scores in Studies 1a & 1b for family and friends & low-sentience animals.*

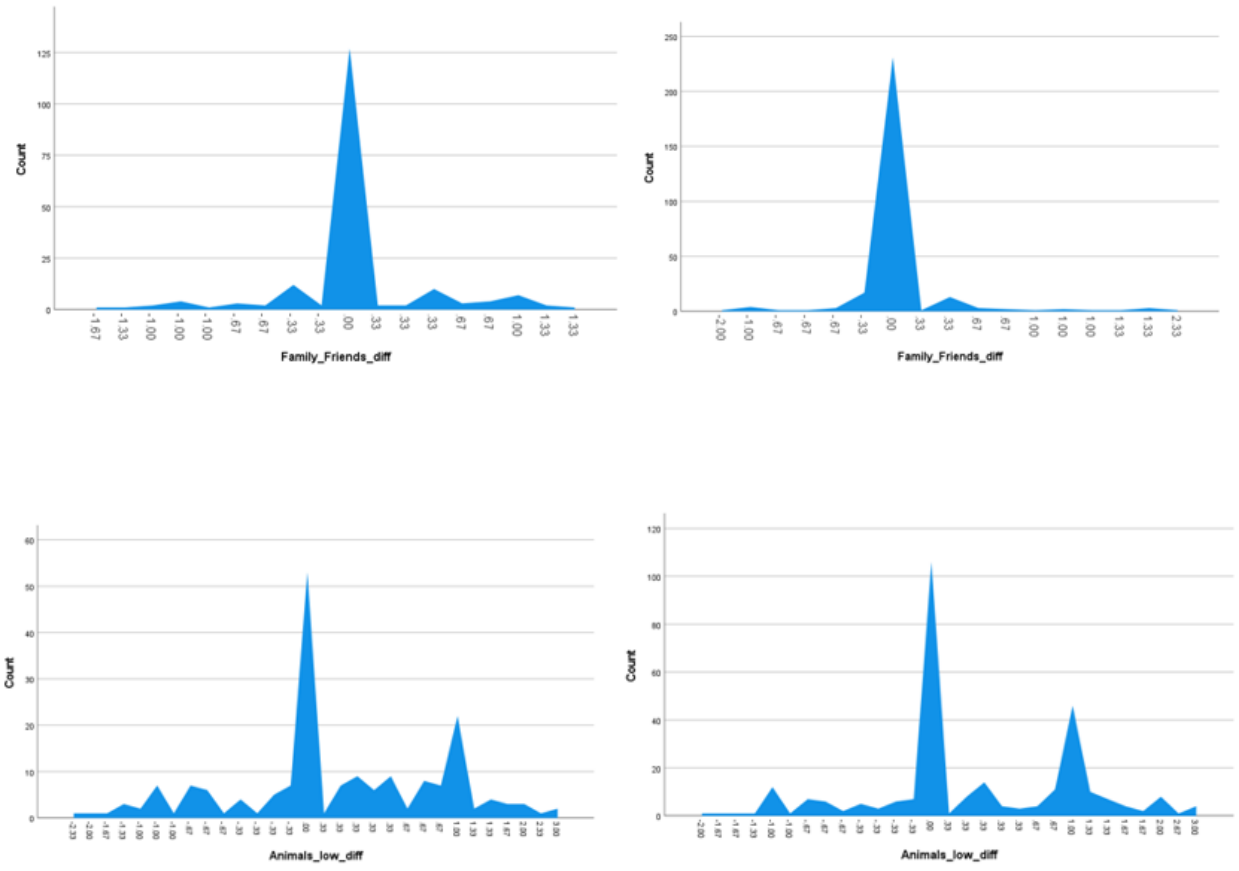

Note. *The X axis shows the amount of discrepancy and the Y axis the number of participants who displayed it. A positive skew is observed in both studies.*

Interim Discussion for Studies 1a & 1b

In sum, across the preliminary online samples recruited in 2018 and 2022, and consistent with prior research on how American samples score on the MES (Kirkland et al., 2022), in Studies 1a and 1b, participants scored about 45 out of the 90 maximum possible points when asked to report their own moral circle. Importantly, people consistently rated their own moral expansiveness as higher compared to that of others for some entity categories, and, to a smaller extent, they also rated their moral expansiveness for these entities as higher than what others would predict for them. The expansiveness score predicted for others had a negative average difference of about 4 points with the expansiveness of the self, the equivalent of over one entity being completely excluded from the moral circle. Therefore, we see support for our initial hypothesis that people would underestimate others' moral expansiveness relative to their own (pre-registered hypothesis 1).

In Study 1a, we see this both in an actual sense (difference between self and other) but also in a relative sense, because people report lower moral expansiveness for how others view them, compared to how they see themselves (difference between self and others for self). Although this latter finding on the self/others-for-self difference did not cleanly replicate in Study 1b for the general expansiveness score (pre-registered hypothesis 2, see Supplemental Analyses in Appendix E for specific Close entity category means potentially driving the MES totals closer together), it deserves a note about how it can potentially inform the rest of our findings. We do see that the categories on which people had discrepant views of how others viewed them, compared to what they reported for themselves, were the same categories in which the greatest self-other differences existed: animals of high and low sentience, plants, and the environment. This can indicate that, to some extent, people's perception that these entities

are morally excluded by others is also reflected on how they think others view them: they expect that others will perceive the same moral inclusivity discrepancy as they have.

Our results clearly identify an entity by type interaction, whereby people are more prone to report the self-other moral inclusivity discrepancy for certain entities that are typically placed further out in the fringes of moral concern (pre-registered hypothesis 3). This is consistent with prior work showing that whether or not pluralistic ignorance (the belief that others hold an opinion different than your own; Katz & Allport, 1931; Prentice & Miller, 1996; Sargent & Newman, 2021) will unfold depends on the specified target in question. For example, in a recent study, Eisner and colleagues (2020) found that Swiss citizens exhibited pluralistic ignorance for a relatively newly debated social target (same-sex female parenting), but participants did not exhibit pluralistic ignorance and instead were accurate about others' views on working mothers, a target whose existence was debated in the past but can be considered, by most, as having reached social consensus as acceptable. Why might we see these differences in our data? There might exist several reasons.

First, and following Eisner and colleagues (2020) research, it is possible that the entities in which we see differences in the moral expansiveness of self and others can sometimes similarly fall under a "gray area" for which social consensus might have yet to be reached (such as with the stigmatized category in the MES measure that includes LGBT individuals, mentally challenged individuals, and refugees). We also see differences, however, in entities that are not "new" in the social landscape - animals, plants, and the environment. Why might those exist? A second possible reason might be related to a kind of fundamental attribution error (Ross, 1977) whereby people perceive these entities (unlike one's friends and family, revered entities, and ingroups) as receiving poor treatment (e.g. exploited and polluted environment, industrialized

farming of animals for human consumption) and thus someone must be treating them poorly, so others must not be as moral expansive towards these entities. This leads to a third potential explanation, a well-documented need for self-enhancement (Epley & Dunning, 2000; Kwan, John, Kenny, Bond, & Robins, 2004; Klein & Epley, 2017).

Reported self-other differences often reflect self-enhancement motives: people think they are more likely than others to eat the way they do because of health reasons, as opposed to less desirable motives, like social pressure (Sproesser, Klusmann, Schupp, & Renner, 2017). Similarly, people think they are less easily persuaded by the media compared to others (Douglas & Sutton, 2004), and the moral domain, specifically, people might be motivated to feel morally superior, and thus report themselves as higher than average others on a host of certain moral traits, when directly asked to do so (Tappin & McKay, 2017). Nevertheless, this doesn't correlate with self-esteem, which complicates the argument about the moral superiority motivation (Tappin & McKay, 2017). To the extent that people don't derive an elevated view of the self when reporting higher moral averages compared to what they predict for others, the question of why people might be motivated to perceive this difference remains unanswered. Both Tappin and McKay (2017), as well as an array of literature on cynicism about the motives of others, pinpoint to the possibility of underestimating others' moral compass, relative to one's own, as an alternative explanation.

In our data, it is possible that people are motivated to see themselves in a positive light, thus overestimating their own expansiveness compared to others. However, there are at least two reasons as to why the self-other discrepancy effect might not be solely attributable to a self-enhancement bias. The first is that if there was a strong motivational goal to self-enhance, people would likely report higher moral inclusivity across *all* entity categories. That is, we

should have seen this applied universally in the data, and not just in select distant entities. If people were hoping to gain reputational benefits we might expect that those would perhaps be easily conferred by self-enhancing across the categories of friends and family, ingroups, and revered entities; morally including such categories over and above others' moral inclusion of the same categories would mean that the self is a superior parent, friend, husband/wife, daughter/son, neighbor, or parishioner to others, However, we do not see self-other differences and thus any evidence of self-enhancement across those categories.

The second reason why the self-other discrepancy effect might not be solely attributable to a self-enhancement bias is that the discrepancy between participants' self-reported moral expansiveness and the one they thought others would predict for them is also discrepant in the same direction, in the same specific categories (pre-registered hypothesis 3). If people are motivated to see themselves in a positive light, then they should also expect others to view them in such a light, which is consistent with recent research on meta-perceptions of immoral behavior, that finds that people's estimations of how others view them in the moral domain is more positive than the reality (Lees, Young, & Waytz, 2022). Nevertheless, in both Studies 1a and 1b, we see that people accurately predict that others don't see them in this way – specifically, the categories in which the self-other discrepancy was the largest (i.e. animals, plants, and the environment) were also the ones were people reported others' perception of them as consistently lower, which means people are, somewhat ironically, accurately tapping into the reality of these discrepancies.

Lack of Access to Information Allows for Cynical Bias to "Creep In"

While we are not entirely equipped to rule out all the possible explanations of the self-other moral inclusivity discrepancy offered above, we can make an educated guess based on

prior literature on how people imagine others' perspectives. I suggest that these differences emerge because of a combination of factors. The groundwork for this discrepancy is potentially laid out by a failure of simulation – people do not have access to others' minds and thus make "educated guesses" for things they are removed from and cannot personally experience (Epley, 2008). This, in turn, gives rise to a general cynical bias that "fills in" these gaps. In other words, people might have less information about how other people morally expand, thus be more prone to inaccuracies, particularly for entities that are further away from the self and therefore harder to simulate for.

Why are people showing a discrepancy when estimating others' moral inclusivity for further entities but not closer entities? One speculation is that we see this effect because people are generally better able to simulate things that are familiar and closer to their own experience (e.g. friends and family) than those that are further away (e.g. animals and the environment) (Preston & De Waal, 2002; Epley, 2008 although see also Malle, 2006 on actor-observer meta-analysis for some conflicting results). Thus, we might see a failure to simulate others' inclusion well for entities for which there is insufficient or inaccurate information for. This harkens back to Chen and colleagues (2022), who found that people underestimated others' pro-environmental behaviors and their approval of those behaviors. In their results, this underestimation was contingent on the observability of the behavior. Behaviors that the authors categorized as belonging to the public sphere (e.g. volunteering for environmental causes, spreading knowledge on environmental protection to family and friends) elicited smaller underestimations compared to behaviors that were categorized into the private sphere (e.g. choosing public transportation or bicycles instead of car, avoid opening the refrigerator door for too long). The observability point is an important one because it raises the possibility that cynicism is contingent on the availability

of information about prosocial behavior. However, this lack of information can lead people to either overestimate others' moral expansion, or underestimate it, relative to the self. In both Studies 1a and 1b, we see consistent underestimation. I propose this discrepancy is due to cynicism – the belief that people are primarily motivated by self-interest (Neumann & Zaki, 2022).

It might be the case that once people have access to information about others' prosociality, cynicism is reduced or goes away completely. On the other hand, scholars have often talked about cynicism as persisting even at the face of disconfirming evidence. In many cases, cynicism has been described as "naive" (Kruger & Gilovich, 1999), or as being "undue" (Critcher & Dunning, 2011), precisely because it can contradict reality even in the eyes of the cynic. For instance, based on participants' performance on a task where they were asked to come up with reasons for big philanthropic donations, Critcher and Dunning (2011) observed that the more participants had the chance to come up with reasons for those donations, the more they generated selfish (but not selfless) explanations, suggesting that even with additional resources that should enable more accurate perception of reality, people do not always dismiss cynical assumptions. Similarly, other work finds that cynicism is positively correlated with endorsement of unsubstantiated claims such as conspiracy theories and paranormal beliefs, constructs that are in opposition and negatively correlated with sound scientific skepticism that emphasizes true facts (Bensley, Watkins, Lilienfeld, Masciocchi, Murtagh, & Rowan, 2022). Thus, the extent to which underestimations of others' prosocial tendencies relative to one's own are a result of lack of access to information or a persistent motivation to maintain beliefs about selfishness, remains unanswered.

My proposed dissertation work aims to test whether these discrepancies are, at least

partially, due to cynicism that “fills in” the gaps that exist because of social distance (see Figure 5 for visual depiction). In order to better understand whether people think others morally expand as much as they themselves do, and whether this is a result of such cynical bias, we need to examine self-other comparisons in the expansion of the moral circle across different levels of social distance, and thus across different levels of access to information about others’ moral landscapes. This can be achieved by re-drawing definitions of “others”. For example, people might be cynical about the moral expansiveness of others if “others” are strangers, but if others are more concretely specified, or even further, are people whose moral landscape participants have access to, such cynical bias might go away (Mastroianni & Gilbert, 2023). However, to the extent that self-other differences persist in a systematic way, even if access to information increases, then this can be an indication that these self-other differences exist for reasons other than uninformed cynical bias (although, in some cases, such bias has been found even in the case of close others, see Zhao & Epley, 2022). An alternative conclusion is that cynicism in the context of moral expansion exists but is, in fact, naive, persistent, and motivated in nature (Kruger & Gilovich, 1999; Critcher & Dunning, 2011), and thus immune to belief updating, even following increased access to information about the moral landscape of others.

Figure 5*: Schematic representation of theoretical framework*

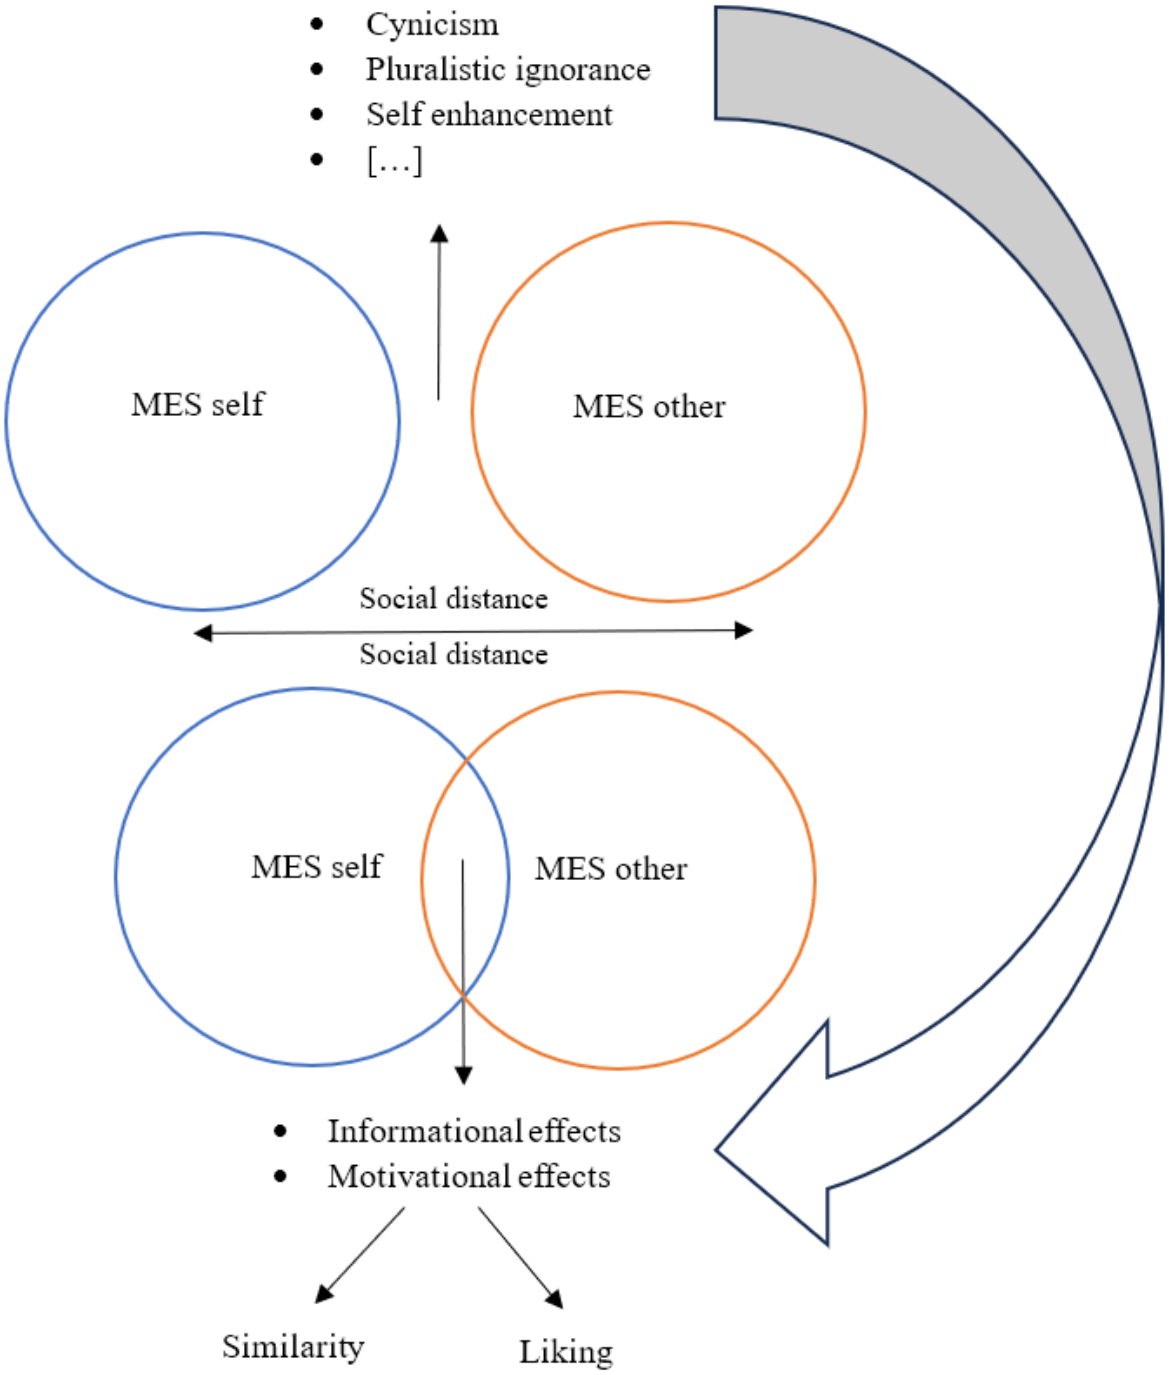

Note. *Visual depiction of proposed theoretical framework for explanation of the self-other discrepancy. The figure outlines that as social distance of self and other increases, so does the discrepancy (top two circles). In the cases where social distance is eliminated (bottom two circles), there could be both informational and motivational explanations for the perceived overlap between self and other.*

Self, Close-Other and Average-Other Comparisons

In the present research, I tested whether access to information about others' moral expansiveness inhibits the moral inclusivity discrepancy, or whether the discrepancy persists even as people have more access 'to another's moral landscape. A very recent example in the literature that employed this approach is the one from Mastroianni and Gilbert (2023) who asked participants to report on the perceived moral decline (i.e. becoming less kind as time goes by) of general others as well as others "in their personal worlds" defined as "all the people with whom you currently interact, in person or otherwise, in your everyday life. This probably includes friends, family members, coworkers, classmates, neighbors, etc." Mastroianni and Gilbert (2023) found that while there was a general perception of moral decline in perceptions of general others, this effect vanished for those whom they defined as close others. Their interpretation (one they admit can be among many others) is exposure to positive (vs. negative) information about the moral behavior of close others. My studies are envisioned to attempt something similar - even though we do not have the data for the moral expansiveness for close others, we expect that the finding will generalize across variations of perceived expansiveness (in the same way it is persistent across variations of perceived moral decline from close others in Mastroianni & Gilbert, 2023).

It's important to note that while Mastroianni and Gilbert's work (2023) proposes a similar mechanism, the dependent variable in is temporal in nature: they find that perceptions of others' morality declines over time, whereas our assessment of morality is agnostic to time and instead is more spatial in nature, with the depictions of the moral landscape expanding outwards in concentric moral circles (Crimston et al., 2016). Mastroianni and Gilbert (2023) find that this moral decline is believed to happen while people acknowledge, at the same time, that the

treatment of marginalized or stigmatized groups in society (e.g. racial minorities) has gotten better. This is in stark contrast to our findings in Studies 1a and 1b, because the categories in which perceived moral expansiveness of others has the largest gap, relative to the self, are precisely the stigmatized or otherwise socially distant ones (e.g. refugee). Thus, while the recent work by Mastroianni and Gilbert (2023) lends support to our hypothesized mechanism about social distance determining whether a perceived self-other gap in morality will surface, our work remains distinct.

In Studies 2, 3, and 4 we removed the barrier of lack of information by asking people to identify and report on a close other's moral landscape, which should be readily accessible. In addition to self and close-other, participants in all 3 studies again reported the moral expansiveness of an average person (note that the wording changed for this perspective from "others" in Studies 1a and 1b, to an "average other" in Studies 2, 3, and 4, but results remained identical). By asking people to report the estimated moral circles of people who are closer to them, we aimed to decrease social distance and potentially increase access to information. We expected that close others' moral circles will be more aligned, if not identical, to the self.

