Title
How the Decision Environment Affects Choices and Judgements

Author
Fridman, Ariel

Publication Date
2023

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA SAN DIEGO

How the Decision Environment Affects Choices and Judgements

A Dissertation submitted in partial satisfaction of the requirements
for the degree Doctor of Philosophy

in

Management

by

Ariel Fridman

Committee in charge:

Professor Wendy Liu, Chair
Professor On Amir
Professor Ayelet Gneezy
Professor Uri Gneezy
Professor Piotr Winkielman

2023

The Dissertation of Ariel Fridman is approved, and it is acceptable in quality and form for publication on microfilm and electronically.

University of California San Diego

2023

DEDICATION

To my family, with love.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

It takes a village to raise a PhD student. This dissertation would not have been possible without the endless guidance and unconditional support of so many. I am deeply grateful to each of the professors, peers, co-authors, friends, and family that have encouraged me along this journey.

First and foremost, I would like to thank my advisor, Wendy Liu. Her inspiration and support brought me to the field of behavioral marketing where I have found an academic home, and she has been a fantastic mentor and advisor – teaching me, among other things, how to refine, test, and position ideas. I would not be where I am today without her.

On Amir has served as a pillar for me – beyond guiding me in the practice of research and serving as a co-author, he has been so generous in the amount of time he has spent mentoring me. I always leave On's office feeling more excited and clear-minded than when I walk in – a testament to his sharp wit and ability to offer clear, precise, and helpful guidance. Beyond our interactions, On has enthusiastically and energetically promoted our work to other researchers in the field, and generously shared his network with me. While I may not share his passion for science fiction, our shared interests in music, comedy, and cool R packages (like parallel processing!) have made it a sheer pleasure and honor to work with him.

Where would I be without the Gneezy's? Without Uri bringing me to UCSD, I'm really not sure. I feel so fortunate to have benefitted from Uri's wisdom and guidance over the years, and I always looked forward to our walks during his negotiation class, and hearing "another one bites the dust" at the end of each course. Ayelet is a gem – her passion, enthusiasm, and love of research are infectious. I'm continually impressed by Ayelet's ability to see and think creatively about the big picture, while simultaneously paying close attention

to details, like the precise phrasing and word choice in writing (e.g., going on "which" hunts). Beyond their brilliance, the kindness that both Ayelet and Uri have shown me over the years is humbling and I am forever grateful.

I am also deeply thankful to Piotr Winkielman – in addition to his deep knowledge of social psychology, which I was lucky enough to benefit from while taking class, his charm and humor have been an enduring source of levity.

In addition to my committee members, several other faculty members have been instrumental in my academic journey. When Rachel Gershon joined the faculty at UCSD, she graciously decided to work with me, despite the fact that we were both junior in the field. Nevertheless, our collaboration blossomed into multiple impactful publications – I'm still amazed by this! During the lockdowns of the COVID-19 pandemic, I always looked forward to our regular Zoom meetings, which provided not only means to advance our research but also a chance to connect with a friend.

Despite the fact that I am not a quantitative marketing student, Karsten Hansen kindly decided to collaborate closely with me on multiple projects, and I have learned so much about the practice of empirical research from him. Karsten has been so generous with his time – I was touched when he offered to meet with me regularly – and I always look forward to our interactions.

The first psychology class I ever took was with Craig McKenzie. I greatly enjoyed hearing the perspective Craig brought, and appreciated his approach to research. Besides being his student, I also learned so much from him by serving as a TA for the course in creativity and innovation that he created.

I also thank the current and former Rady faculty members who have provided me opportunities to learn and have given me valuable research advice. Thank you to Uma Karmarkar, Kanishka Misra, Kenneth Wilbur, Robert Sanders, Sally Sadoff, Marta Serra-Garcia, Charles Sprenger, Pamela Smith, Christopher Oveis, and Zhe Zhang. Finally, I am grateful for Dean Ordóñez, who has always been such a supportive leader.

Thank you to the current and former PhD students who have been supportive and encouraging colleagues and friends. I want to give a special shout out to Prabhanjan Didwania, Junxiong Gao, Jessica Kim, and Yida Peng who started the PhD program with me. It's been so amazing to all grow together. I also want to highlight the role Giulia Maimone has played as a constant source of optimism, cheer, and insight through all the highs and lows – having you as an officemate and friend has been a gift. Thank you to Meenakshi Balakrishna, Brianna Chew, Katie Hillegass, Olivia Jurkiewicz, Michelle Kim, Seung Kim, Xiaofeng Liu, Amanda Nachman, Kohei Hayashida, John-Henry Pezzuto, Carolina Raffaelli, Heather Romero-Kornblum, Anindo Sarkar, Gal Smitizsky, Shuang Wu, Paul Wynns, and Jean Zhang for support, advice on research, and for being cheerleaders throughout the journey. Lastly, thank you to former Rady PhD students – Kristen Duke and Alicea Lieberman – who have continued to be fantastic role models.