Studies 2, 3, and 4, while high in ecological validity given participants' direct link with a real-world referent, still conflate a number of explanations for a self-other alignment, and thus an effect could not be solely attributable to increased access to information. First, people gravitate towards people with whom they share similar moral values (Byrne, 1961; Philipp-Muller, Wallace, Sawicki, Patton, & Wegener, 2020; Brown, 2020), thus it would be unsurprising if their moral landscapes looked similar. For this reason, Study 2 additionally employed a similarity manipulation referring to the close-other person participants reported on, whereby participants either reflected on being similar or dissimilar from their close others.

Second, people like those close to them more than general others, and this liking can color their motivations to paint them as highly moral (Bocian, Baryla, Kulesza, Schnall, & Wojciszke, 2018; Lee & Holyoak, 2020) which, in this context, could translate to seeing them as highly morally expansive. In order to address these concerns, and to better understand whether access to information determines whether the self-other moral inclusivity discrepancy will occur, we also conducted Study 3, where the manipulation targeted liking, and participants either reflected on reasons they liked or disliked their close others. If perceptions of others' moral expansiveness significantly change as a function of either feeling more similar to them (Study 2) or liking them more (Study 3) compared to the average person, this would speak to perceptions of moral expansiveness as not primarily relying on informational access, but instead prone to interpersonal motivation effects. In other words, to the extent that people's predictions of their close others' expansiveness changes as a function of how much they are encouraged to feel similar to them, or reminded of reasons why they like them, this would refute the prediction that informational effects are primarily at play.

Finally, for those targets that social distance cannot be eliminated (strangers) and thus stereotypes are more prone to be at work (Pratto & Bargh, 1991; Rubinstein, Jussim, & Stevens, 2018), correcting cynical assumptions can buffer against this discrepancy. In Study 4, we directly attempted to manipulate people's assumptions about cynicism in the world, expecting that this will increase only the perceived moral expansion of others', but not the self's or their close others' moral expansiveness, for which participants already have informational access to. That is, while we can decrease the perceived social distance of general "others" by asking people to bring to mind a specific close other (e.g. a family member or friend) whose moral landscape would be more readily accessible, the puzzle still remains about how people can be made less cynical about

general others whom they have no relation to. In these cases, changing the cynical stereotypes that people hold is one avenue to address the self-other moral inclusivity discrepancy for unknown others, which Study 4 attempted to do. Thus, this set of studies will attempt to test whether the discrepancy between self and other is both a) dependent on social distance and b) is a result of cynical assumptions that people hold.

Chapter 3: Study 2 - MES of Similar vs. Dissimilar Close Others

Study 2

The aim of Study 2 was to examine one potential antecedent of the self-other discrepancy observed in the 2018 and 2022 samples, specifically access to information as afforded by social distance, and moderation by felt similarity. Because we would expect that people will be more familiar with a close other's moral circle than with a stranger's, we expected that the moral inclusivity discrepancy people showed in Studies 1a and 1b will be reduced or eliminated for the moral circles people report for those close to them, but that it would persist for average others (see Figure 5). We also expected that we would successfully reduce felt similarity to close others in the dissimilarity condition, but that moral expansiveness scores for close others would remain unchanged across the dissimilarity, similarity, and control conditions, because access to information about close other's moral landscapes remains unchanged. We additionally expected to observe significant negative correlations between three separate cynicism measures and all MES scores. Our hypotheses for Study 2 were pre-registered on the Open Science Framework (access at https://osf.io/3su2h) and consisted of the following:

1) For distant entities (but not proximal entities) in the moral circle, we expect that the moral expansiveness score (MES) for the self will be higher than the one participants will report for average others' MES.
2) For both distant and proximal entities in the moral circle, we expect that the moral expansiveness score (MES) for the self will not be significantly different from the one participants will report for close other's MES under any of the close-other manipulations (similarity, dissimilarity, control).

3) For distant entities (but not proximal entities) in the moral circle, we expect that the moral expansiveness score (MES) for close others will be higher than the one participants will report for average others' MES under all of the close-other manipulations (similarity, dissimilarity, control).
4) For both distant entities and proximal entities in the moral circle, we expect that the moral expansiveness score (MES) for close others will not be significantly different across the three between-subjects manipulations (similarity, dissimilarity, control).
5) We expect that participants assigned to reflect on being dissimilar to their close other will report lower felt similarity with their close other, compared to participants assigned to the similarity-close other and control-close other conditions.
6) We expect that three measures of cynicism (subscale from Cook & Medley, 1954; Chowdhury & Fernando, 2014; competitive primal world beliefs, Clifton et al, 2019) will correlate negatively with all MES scores (self, close others, average others).

Participants

Based on the effects found in Studies 1a and 1b and the expectation for some attrition, Study 2 aimed to recruit at least N=350 participants online on CloudResearch.com (Litman, Robinson & Abberbock, 2017). We restricted our search to those residing within the United States. Following exclusions of participants with incomplete data on the MES, nonsensical responses to the open-ended prompts of close-other identification and similarity/dissimilarity, or who failed the attention check, we obtained a final sample of N=298. A sensitivity analysis for a repeated measures ANOVA with a within-between interaction using G* Power software (Faul, Erdfelder, Lang, & Buchner, 2007) revealed that with a sample size of N=298, Study 2 can detect an effect as small as f = .102. Participants were assigned to between-subjects conditions using

the "evenly-present elements" feature in Qualtrics software. The sample was fairly evenly distributed across the three between-subject conditions of close-other similarity (N=98), close-other dissimilarity (N=91), and control (N=100), and the vast majority of attrition occurred in the introductory stages of the survey, before participants saw any content that was differentiated by condition assignment.

In Study 2, the sample's mean age was M_{age} = 40.65 (SD=11.25) years, 55% of participants identified their gender as male, 43.9% as female, and 1% as other. For their ethnicity, 73% of participants identified as White or Caucasian, 11.8% as Black or African American, 5.9% as Asian or Pacific Islander, 2.1% as Hispanic or Latino, 6.8% reported a combination of ethnicity backgrounds, and 0.3% chose "other" as their ethnicity category. The sample was near the mid-point on a 7-point conservativism scale (M=3.47, SD=1.81) ranging from Extremely Liberal (1) to Extremely Conservative (7) and appeared to be slightly to somewhat religious (M=2.40, SD=1.42) on a 5-point religiosity scale, ranging from Not at All Religious (1) to Extremely Religious (5).

Measures

Participants were asked to fill out the Moral Expansiveness Scale (Crimston et al., 2016) and following the average other and close other MES perspectives, participants were also asked to report how similar they felt towards and how much they liked 1) the average person and 2) their identified close other, respectively, on a scale from 1 (Not at all) to 7 (Extremely). In a multiple-choice question, participants were also asked to report the relationship they had with their identified close-other (i.e. friend, family member, partner, colleague, or other category). Participants were also asked to complete three separate cynicism measures which included a) the cynicism subscale from Cook & Medley's hostility scale (1954), example item: "*I think*

most people would lie to get ahead."; b) the cynicism scale adapted from Wrightsman's philosophies of human nature scale (1964), example item: "*People pretend to care more about one another than they really do*." ; and c) the cooperative vs. competitive primal world beliefs scale from Clifton et al., (2019), example item: "*For all life—from the smallest organisms, to plants, animals, and for people too—everything is a cut-throat competition*." (For complete scale measures refer to Appendix D).

The measures have been widely validated (Smith & Frohm, 1985; Alsaid, Li, Chiou & Lee, 2023; Clifton & Yaden, 2021), and although the target construct from the use of these scales in psychological and market research varied from attempts to measure distrust, to belief in a "jungle" or "competitive" world, they all have also been categorized as "cynicism" (Choy, Eom & Li, 2021; Chowdhury & Fernando, 2014; Meehan, Zeigler-Hill, & Shackelford, 2022). We thus included these measures to broadly gauge the relationship between the MES and the construct of cynicism (the belief that people are primarily motivated by self-interest). (For additional exploratory measures and analyses included in Study 2 refer to Supplemental Analyses in Appendix E).

Procedure

After participants consented to take part in the study, they were given the same instructions and information pertaining to the MES measure as with Studies 1a and 1b, and all participants were similarly asked to fill out the MES measure from three perspectives. Two of the perspectives were identical to Studies 1a and 1b, such that participants were asked to complete the MES from their own perspective, and from the perspective of an average other. However, for the third perspective that asked participants to fill out the MES from the perspective of a close other, participants were first asked to bring to mind and type in the name

of a close other, and were subsequently assigned to three between-subjects conditions: the similarity condition, in which they were asked to reflect and write 2-3 sentences on what makes them similar to their identified close-other; the dissimilarity condition, in which participants were asked to reflect and write 2-3 sentences on what makes them dissimilar to their identified close-other, or the control condition, which had no similarity manipulation and instead asked participants to identify their close other and then proceed to fill out the MES (see Appendix B for condition prompts). After participants filled out these three perspectives, they moved on to complete the main cynicism and other exploratory measures, filled out demographics, and lastly, they were debriefed.

Results

For their close other, 54.3% of participants identified a friend (N=157), 23.5% identified a partner (N= 68), 18.3% a family member (N= 53), 2.8% a colleague (N=8) and 1% of participants chose someone belonging to an "other" category (N=3).

We conducted a repeated measures ANOVA test on the 7-point liking scale, F (1, 288) = 633.21, p<.001, partial η^2 =.687, which confirmed that, on average, the identified close others were liked (M=6.29, SD=0.96, 95% CI [6.18, 6.40] significantly more compared to average others (M=4.27, SD=1.17, 95% CI [4.16, 4.41], (M_{diff} = 2.02, p<.001, 95% CI [1.86, 2.18]). A separate test on the 7-point similarity scale, F (1, 288) = 130.41, p<.001, partial η^2 =.312, additionally confirmed that close others were also perceived as more similar to the participant (M=5.11, SD=1.38, 95% CI [4.95, 5.27], compared to average others (M=3.90, SD=1.54, 95% CI [3.72, 4.08], (M_{diff} = 1.21, p<.001, 95% CI [1.00, 1.42]).

Prior to comparing the self, average other, and close other MES overall average scores, we checked whether 1) the similarity manipulation was effective in reducing felt similarity with

close-others (pre-registered hypothesis 5) and whether liking scores were also impacted by the manipulation, and 2) the same manipulation resulted in any significant changes on the predicted scores of MES for close others across the control, similarity, and dissimilarity between-subjects conditions (pre-registered hypothesis 4).

Felt Similarity for Close-Other Across Similarity Conditions. A one-way ANOVA test with the felt similarity score towards close others as the dependent variable and condition assignment as the independent variable revealed that we successfully manipulated felt similarity towards close others, with a significant main effect of condition, F (2, 286) = 33.71, p<.001, partial η^2 =.191. Post-hoc tests using the Bonferroni correction revealed a significant difference on the 1–7-point similarity scale (M_{diff} = -0.83, p<.001, 95% CI [-0.40, -1.27]) between the similarity ratings completed for those in the dissimilarity condition (M=4.32, SD=1.48) and those in the control condition (M=5.15, SD=1.25) such that people in the dissimilarity condition felt significantly less similar to their identified close others compared to those in the control condition. There was also a significant difference (M_{diff} = 0.66, p<.001, 95% CI [0.23, 1.08]) between those in the similarity condition (M=5.81, SD=0.97) and those in the control condition such that people in the similarity condition felt more similar to their identified close others, compared to those in the control condition. The results supported pre-registered hypothesis 5.

Liking for Close-Other Across Similarity Conditions. We repeated the one-way ANOVA test, this time with the liking score towards close others on the 1–7-point liking scale as the dependent variable, and condition assignment as the independent variable. The results revealed that the similarity manipulation had a significant main effect on liking scores, F (2, 286) = 4.08, p= .018, partial η^2 =.028. Post-hoc tests using the Bonferroni correction showed that the small difference (M_{diff} = -0.35, p<.031, 95% CI [-0.68,-0.02]) lied in the comparison between

those in the control condition (M=6.16, SD=1.03) and those in the similarity condition (M=6.51, SD=-0.85) such that people in the similarity condition reported liking their identified close others more compared to those in the control condition. The liking scores for close others between those in the control condition and those in the dissimilarity condition were virtually indistinguishable (M_{diff} = -0.03, p= 1.00, 95% CI [-.36, .31]), with participants in the dissimilarity condition reporting liking their close others the same as in the absence of a manipulation, indicating that the dissimilarity manipulation had no effect on liking scores.

Perceived MES Score for Close-Other Across Similarity Conditions. A one-way ANOVA test with the perceived MES score of close others as the dependent variable and condition assignment as the independent variable revealed no main effect of condition, F (2, 286) = 0.22, p= .802, partial η^2 =.002, indicating that people across the similarity ($MES_{sim=}$ 42.15, SD=12.18), dissimilarity ($MES_{dissim=}$ 42.64, SD=14.36), and control ($MES_{control}$= 41.37, SD=13.47) conditions did not perceive or report their close-others' moral expansiveness any differently as a function of the similarity condition they were assigned to. This was true even when we ran a two-way Condition x Entity type interaction, to test the effect of condition on the type of entity that the MES score was composed of, checking whether condition had a differential effect on the perceived MES of close others for proximal entities versus the perceived MES of close others for distant entities, which it did not, F (2, 286) = .009, p= .992, partial η^2 =.000.[2] The results supported pre-registered hypothesis 4.

[2] We found a small significant correlation between similarity to close others and close others' perceived MES score, r=.157, p=.008. We used Z-scores at high (+1 SD) and low levels (-1 SD) of felt similarity to test whether the gap between one's own expansiveness and that of their close others differs as a function of how similar participants felt to their close others, which was not significant, F (1, 287) = 0.96, p= .329, partial η^2 =.003. We obtained a similar correlation between the liking rating for close others and their perceived MES scores, r=.154, p=.009, thus we repeated the analysis at ± 1SD relative to the mean of the liking score, and we obtained the same results on a non-significant interaction, F (1, 287) = 3.64, p= .057, partial η^2 =.013.

MES Comparison Between Self, Average Other, and Close Other. Given that we obtained no differences in the perceived MES score for close others as a function of condition, we collapsed the three MES scores for close others (MES_{sim}, MES_{dissim}, $MES_{control}$) into one score (MES_{close}) to conduct a repeated-measures test of the three within-subjects perspectives of MES self, MES average other and MES close other (collapsed). A repeated measures ANOVA with a Greenhouse-Geisser correction revealed a significant difference between the three MES perspectives, F (1.99, 572.99) = 22.87, p<.001, partial η^2 =.074. Post hoc tests using the Bonferroni correction revealed a significant difference (M_{diff} = -4.63, p<.001, 95% CI [-6.34, -2.93]) between the scores for the MES of self (M_{self}=43.60, SE=.73) and the scores for the MES of average others (M_{avrg}=38.97, SE=.68) such that people perceive average others as less morally expansive compared to themselves. As with Studies 1a and 1b, this difference (M_{diff} = -4.63) was the equivalent of over one entity being perceived as completely excluded from others' moral circles, compared to that of the self. There was also a significant difference (M_{diff} = -3.07, p<.001, 95% CI [-4.78, -1.36]) between the scores for the MES of average others and the MES scores of close others ($M_{close=}$42.03, SE=.78) such that people perceived their close others to be more morally expansive than the average person.

MES Self, Average, and Close Other Across Proximal and Distant Entities. Because our higher-level comparisons are crossing the between-subjects factor for the close-other perspective (Condition: similarity, dissimilarity, control), and the within-subjects factors of MES perspective (self, average other, close other) and Entity type (proximal, distant), we opted to conduct these analyses using a linear mixed-effects model to account for the fact that the MES perspective for self and average other did not receive a similarity manipulation (i.e. the design was not fully crossed between all factors). As such, a linear mixed-effects model with one

between-subjects factor (condition) and two within-subjects factors (MES perspective, Entity type) cross-tested all possible interactions between the factors, treating the cells that did not receive the condition treatment as redundant, and thus set to 0. The results that follow report Type III tests of fixed effects.

Before we examined interactions we confirmed that, as with Studies 1a and 1b, there was a main effect of entity type, F (1, 1626.23) = 2556.24, $p<.001$, such that across all conditions and MES perspectives, out of the maximum 3 points an entity could receive if placed in the innermost part of the MES, participants reported of themselves and predicted for others more moral inclusivity (M_{diff} = 1.13, $p<.001$, 95% CI [1.09, 1.17]) of proximal entities (average of family and friends, ingroup, and revered entities) ($M_{proximal}$=2.18, SE=.01), than of distant entities (average of outgroups, stigmatized entities, animals of high sentience, animals of low sentience, plants, the environment, and villains) ($M_{distant}$=1.05, SE=.02). (Note that for Studies 2-4 that included the close-other perspective, the first group of entities will be referred to as "proximal" as opposed to "close" to avoid confusion).

Note that, as mentioned earlier, because condition only targeted the close-other MES perspective, the cells that did not receive the condition treatment (MES self, MES average other) were treated as redundant by the linear mixed-effects model and thus set to 0. As such, we were able to test for the Condition (similarity, dissimilarity, control) x MES perspective (self, average other, close other) x Entity type (proximal, distant) interaction, which not significant, F (4, 1052.34) = 0.17, p=.953, and there was no Condition (similarity, dissimilarity, control) x Entity type (proximal, distant) significant interaction, F (2, 1626.23) = 0.39, p=.675.

After confirming that condition had no effect on MES scores, we moved on to examine the rest of the interactions of interest, namely whether there was a significant MES perspective

(self, average other, close other) x Entity type (proximal, distant) interaction as with Studies 1a and 1b (pre-registered hypotheses 1-3). Indeed, we replicated the significant interaction in the same fashion, F (2, 1052.34) = 20.00, p<.001, such that proximal entities were perceived to be equally morally included in all perspectives - self's ($M_{self\_proximal}$=2.14, SE=.03, 95% CI [2.09, 2.19]), close others' ($M_{close\_proximal}$=2.17, SE=.03, 95% CI [2.12, 2.22]), average others' ($M_{avrg\_proximal}$=2.22, SE=.02, 95% CI [2.17, 2.26]), in contrast to distant entities, which were perceived to be included less in the moral landscapes of average others ($M_{avrg\_distant}$=0.91, SE=.03, 95% CI [0.85, 0.96]), compared to the moral landscapes of the self ($M_{self\_distant}$=1.16, SE=.03, 95% CI [1.10, 1.22]), and compared to the moral landscapes of close others ($M_{close\_distant}$=1.07, SE=.03, 95% CI [1.01, 1.39]). As predicted, distant entities in the moral landscape of the self were not placed any differently compared to where they were perceived to belong in the moral landscape of close others. The results supported pre-registered hypotheses 1-3.

MES Associations with Cynicism Measures. All three cynicism scales displayed excellent reliability, and the three correlated at levels of r = .59 (p<.001) and above with each other, indicating that, although they were not identical operationalizations of cynicism, all three were tapping into the same construct (see Table 3 for Cronbach's alphas, descriptives, and correlations). The sample appeared moderately cynical based on agreements to cynical statements rated on a 1-7–scale (1 - not at all, 7 - completely true), and levels of cynicism across all three measures did not differ as a function of the similarity condition participants were assigned to (all ps>.136). We examined whether cynicism correlated negatively with all three MES perspectives (self, close other, average other) (pre-registered hypothesis 6). All cynicism measures showed small significant negative correlations with the MES self score, and two (out

of three) cynicism measures correlated negatively with the perceived MES scores of close others (when collapsed across the similarity, dissimilarity, and control conditions). However, unlike our predictions, there was no relationship between the cynicism measures and the perceived MES score for average others. (For a lengthier discussion and analysis of the cynicism constructs see aggregate section following the interim discussion of Study 3). The results partially supported pre-registered hypothesis 6.

Table 3: *Reliabilities, Means, & Pearson's correlations between variables in Study 2*

Variable	α	*M* (SD)	1	2	3	4	5	6	7	8	9	10
1. Cynicism (Cook & Medley, 1954)	.890	*M*=4.18 (1.24)	--									
2. Cynicism (Wrightsman,1964)	.907	*M* =4.46 (1.38)	.82***	--								
3. Comp vs. Coop beliefs (Clifton et al., 2019)	.812	*M* =3.76 (1.38)	.68***	.59***	--							
4. MES self	.907	*M* =43.60 (12.34)	-.15*	-.12*	-.20**	--						
5. MES close other	.918	*M* =42.03 (13.31)	-.13*	-.09	-.16**	.61**	--					
6. MES average other	.904	*M* =38.97 (11.54)	-.03	-.04	-.07	.49***	.54***	--				
7. Similar to close other	--	*M* =5.11 (1.38)	-.05	-.009	-.05	.12*	.16**	.12*	--			
8. Similar to average other	--	*M*=3.90 (1.56)	-.00	-.06	-.07	.04	.02	.26***	.24***	--		
9. Like close other	--	*M*=6.29 (0.96)	-.11	-.05	-.11	.06	.15**	.03	.54***	.16**	--	
10. Like average other	--	*M*=4.27 (1.17)	.01	-.04	-.03	.18**	.15*	.27***	.20***	.73***	.19**	--

Note 1. *All close-other relevant variables (MES close other, Similarity, Liking) are collapsed across the three between-subjects similarity conditions.*

Note 2. ** p<.05, ** p<.01, *** p<.001*

Interim Discussion for Study 2

In Study 2, participants exhibited the moral inclusivity discrepancy that we observed in Studies 1a and 1b between their own moral expansiveness scores, and that they predicted for average others for distant entities (pre-registered hypothesis 1). As we expected, this was not the case for the predicted moral expansiveness scores for those participants identified as close to them, whose predicted scores matched the participants themselves for both close and distant entities in the moral circle (pre-registered hypothesis 2). When comparing the perceived moral landscapes of those that participants identified as being close to them compared to those of average others, close others were perceived as more morally expansive than the average person when it came to distant entities (pre-registered hypothesis 3).

Critically, and as we predicted, despite the fact that the more similar participants felt to their close others, the higher they reported their close others' moral expansiveness to be as evidenced by a small significant correlation (r=.16, p<.010), they do not report their close others' MES scores (for distant or for proximal entities) any differently when assigned to a dissimilarity condition (pre-registered hypothesis 4) that successfully managed to decrease felt similarity with close others (pre-registered hypothesis 5). However, we also (unsurprisingly) obtained a ceiling effect when asking participants how much they like their close others, thus we turned our attention to a more targeted, liking-relevant manipulation, to rule out the possibility that participants' reports of their close others' moral expansiveness is influenced by how much they like them.

Chapter 4: Study 3 – MES of Liked vs. Disliked Close Others

Study 3

The aim of Study 3 was to test whether, aside from similarity, liking alone cannot account for the elimination of the moral expansiveness discrepancy in the case of close-others. In other words, we wanted to rule out the possibility that people make attributions of moral expansiveness based on how much they like their close others. We expected that we would successfully reduce liking for close others in the disliking condition, but that moral expansiveness scores for close others would remain unchanged across the disliking, liking, and control conditions, because access to information about close other's moral landscapes should remain constant. As with Study 2, we additionally expected to observe significant negative correlations between three separate cynicism measures and all MES scores. Our hypotheses for Study 3 were pre-registered on the Open Science Framework (access at https://osf.io/yut6s) and consisted of the following:

1) For distant entities (but not proximal entities) in the moral circle, we expect that the moral expansiveness score (MES) for the self will be higher than the one participants will report for average others' MES.
2) For both distant and proximal entities in the moral circle, we expect that the moral expansiveness score (MES) for the self will not be significantly different from the one participants will report for close other's MES under any of the close-other manipulations (liking, disliking, control).
3) For distant entities (but not proximal entities) in the moral circle, we expect that the moral expansiveness score (MES) for close others will be higher than the one participants will report for average others' MES under all of the close-other

manipulations (liking, disliking, control).

4) For both distant entities and proximal entities in the moral circle, we expect that the moral expansiveness score (MES) for close others will not be significantly different across the three between-subjects manipulations (liking, disliking, control).
5) We expect that participants assigned to reflect on reasons for disliking or feeling annoyed with their close other will report lower felt liking for their close other, compared to participants assigned to the liking-close other and control-close other conditions.
6) We expect that three measures of cynicism (subscale from Cook & Medley, 1954; Chowdhury & Fernando, 2014; competitive primal world beliefs, Clifton et al, 2019) will correlate negatively with all MES scores (self, close others, average others).

Participants

As with Study 2 that used a between-subjects component, and the expectation for some moderate attrition, Study 3 aimed to recruit at least N=450 participants. We recruited our participants from an undergraduate student sample in a large public university in the Northeast of the United States. Following exclusions of participants with incomplete data on the MES, nonsensical responses to the open-ended prompts of close-other identification and liking/disliking, or who failed the attention check, we obtained a final sample of N=421. A sensitivity analysis for a repeated measures ANOVA with a within-between interaction using G* Power software (Faul, Erdfelder, Lang, & Buchner, 2007) revealed that with a sample size of N=421, Study 3 can detect an effect as small as f = .086. Participants were assigned to between-subjects conditions using the “evenly-present elements” feature in Qualtrics software.

The sample was fairly evenly distributed across the three between-subject conditions of close-other liking (N=135), close-other disliking (N=144), and control (N=142), and the vast majority of attrition occurred in the introductory stages of the survey, before participants saw any content that was differentiated by condition assignment.

In Study 3, the sample's mean age was M_{age} = 18.88 (SD=1.87) years, 77.7% of participants identified their gender as female, 20.9% as male, 1.2% as other, and 0.2% did not provide gender information. For their ethnicity, 65.6% of participants identified as White or Caucasian, 8.1% as Asian or Pacific Islander, 6.9% as Hispanic or Latino, 5.2% as Black or African American, 12.8% reported a combination of ethnicity backgrounds, and 1.2% chose "other" as their ethnicity category. The sample was near the mid-point on a 7-point conservativism scale (M=3.66, SD=1.51) ranging from Extremely Liberal (1) to Extremely Conservative (7) and was slightly to somewhat religious (M=2.42, SD=1.10) on a 5-point religiosity scale, ranging from Not at All Religious (1) to Extremely Religious (5).

Measures

All measures were identical to Study 2. Participants filled out the Moral Expansiveness Scale (Crimston et al., 2016) alongside liking and similarity ratings for average and close others on a scale from 1(Not at all) to 7 (Extremely), and reported the relationship they had with their close others (i.e. friend, family member, partner, colleague, or other category). Participants were also again asked to complete three separate cynicism measures (cynicism subscale from Cook & Medley's hostility scale, 1954; cynicism scale adapted from Wrightsman 1964; cooperative vs. competitive primal world beliefs from Clifton et al., 2019), as well as other exploratory measures and demographics before they were debriefed. (For additional exploratory measures and analyses included in Study 3 refer to Supplemental Analyses in Appendix E).

Procedure

After participants consented to take part in the study, they were given instructions and information pertaining to MES measure in the same fashion as with Studies 1a, 1b and 2. As with Study 2, after completing the self and average other perspectives, they were then assigned to three between-subjects conditions. In the liking condition, after identifying a close-other, participants were asked to reflect and write 2-3 sentences on what makes them like their identified close-other. In the disliking condition, participants were asked to reflect and write 2-3 sentences on what makes them sometimes dislike or feel annoyed with their identified close-other. Finally, participants assigned to the control condition received no liking manipulation; instead, they solely identified their close other and then proceeded to fill out the MES from their close other perspective (see Appendix B for condition prompts). After completing all three MES perspectives, participants again completed the three cynicism measures, filled out other exploratory measures and demographic questions, and were debriefed.

Results

In contrast to Study 2 where most participants identified a friend, in Study 3 over half of participants, at 50.4% (N=212) identified a family member as their close other, followed by 34.7% who identified a friend (N=146), 13.5% who identified a partner (N= 57), 0.2% a colleague (N=1) and 1.2% of participants chose someone belonging to an "other" category (N=5).

We conducted the same repeated measures ANOVA tests on the 7-point liking and similarity scales, which replicated that, on average, the identified close others were liked (M=6.69, SD=0.67, 95% CI [6.63, 6.75] significantly more compared to average others

(M=4.29, SD=0.99, 95% CI [4.19, 4.38], (M_{diff} = 2.40, p<.001, 95% CI [2.29, 2.51]), F (1, 420) =1884.28, p<.001, partial η^2 =.818, and were also perceived as more similar to the participant (M=5.40, SD=1.26) [95% CI [5.28, 5.52], compared to average others (M=3.91, SD=1.24) [95% CI [3.80, 4.03], (M_{diff} = 1.50, p<.001, 95% CI [1.34, 1.64]), F (1, 420) = 401.35, p<.001, partial η^2 =.489.

As with Study 2, before comparing the self, average other, and close other MES overall average scores, we checked whether 1) the liking manipulation was effective in reducing the liking of close-others (pre-registered hypothesis 5), and 2) the same manipulation resulted in any significant changes on the predicted scores of MES for close others across the control, liking, and disliking between-subjects conditions (pre-registered hypothesis 4).

Liking for Close-Other Across Liking Conditions. Although the vast majority of participants who got assigned to the disliking between-subjects condition were able to come up with reasons they disliked or sometimes got annoyed with their identified close other, there was a handful of participants (N=6) who expressed complete inability to do so (example response "*Honestly there is nothing that she does that makes me feel dislike or feel annoyed*."). There was also another set of participants (N=6) who communicated that they found it difficult to do so but came up with a negative attribute (example response: "*It is honestly quite hard to find reasons to dislike her. I guess the main thing would be disagreements based on politics, but even then they are few and far between.* […]"). We included the latter set of participants who were still able to come up with a reason in the between-subjects test for the main effect of condition on liking scores, but exploratorily excluded the first group who were completely unable to list a reason. We tested the main effect of condition both with and without this group in the dataset.

Using the entire dataset without exclusions, a one-way ANOVA test with the liking score

for close others as the dependent variable and condition assignment as the independent variable revealed that we did not manage to successfully manipulate liking towards close others. There was no significant main effect of condition, F (2, 418) = 2.95, p= .053, partial η^2 =.014, with means across the disliking (M=6.58, SD=0.82), liking (M=6.77, SD=0.57), and control (M=6.72, SD=0.56) conditions all displaying a near ceiling effect on the 1-7-point liking scale, which was consistent with the liking means we obtained in Study 2.

When repeating the analysis but excluding those in the disliking condition who were unable to list reasons for why they disliked their close others (N=6), the one-way ANOVA test showed a small, significant main effect of condition, F (2, 412) = 3.48, p= .032, partial η^2 =.017. As expected, the difference was driven by the now slightly lower mean of the disliking condition (M=6.57, SD=0.84), which was significantly different only from the mean liking score of those in the liking condition, but not from those in the control condition (means across these two remain the same as with the analyses above, given no exclusions in these two conditions). We decided to use a conservative approach and retain all participants in the disliking condition in the remaining analyses, given that all went through the process of attempting to identify reasons for why they disliked their close other. We therefore considered the inability to do so as part of the heterogeneous response to the manipulation. Thus, the results failed to support pre-registered hypothesis 5.

Felt Similarity for Close-Other Across Liking Conditions. A one-way ANOVA test with the felt similarity score towards close others as the dependent variable and condition assignment as the independent variable showed no significant main effect of condition, F (2, 428) = 2.42, p= .090, partial η^2 =.011, suggesting that manipulating liking for close others has no effect on perceived similarity with close others.

Perceived MES Score for Close-Other Across Liking Conditions. We repeated a one-way ANOVA test with the perceived MES score of close others as the dependent variable and condition assignment as the independent variable. As with Study 2, the test revealed no main effect of condition, F (2, 418) = 0.02, p= .980, partial η^2 =.000, indicating that people across the liking (MES_{like}= 41.26, SD=13.07), disliking ($MES_{dislike}$= 40.98, SD=24.49), and control ($MES_{control}$= 41.01, SD=12.24) conditions did not differentially report their close-others' moral expansiveness based on the liking condition they were assigned to. Again, condition had no effect even when we tested the interaction between the effect of condition on the type of entity that the MES score was composed of, seeing no differential impact on the perceived moral inclusivity of proximal versus distant entities for close others, F (2, 418) = 1.01, p= .367, partial η^2 =.005.[3] The results supported pre-registered hypothesis 4.

MES Comparison Between Self, Average Other, and Close Other. Although our manipulation check did not confirm that we successfully manipulated liking for close others, the lack of differences in the MES scores as a function of the liking condition, and the lack of correlation between the liking rating and the reported MES scores for close others across conditions, allowed us to proceed with conducting a repeated-measures test of the three within-subjects perspectives of MES self, MES average other, and again with the collapsed MES close

[3] We found no correlation between liking for close others, and close others' perceived MES score, r=.016, p=.743, nor a significant moderation of the MES self and MES close other scores at high (+1 SD) and low levels (-1 SD) of the liking rating, F (1, 419) = 0.37, p= .542, partial η^2 =.001. However, we obtained a similar correlation as with Study 2 on how similar participants felt to their close others, and the perceived moral expansiveness of their close others, r=.157, p=.001. This time, the moderation analysis by similarity was significant, F (1, 419) = 11.10, p<.001, partial η^2 =.026. At high levels of felt similarity with close others, participants had no difference (M_{diff} = 0.16, p=.822, 95% CI, [-1.26, 1.58]) in the MES scores they reported of themselves (M=43.22, SE=.78), and what they predicted for their close others (M=43.06, SE=.86) but at low levels of felt similarity, participants had a higher gap (M_{diff} = 3.57, p<.001, 95% CI, [2.15, 4.99]), such that their own MES score (M=42.67, SE=.78) was significantly higher than what they predicted for close others (M=39.11, SE=.86).

other scores (MES_{like}, $MES_{dislike}$, $MES_{control}$) forming one score (MES_{close}).

A repeated measures ANOVA with a Greenhouse-Geisser correction revealed a significant difference between the three MES perspectives, F (1.97, 828.03) = 23.23, p<.001, partial η^2 =.052. Post hoc tests using the Bonferroni correction revealed a significant difference (M_{diff} = -3.70, p<.001, 95% CI [-4.99, -2.41]) between the scores for the MES of self (M_{self}=42.95, SE=.55) and the scores for the MES of average others (M_{avrg}=39.25, SE=.55) such that people perceive average others as less morally expansive compared to themselves. Similar to Studies 1a, 1b, and 2, this difference (M_{diff} = -3.70) was the equivalent of over one entity being perceived as completely excluded from others' moral circles, compared to that of the self.

As with Study 2, there was also a smaller but significant difference (M_{diff} = -1.83, p= .004, 95% CI [-3.21, -.46]) between the scores for the MES of average others and the MES scores of close others ($M_{close=}$41.08, SE=.61) such that people perceived their close others to be more morally expansive than the average person. We also, however, found an additional difference of about equal size between the MES scores of the self and that of close others, such that participants perceived themselves as more morally expansive than their close others.

MES Self, Average, and Close Other Across Proximal and Distant Entities. As with Study 2, in order to account for the fact that the MES perspective for self and average other did not receive a liking manipulation, we opted to conduct the higher-order interaction analyses to test our main predictions (pre-registered hypotheses 1-3) using a linear mixed-effects model, that treated the cells that did not receive the condition treatment as redundant, with one between-subjects factor (condition) and two within-subjects factors (MES perspective, Entity type). The results that follow report Type III tests of fixed effects.

Before we examined interactions, we confirmed that, as with Studies 1a, 1b, and 2, there

was a main effect of entity type, F (1, 2357.28) = 5429.68, p<.001, such that across all conditions, out of the maximum 3 points an entity could receive if placed in the innermost part of the MES, participants reported of themselves and predicted for others more moral inclusivity (M_{diff} = 1.23, p<.001, 95% CI [1.19, 1.26]) of proximal entities (average of family and friends, ingroup, and revered entities) ($M_{proximal}$=2.23, SE=.01), than of distant entities (average of outgroups, stigmatized entities, animals of high sentience, animals of low sentience, plants, the environment, and villains) ($M_{distant}$=1.00, SE=.01).

As expected, similar to Study 2's interactions with the similarity condition and following the non-treated cells (MES self, MES average other) being set to 0, the highest-order Condition (liking, disliking, control) x MES perspective (self, average other, close other) x Entity type (proximal, distant) interaction was not significant, F (4, 1515.76) = 0.09, p=.986 , nor was the Condition (liking, disliking, control) x Entity type (proximal, distant) interaction, F (2, 2357.28) = 0.99, p=.373.

The MES perspective (self, average other, close other) x Entity type (proximal, distant) interaction was again significant, F (2, 1515.76) = 17.90, p<.001, such that proximal entities were perceived to be equally morally included in all perspectives – self's ($M_{self\_proximal}$=2.21, SE=.02, 95% CI [2.18, 2.25]), close others' ($M_{close\_proximal}$=2.22, SE=.02, 95% CI [2.18, 2.2]), average others' ($M_{avrg\_proximal}$=2.25, SE=.02, 95% CI [2.22, 2.29]) but distant entities were perceived to be less included in average others' moral landscapes ($M_{avrg\_distant}$=.90, SE=.02, 95% CI [0.86, 0.95]), compared to the moral landscape of both the self ($M_{self\_distant}$=1.10, SE=.02, 95% CI [1.06, 1.14]) and that of close others ($M_{close\_distant}$=1.01, SE=.02, 95% CI [0.96, 1.05]). These results supported pre-registered hypotheses 1 and 3. However, this time, there was also a significant difference between the MES score for distant entities for self and the predicted MES

score for distant entities for close others such that, close others were also perceived to be less expansive compared to the self for distant entities, although to a smaller extent compared to average others. This result failed to support pre-registered hypothesis 2.

MES Associations with Cynicism Measures. The three cynicism scales displayed lower reliability compared to Study 2, but still ranging from good to excellent (see Table 4 for Cronbach's alphas, descriptives, and correlations). The three scales correlated at levels between r = .44 and r = .68, which was lower than Study 2's correlations between the cynicism measures, but still indicated substantial overlap between the three scales. The sample appeared moderately cynical on a 1-7 scale (1 – not at all, 7- completely true). As with Study, we again examined whether cynicism correlated negatively with the three MES perspectives (self, close other, average other) (pre-registered hypothesis 6). Contrary to our expectations, the cynicism measures did not show the expected negative relationships with the MES self, close other, and average other scores, except for the adapted Wrightsman (1964) measure correlating negatively with the MES close other score at r=-.14, p=.005. These results failed to support pre-registered hypothesis 6. (For a lengthier discussion and analysis of the cynicism constructs see aggregate section following the interim discussion of Study 3).

Table 4: *Reliabilities, Means, & Pearson's correlations between variables in Study 3*

Variable	α	M (SD)	1	2	3	4	5	6	7	8	9	10
1. Cynicism (Cook & Medley, 1954)	.820	*M*=3.89 (1.02)	--									
2. Cynicism (Wrightsman,1964)	.811	*M* =4.81 (1.11)	.68***	--								
3. Comp vs. Coop beliefs (Clifton et al., 2019)	.697	*M* =3.52 (1.12)	.55***	.44***	--							
4. MES self	.907	*M* =42.95 (11.34)	-.02	-.06	-.04	--						
5. MES close other	.917	*M* =41.08 (12.57)	-.06	-.14**	-.05	.61***	--					
6. MES average other	.910	*M* =39.25 (11.28)	.06	.01	.03	.52***	.52***	--				
7. Similar to close other	--	*M* =5.40 (1.26)	-.14**	-.05	-.10*	.02	.16**	.06	--			
8. Similar to average other	--	*M*=3.91 (1.24)	-.16**	-.14**	-.12*	-.15**	-.01	.11*	.26***	--		
9. Like close other	--	*M*=6.69 (0.67)	-.14**	-.02	-.08	-.01	.02	.03	.44***	.08	--	
10. Like average other	--	*M*=4.29 (0.99)	-.25***	-.18***	-.22***	-.03	-.00	.08	.19***	.53***	.11*	--

Note 1. *All close-other relevant variables (MES close other, Similarity, Liking) are collapsed across the three between-subjects similarity conditions.*

Note 2. ** p<.05, ** p<.01, *** p<.001*

Interim Discussion for Study 3

The results of Study 3 replicated the finding that the moral expansiveness of the self for distant entities is greater than that of average others, and that close others are perceived as more morally expansive than average others. However, Study 3 additionally found that participants' close others were also seen as less morally expansive compared to the self on distant entities, which contradicts the lack of differences between self and close other that we obtained in Study 2, although the difference was smaller than the self-average other gap. Importantly, we again obtained ceiling effects on the liking responses people gave for their close others across all the liking, disliking, and control conditions (above a 6 on a 1-7 Likert scale) and, although the similarity manipulation in Study 2 seemed to increase liking scores for close others, the direct manipulation of liking failed to impact how much close others are liked.

There are at least two reasons as to why we were unable to successfully manipulate liking towards close others in Study 3. The first and most straightforward reason is that given the centrality of close others in someone's life, it would and should be hard to be able to change someone's (largely) positive feelings for a close other person in their life. Thus, it is perhaps not surprising that we were able to only slightly shift liking for close others downwards, only when we exploratorily removed those who found it hard and refused to cite reasons for which they disliked or sometimes felt annoyed with their close others.

The second reason as to why we might have been unsuccessful in reducing liking for close others has to do with the nature of Study 3's sample, which used university college students, with an average age of about 18 years (M_{age} = 18.88), and a very small age range, (SD=1.87) compared to Study 2 which used an online sample with an average age of about 40 (M_{age} = 40.65) and a much larger age range (SD=11.25). Indeed, the student sample, likely due

to their young age, were nearly three times as likely to identify a family member (often parent), as their close other (identified a parent in 50.4% of the cases), compared to online participants in Study 2 (identified a parent only in 18.3% of the cases). Our manipulation of encouraging people to reflect on reasons as to why they dislike or sometimes feel annoyed with their close others might have had less of a chance to be effective in the case of family members, particularly parents, compared to other close others who might be more relevant later on in one's life, such as their partners, friends, or co-workers, although this remains speculative.

Beyond liking, Study 3 replicated Study 2's finding that the greater the similarity one perceives in their close other, the higher they perceived their close others' moral expansiveness score to be. Importantly, unlike Study 2 where similarity was directly manipulated, in Study 3 where similarity was not the forefront of the manipulation, we found that the less similar someone perceived their close other to be to them, the larger the gap between their moral expansiveness scores and the scores they perceived their close others to have. This suggests that even though our similarity manipulation in Study 2 might have been effective in shifting similarity scores, it still might not have captured naturally-occurring differences in how similar people feel to their close ones, that might lead them to report their close others' moral expansiveness as more or less similar to themselves.

More interestingly, because over half of the participants in Study 3 identified a family member (often a parent) as their close other, this sample was primarily reporting their parents' moral expansiveness. This, in conjunction with the fact that we replicated the discrepancy we typically find between the MES of the self and the perceived MES of average others in distant entities in the case of close others, presents an additional interpretation of the similarity moderation results. It is possible that the student sample would generally think of their parents as

both more different and less expansive due to them being older and potentially not as liberal - e.g. some students explicitly said that their close others had "outdated beliefs" or “made the "occasional political off-comment" when asked to find reasons why they disliked them in the open-ended prompt. In other words, this sample, due to its young and limited age range and their increased likelihood to report on a parent figure as their close-other, might pose a unique case of when similarity is particularly relevant as a moderator. This might be due to exacerbated differences between themselves and their close others that due to perceived generational ideological gaps, are uniquely associated with the contents of the distant entity categories of the MES.