Finally, I would like to express my deepest appreciation and gratitude to my loving family for their unwavering support throughout my life. Words cannot capture how grateful I am to them. Their dedication has been the driving force behind my academic success, and I would not be where I am today without them. I want to thank my parents, Rachel and Moshe, and sister, Adie, for their unconditional love and support. My parents provided a foundational pillar that has been instrumental to my achievements, and my sister helped me recognize from an early age my love of teaching, often reminding me how I taught her to read (though I'm sure this isn't entirely true). I am also deeply grateful to my extended family – your words of encouragement, excitement, and cheer for this latest endeavor have meant the world to me. To my daughter Aviva, you have been my biggest source of inspiration and motivation, and I'm so excited to watch you grow. Lastly, I would like to express my heartfelt appreciation to my wife, Lotem, who has been by my side every step of this journey. From encouraging me to pursue a PhD and moving to San Diego with me, to countless conversations about research and edits to my manuscripts, her love, understanding, and patience have carried me through to this point and I am forever grateful. I thank you from the bottom of my heart.

Chapter 1, in full, has been submitted for publication of the material. Ariel Fridman, On Amir, and Karsten Hansen. The dissertation author was the primary investigator and author of this paper.

Chapter 2, in full, is a reprint of previously published material as it appears in the Proceedings of the National Academy of Science, 119(49), e2215633119, Rachel Gershon and Ariel Fridman. The dissertation author was one of the primary investigators and authors of this paper.

Chapter 3, in full, is a reprint of previously published material as it appears in Scientific Reports 12, 4886, Ariel Fridman, Rachel Gershon, and Ayelet Gneezy. The dissertation author was the primary investigator and author of this paper.

VITA

2015	Bachelor of Arts, Economics, University of California, Berkeley
2017-2023	Teaching Assistant, Rady School of Management, University of California San Diego
2023	Doctor of Philosophy, Management, University of California San Diego

PUBLICATIONS

Fridman, Ariel* and Rachel Gershon* (2023), "Reply to Kasper et al.: Identity threat is lower when supporting an opposing group member than an opposing group," Proceedings of the National Academy of Sciences, 120(12), e2301641120.

Gershon, Rachel* and **Ariel Fridman*** (2022), "Individuals prefer to harm their own group rather than help an opposing group," Proceedings of the National Academy of Sciences, 119(49), e2215633119.

Fridman, Ariel, Rachel Gershon, and Ayelet Gneezy (2022), "Increased Generosity under COVID-19 Threat," Scientific Reports, 12, 4886.

Fridman, Ariel, Rachel Gershon, and Ayelet Gneezy (2021), "COVID-19 and Vaccine Hesitancy: A Longitudinal Study," PLOS One, 16(4), e0250123.

*denotes equal contribution

ABSTRACT OF THE DISSERTATION

How the Decision Environment Affects Choices and Judgements

by

Ariel Fridman

Doctor of Philosophy in Management

University of California San Diego, 2023

Professor Wendy Liu, Chair

This dissertation comprises three papers examining how the environment or context in which a decision is made affects choices.

Chapter 1 investigates how the presence of seemingly irrelevant alternatives in an assortment can systematically affect people's choices. We analyze a large dataset of real-world purchase decisions in an online marketplace and find evidence for the asymmetric dominance effect, whereby the inclusion of an option that is inferior to another one in the assortment leads to a preference shift towards the superior option. This work sheds light on where, when, and why this effect occurs. We identify a novel moderator of the effect – the ability to sample or experience the options available – and find evidence consistent with a perceptual mechanism underlying the effect.

Chapter 2 explores how people make allocations between themselves and others in highly polarized environments (e.g., abortion, gun control, political parties) using a novel paradigm. In two large nationally representative samples, we asked participants to make lose-lose decisions: either subtract funds from their side of an issue or add funds to the opposing side. Strikingly, individuals were so averse to supporting opposing groups that they preferred to enact equivalent or greater financial harm to their own group instead – a preference that cannot be explained by existing theories. Instead, this work highlights the important role that identity plays in decision-making in polarized environments, underscoring the psychological barriers that impede the advancement of important causes.