Interim Discussion Across Studies 2 & 3

Studies 2 and 3 tested whether the MES gap between self and other disappears in the case of close others even when participants are given reasons to dislike or feel dissimilar to them. Although we found no differences between the moral expansiveness participants reported of themselves and their close others in Study 2, we did find that participants reported themselves as more expansive than their close others in Study 3, which might reflect naturally-occurring differences between participants and their close others.

Precisely because MES self and MES close other are comparably expansive relative to average others, increasing or reducing cynical beliefs should primarily impact average others’ perceived moral expansiveness. We followed a fairly standard process whereby we first attempted to establish an association between moral expansiveness and cynicism in Studies 2 and 3, which we report on in the sections below, and subsequently manipulated cynicism in an attempt to shift responses on moral expansiveness, which Study 4 aimed to do.

MES Self-Average Other Gap at Different Levels of Cynicism

While we used other exploratory measures included in Studies 1a and 1b as our closest approximation to cynicism (see Supplementary Analyses in Appendix E), Studies 2 & 3 included additional cynicism measures to directly tap into the construct (see negative and non-significant correlations between MES perspectives and cynicism measures in Tables 3 & 4). Specifically, as mentioned earlier, we included measures by Cook and Medley (1954), Wrigthsman (1964), and Clifton and colleagues (2019). We used Z scores to differentiate between those low (-1SD) and high (+1SD) on the three scales and observed how the moral expansiveness discrepancy we found across all studies between self and average other behaves at different cynicism levels. The results painted a somewhat consistent picture, but contrary to what we initially expected, which was that the discrepancy would be higher at higher levels of cynicism (accounted for by a substantial reduction in the perceived moral expansiveness of others).

The cynicism scale by Wrightman (1964) as adapted by Chowdhury and Fernando (2014), had no significant interaction between MES perspective (self vs. average other) and cynicism (high vs. low) in Study 2, F (1, 287) = 1.90, p=.169, partial η^2 =.007, or in Study 3, F (1, 418) = 2.16, p=.143, partial η^2 =.005, and the perceived MES scores for average others do not show a substantial decline at high levels of this measure in either of the two studies (see Figure 6 for the plotted MES perspective x Wrightman cynicism scale moderation across Studies 2 and 3).

Figure 6: *MES perspective x Wrigthsman cynicism interaction test in Studies 2 & 3*

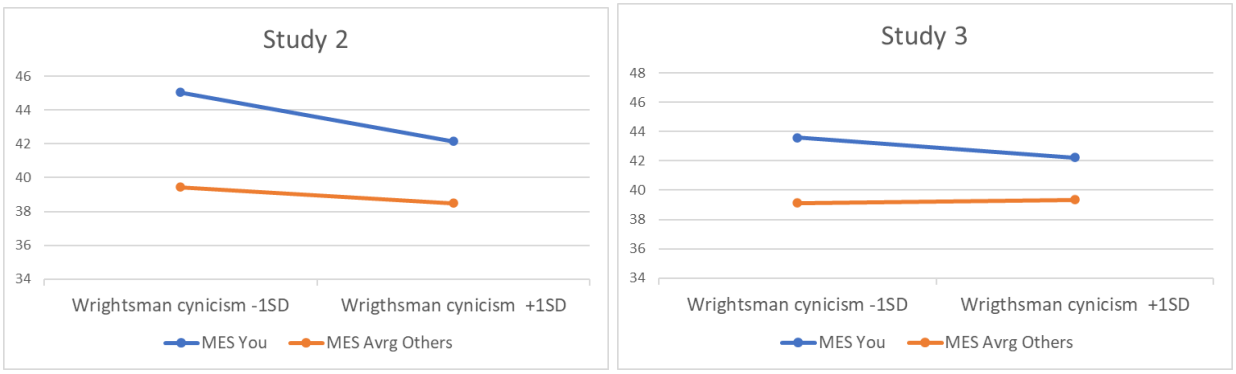

Note. *Plotted estimated marginal means at ± 1SD (out of the maximum possible MES total of 90 points - Y axis).*

When using the Cook and Medley (1954) cynicism scale as moderator, we see a significant interaction between MES perspective and cynicism in Study 2, F (1, 287) = 4.39, p=.037, partial η^2 =.015, but no moderation in Study 3, F (1, 418) = 2.53, p=.112, partial η^2 =.006 (see Figure 7 for the plotted MES perspective x Cook & Medley cynicism scale moderation across Studies 2 and 3). Again, the pattern of results seemed to be such that the self-average other gap in expansiveness was greater at low levels, as opposed to high levels of cynicism, and significantly so in Study 2.

Figure 7: *MES perspective x Cook & Medley cynicism interaction test in Studies 2 & 3*

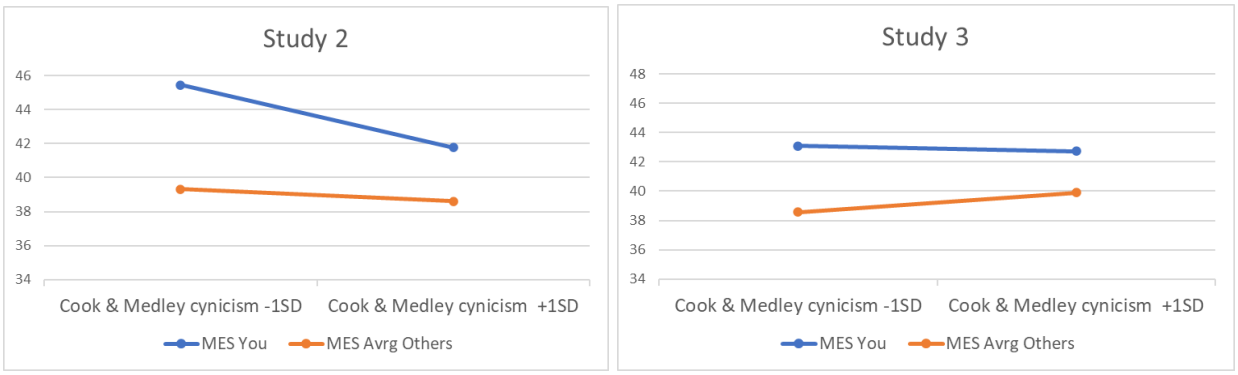

Note. *Plotted estimated marginal means at ± 1SD (out of the maximum possible MES total of 90 points - Y axis).*

Finally, when using the Clifton and colleagues (2019) cooperative vs. competitive belief cynicism scale as moderator, we again see a significant interaction between MES perspective and cynicism in Study 2, F (1, 287) = 4.77, p=.030, partial η^2 =.016, but no moderation in Study 3, F (1, 418) = 2.48, p=.116, partial η^2 =.006 (see Figure 8 for the plotted MES perspective x Clifton cynicism scale moderation across Studies 2 and 3).

Figure 8: *MES perspective x Clifton cynicism interaction test in Studies 2 & 3*

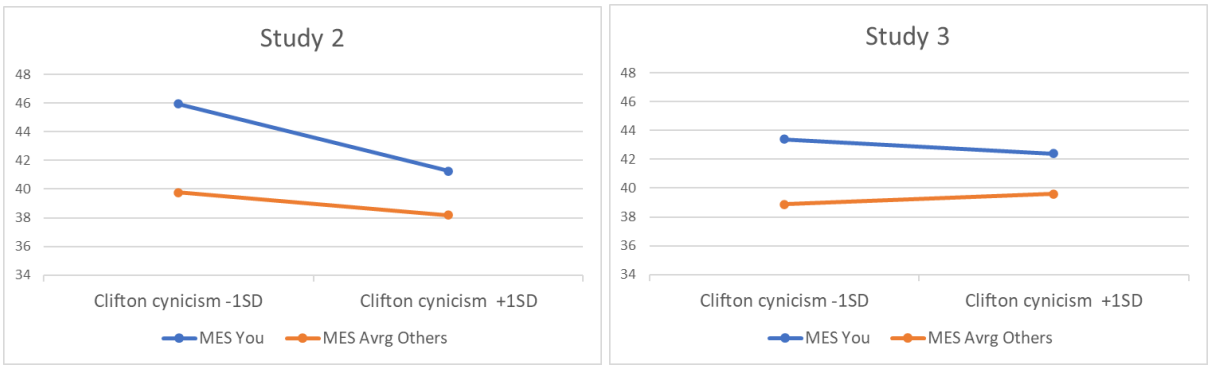

Note. *Plotted estimated marginal means at ± 1SD (out of the maximum possible MES total of 90 points - Y axis).*

In addition to the greater gap that appears between self and others when people are least cynical (which is the opposite of what we expected), what becomes evident from Figures 5-7, is that one's own moral expansiveness score presents more movement across cynicism levels, whereas the perception of average others' MES seems to be consistently stable. Overall, although Study 2 seemed to come closer to the pattern of results we were expecting compared to Study 3, whereby both one's own moral expansiveness and what they estimated for average others lowered at high levels of cynicism (see Study 2 graphs across all 3 cynicism moderators), the movement within the prediction of average others' expansiveness scores continued to be minimal. This indicates that people's own expansiveness scores were more related to how cynical they were, compared to what they reported other's expansiveness to be, which replicated results we found with other individual-trait measures that were relevant to prosociality and the MES (see Supplemental Analyses in Appendix E for fear of compassion

results).

While the reason for this could be psychometric in nature given one's own scores on any given measure (in this case, cynicism) would correlate more strongly with self-reports compared to estimations for others, we still expect that people high on cynicism should have more constricted moral circles both for the self *and* in their perceptions of others. This expectation is shaped by research showing that people who believe in true altruism are also more prosocial themselves (Carlson & Zaki, 2022), as well as research on prosocial orientation showing that people who are more prosocial tend to have a more balanced view of others' prosociality, than those with non-prosocial orientations (Van Lange, 1992; Critcher & Dunning, 2011). Given our cynicism moderation results, and in juxtaposition with this literature, we decided to employ our direct cynicism manipulation in Study 4 using Clifton and colleagues' (2019) competitive vs. cooperative scale that showed both the strongest correlations with overall MES scores (see Tables 2 and 3), and the greatest decline at high levels of the measure, in MES average other scores (see Figure 8). Thus, we persisted in targeting a greater decrease of the perceived MES average other score, since it has so far remained relatively resistant to correlational movement across higher levels of cynicism.

Chapter 5: Study 4 - Manipulating Cynicism

Study 4

Prior research using the construct of moral expansiveness, and the MES measure more specifically (Crimston et al., 2016) has negatively linked constructs related to cynicism with moral expansion. For example, across 36 countries, disagreement with the statement "*I completely trust most other people*" that is reflecting generalized trust is associated with smaller moral circles (Kirkland et al., 2022), as is agreement with statements like "*Being too compassionate makes people soft and easy to take advantage of*" which reflects fear of compassion (Crimston, Blessing, Gilbert, & Kirby, 2022) that we measured and replicated in all earlier studies as well (see Appendix E for Supplemental Analyses). This prior work has thus made connections between cynical assumptions and mindsets and the moral circle, providing the ground on which these assumptions can be flipped or manipulated, which Study 4 attempted to do.

As discussed earlier, cynical beliefs as an antecedent for prosociality-relevant outcomes (in this case, moral expansiveness) might be particularly relevant when people do not have enough information about others (Pratto & Bargh, 1991; Rubinstein, Jussim, & Stevens, 2018), and thus we expect cynicism to be most evidently at work in the case of an unknown average other. Because we would expect that people will be more familiar with a close other's moral circle than with an average person's, we expected that in Study 4, manipulating a cynical outlook prior to filling out the MES, will have the strongest impact on the perceived moral circles of unknown others, but less of an impact on the circles people report for their close others (or their own).

Study 4 aimed to directly manipulate cynical assumptions, to test whether the moral

landscape of general others, but not primarily that of close others for which information already exists, is prone to updating. In Study 4 we expected that participants in the anti-cynicism condition will no longer exhibit a gap between their own expansiveness scores and that of general others in the distant categories of the moral circle (stigmatized, villains, animals of high-sentience, animals of low-sentience, plants, and the environment), but that this gap will persist both in the cynicism and the control conditions. Our hypotheses for Study 4 were pre-registered on the Open Science Framework (access at https://osf.io/8fu5v) and consisted of the following:

1) We expect a significant main effect of condition on the main cynicism measure (competitive world beliefs, Clifton et al, 2019), such that participants assigned to the anti-cynicism condition will report the lowest belief in a competitive world, compared to participants in the control and pro-cynicism conditions.
2) We expect a significant Condition (pro-cynicism vs. anti-cynicism vs. control) x MES perspective (self vs. average other) x Entity type (proximal vs. distant) interaction, such that the moral expansiveness score (MES) for the self will be higher than the one participants will report for average others' MES in the control, and pro-cynicism conditions, but not in the anti-cynicism condition.
3) We expect a null Condition (pro-cynicism vs. anti-cynicism vs. control) x MES perspective (self vs. close other) x Entity type (proximal vs. distant) interaction, such that for both distant and proximal entities in the moral circle, we expect that the moral expansiveness score (MES) for the self will not be significantly different from the one participants will report for close other's MES under any of the conditions (control, pro-cynicism, anti-cynicism).

Participants

Based on the sampling approach of prior studies for a within-between subjects design, Study 4 aimed to recruit at least N=400 participants. We recruited our participants on the online recruitment platform, Prolific, and restricted our search to those residing within the United States. Following exclusions of participants with incomplete data on the MES, nonsensical responses to the open-ended prompts of close-other identification, or who failed the attention check, we obtained a final sample of N=378. A sensitivity analysis for a repeated measures ANOVA with a within-between interaction using G* Power software (Faul, Erdfelder, Lang, & Buchner, 2007) revealed that with a sample size of N=378, Study 4 can detect an effect as small as f = .091. The sample was fairly evenly distributed across the three between-subject conditions of cynicism (N=127), anti-cynicism (N=125), and control (N=126), and attrition was not disproportionate in any single condition assignment.

In Study 4, the sample's mean age was M_{age} = 37.32 (SD=12.91) years, 52.6% of participants identified their gender as male, 45.5% as female, and 1.9% as other. For their ethnicity, 67.2% of participants identified as White or Caucasian, 9.3% as Black or African American, 7.9% as Asian or Pacific Islander, 6.1% as Hispanic or Latino, 0.3% as Native American, 8.7% reported a combination of ethnicity backgrounds, and 0.8% chose "other" as their ethnicity category. The sample leaned liberal on a 7-point conservativism scale (M=3.08, SD=1.68) ranging from Extremely Liberal (1) to Extremely Conservative (7) and was just slightly religious (M=1.99, SD=1.25) on a 5-point religiosity scale, ranging from Not at All Religious (1) to Extremely Religious (5).

Measures

Given that existing attempts to manipulate cynicism in the literature were given in the

context of financial incentives and monetary rewards (e.g. by giving people feedback on others' trustworthiness in economic games; Fetchenhauer & Dunning, 2010), in contexts of institutional trust (e.g. increase trust in the police; Spadaro, Gangl, Van Prooijen, Van Lange, & Mosso, 2020), or in the form of a longitudinal multi-year intervention (Stavrova & Ehlebracht, 2018), we created stimuli to be more targeted to the purposes and survey length of this study. All cynicism prompts and materials for Study 4 can be found in Appendix C.

As with prior studies, participants filled out the Moral Expansiveness Scale (Crimston et al., 2016), and identical to Studies 2 and 3, they reported liking and similarity ratings for average and close others on a scale from 1(Not at all) to 7 (Extremely), as well as the relationship they had with their close others (i.e. friend, family member, partner, colleague, or other category). Given that we observed the most consistent and strongest correlations between the MES (across the self and close other perspectives) on the competitive versus cooperative world beliefs scale (Clifton et al., 2019) in Study 2, as well as the greatest movement in predictions of MES average-other scores across different levels on the scale, we opted to use this scale as our measure of cynicism, which also served as a manipulation check for the effect of condition.

Procedure

Participants were first assigned to one of three between-subjects conditions using the "evenly-present elements" feature in Qualtrics software – the cynicism condition, the anti-cynicism condition, and a control condition, where participants received no additional information prior to completing the MES. Participants assigned to the cynicism condition read a summary ostensibly from a popular psychology magazine article, discussing research showing that human nature is generally much more selfish and competitive towards others and the world than we have assumed. Participants in the anti-cynicism condition read a same-length summary,

discussing research showing that human nature is generally kind towards others and the world, and that people care more about the world around them than we assume they do. Participants in the control condition did not receive any information and instead proceeded directly to the MES measures. In both the cynicism and anti-cynicism prompts, the cited research was real and linked references were provided at the end (see Appendix C for full prompts presented to participants). All participants were subsequently asked to fill out the Moral Expansiveness Scale (Crimston et al., 2016) in the same way as with Studies 2 and 3 - for themselves, for a close other that they were asked to bring to mind, and for an average other, in counterbalanced order. Following all MES perspectives, participants filled out the (Clifton et al., 2019) cynicism manipulation check, reported demographics, and were debriefed.

Results

For their close other in Study 4, 41% of participants (N=155) identified a friend, another 31.5% identified a partner (N=119), 26.7% identified a family member (N= 101), and 0.8% of participants chose someone belonging to an "other" category (N=3).

The repeated measures ANOVA tests on the 7-point liking and similarity scales, replicated that, on average, the identified close others were liked (M=6.51, SD=0.75, 95% CI [6.44, 6.59] significantly more compared to average others (M=4.04, SD=1.12, 95% CI [3.93, 4.15], (M_{diff} = 2.47, p<.001, 95% CI [2.35, 2.59]), F (1, 377) = 1718.85, p<.001, partial η^2=.820, and were also perceived as more similar to the participant (M=5.19, SD=1.23, 95% CI [5.06, 5.31]), compared to average others (M=3.64, SD=1.24, 95% CI [3.52, 3.77]), (M_{diff} = 1.55, p<.001, 95% CI [1.40, 1.69]), F (1, 377) = 444.22, p<.001, partial η^2 =.541.

We confirmed that the competitive versus cooperative beliefs scale (Clifton et al., 2019) displayed good reliability and the expected negative relationships with the MES self, and MES

close other scores that we also observed in Study 2, and this time also with the perceived MES average other scores (see Table 5 for Cronbach's alphas, descriptives, and correlations). The sample appeared moderately cynical on the same 1-7 scale (1 – not at all, 7 – completely true).

Table 5: *Reliabilities, Means, & Pearson's correlations between variables in Study 4*

Variable	α	M (SD)	1	2	3	4	5	6	7	8
1. Comp vs. Coop beliefs (Clifton et al., 2019)	.865	*M* =3.48 (1.38)	--							
2. MES self	.936	*M* =42.97 (12.72)	-.27***	--						
3. MES close other	.938	*M* =40.01 (13.04)	-.20***	.65***	--					
4. MES average other	.935	*M* =36.51 (11.95)	-.16**	.56***	.51***	--				
5. Similar to close other	--	*M* =5.19 (1.23)	-.19***	.05	.19***	.06	--			
6. Similar to average other	--	*M*=3.64 (1.24)	.15**	.13*	.08	.22***	.33***			
7. Like close other	--	*M*=6.51 (0.75)	-.14**	.08	.18***	.10	.42***	.21***	--	
8. Like average other	--	*M*=4.04 (1.12)	-.19***	.17***	.17***	.31***	.31***	.68***	.28***	--

Note 1. *The competitive vs. cooperative scale mean is collapsing across the three between-subjects conditions.*

Note 2. * *p*<.05, ** *p*<.01, *** *p*<.001

Following these, we moved on to check the main analyses of interest, namely 1) whether the cynicism manipulation was effective in reducing self-reported belief in cynicism on the competitive vs. cooperative world belief subscale (Clifton et al., 2019), (pre-registered hypothesis 1) and 2) whether the cynicism manipulation was effective in reducing actual MES scores across any of the MES perspectives, but particularly in the perceived moral expansiveness of average others (pre-registered hypotheses 2 & 3).

Belief in Cynicism Across the Three Cynicism Conditions. Contrary to our expectations, a one-way ANOVA test with the cynicism score on the cooperative vs. competitive world belief subscale as the dependent variable and condition assignment as the independent variable revealed no main effect of condition, F (2, 375) = 0.57, p= .567, partial η^2 =.003, ($M_{control}$=3.50, SD=1.25; M_{cyn}=3.57, SD=1.46; $M_{anti\text{-}cyn}$=3.38, SD=1.42) indicating that we did not successfully manipulate belief in cynicism as measured by this particular scale. This result failed to support pre-registered hypothesis 1.

MES Self, Close, and Average Other Across the Three Cynicism Conditions (Totals). Even though condition had no effect on the cooperative vs. competitive world belief subscale (Clifton et al., 2019) that we used as a manipulation check, in a repeated-measures ANOVA test with a Greenhouse-Geisser correction where we used the overall totals of each MES perspective as the dependent variable (maximum possible score of 90 points), condition interacted with the type of MES perspective, F (3.91, 733.24) = 3.36, p=.010, partial η^2 =.018. Specifically, and consistent with what we expected, condition had no effect (all pairwise comparisons p=.100) on the self-reported moral expansiveness of the self ($M_{self\_control}$=43.71, SE=1.14; M_{self_cyn}=42.40, SE=1.13; $M_{self\_anti\text{-}cyn}$=42.79, SE=1.14) or the perceived moral expansiveness of close others (all pairwise comparisons p=.100) ($M_{close\_control}$=39.63, SE=1.16;

M_{close_cyn}=39.52, SE=1.16; $M_{close\_anti\text{-}cyn}$=40.88, SE=1.17), but it did significantly affect the perceived moral expansiveness of average others ($M_{avrg\_control}$=37.59, SE=1.05; M_{avrg_cyn}=33.57, SE=1.05; $M_{avrg\_anti\text{-}cyn}$=38.42, SE=1.06). Specifically, within the perceived moral expansiveness of average-other scores, there were notable significant reductions between those in the cynicism condition and those in the control (M_{diff} =-4.02, p=.021, 95% CI [-7.59, -0.4]) and anti-cynicism (M_{diff} = -4.86, p=.004, 95% CI [-8.43,-1.28]) conditions, suggesting that participants' perceptions of the MES scores of average others were only affected by the cynicism condition, but not the anti-cynicism condition.

MES Self, Average, and Close Other for Proximal and Distant Entities Across the Three Cynicism Conditions. Although the design of Study 4 was fully crossed (all MES perspectives received a cynicism-relevant manipulation or control), for the remainder of the analyses that added the entity type as the additional factor, we conducted both a repeated-measures within-between subjects ANOVA with a Greenhouse-Geisser correction, as well as a linear mixed-effects model to better account for compound symmetry (the assumption that there is constant variance and constant covariance between observations), which the ANOVA analysis did not test for, and instead assumed (Bruin, 2006). Following pre-registration plans for this study, the following analyses report both and we note where the results of the two analytical approaches might diverge.

Before we examined interactions with entity type, a repeated-measures within-between subjects ANOVA with a Greenhouse-Geisser correction confirmed that, as with Studies 1a, 1b, 2, and 3, there was a main effect of entity type, F (1, 375.00) = 4314.22, p<.001, partial η^2 =.920, such that across all conditions, out of the maximum 3 points an entity could receive if placed in the innermost part of the MES, participants reported of themselves and predicted for others more

moral inclusivity (M_{diff} = 1.16, p<.001, 95% CI [1.12, 1.19]) of proximal entities (average of family and friends, ingroup, and revered entities) ($M_{proximal}$=2.14, SE=.02), than of distant entities (average of outgroups, stigmatized entities, animals of high sentience, animals of low sentience, plants, the environment, and villains) ($M_{distant}$=0.98, SE=.02). The linear mixed-effects model confirmed the main effect of entity type, F (1, 2058.33) = 3980.46, p<.001.

We moved on to check the highest-order Condition (cynicism, anti-cynicism, control) x MES perspective (self, close other, average other) x Entity type (proximal, distant) interaction, which was significant under the repeated-measures within-between subjects ANOVA with a Greenhouse-Geisser correction, F (3.78, 708.93) = 2.99, p=.021, partial η^2 =.016, and remained significant with a Huynh-Feldt correction, F (3.82, 716.21) = 2.99, p=.020, partial η^2 =.016, but not with the Lower-Bound correction, F (2.00, 375.00) = 2.99, p=.052, partial η^2 =.016. The linear mixed-effects model did not confirm the Condition x MES perspective x Cynicism condition interaction as significant, F (4, 1434.08) = 0.79, p=.533.

Given that we obtained significance with the within-between subjects ANOVA with the Greenhouse-Geisser and Huynh-Feldt corrections, we decomposed the interaction, seeing that the only differences between conditions stemmed from distant entities being perceived to be less included in average others' moral landscapes in the cynicism condition ($M_{cyn\_avrg\_distant}$=0.68, SE=.04, 95% CI [0.60, 0.77]), compared to the other two conditions. Specifically, the cynicism condition differed from the anti-cynicism condition ($M_{anti\text{-}cyn\_avrg\_distant}$=.89, SE=.04, 95% CI [0.81, 0.98]) (M_{diff} = -0.21, p<.001, 95% CI [-0.35, -0.07]) as well as the control condition ($M_{contr\_avrg\_distant}$=.87, SE=.04, 95% CI [0.79, 0.95]) (M_{diff} = -0.16, p= .005, 95% CI [-0.33, -0.04]). No differential effect of condition was observed on the self or close other perspectives' inclusion of either proximal or distant entities (all ps>.686) (see Figure 9 for the plotted

interaction).

Figure 9: *Condition x MES perspective x Entity type interaction in Study 4*

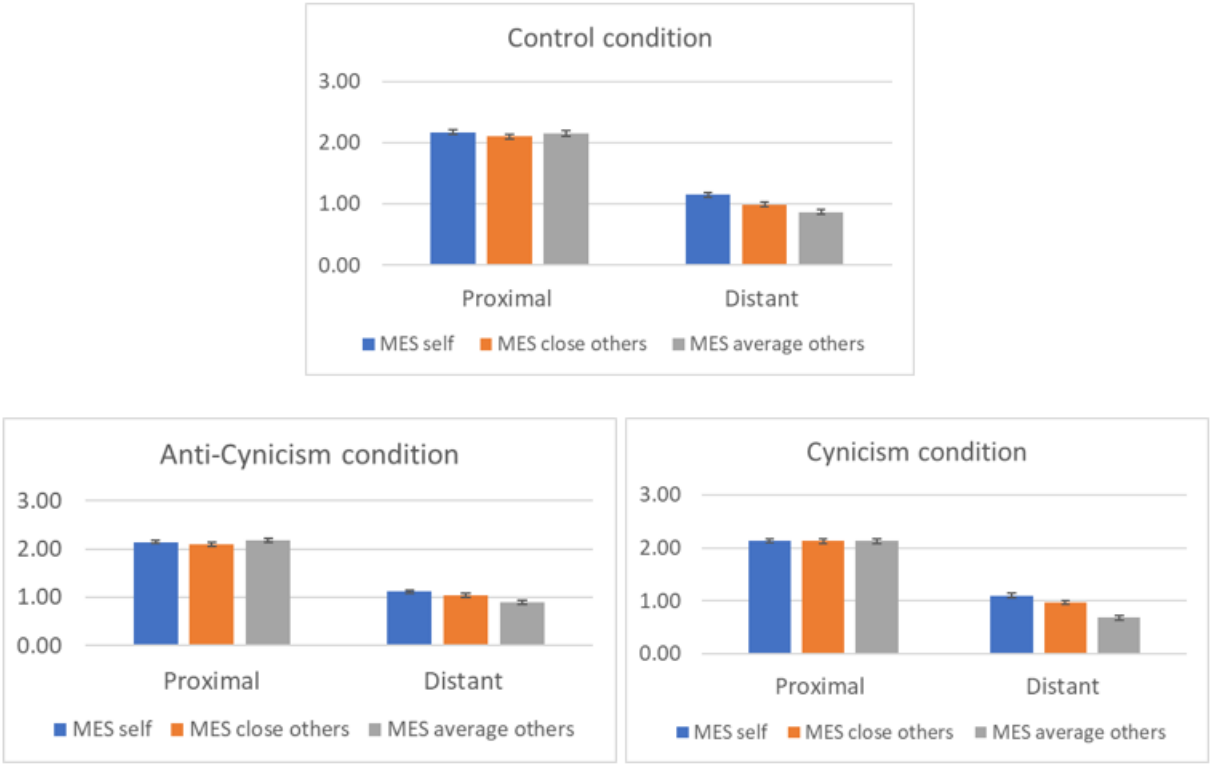

Note. *Plotted estimated marginal means. Error bars indicate standard error of the means.*

When breaking down the ANOVA interaction to compare the three MES perspectives across conditions, we confirmed that we failed to eliminate or lessen the gap between the moral expansiveness of the self and the one participants perceived for average others, as the discrepancy for distant entities among the two perspectives persisted in the anti-cynicism condition (M_{diff} = -0.26, p<.001, 95% CI [-0.33, -0.12]) and, as was expected, in the cynicism (M_{diff} = -0.42, p<.001, 95% CI [-0.53, -0.31]) and control (M_{diff} = -0.28, p<.001, 95% CI [-0.39, -0.18]) conditions. Thus, although one statistical approach to the results supported the Condition (pro-cynicism vs. anti-cynicism vs. control) x MES perspective (self vs. average other) x Entity

type (proximal vs. distant) significant interaction that we expected to find (pre-registered hypothesis 2), the direction of the effect was the opposite of what we predicted, with the MES discrepancy widening in the cynicism condition, as opposed to shrinking in the anti-cynicism condition. Thus, these results failed to support pre-registered hypothesis 2.

When comparing the self and close other perspectives for distant entities across the three conditions under the same interaction, we see the same moral inclusivity discrepancy we expected for the average other perspective. Specifically, close others were seen as being less inclusive of distant entities compared to the self in the cynicism (M_{diff} = -0.13, p=.002, 95% CI [-0.22, -0.04]) and control (M_{diff} = -0.17, p<.001, 95% CI [-0.26, -0.07]) conditions, but not in the anti-cynicism condition (M_{diff} = -0.07, p=.193, 95% CI [-0.16, 0.02]). These results failed to support pre-registered hypothesis 3, which expected no differences between the self and close other perspectives across the three cynicism conditions.

MES Self, Average, and Close Other Across the Three Cynicism Conditions (Entity-level). We further decomposed the lower-order Condition (cynicism, anti-cynicism, control) x MES (self, close other, average other) perspective interaction, which was significant with the Greenhouse-Geisser correction, F (3.95, 739.83) = 2.90, p=.022, partial η^2 =.015, and the Huynh-Feldt correction, F (3.99, 747.68) = 2.90, p=.021, partial η^2 =.015, but not the Lower-Bound correction, F (2.00, 375.00) = 2.90, p=.056, partial η^2 =.015. The linear mixed-effects model did not confirm the MES perspective x Cynicism condition interaction as significant, F (2, 2058.33) = 2.51, p=.081.

Given that we obtained significance with the within-between subjects ANOVA with the Greenhouse-Geisser and Huynh-Feldt corrections, we decomposed the Condition x MES perspective interaction, seeing, again, that the only difference between conditions stemmed from

average others being perceived to be less morally inclusive in the cynicism condition (M_{cyn_avrg}=1.41, SE=.03, 95% CI [1.35, 1.47]), compared to the anti-cynicism condition ($M_{anti\text{-}cyn\_avrg}$=1.54, SE=.03, 95% CI [1.48, 1.60]) (M_{diff} = -0.13, p=.013, 95% CI [-0.24, -0.02]), but the difference with the control condition did not reach significance this time (M_{contr_avrg}=1.51, SE=.03, 95% CI [1.45, 1.57]) (M_{diff} = -0.10, p= .080, 95% CI [-0.21, 0.01]). Condition had no other unique effect on the reports of moral expansiveness of the self or close other perspectives (all ps=1.000).

MES Comparison Between Self, Average Other, and Close Other Perspectives. In a separate analysis, we collapsed across the three cynicism conditions and re-tested the average differences between the three MES perspective totals to check whether the pattern with the main effect of perspective (MES self ≥ MES close other > MES average other) replicates, as with Studies 2, and 3. A repeated measures ANOVA with a Greenhouse-Geisser correction showed the same significant difference between the three MES perspectives as we obtained in Studies 1a, 1b, 2, and 3, F (1.95, 735.37) = 57.14, p<.001, partial η^2 =.132. Post hoc tests using the Bonferroni correction showed a significant difference (M_{diff} = 6.46, p<.001, 95% CI [4.98, 7.93]) between the scores for the MES of self (M_{self}=42.97, SE=.65) and the scores for the MES of average others (M_{avrg}=36.51, SE=.62) such that people perceive average others as less morally expansive compared to themselves by over two entities.

This increase in the mean difference between the self and average other scores (in Studies 1a, 1b, 2, and 3, the difference was the equivalent of a little over one entity while in Study 4, it was a little over two) is due the effect of the condition, which had a unique, reducing, impact only on the perceived expansiveness scores of average others. As with Studies 2 and 3, there was also a significant difference (M_{diff} = -3.49, p<.001, 95% CI [-5.03, -1.95]) between the scores for

the MES of average others and the MES scores of close others ($M_{close=}$40.01, SE=.67) such that people perceived their close others to be more morally expansive than what they perceived the average person to be, by the equivalent of a little over one entity. Lastly, we also replicated the additional smaller difference we found in Study 3, between the MES scores of the self and that of close others, such that participants perceived themselves as more morally expansive than their close others (M_{diff} = 2.96, p<.001, 95% CI [1.62, 4.30]).

MES Self - MES Close Other Comparison Across Low and High Felt Similarity. Study 4 did not manipulate similarity or liking for close others, and there were no differences in felt similarity, F (2, 375) = 0.58, p=.562, partial η^2 =.003, or liking, F (2, 375) = 0.80, p=.449, partial η^2 =.004, as a function of the cynicism condition participants were assigned to. Nevertheless, we wanted to check associations of the MES discrepancy with similarity and liking again, given the positive relationship we observed between similarity and the perceived MES scores of close others in Studies 2 and 3, and the significant moderation by felt similarity in Study 2.

As with Studies 2 and 3, we again obtained a significant correlation between how similar participants felt to their close others, and the perceived moral expansiveness of their close others, r=.19, p<.001 (see Table 5 for relevant correlations). The moderation of the difference between MES self and MES close other, at high (+1 SD) and low levels (-1 SD) of felt similarity was again significant, F (1, 376) = 11.13, p<.001, partial η^2 =.029. At high levels of felt similarity with close others, participants had no difference (M_{diff} = 1.13, p=.149, 95% CI [-0.40, 2.66]) in the MES scores they reported of themselves (M=43.64, SE=.93), and what they predicted for their close others (M=42.51, SE=.93). However, at low levels of felt similarity, participants reported being more morally expansive (M=42.30, SE=.93) compared to their close others

(M=37.50, SE=.93), (M_{diff} = 4.80, p<.001, 95% CI [3.27, 6.33]), replicating the result we obtained in Study 3.

MES Self - MES Close Other Comparison Across Low and High Liking. Even though we only saw a small significant correlation between liking and the perceived MES score for close others in Study 2, but no correlation in Study 3, and no moderation of the MES self-MES close other comparison by liking score in any of the two studies, we conducted the same analyses in Study 4 to test these relationships with an additional sample. We found the same significant association as with Study 2, between how much participants reported liking their close others, and how morally expansive they perceived them to be, r=.18, p<.001. This time, the liking score additionally significantly moderated the MES self-MES close other difference, F (1, 376) = 5.98, p=.015, partial η^2 =.016, in the same way that the similarity score did in Studies 3 and 4 - participants who reported liking their close other more had a smaller difference (M_{diff} = 1.61, p=.041, 95% CI [0.67, 3.15]) between their own MES scores (M=43.97, SE=.92), and what they predicted for their close others (M=42.36, SE=.94). However, at low levels of liking, participants reported being more morally expansive (M=41.97, SE=.92) than their close others (M=37.65, SE=.94), (M_{diff} = 4.32, p<.001, 95% CI [2.78, 5.86]).

Interim Discussion for Study 4

Although our results were mixed due to the assumptions of the linear mixed-effects model determining non-significance on the highest-order and some of the lower-order interactions, we have some evidence from at least one statistical approach that the cynicism manipulation in Study 4 successfully managed to decrease the perceived moral expansiveness of average others, but not that of close others, or that of the self for distant entities. This is consistent with our prediction that, because people already have access to their close others' (and

their own) moral landscapes, there is little movement that a cynical (or anti-cynical) outlook could have resulted to. In other words, we confirmed that for those who participants don't know, perceptions of their moral expansiveness are prone to informational biases relevant to cynicism. The competitive vs. cooperative world belief subscale by Clifton and colleagues (2019), however, failed to confirm that cynicism as measured by the specific scale was successfully manipulated, which is inconsistent with the effect we saw on the moral expansiveness measure itself.

Some additional methodological choices that might have impacted the fact that we obtained differences as a function of the condition on the MES scores but not on the cynicism manipulation measure, is that we presented the latter after participants went through all 3 perspectives of the MES, by which point the effect might have faded. Should we have put the scale right after the manipulation, we could have captured more of what the manipulation was intending to do, which seemed to have impacted the MES in the expected ways. The other possible reason is that the cynicism manipulation and what we used as a manipulation check, although both targeting cynicism as a construct, could have diverged in how they were perceived by participants. For example, the competitive vs. cooperative cynicism measure made broad philosophical statements about life being a cut-throat competition, but did not necessarily mention specific, tangible, human behaviors that were documented in the cynicism manipulation prompts.

Although we were able to see movement on the moral expansiveness scale as a function of the cynicism manipulation, we were not able to shift perceived moral expansiveness scores for average others upwards, even with a comparable anti-cynicism manipulation sharing true facts about human kindness. There may be at least two reasons for why we weren't able to do so. The

first is that people might be predisposed to believe information that confirms people's bad intentions, as opposed to good ones. This might serve an adaptive purpose preparing one for and preventing one from experiencing harm in the context of social interactions (Baumeister, Bratslavsky, Finkenauer, & Vohs, 2001).

The second reason might be that prosociality-relevant inductions tend to resemble people's–baselines - for example, earlier work has found that encouraging participants to empathize with a suffering target does not result in any additional increases beyond what participants already do without any instructions (McAuliffe, Carter, Berhane, Snihur, & McCullough, 2020), but encouraging participants to remain detached does reduce their empathy for suffering. It is possible that something similar is happening in–this data - people might already be reaching a "maximum" of how much they can be convinced that the world is a kind place, but asking people to walk back from that assumption is far easier.

Chapter 6: General Discussion

Recap of Research Aims

Across the literature, there is a well-documented underestimation of prosociality happening across domains and contexts (blood donation - Miller & Ratner, 1998; Brick et al., 2021; everyday requests for help and environmental behavior - Kogut & Beyth-Marom, 2008; Chen, Wan, & Yang, 2022; negotiations - Mason, Wiley, & Ames, 2018; everyday requests for help - Zhao & Epley, 2022; climate policy support - Sparkman, Geiger, & Weber, 2022; COVID-19 collective health preventative behavior - Graupensperger, Lee, & Larimer, 2021; same-sex marriage approval - Eisner, Turner-Zwinkels, & Spini, 2021). My work has attempted to zoom out of individual helping and cooperative behaviors and intentions, and ask whether the effect is broader, that is, whether we underestimate others' moral landscape en large. Further, my dissertation studies aimed to discern whether this discrepancy can be traced back to cynical biases that result from gaps in simulation, found when social distance is increased, whether those biases are reversible when such distance is eliminated, and whether this same discrepancy can be eliminated even for strangers when participants are reminded that we tend to underestimate the goodness in others.

Our initial Studies 1a and 1b have lent support to our initial prediction that moral landscapes are imagined as more constricted for others compared to that of the self, and have further specified that this underestimation might be contingent on social distance between self and other. My proposed dissertation studies aimed to establish that we see this because people's simulation for things that are further away are more prone to inaccuracies (Epley, 2008) because people have to "guess" with limited availability of information. In order to test whether restoring this lack of access to information eliminates the moral inclusivity discrepancy, we asked

participants to imagine the moral circle of those whose moral landscapes they have access to (close others).

We predicted that the moral circle of close others will not look significantly different than the self, even when we manipulated similarity for close others (Study 2), and liking for close others (Study 3), two factors that we would expect could be responsible for bringing the MES self scores and MES close others scores closer together. Study 4 attempted to test a cynical bias as an explanation of the moral inclusivity discrepancy by directly exposing participants to an anti-cynicism prompt prior to the completion of the moral expansiveness measure. We expected that this prompt would "correct" people's predicted moral expansiveness scores for strangers (average others), but this effect would be muted for the moral expansiveness scores they predicted for their close others.

Main Takeaways

Did the Perceived Moral Expansiveness of Close others (More Closely) Match That of the Self?

While Study 2 did not find any differences between the moral expansiveness participants reported for themselves and what they predicted for their close others, this difference did appear in the more highly-powered Studies 3 and 4, showing the same moral inclusivity discrepancy we observed in the comparisons between one's own moral expansiveness and their perception of the average person's, concentrated again in the category of distant entities. In other words, although participants did have greater informational access to their close others' moral landscapes, they still thought that their closer others were less inclusive to entities that are typically stigmatized or marginalized. It's important to note that while this difference appeared, there was still a greater overlap between one's own moral inclusivity and their close others', and the moral landscapes of

close others were still, in the minds of our participants, significantly more inclusive compared to those of the average person. This could suggest a few possible explanations for the existence of the discrepancies.

First, participants are privy to their own internal experiences, preferences, and ideals, and their reports of others are always going to come from a third-person vantage point. Thus there is a possibility that reports of others' experiences (whether close or general others) might correlate with each other more strongly than with one's own. Beyond this, our findings might be complimenting a long line of previous work finding that such "empathy gaps" result in underestimations or "dampening" of experiences, whereby people systematically underreport both the visceral (Loewenstein, 1996) and the cognitive (Van Boven, Dunning, & Loewenstein, 2000; Faro & Rottenstreich, 2006) experiences of others. However, this need not negate a motivated account for our results.

To the extent that we found accurate (i.e. reflecting the sample average for the MES self perspective) estimations for both close and average others on entities typically placed proximal to one's moral circle - family and friends, ingroups, and revered individuals, we can assume that participants have cognitively calibrated to reach what they think is plausible for both proximal and distant entities (Epley, Keysar, Van Boven, & Gilovich, 2004). Thus, if the moral inclusivity discrepancy persists, to some extent, even in the cases of close others where calibration should be optimized, it can signal that its existence is independent of access to information and instead might be more motivated in nature: people's cynical intuitions might span even those closest to them (Zhao & Epley, 2022).

Did Perceived Similarity With and Liking for Close-Other Affect the Difference Between MES Self-MES Close Other?

Studies 2 and 3 attempted to answer the question of whether moral expansiveness between self and close others resembles one another merely because participants feel more similar to or like their close others more, compared to average others, and not because they have informational access to their close others' moral worlds. Study 2 was successful in shifting participants' perceived similarity to their close other following our 2-minute similarity manipulation, without, however, shifting their perceived moral expansiveness accordingly. While that could suggest that felt similarity is unrelated to perceptions of moral expansiveness, it would be hasty to accept a null interaction without examining naturally-occurring similarity that participants perceive between themselves and their reported close other individual, that could still be influencing their MES ratings.

Indeed, the more similar participants perceived their close others to be, the higher they perceived their moral expansiveness to be across all studies (Study 2, r=.16, p<.010; Study 3, r=.16, p<.010; Study 4, r=.19, p<.001). We saw this precise pattern in Studies 3 and 4, whereby the more similar participants felt to their close others, the closer their MES scores were to each other. Thus, we can conclude that naturally-occurring similarity can and does play a role in the closer overlap we hypothesized and saw between participants and their view of their close others. This also suggests that there could also be intergroup effects at play. Although we have attempted to keep the description of "others" as devoid of context as possible to control for motivational effects, we still have assumed that close others will be, more often than not, ingroup members (in various dimensions relevant to the participant). However, this remains an assumption, and to the extent that we have evidence for naturally-occurring perceived similarity bringing the self and

close other imagined moral expansiveness closer together, there is a very real possibility that we are seeing intergroup effects at play.

Unsurprisingly, it proved difficult to get participants to dislike their close others in Study 3, in order to be able to conclude that liking is not a deciding factor in the moral expansiveness estimations for close others. When excluding a handful of participants who found it difficult to list reasons for why they disliked their close others, the manipulation check confirmed that those in the liking and those in the disliking condition substantially differed in their liking ratings as we had hoped to achieve, but the difference was small, and ratings across all conditions still reached ceiling effects (above a 6 on a 7-point Likert scale).

As with the similarity relationship to MES, in the case of liking too, the more participants liked their close others, the higher they perceived their moral expansiveness to be across two out of the three studies (Study 2, r=.15, p<.010; Study 4, r=.18, p<.001) and in Study 4, the more participants liked their close others, the closer their MES scores were to each other. Although we ultimately did not find any differences in moral expansiveness as a function of the liking condition in Study 3, the existing correlations and liking moderation in Study 4 potentially support earlier findings showing that people are motivated to see those they like as highly moral (Bocian, Baryla, Kulesza, Schnall, & Wojciszke, 2018; Lee & Holyoak, 2020). It's important to note, however, that in Study 3 which used a university student sample, we did not see a relationship between how much participants liked their close others, and their estimation of their close others' moral expansiveness (Study 3, r=.02, p>.050), possibly presenting a case for close others (often parents) that can be liked, but perhaps due to generational differences (e.g. see link between age and conservatism; Cornelis, Van Hiel, Roets, & Kossowska, 2009) are not necessarily thought of as more morally expansive.

It's important to note that even at low levels of similarity and liking where the gap between one's own moral expansiveness and that of close others was significantly larger, the nature of the gap was such that participants in Studies 3 and 4 still perceived their close others as *less* morally expansive than themselves. It is possible that someone might perceive themselves as less similar to a close other or to like them less, but still perceive them as more (compared to less)–expansive - however, this did not seem to be the case in our data. In other words, when participants felt less similar to or liked their close other less, they always underestimated their expansiveness relative to their own. Although this may be due to a variety of reasons (e.g. cynicism, empathy gaps, differentiation/distinctiveness needs), it generally suggests that other systemic factors might be at play.

Did Manipulating Cynicism Uniquely Affect the Perception of MES Scores for Average Others?

Across the studies, and prior to attempting to manipulate cynicism, we found correlational evidence that when people are least cynical, the gap between their expansiveness and what they predict for average others is wider. This was because one's own moral inclusivity is higher as cynicism levels remain low, but estimations of average other's inclusivity remain constant and relatively immune to cynicism-relevant individual differences. In other words, in many of the studies, we see that whether or not one holds a more cynical outlook has little impact on how they imagine the moral landscapes of others, and instead is more relevant for their own, or their perception of their close other's moral landscapes (see relevant correlations in Tables 3 and 4). While the reason for this could be psychometric in nature – people's own scores correlate better with any given measure that is self-report, compared to an estimate they give for others, it does put into question our initial prediction that cynicism was the main reason for the

moral inclusivity discrepancy we observed across all studies.

The more important question, after identifying that perceptions of others' expansiveness remain consistently unmoved, is perhaps not the gap between self and average other in expansiveness, but critically, the question of how to move the expansiveness perceptions of others towards distant entities in the moral circle, upwards. The results of Study 4 point to the possibility that cynicism-relevant content might still be one avenue to move perceptions of others' moral expansiveness towards distant entities. Although the results on the effects of the manipulation differed depending on the statistical approach taken, and cynicism as measured in the manipulation check did not converge with the observed changes on the perceived moral expansiveness of average others, we did see a substantial decrease in the perception of the average person's moral inclusivity of distant entities in the moral circle.

Specifically, a 1-minute passage designed to manipulate cynicism (see Appendix C) managed to recreate the same gap (of the equivalent of a one-entity exclusion) that we found naturally-occurring in all previous studies between the moral expansiveness scores of the self and those of average others. Critically, and as we predicted, one statistical interpretation of the data shows this effect to be unique to the perceptions of average others, but not close others, or the self, even though correlations in earlier studies (see Tables 3 and 4) suggested that trait cynicism was more strongly correlated with one's own scores and those they predicted for their close others. This potentially lends support to our initial theory that because participants already have access to their own moral landscapes and those of their close others, information that "fills the gaps" for others out in the world, broadly speaking, is more impactful in how they imagine their inclusivity for entities that are most removed from the moral circle.

Although a handful of recent studies have managed to shift one's own moral

expansiveness following experiences of awe induction (Song, Klebl, & Bastian, 2023), or compassion training (Kirby, Hoang, & Crimston, in press), to our knowledge, these results present the first successful attempt to shift perceptions of *others'* moral expansiveness in an experimental setting. However, perceptions of moral inclusivity for distant entities among average others was shifted downwards in the cynicism condition, but not upwards in the anti-cynicism condition, the latter being what we were hoping to achieve. Specifically, while we thought that the cynicism and control conditions would look similar to each other because we hypothesized people's default (control) reflects a cynical reality, and thus predicted movement only on the anti-cynicism condition, we were proven wrong. It's possible that cynicism runs so deep, that even small cues can activate or re-confirm existing biases. As mentioned in the interim discussion of Study 4, this readiness to process information relevant to people's bad intentions might serve an adaptive purpose of preventing harm in the context of social interactions (Baumeister, Bratslavsky, Finkenauer, & Vohs, 2001).

Importance of Contribution & Big-Picture Discussion

Overall, our data establish a few "firsts" in the understanding of how people imagine the moral landscapes of others. First, we have robust evidence across five samples that those landscapes are systematically imagined as more constricted relative to the self, conceptually replicating numerous other papers finding self-other differences in various domains of prosociality (Miller & Ratner, 1998; Brick et al., 2021; Kogut & Beyth-Marom, 2008; Chen, Wan, & Yang, 2022; Mason, Wiley, & Ames, 2018; Zhao & Epley, 2022; Sparkman, Geiger, & Weber, 2022; Graupensperger, Lee, & Larimer, 2021). Secondly, we also established that when those moral landscapes of others are imagined, they are not constricted universally, rather; they seem to shrink when it comes to people and entities in the world for which we perhaps have yet

to reach social consensus on their deserved moral inclusivity (Eisner, Turner-Zwinkels, & Spini, 2021). Thirdly, we show that both informational and motivational effects can be at play when that imagined moral constriction occurs, which replicates adjacent work on the perception of moral decline (Mastroianni & Gilbert, 2023) when those imagined others fall at different places on the spectrum of social distance. Finally, we present a first brief attempt to shift the perceptions of moral expansion in others by manipulating assumptions that people hold about kindness in others, and pave the way for future work to continue uncovering the reasons for this imagined moral discrepancy.

Limitations and Future Directions

Although this dissertation makes important contributions to the literature, it is important to discuss its limitations. First, while we were able to conceptualize "others" as other people in the study in Studies 1a and 1b, we have no accurate way of knowing whether the close others' moral expansiveness scores that participants will report (Studies 2, 3, & 4), accurately reflect reality. This presents a legitimate barrier to making accuracy claims without having access to close others' actual moral expansiveness scores, although work has shown that second-hand reports can be trusted in describing a close other's traits accurately (Vazire & Mehl, 2008; Vazire & Carlson, 2011). Since our attempt to manipulate cynicism (Study 4) was not predicted to shift the perceived scores of close others, (and, in fact, it did not), we can be less concerned about this particular limitation. However, we can't speak confidently of a moral inclusivity "error".

A second limitation is about the scope of our claims for informational versus motivational mechanisms for understanding the moral expansiveness discrepancy. While we have tried to argue for an informational mechanism for the discrepancy using a close-others design that made availability of information on one's moral landscape possible (Study 2, trying

to rule out felt similarity as an alternative explanation; Study 3, trying to rule out liking as an alternative explanation), it was not possible to rule out liking for close others as a potential explanation for the similarity of self and close-other MES scores, given the results of Study 3. Because people are expected to exhibit a desire to morally elevate close others (Bocian, Baryla, Kulesza, Schnall, & Wojciszke, 2018; Lee & Holyoak, 2020), motivational effects associated with reporting higher moral expansiveness could still be at play. Thus, we cannot decouple informational effects from motivational effects in the case of close others, at least when it comes to liking. As such, our methodological choice implies a trade-off between ecological validity of imagining others whose landscapes can only be accessed through this kind of close personal connection, and the inability to claim purely informational effects on moral expansiveness because of the motivational "baggage" that comes precisely due to the nature of this connection.

A third limitation is that although our designs steered clear of explicitly inducing group membership motivations by describing general or average "others" without any additional prompting, it is of course not unreasonable to think that those others could still be thought of as outgroups, in some respect. The fact that naturally-occurring similarity in imagined close others still seemed to bring the moral expansiveness of self and close other closer together in some of our studies, suggests that even close others can be thought of as socially close but potentially still belonging to some outgroup category. Relatedly, to the extent that participants saw themselves as belonging to any group categories referenced in the moral expansiveness measure itself (e.g. stigmatized groups), could have also introduced group-relevant motivations that mattered for how the moral landscapes of others were imagined. Although some attempts have been made to try and understand how moral expansiveness is explained by who expands and towards which entities (Jaeger & Wilks, 2023), future work should further unpack how and when intergroup

motivations are relevant in how the moral landscapes of others are imagined.

A third limitation is that our studies did not take into consideration variation of moral expansion across different sub-categories, which is a consideration that prior research has deemed relevant (Rottman, Crimston, & Syropoulos, 2021). That is, while people can, on average, have the same moral expansiveness score, and thus appear equally expansive, some might prioritize moral inclusion of environmental entities more than others, a phenomenon of "tree-huggers" versus "human-lovers" according to Rottman and colleagues (2021). Given that the moral inclusivity discrepancy is contained primarily within those entities in which this variation appears, it is important that next steps of this work examine variation in people's inclusion of entity subcategories, and the extent to which our effects are contained within people who do prioritize nature over humans. Although the self-general other discrepancy that we find in all 5 studies should not be dependent on whether one is more expansive towards nature or not, it is possible that people within that category exaggerate this discrepancy even further, because of heightened concerns for environmental protection and conservation (Rottman, Crimston, & Syropoulos, 2021).

Fourth, while the measure of moral expansiveness that we used is well-validated and well-correlated with constructs of prosocial intentions and behavior (Crimston et al., 2016), as well as widely used as a predictor and a dependent variable, especially in recent years (Rottman et al., 2021; Kirkland et al., 2022; Song, 2023) it is one of many operationalizations of moral expansion that exist in the literature (for other examples see Reed & Aquino, 2003; Waytz, Iyer, Young, Haidt, & Graham, 2019; Yudkin, Prosser, Heller, McRae, Chakroff & Crockett, 2022). Thus, our findings remain, to some extent, measure-specific and future work should attempt to replicate the self-other discrepancy effect across different operationalizations of the construct of

moral expansiveness.

Finally, given that our manipulation intending to shift cynicism did have an effect on perceptions of moral expansiveness for others, but not on a cynicism manipulation check, it is unclear whether what we manipulated was indeed, cynicism. It also remains unclear if the shift we saw on moral expansiveness is a long-lasting or sustainable one, thus our ability to claim that we have produced an intervention that can upregulate (but not downregulate) cynicism remains limited, for these reasons. This is further complicated by the challenge of defining a construct as broad as cynicism, which has sometimes been described as persisting even at the face of disconfirming evidence, earning it the adjective of "naive" (Kruger & Gilovich, 1999), or it being "undue" (Critcher & Dunning, 2011), suggesting that it might reflect a stable dispositional trait, although in other work it has been described as a belief (Stavrova & Ehlebracht, 2019; Hilbig, Moshagen, Thielmann, & Zettler, 2022). If cynicism is a stable trait, for example, attempts to shift it in short experimental surveys like ours might remain unsuccessful, and designs that are more longitudinal in nature might be required (Stavrova & Ehlebracht, 2018). Future work should develop and calibrate multiple possibilities for manipulating cynicism depending on the construct definition.

The cynicism explanation is only one of many that could be explaining the self-other moral expansiveness discrepancy. There are also explanations that have not been tested in this set of studies such as self-enhancement (Epley & Dunning, 2000; Kwan, John, Kenny, Bond, & Robins, 2004; Klein & Epley, 2017), or pluralistic ignorance (Katz & Allport, 1931; Prentice & Miller, 1996; Sargent & Newman, 2021). Given that cynicism can develop as a response to repeated exposure to negative news and events (Buchanan & Sandstrom, 2023), as well as the deliberate seeking out and paying attention to such information (Mastroianni & Gilbert, 2023),

pluralistic ignorance (the belief that others hold opinions that are different from the self's) and other explanations need not to be completely separated from cynicism. Thus, additional explanations for the moral inclusivity discrepancy we observed need not to be completely separated from cynicism. Future work can test cynicism against, or complimentary to other explanations.

Conclusion

People exhibit a moral inclusivity discrepancy when imagining the moral landscape of others. In this dissertation, I aimed to map and understand where this discrepancy is coming from and whether it is the result of "unfilled gaps" in the place of which people cynically assume moral constriction. The self-other discrepancy gap can be due to multiple reasons beyond just a cynical stance in life. But we found some mixed evidence that a short intervention intending to manipulate cynicism did uniquely lower the perception of moral expansiveness for other people. Although the nature, strength, and duration of the manipulation all warrant further examination in future work, our results speak to the potential of one possible avenue to do so. This work furthers our understanding of how we perceive prosociality in others and needs to be taken into consideration when attempting to change perceptions of and for a more morally inclusive society and world.

References

Allum, N. (2011). What makes some people think astrology is scientific?. *Science Communication, 33*(3), 341-366.

Alsaid, A., Li, M., Chiou, E. K., & Lee, J. D. (2023). Measuring trust: A text analysis approach to compare, contrast, and select trust questionnaires. *Frontiers in Psychology*, 14.

Barefoot, J. C., Dodge, K. A., Peterson, B. L., Dahlstrom, W. G., & Williams, R. B. (1989). The Cook-Medley hostility scale: item content and ability to predict survival. *Psychosomatic medicine*.

Batista, R. M., Sussman, A. B., & Trueblood, J. S. (2023). Self-other differences in perceptions of wealth. *Journal of Experimental Social Psychology, 104*, 104420.

Baumeister, R. F., Bratslavsky, E., Finkenauer, C., & Vohs, K. D. (2001). Bad is stronger than good. *Review of General Psychology*, *5*(4), 323-370.

Bensley, D. A., Watkins, C., Lilienfeld, S. O., Masciocchi, C., Murtagh, M. P., & Rowan, K. (2022). Skepticism, cynicism, and cognitive style predictors of the generality of unsubstantiated belief. *Applied Cognitive Psychology, 36*(1), 83-99.

Bocian, K., Baryla, W., Kulesza, W. M., Schnall, S., & Wojciszke, B. (2018). The mere liking effect: Attitudinal influences on attributions of moral character. *Journal of Experimental Social Psychology, 79*, 9-20.

Brick, C., Fillon, A., Yeung, S. K., Wang, M., Lyu, H., Ho, J. Y. J., ... & Feldman, G.

(2021). Self-interest is overestimated: Two successful pre-registered replications and extensions of Miller and Ratner (1998). *Collabra: Psychology, 7*(1), 23443.

Brown, A. (2020, August 28). Most Democrats who are looking for a relationship would not consider dating a Trump voter. Pew Research Center. https://pewrsr.ch/3bwR7aK

Bruin, J. (2006). Newtest: command to compute new test. UCLA: Statistical Consulting Group. https://stats.oarc.ucla.edu/stata/ado/analysis/.

Buchanan, K., & Sandstrom, G. M. (2023). Buffering the effects of bad news: Exposure to others' kindness alleviates the aversive effects of viewing others' acts of immorality. *Plos one, 18*(5), e0284438.

Burger, J. M., Messian, N., Patel, S., Del Prado, A., & Anderson, C. (2004). What a coincidence! The effects of incidental similarity on compliance. *Personality and Social Psychology Bulletin, 30*(1), 35-43.

Byrne, D. (1961). Interpersonal attraction and attitude similarity. *The Journal of Abnormal and Social Psychology, 62*(3), 713-715.

Carlson, R. W., & Zaki, J. (2022). Belief in altruistic motives predicts prosocial actions and inferences. *Psychological Reports, 125*(4), 2191-2212.

Chowdhury, R. M., & Fernando, M. (2014). The relationships of empathy, moral identity and cynicism with consumers' ethical beliefs: The mediating role of moral disengagement. *Journal of Business Ethics, 124*, 677-694.

Chen, S., Wan, F., & Yang, S. (2022). Normative misperceptions regarding pro-environmental behavior: Mediating roles of outcome efficacy and problem awareness. *Journal of Environmental Psychology, 84*, 101917.

Choy, B. K., Eom, K., & Li, N. P. (2021). Too cynical to reconnect: Cynicism moderates the effect of social exclusion on prosociality through empathy. *Personality and Individual Differences, 178*, 110871.

Clifton, J. D., Baker, J. D., Park, C. L., Yaden, D. B., Clifton, AI., Terni, P., ... & Seligman, M. E. (2019). Primal world beliefs. *Psychological Assessment, 31*(1), 82.

Clifton, J. D., & Yaden, D. B. (2021). Brief measures of the four highest-order primal world beliefs. P*sychological Assessment, 33*(12), 1267.

Cook, W. W., & Medley, D. M. (1954). Proposed hostility and Pharisaic-virtue scales for the MMPI. *Journal of Applied Psychology, 38*(6), 414–418.

Cornelis, I., Van Hiel, A., Roets, A., & Kossowska, M. (2009). Age differences in conservatism: Evidence on the mediating effects of personality and cognitive style. *Journal of Personality, 77*(1), 51-88.

Crimston, C. R., Bain, P. G., Hornsey, M. J., & Bastian, B. (2016). Moral expansiveness: Examining variability in the extension of the moral world. *Journal of Personality and Social Psychology, 111*(4), 636–653.

Crimston, C. R., Blessing, S., Gilbert, P., & Kirby, J. N. (2022). Fear leads to suffering: Fears of compassion predict restriction of the moral boundary. *British Journal of*

Social Psychology, 61(1), 345-365.

Critcher, C. R., & Dunning, D. (2011). No good deed goes unquestioned: Cynical reconstruals maintain belief in the power of self-interest. *Journal of Experimental Social Psychology, 47*(6), 1207-1213.

Davis, M. H. (1983). Measuring individual differences in empathy: Evidence for a multidimensional approach. *Journal of Personality and Social Psychology, 44*(1), 113-126.

Douglas, K. M., & Sutton, R. M. (2004). Right about others, wrong about ourselves? Actual and perceived self-other differences in resistance to persuasion. B*ritish Journal of Social Psychology, 43*(4), 585-603.

Eisner, L., Spini, D., & Sommet, N. (2020). A contingent perspective on pluralistic ignorance: When the attitudinal object matters. *International Journal of Public Opinion Research, 32*(1), 25-45.

Eisner, L., Turner-Zwinkels, F., & Spini, D. (2021). The impact of laws on norms perceptions. *Personality and Social Psychology Bulletin, 47*(7), 1071-1083.

Epley, N. (2008). Solving the (real) other minds problem. *Social and Personality Psychology Compass, 2*(3), 1455-1474.

Epley, N., & Dunning" D. (2000). Feeling" holier than thou": are self-serving assessments produced by errors in self-or social prediction?. *Journal of Personality and Social Psychology, 79*(6), 861.

Epley, N., Keysar, B., Van Boven, L., & Gilovich, T. (2004). Perspective taking as egocentric anchoring and adjustment. *Journal of Personality and Social Psychology, 87*(3), 327.

Everett, J. A., Faber, N. S., Savulescu, J., & Crockett, M. J. (2018). The costs of being consequentialist: Social inference from instrumental harm and impartial beneficence. *Journal of Experimental Social Psychology, 79*, 200-216.

Faro, D., & Rottenstreich, Y. (2006). Affect, empathy, and regressive mispredictions of others' preferences under risk. *Management Science, 52*(4), 529-541.

Faul, F., Erdfelder, E., Lang, A. G., & Buchner, A. (2007). G* Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods, 39*(2), 175-191.

Fetchenhauer, D., & Dunning, D. (2010). Why so cynical? Asymmetric feedback underlies misguided skepticism regarding the trustworthiness of others. *Psychological Science, 21*(2), 189-193.

Geber, S., Baumann, E., Czerwinski, F., & Klimmt, C. (2021). The effects of social norms among peer groups on risk behavior: A multilevel approach to differentiate perceived and collective norms. *Communication Research, 48*(3), 319-345.

Gilbert, P., McEwan, K., Matos, M., & Rivis, A. (2011). Fears of compassion: Development of three self-report measures. *Psychology and Psychotherapy: Theory, research and practice, 84*(3), 239-255.

Graupensperger, S., Lee, C. M., & Larimer, M. E. (2021). Young adults underestimate how well peers adhere to COVID-19 preventive behavioral guidelines. *The Journal of Primary Prevention*, *42*(3), 309-318.

Hilbig, B. E., Moshagen, M., Thielmann, I., & Zettler, I. (2022). Making rights from wrongs: The crucial role of beliefs and justifications for the expression of aversive personality. *Journal of Experimental Psychology: General, 151*(11), 2730.

Jaeger, B., & Wilks, M. (2023). The relative importance of target and judge characteristics in shaping the moral circle. *Cognitive Science, 47*(10), e13362.

Kirby, J.N., Hoang, A., & Crimston, C.R. (in press). A brief compassion focused therapy intervention can increase moral expansiveness: A randomized controlled Trial. *Mindfulness*. doi.org/10.1007/s12671-023-02300-7.

Kirkland, K., Crimston, C. R., Jetten, J., Rudnev, M., Acevedo-TrianaI., Amiot, C. E., ... & Bastian, B. (2022). Moral expansiveness around the world: The role of societal factors across 36 countries. *Social Psychological and Personality Science*, 19485506221101767.

Klein, N., & Epley, N. (2017). Less evil than you: Bounded self-righteousness in character inferences, emotional reactions, and behavioral extremes. *Personality and Social Psychology Bulletin, 43*(8), 1202-1212.

Kogut, T., & Beyth-Marom, R. (2008). Who helps more? How self-other discrepancies influence decisions in helping situations. *Judgment and*

Decision Making, 3(8), 595.

Kruger, J., “ Iich,”T. (1999). " Naive cynicism" in everyday theories of responsibility assessment: On biased assumptions of bias. *Journal of Personality and Social Psychology, 76*(5), 743.

Kwan, V. S., John, O. P., Kenny, D. A., Bond, M. H., & Robins, R. W. (2004). Reconceptualizing individual differences in self-enhancement bias: an interpersonal approach. *Psychological review, 111*(1), 94.

Laham, S. M. (2009). Expanding the moral circle: Inclusion and exclusion mindsets and the circle of moral regard. *Journal of Experimental Social Psychology, 45*(1), 250-253.

Law, K. F., Campbell, D., & Gaesser, B. (2022). Biased benevolence: The perceived morality of effective altruism across social distance. *Personality and Social Psychology Bulletin*, *48*(3), 426-444.

Lee, J., & Holyoak, K. J. (2020). “But he’s my brother”: The impact of family obligation on moral judgments and decisions. *Memory & cognition, 48,* 158-170.

Lees, J., Young, L., & Waytz, A. (2022).’Morally questionable actors' meta-perceptions are accurate but overly positive. *Journal of Experimental Social Psychology, 102,* 104371.

Leviston, Z., & Uren, H. V’ (2020). Overestimating one's “green” behavior: Better-than-

average bias may function to reduce perceived personal threat from climate change. *Journal of Social Issues, 76*(1), 70-85.

Lindsey-Hall, K. K., Jaramillo, S., Baker, T. L., & Bachrach, D. G. (2023). An examination of frontline employee–customer incidental similarities in service failure and recovery contexts. *Psychology & Marketing*, 1– 14.

Litman, L., Robinson, J., & Abberbock, T. (2017). TurkPrime. com: A versatile crowdsourcing data acquisition platform for the behavioral sciences. *Behavior Research Methods, 49*(2), 433-442.

Loewenstein, G. (1996). Out of control: Visceral influences on behavior. *Organizational Behavior and Human Decision Processes, 65*(3), 272-292.

Malle, B. F. (2006). The actor-observer asymmetry in attribution: a (surprising) meta-analysis. *Psychological Bulletin, 132*(6), 895–919.

Mason, M. F., Wiley, E. A., & Ames, D. R. (2018). From belief to deceit: H'w expectancies about others' ethics shape deception in negotiations. *Journal of Experimental Social Psychology, 76*, 239-248.

Mastroianni, A. M., & Gilbert, D. T. (2023). The illusion of moral decline. *Nature*, 1-8.

McAuliffe, W. H., Carter, E. C., Berhane, J., Snihur, A. C., & McCullough, M. E. (2020). Is empathy the default response to suffering? A meta-analytic evaluation of perspective taking's effect on empathic concern. *Personality and Social Psychology Review, 24*(2), 141-162.

McManus, R. M., Kleiman-Weiner, M., & Young, L. (2020). What we owe to family:

The impact of special obligations on moral judgment. *Psychological Science, 31*(3), 227-242.

McManus, R. M., Young, L., & Sweetman, J. (2023). Psychology is a property of persons, not averages or distributions: Confronting the group-to-person generalizability problem in experimental psychology. *Advances in Methods and Practices in Psychological Science, 6*(3), 25152459231186615.

Meehan, M. K., Zeigler-Hill, V., & Shackelford, T. K. (2022). Dark personality traits and anti-natalist beliefs: The mediating roles of primal world beliefs. *Philosophical Psychology*, 1-23.

Miller, D. T., & Ratner, R. K. (1998). The disparity between the actual and assumed power of self-interest. *Journal of Personality and Social Psychology, 74*(1), 53.

Neumann, E., & Zaki, J. (2022). Toward a social psychology of cynicism. *Trends in Cognitive Sciences*, *27*(1), 1-3.

Opotow, S. (1990a). Moral exclusion and injustice: An introduction. *Journal of Social Issues, 46*(1), 1-20.

Opotow, S. (1990b). Deterring moral exclusion. *Journal of Social Issues, 46*(1), 173-182.

Opotow, S., & Weiss, L. (2000). New ways of thinking about environmentalism: Denial and the process of moral exclusion in environmental conflict. *Journal of Social Issues, 56*(3), 475-490.

Paolacci, G., Chandler, J., & Ipeirotis, P. G. (2010). Running experiments on amazon mechanical turk. *Judgment and Decision making, 5*(5), 411-419.

Philipp-Muller, A., Wallace, L. E., Sawicki, V., Patton, K. M., & Wegener, D. T. (2020). Understanding when similarity-induced affective attraction predicts willingness to affiliate: An attitude strength perspective. *Frontiers in psychology, 11*, 1919.

Pratto, F., & Bargh, J. A. (1991). Stereotyping based on apparently individuating information: Trait and global components of sex stereotypes under attention overload. *Journal of Experimental Social Psychology, 27*(1), 26-47.

Prentice, D. A., & Miller, D. T. (1996). Pluralistic ignorance and the perpetuation of social norms by unwitting actors. In *Advances in experimental social psychology* (Vol. 28, pp. 161-209). Academic Press.

Preston, S. D., & De Waal, F. B. (2002). Empathy: Its ultimate and proximate bases. *Behavioral and brain sciences, 25*(1), 1-20.

Reed II, A., & Aquino, K. F. (2003). Moral identity and the expanding circle of moral regard toward out-groups. *Journal of Personality and Social Psychology, 84*(6), 1270–1286.

Ross, L. (1977). The intuitive psychologist and his shortcomings: Distortions in the attribution process. In *Advances in experimental social psychology* (Vol. 10, pp. 173-220). Academic Press.

Rottman, J., Crimston, C. R., & Syropoulos, S. (2021). Tree-huggers versus human-

lovers: Anthropomorphism and dehumanization predict valuing nature over outgroups. *Cognitive Science, 45*(4), e12967.

Rubinstein, R. S., Jussim, L., & Stevens, S. T. (2018). Reliance on individuating information and stereotypes in implicit and explicit person perception. *Journal of Experimental Social Psychology, 75*, 54-70.

Sabucedo, J. M., Dono, M., Alzate, M., & Seoane, G. (2018). The importance of protesters' morals: Moral obligation as a key variable to understand collective action. *Frontiers in Psychology, 9*, 418.

Sargent, R. H., & Newman, L. S. (2021). Pluralistic ignorance research in psychology: A scoping review of topic and method variation and directions for future research. *Review of General Psychology, 25*(2), 163-184.

Singer, P. (1981). *The expanding circle*. Oxford: Clarendon Press.

Smith, T. W.,'& Frohm, K. D. (1985). What's so unhealthy about hostility? Construct validity and psychosocial correlates of the Cook and Medley Ho scale. *Health Psychology, 4*(6), 503.

Song, J. Y., Klebl, C., & Bastian, B. (2023). Awe promotes moral expansiveness via the small-self. *Frontiers in Psychology, 14*, 399.

Spadaro, G., Gangl, K., Van Prooijen, J. W., Van Lange, P. A., & Mosso, C. O. (2020). Enhancing feelings of security: How institutional trust promotes interpersonal trust. *PloS one, 15*(9), e0237934.

Sparkman, G., Geiger, N., & Weber, E. U. (2022). Americans experience a false social

reality by underestimating popular climate policy support by nearly half. *Nature Communications, 13*(1), 1-9.

Sproesser, G., Klusmann, V., Schupp, H. T., & Renner, B. (2017). Self-other differences in perceiving why people eat what they eat. *Frontiers in psychology*, *8*, 209.

Stavrova, O., & Ehlebracht, D. (2018). Education as an antidote to cynicism: A longitudinal investigation. *Social Psychological and Personality Science, 9*(1), 59-69.

Stavrova, O., & Ehlebracht, D. (2019). The cynical genius illusion: Exploring and debunking lay beliefs about cynicism and competence. *Personality and Social Psychology Bulletin, 45*(2), 254-269.

Tappin, B. M., & McKay, R. T. (2017). The illusion of moral superiority. *Social Psychological and Personality Science*, *8*(6), 623-631.

Van Boven, L., Dunning, D., & Loewenstein, G. (2000). Egocentric empathy gaps between owners and buyers: misperceptions of the endowment effect. *Journal of Personality and Social Psychology, 79*(1), 66.

Van Lange, P. A. (1992). Confidence in expectations: A test of the triangle hypothesis. *European Journal of Personality, 6*(5), 371-379.

Vazire, S., & Carlson, E. N. (2011). Others sometimes know us better than we know ourselves. *Current Directions in Psychological Science, 20*(2), 104-108.

Vazire, S., & Mehl, M. R. (2008). Knowing me, knowing you: the accuracy and unique predictive validity of self-ratings and other-ratings of daily behavior. *Journal of Personality and Social Psychology, 95*(5), 1202.

Wang, Y. A., & Todd, A. R. (2021). Evaluations of empathizers depend on the target of empathy. *Journal of Personality and Social Psychology, 121*(5), 1005–1028.

Waytz, A., Iyer, R., Young, L., Haidt, J., & Graham, J. (2019). Ideological differences in the expanse of the moral circle. *Nature Communications, 10*(1), 1-12.Wrightsman Jr, L. S. (1964). Measurement of philosophies of human nature. *Psychological Reports, 14*(3), 743-751.

Yu, Z., Guindani, M., Grieco, S. F., Chen, L., Holmes, T. C., & Xu, X. (2022). Beyond t test and ANOVA: applications of mixed-effects models for more rigorous statistical analysis in neuroscience research. *Neuron, 110*(1), 21-35.

Yudkin, D. A., Prosser, A. M., Heller, S. M., McRae, K., Chakroff, A., & Crockett, M. J. (2022). Prosocial correlates of transformative experiences at secular multi-day mass gatherings. *Nature communications, 13*(1), 2600.

Zhao, X., & Epley, N. (2022). Surprisingly happy to have helped: How underestimating prosociality creates a misplaced barrier to asking for help. *Psychological Science, 33*(10), 1708-1731.

Appendix A: Moral Expansiveness Measure Used in Studies 1a, 1b, 2, 3, & 4

Moral Circles

People sometimes talk about **'circles of moral concern'**. These circles are simple ways to make sense of the levels of moral consideration we have for different entities (e.g., people, animals, and the environment).

Where we place these entities within our moral circles is important and has direct consequences for how we treat them. For example, you might have close family or friends that are central to your moral world, this means you would be willing to make personal sacrifices for them. However, if we do not include an entity within our moral circles, this means we do not believe they are deserving of moral care and consideration, and wouldn't want to make sacrifices for them.

On the following tasks you are given the opportunity to organize a range of entities and place them **within your own moral circles** that reflect your individual views and feelings. You will also have the opportunity to reflect how entities fit **within other's moral circles**, as well as how **others think entities fit within your moral circles**.

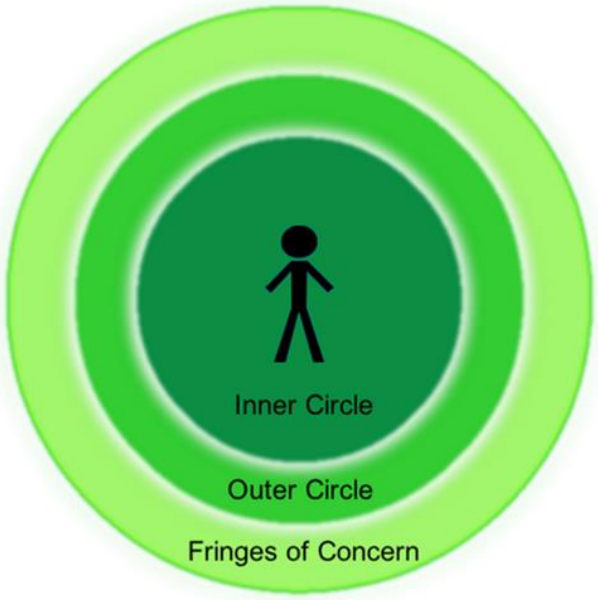

Inner Circle of Moral Concern: These entities deserve the **highest level of moral concern and standing.** You have a moral obligation to ensure their welfare and feel a sense of personal responsibility for their treatment.

Outer Circle of Moral Concern: These entities deserve **moderate moral concern and standing.** You are concerned about their moral treatment; however, your sense of obligation and personal responsibility is greatly reduced.

Fringes of Moral Concern: These entities deserve **minimal moral concern and standing**, but you are not morally obligated or personally responsible for their moral treatment.

Outside the Moral Boundary: These entities deserve **no moral concern or standing.** Feeling concern or personal responsibility for their moral treatment is extreme or nonsensical.

MES Entity List (United States)
Family/Friends
Family member
Close friend
Partner/spouse
In-group
American citizen
Somebody from your neighborhood
Co-worker
Out-group
Foreign citizen
Member of opposing political party
Somebody with different religious beliefs
Revered
U.S. President (position not specific individual)
U.S. Soldier
Charity worker
Stigmatized
Homosexual
Mentally challenged individual
Refugee
Villains
Murderer
Terrorist
Child molester
Animals high-sentient
Chimpanzee
Dolphin
Cow
Animals low-sentient
Chicken
Fish
Bee
Plants
Redwood tree
Apple tree
Rose bush
Environment
Coral reef
Old-growth forest
Grand Canyon National Park

Appendix B: Studies 2 & 3 Similarity & Liking manipulation materials

Next, you will be completing this version of the task from the perspective of a person who is close to you. To do this, first identify this person. It can be a friend, a family member, a partner, or anyone else who ’ou are close with. Once you've chosen the person, type their name below, and click continue.

Study 2

Now, think of (name of close-other). What makes you similar to / different from this person? Take a few minutes to write 2-3 sentences, identifying as many similarities / differences between you and (name of close-other) as possible, focusing on things such as personality, character, preferences, likes, opinions, etc.
Please write for at least 2 minutes. After 2 minutes have elapsed, you will be able to proceed by clicking continue.

Study 3

Now, think of (name of close-other). What makes you like / sometimes dislike or feel annoyed with this person? Take a few minutes to write 2-3 sentences, identifying as many reasons as possible as to why you like / occasionally dislike or feel annoyed with (name of close-other), focusing on things such as personality, character, preferences, likes, opinions, etc.
Please write for at least 2 minutes. After 2 minutes have elapsed, you will be able to proceed by clicking continue.

Appendix C: Study 4 Cynicism Manipulation Materials

Prompt given to participants in the Anti-Cynicism condition:

Please take the next 1-2 minutes to carefully read the following excerpt from a Psychology Today article. When you are ready to proceed, click on the "continue" button that will appear at the end.

Recent research trends show that human beings are inherently kind towards others, and that we often underestimate how much people care about each other and the world around them. Over the past decade, various researchers have conducted numerous studies on this topic, and the results are clear: people are wired to be cooperative.

A study[1] conducted by social psychologists showed that when asking people how kind, nice, honest or good others are, participants tended to underestimate what others generally report, even though they themselves reported being treated with kindness by others in their lives.
In another study[2] by researchers focusing on evolutionary dynamics, participants were asked to play a game where they could earn money by taking resources from others, or by sharing resources with others. The results showed that when participants were asked to make these decisions quickly, they gave away more resources instead of keeping them to themselves, indicating that our natural inclination is to be giving and helpful.

This cooperative behavior is not just constrained towards humans.

Another study[3] by environmental psychologists showed that even though 66-80% of Americans support policies aimed at mitigating climate change, the percentage of support Americans estimate is at 37–43%. This, combined with rapidly rising trends in plant-based diets and a heightened adoption of eco-friendly lifestyles suggest that humans are less self-centered than we have assumed.

So why do we often assume the worst in others? This may be due to the fact that we are often exposed to negative news and events, which can skew our perception of the world. However, as more research is conducted, it is becoming clear that people are not as selfish or cynical as we thought.

References to cited research:

[1] Mastroianni, A. M., & Gilbert, D. T. (2023). The illusion of moral decline. Nature, 1-8.

[2] Rand, D. G., Greene, J. D., & Nowak, M. A. (2012). Spontaneous giving and calculated greed. Nature, 489(7416), 427-430.

[3] Sparkman, G., Geiger, N., & Weber, E. U. (2022). Americans experience a false social reality by underestimating popular climate policy support by nearly half. Nature communications, 13(1), 4779.

Prompt given to participants in the Pro-Cynicism condition:

Please take the next 1-2 minutes to carefully read the following excerpt from a Psychology Today article. When you are ready to proceed, click on the "continue" button that will appear at the end.

Recent research trends show that human beings are inherently self-centered in relation to others, and that we often underestimate how much people care about themselves instead of the world around them. Over the past decade, social psychologists have conducted numerous studies on this topic, and the results are clear: people are wired to be competitive.

A study[1] conducted by developmental psychologists with infants as participants, showed that even at an early age, individuals may showcase signs of selfish inclinations by preferring to accept toys overwhelmingly from those that have favored them earlier in the experiment, or who otherwise look similar to them.

In another study[2] conducted by economic scientists, a behavioral economic experiment showed that participants often prioritized their own gain over collective welfare when they could see how much money they could earn by not cooperating with a fellow player, compared to when that information remained hidden.

This competitive behavior is not just constrained towards humans.
Another study[3] shows that even though 58% of Americans think that state officials are doing too little to address climate change, an even larger percentage of 68% believes that oil, coal, and natural gas should still continue to be used as main sources of energy. This, combined with rapidly rising trends in meat-based diets and a heightened adoption of consumer-focused lifestyles suggest that humans are more self-centered than we have assumed.

So why do we often assume the best in others? This may be due to the fact that we are often exposed to positive news and events, which can skew our perception of the world. However, as more research is conducted, it is becoming clear that people are just as selfish or as cynical as we thought.

References to cited research:

[1] Wynn, K., Bloom, P., Jordan, A., Marshall, J., & Sheskin, M. (2018). Not noble savages after all: Limits to early altruism. Current Directions in Psychological Science, 27(1), 3-8.

[2] Burton-Chellew, M. N., & Guérin, C. (2022). Self-interested learning is more important than fair-minded conditional cooperation in public-goods games. Evolutionary Human Sciences, 4, e46.

[3] Pew Research Center. (2023, August 9). Americans' views of climate change in 8 charts | Pew Research Center. https://www.pewresearch.org/short-reads/2023/08/09/what-the-data-says-about-americans- views-of-climate-change/

Appendix D: Cynicism & Fear of Compassion Measures Used in Studies 1-4

Fear of Compassion scale (Gilbert, Mc–wan, Matos, & Rivis, 2011) - Studies 1a, 1b, 2, & 3

1. Being too compassionate makes people soft and easy to take advantage of.
2. People will take advantage of you if you are too forgiving and compassionate.
3. I fear that being too compassionate makes people an easy target.
4. I fear that if I am compassionate, some people will become too dependent on me.
5. People will take advantage of me if they see me as too compassionate.
6. I worry that if I am compassionate, vulnerable people can be drawn to me and drain my emotional resources.
7. Being compassionate towards people who have done bad things is letting them off the hook.
8. There are some people in life who don't deserve compassion.
9. For some people I think discipline and proper punishments are more helpful than being compassionate to them.
10. People need to help themselves rather than waiting for others to help them.

Cynicism subscale (hostility subscale; Cook & Medley, 1954) - Studies 2 & 3

1. I think most people would lie to get ahead.
2. Most people are honest chiefly through fear of getting caught.
3. Most people will use somewhat unfair means to gain profit or any advantage rather than to lose it.
4. No one cares much what happens to you.
5. It is safer to trust nobody.
6. Most people make friends because friends are likely to be useful to them.
7. Most people inwardly dislike putting themselves out to help other people.
8. I commonly wonder what hidden reason another person may have for doing something nice to me.

Cynicism scale (adopted from Wrightsman 1964) - Studies 2 & 3

1. If most people could get into a movie without paying and be sure that they would not be seen, they would do it.
2. Most people would tell a lie if they could gain by it.
3. People claim that they have ethical standards regarding honesty and morality, but few people stick to them when the chips are down.
4. People pretend to care more about one another than they really do.
5. Most people are not really honest for a desirable reason; they are afraid of getting caught.

Cooperative vs. Competitive Primal World Beliefs (Clifton et al., 2019) - Studies 2, 3, & 4

1. Instead of being cooperative, life is a brutal contest where you got to do whatever it takes to survive.
2. For all life—from the smallest organisms, to plants, animals, and for people too—everything is a cut-throat competition.
3. Instead of being cooperative, the world is a cut-throat and competitive place.
4. The world runs on trust and cooperation way more than suspicion and competition. (R)

Appendix E: Supplementary analyses

Study 1a

Moral Expansiveness of Entity categories for Self vs. Others

In order to detect differences across the three perspectives for each entity category, we conducted repeated measures ANOVAs for all ten entity categories, comparing the self, others, and others for self perspectives. A repeated measures ANOVA with a Greenhouse-Geisser correction revealed that there were no significant differences between the three perspectives for friends and family, F (1.99, 368.07) = 0.15, p= .861, partial η^2 =.001, ingroups, F (1.98, 365.73) =0.65, p= .522, partial η^2 =.003, and revered entities, F (1.97, 363.62) = 0.18, p= .830, partial η^2=.001. Nevertheless, the repeated measures ANOVA was significant in almost all of the remaining entities categories, and pairwise comparisons between the three perspectives with a Bonferroni correction revealed a consistent difference between the moral expansiveness scores participants reported for themselves, and those they predicted for other people.

For stigmatized groups, F (1.91, 353.17) = 7.90, p<.001, partial η^2 =.041, out of the maximum possible 3 points, participants rated themselves as more morally inclusive compared to others (M_{diff} = 0.20, p= .002, 95% CI [.063, .335]), and this was also the case for outgroup entities, F (1.99, 367.74) = 4.44, p= .013, partial η^2 =.023, (M_{diff} = 0.14, p= .014, 95% CI [.022, .261]), animals of high sentience, F (1.93, 357.80) = 13.89, p<.001, partial η^2 =.070, (M_{diff} =0.28, p<.001, 95% CI [.142, .417]) animals of low sentience, F (1.98, 366.97) = 6.06, p= .003, partial η^2 =.032, (M_{diff} = 0.20, p= .003, 95% CI [.055, .336]), plants, F (1.99, 369.91) = 8.43, p<.001, partial η^2 =.044, (M_{diff} = 0.22, p<.001, 95% CI [.088, .342]), and the environment, F (1.97, 364.93) = 19.16, p<.001, partial η^2 =.094, (M_{diff} = 0.35, p<.001, 95% CI [.207, .495]). The repeated measures ANOVA for the villain category was significant, F (1.99, 368.39) = 3.36, p=.036, partial η^2 =.018, but the means had floor effects, as is typical for this category (Crimston et al., 2016). Although the direction of the means was the same as with the previous comparisons, there was no statistical difference between self and other for the villain category (M_{diff} = 0.09, p=.121, 95% CI [-.015, .191]).

Moral Expansiveness of Entity categories for Self vs. Others' predictions for self

The paired comparisons between the scores for participants own moral inclusivity and the one they thought others would predict for them (e.g. meta-perceptions about their moral inclusivity), however, is of interest, because it can give us a relative sense of whether participants accurately tap into the self-other discrepancy pattern (i.e. by predicting that others would see them as less morally inclusive than they actually are). The pairwise comparisons followed a similar, but weaker, pattern, such that people generally reported higher expansiveness scores than what they thought others would predict for them in the direction of most of the means, but the mean difference was only statistically significant for the entity categories with the largest self-other discrepancy effects: the environment, (M_{diff} = 0.19, p= .002, 95% CI [.207, .495]), and animals of high sentience (M_{diff} = 0.16, p=.004, 95% CI [.042, .277]), with trending effects in the expected direction for animals of low sentience (M_{diff} = 0.12, p= .076, 95% CI [-.009, .252],).

Study 1b

Moral Expansiveness of Entity categories for Self vs. Others

We repeated these comparisons for all ten entity categories across the three MES perspectives for Study 1b. As with Study 1a, there were no significant differences between the three perspectives for friends and family, F (1.75, 497.91) = 0.61, p=.520, partial η^2 =.002, and ingroups, F (1.98, 564.46) = 0.22, p= .800, partial η^2 =.001. Among revered entities, there were significant differences, F (1.99, 567.18) = 8.78, p<.001, partial η^2 =.030, in a direction opposite to the one that we most frequently found, such that participants reported others' inclusivity for revered entities as greater than their own (M_{diff} = -0.14, p<.001, 95% CI [-.219, -.053]). For the rest of the category comparisons between self and other, results were identical to Study 1a such that, out of the maximum possible 3 points an entity could receive, participants underestimated others' moral inclusivity relative to their own for outgroups, F (1.98, 563.60) = 9.21, p<.001, partial η^2 =.031, (M_{diff} = 0.18, p<.001, 95% CI [.063, .252]) and stigmatized human entities, F (1.94, 553.41) = 8.67, p= <.001, partial η^2 =.30, (M_{diff} = 0.14, p=.002, 95% CI [.041, .239]) compared to their own, as well as for the non-human categories of animals of high sentience, F (1.91, 544.99) = 20.27, p<.001, partial η^2 =.066, (M_{diff} = 0.27, p<.001, 95% CI [.161, .387]), animals of low sentience, F (1.94, 551.64) = 27.66, p<.001, partial η^2 =.088, (M_{diff} = 0.32, p<.001, 95% CI [.210, .438]), plants, F (1.86, 530.48) = 21.66, p<.001, partial η^2 =.071, (M_{diff} = 0.27, p<.001, 95% CI [.160, .379]), and the environment, F (1.92, 547.31) = 28.26, p<.001, partial η^2 =.090, (M_{diff} = 0.32, p<.001, 95% CI [.208, .430]). As with Study 1a, the villain category showed floor effects and had no significant differences between the self and other perspectives, F (1.98, 564.51) = 0.18, p= .832, partial η^2 =.001.

Moral Expansiveness of Entity categories for Self vs. Others' predictions for self

We repeated analyses of meta-perceptions for the self across the ten entity categories in Study 1b. The differences between the moral expansiveness scores for the self and the scores others would predict for the self in Study 1b appeared in the revered entity (M_{diff} =- 0.10, p=.011, 95% CI [-.219, -.053]) in the same direction as the self-other difference reported above, such that participants thought others would predict they would be more inclusive for revered entities than they actually were. As with Study 1a, differences between self and what others would predict for the self appeared consistently in all the non-human categories of animals of high sentience (M_{diff} = 0.13, p= .004, 95% CI [.032, .219]), animals of low sentience (M_{diff} = 0.15, p<.001, 95% CI [.054, .249]), plants (M_{diff} = 0.10, p= .018, 95% CI [.013, .185]), and the environment, (M_{diff} = 0.14, p=.001, 95% CI [.044, .229]).

MES Discrepancy Scores vs. Compassion Discrepancy Scores

To complete the development and investigation of the nature of the MES discrepancy score, we also wanted to compare and link it to other relevant constructs, given the existing multiple discrepancies in the perception of prosocial behavior that were discussed earlier in the introduction. Specifically, in both studies 1a and 1b, we asked participants to report how compassionate they thought they were, and how compassionate they thought others were, as well as how much they themselves feared compassion (based on fear of compassion scale from

Gilbert, McEwan, Matos, & Rivis, 2011).Exploratory pre-registered hypotheses and uses for these additional variables for Study 1b can be found at https://osf.io/kmqpd.

To that end, we created another discrepancy score using two additional items from Studies 1a and 1b that assessed self and other perceptions of compassion (*How compassionate do you consider yourself to be?* and *How compassionate do you consider others to be?*, the latter subtracted from the former), finding a similar discrepancy with the compassion perceptions. The MES discrepancy score and the Compassion discrepancy score (compassion self - compassion other) are, as expected, significantly positively correlated (r=.292, p<.001 in Study 1a, and r=.343, p<.001 in Study 1b), and the size of the correlation indicates that the two discrepancies are related but distinct variables.

In order to further understand how the MES discrepancy behaves as a variable in predicting compassionate outcomes, we report here its relationship with some additional dependent variables included in Studies 1a and 1b that referred to homeless targets whose pictures participants saw after completing the MES. After viewing these pictures, participants were asked, among others, how much compassion they were feeling in that moment, and subsequently how much they wished they could reduce the suffering of the homeless, and the extent to which they would support a move of the homeless in a shelter in their own neighborhood.

A multiple regression model predicting state compassion after viewing the homeless targets using the compassion discrepancy and the MES discrepancy as separate predictors revealed that both the compassion t =4.122, p<.001 and MES discrepancies t =3.069, p=.002 predicted a significant amount of variance in the model in Study 1a. In Study 1b the compassion discrepancy predicted state compassion for homeless targets t =7.854, p<.001, whereas the MES discrepancy did not t =1.812, p=.071.

However, both discrepancies predicted (to a smaller extent) the behavioral intention measure of wishing to reduce the homeless suffering in Study 1a (MES discrepancy, t =3.696, p<.001, Compassion discrepancy, t =2.215, p=.028) and in Study 1b (MES discrepancy, t =3.703, p<.001, Compassion discrepancy, t =3.648, p<.001). Supporting the move of the homeless in the participant's neighborhood was significantly predicted only by the MES discrepancy in Study 1a (MES discrepancy, t =2.765, p=.006, Compassion discrepancy, t =1.218, p=.225), and by both discrepancies in Study 1b (MES discrepancy, t =4.395, p<.001, Compassion discrepancy, t =2.618, p=.009).

MES Scores and Associations with Fear of Compassion and Moral Obligation Measures

Studies 1a and 1b did not directly assess cynicism as a trait measure, nor did they establish a link between cynicism as a construct and moral expansiveness. Although fearing compassion and being cynical are not the same thing, it is the measure included in Studies 1a and 1b that bears the closest resemblance to the construct of cynicism (defined as the belief that people are primarily motivated by self-interest). Prior work has found a negative correlation between fearing compassion and moral expansiveness (Crimston et al., 2021), as well as a negative correlation between generalized trust ratings and moral expansiveness (Kirkland et al., 2021), both of which are relevant (but not exact) measures of cynicism. In the following paragraphs, we report on the correlation between the fear of compassion scale (see Appendix D for full list of items) that was included in Studies 1a and 1b (see also Crimston et al., 2021) and the MES scores for self and other.

The fear of compassion scale (Gilbert, McEwan, Matos, & Rivis, 2011), while including some items on beliefs about others' compassion (e.g. "People will take advantage of you if you are too forgiving and compassionate"), it mostly assesses beliefs about the self being compassionate towards others (e.g. "*I fear that if I am compassionate, some people will become too dependent on me*", "*People will take advantage of me if they see me as too compassionate*"). Indeed, while we do not see any relationship between fear of compassion and the MES scores for the self and other perspectives in Study 1a (with MES_{self} r=. -111, p=.132, with MES_{others} r=.086, p=.243), when checking whether people differ in their scoring of MES across self and others based on their fear of compassion scores in Study 1a, we see a significant interaction, F (1, 184) = 9.34, p=.003, partial η^2 =.048. For those who considerably fear compassion (defined as +1SD above the mean fear of compassion score), there is no meaningful difference between the moral expansiveness they report for themselves (MES_{self}=43.08) and that of others (MES_{others}=41.28), (pairwise comparison p=.114). There is, however, a significant difference between self (MES_{self}=45.87) and others (MES_{others}=39.17) for those who do not fear compassion (defined as -1SD above the mean fear of compassion score), (pairwise comparison p<.001).

In Study 1b we see a correlation between fear of compassion with MES_{self} r=-.220, p<.001, but not with MES_{others} r=.054, p=.366. The significant interaction found in Study 1a replicates in Study 1b (p<.001), F (1, 284) = 26.68, p<.001, partial η^2 =.086, with the pattern in Study 1b additionally indicating that the difference between MES self and other for those who do not fear compassion is driven by one's own high moral expansiveness mean (MES_{self}=46.83) compared to what they think of others (MES_{others}=39.13), (pairwise comparison p<.001). For those who do fear compassion, on the other hand, their moral expansion is virtually indistinguishable from that they report for others (MES_{self}=41.01; MES_{others}=40.51), (pairwise comparison p=.616).

In other words, to the extent that fear of compassion is our closest (yet imperfect) proxy to cynicism, we see that those who score higher on fear of compassion actually show no self-other discrepancy on moral expansion, whereas those who do not fear compassion show a rather large gap. This seems to be primarily because, at least in the more highly powered Study 1b, fear of extending compassion seems to be associated with changes in one's own moral expansion (people who fear compassion do not extend their moral circle as much as people who don't) more so than what people predict for others' moral expansion (which doesn't seem to be moving around much based on fear of compassion). In a nutshell, this seems to indicate that the gap between self and others in MES scores might be due to motivational fluctuations of different kinds that seem to relate to the self, not so much to others.

Studies 2 and 3 replicated this pattern: those who were the least fearful of compassion showed a greater gap between their own MES scores, compared to those most fearful of compassion. Specifically, in Study 2, when using the fear of compassion (FC) as moderator of the difference between the MES score for self and perceived MES score for average others, the MES perspective x Fear of Compassion interaction was significant, F (1, 287) = 12.49, p<.001, partial η^2 =.042, replicating the pattern of the pairwise comparison that we saw in Studies 1a and 1b. This was the case with the Study 3 data as well - the MES perspective x Fear of Compassion moderation was significant, F (1, 419) = 9.46, p=.002, partial η^2 =.022. (See Supplementary Figure 1 for the plotted MES perspective x Fear of Compassion moderation across Studies 1a, 1b, 2, and 3).

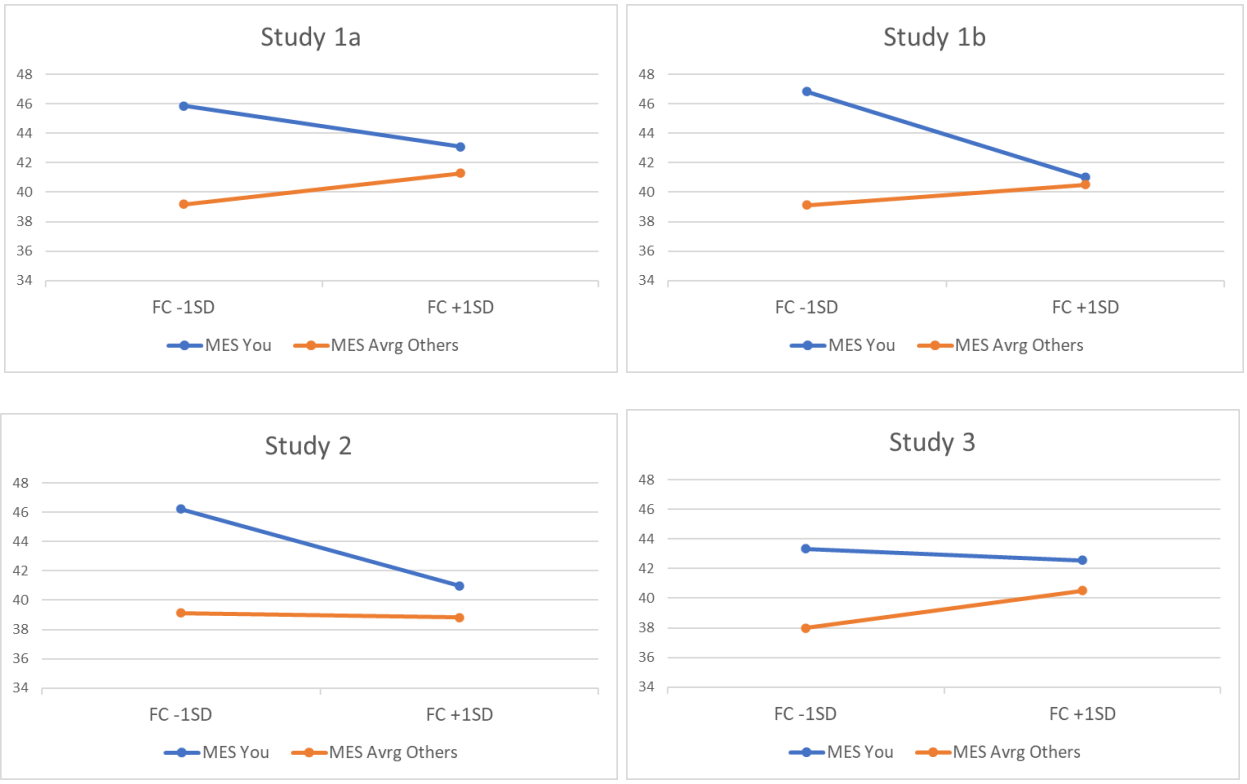

Supplementary Figure 1: Estimated marginal means (out of the maximum possible MES total of 90 points - Y axis) for the MES perspective (self, average others) x Fear of Compassion (+1SD vs. -1SD) interaction test across Studies 1a, 1b, 2, and 3.

This seems to also be the case for the moral obligation scale (adapted from Sabucedo, Dono, Alzate, & Seoane, 2018) we included in an exploratory fashion in Study 1b (example item: "Reducing the suffering of others constitutes a moral obligation to oneself"). Moral obligation correlates significantly with the moral expansiveness score for the self (r=.354, p<.001), and with how moral expansive others view the self (r=.233, p<.001), but there is no relationship between moral obligation and how moral expansive others are perceived to be (r=.083, p=.160). In addition, those scoring high vs. low on moral obligation differentially report a gap between the scores for MES self and MES others, as indicated by a significant interaction, F (1, 284) = 26.88, p<.001, partial η^2 =.086. Specifically, those scoring high on moral obligation have a higher discrepancy between their MES self score (MES_{self}=48.60) and that they report for others (MES_{others}=40.89) (pairwise comparison p<.001), compared to those who score low on moral obligation and who, instead, show comparable moral expansiveness between themselves (MES_{self}=39.23) and what their report for others (MES_{others}=38.75), (pairwise comparison p=.625). Therefore, we do see evidence of the predictive ability of the MES discrepancy between self and other for prosocial responses, and although the results for this are mixed across Studies 1a and 1b, there is some promise of unique predictive ability of the MES discrepancy as well (relative to the compassion discrepancy) for more behavioral intentions relevant to compassion (e.g. supporting the move of the homeless into one's neighborhood).

Thus, when consulting secondary data in Studies 1a and 1b, we see two relevant moderators of the MES - fear of compassion and moral obligation, such that those not fearing

compassion and feeling a higher sense of moral obligation to reduce suffering in the world, report being more morally expansive compared to others. Both the fear of compassion and moral obligation scale items have the self as the focus in nearly all items (e.g. “Reducing the suffering of others would make me feel proud of myself”) and thus the significant association with MES self (but not other) that we find in Study 1b makes methodological sense.

As mentioned above, those higher on fear of compassion tend to report lower moral expansiveness scores overall (descriptively in Study 1a, and significantly lower in Study 1b), replicating prior work (Crimston et al., 2021). This seems to eliminate the MES discrepancy due to a lower gap between a low MES score for self, and a comparably low score for others. In other words, the motivation driving one away from compassion and thus from expressing moral inclusion seems mostly relevant to the self MES score (significantly so in Study 1b), and not that of others. Thus, both the fear of compassion scale (Studies 1a and 1b) and the moral obligation scale (Study 1b) seem to matter for motivations that are relevant to self, but seem to be doing little to beliefs about others’ moral expansion.

We think that this might be due to at least two reasons. First, the issue might be methodological in nature - as mentioned above, both fear of compassion scale items and particularly the moral obligation scale items revolve around and explicitly mention the self, not others, which makes it harder to see an association between the measures and the MES other condition. Second, it is possible that people who do not fear compassion, and those who feel a heightened sense of moral obligation to reduce suffering in the world are not cynics per se, but earnest moralists, with wider moral circles than the average other.

Eliana Hadjiandreou
Vita

EDUCATION

May 2024 *(anticipated)*	Pennsylvania State University, University Park, PA Doctorate (PhD), Social Psychology, Advisor: Dr. Daryl Cameron
Sept 2022 - June 2023	Stanford University, Palo Alto, CA Visiting Fellowship, Stanford Social Neuroscience Lab, Advisor: Dr. Jamil Zaki
December 2020	Pennsylvania State University, University Park, PA Master of Science (M.S.) in Psychology, Advisor: Dr. Daryl Cameron
May 2016	Clark University, Worcester, MA Bachelor of Arts (B.A.), *Summa Cum Laude* with High Honors in Psychology

PUBLISHED PEER-REVIWED PAPERS

Hadjiandreou, E. & Cameron, C. D. (2022). Adversity-based identities drive social change. *Trends in Cognitive Sciences.* [access]

PUBLISHED BOOK CHAPTERS

Hadjiandreou, E. (2016). Average as optimum. In *Ethical Ripples of Creativity and Innovation* (pp. 174-181). Palgrave Macmillan, London. [access]

TEACHING ASSISTANSHIPS

Jan. 2022 – May 2022	Social and Personality Development, Instructor: Dr. Nhi Thai	Penn State
Aug. 2021 – Dec 2021	Basic Research Methods in Psychology, Instructor: Dr. Nicholas Pearson	Penn State
June 2021 – July 2021	Center for the Science of Moral Understanding Summer Series (virtual)	UNC Chapel Hill
Dec. 2020 – May 2021	Intro to Well-Being & Positive Psychology, Instructor: Dr. Nhi Thai	Penn State
Jan. 2020 – Dec 2020	Intro to Psychology, Instructor: Dr. Joshua Wede	Penn State
Jan. 2020 – May 2020	Moral Psychology, Instructor: Dr. Daryl Cameron	Penn State
Aug. 2019 – Dec. 2019	Intro to Social Psychology, Instructor: Dr. Nicholas Pearson	Penn State
Aug. 2013 – Dec. 2013	First-Year Psychology Seminar, Instructor: Dr. Michael Bamberg	Clark University

SERVICE TO THE FIELD

Ad-hoc review

Journal of Experimental Psychology: General
Journal of Social and Political Psychology
Emotion (co-review)
Current Psychology
Frontiers in Psychology (retracted invitation due to PhD requirement)
Clark Scholarly Undergraduate Research Journal