Finally, chapter 3 reveals how financial generosity is affected by crises, specifically the COVID-19 pandemic. While psychological theories have supported the possibility that both increased selfishness and increased generosity may emerge in these situations, we find convergent evidence from both a dataset capturing real-world donations and a 6-month longitudinal study that individuals exhibited greater financial generosity when their county experienced COVID-19 threat – a silver lining amidst the tragedy of the pandemic.

Together, these three chapters demonstrate the important role that the environment of a choice can play in systematically affecting decision-making, in situations ranging from purchase decisions to charitable giving.

CHAPTER 1

DOMINANCE EFFECTS IN THE WILD

Ariel Fridman, On Amir, and Karsten Hansen

Rady School of Management, University of California, San Diego, La Jolla, CA, 92093, USA

Abstract

Dominance is the strongest form of preference relations that renders one alternative clearly preferred, and the other, well, irrelevant. An extensive literature has studied dominance effects using stylized lab experiments and found that, surprisingly, people's preferences among options can depend on the presence of irrelevant options in the choice set. However, null results in some recent lab studies and lack of real-world evidence call into question whether, when, and how the effect exists. We identify an important moderator for the dominance effect – preference uncertainty – and test it in both a real marketplace for digital freelance services and a lab experiment. Further, consistent with a perceptual mechanism for dominance effects, we also find evidence for several additional moderators that help explain how the effect works, such as the count of dominated alternatives and the magnitude of dominance. This work is the first to be able to use consequential field data to shed light on when and why dominance effects occur, with implications for marketers, choice architects, user interface designers, and policymakers.

Introduction

A substantial amount of research across a variety of fields, including economics, psychology, marketing, organizational behavior, and law, have demonstrated that choices are susceptible to contextual influences (Kahneman and Tversky 1984, Tetlock 1985, Tversky and Simonson 1993, Kelman et al. 1996, Thaler 1999, Mailath and Postlewaite 2003, Griffin et al. 2005). Context has been defined broadly as any factor that has the potential to shift the choice outcome by altering the process by which the decision is made (Thomadsen et al. 2018). One of the most studied context effects is the asymmetric dominance effect (also known as the attraction effect or decoy effect; here referred to simply as the "dominance effect"). It has been demonstrated in a wide variety of settings: choices for consumer goods (Huber et al. 1982, Huber and Puto 1983, see Heath and Chatterjee 1995 for a review), gambles (Huber et al. 1982), medical decisions (Schwartz and Chapman 1999), basic perceptual tasks (Trueblood et al. 2013), and animal behavior (Bateson et al. 2002, Shafir et al. 2002, Schuck-Paim et al. 2004, Lea and Ryan 2015). However, recent null findings in the literature (Frederick et al. 2014, Yang and Lynn 2014) call into question whether, when, and how the effect exists. We attempt to answer these questions, and also provide evidence for an important moderator for the effect – the ability to sample or experience alternatives. We explicitly test this moderator in a real-world dataset involving consequential choices from fiverr.com (for purchases of services like logo design, translation, etc.) and in a laboratory experiment. Our work is among the first to show that the dominance effect can indeed be observed outside the lab, and importantly, we also describe when and where these effects are likely to be found.

Dominance of one alternative over the other happens when the alternative is clearly better than the other on all attributes (strong dominance), or on some of its attributes but not weaker on any (weak dominance). As such, dominance is the purest and simplest of preference relations, and should render the dominated alternative irrelevant for choice. Research into the effects of context identified a violation of this simple logic in the case of asymmetrically dominated alternatives. Originally demonstrated by Huber et al. (1982), given a choice set of two alternatives, the addition of a third dominated alternative that is clearly inferior to one of the existing alternatives (but not the other) can result in a preference shift toward the alternative that dominates it. Consider the following real-world example involving collectible coins (Figure 1.1). While coins A, B, and C are different, coins A and A' are identical. However, the price of coin A' is greater, and the seller has a lower average rating and number of ratings. Therefore, the coin A' is dominated by coin A. While coin A' is effectively an irrelevant alternative in this assortment, the dominance effect suggests that its presence matters. Specifically, the effect occurs if coin A is more likely to be purchased when A' is present (vs. absent) in the assortment. The preference shifts induced by the dominance effect violate a classic tenet of rational decision theory – the Independence of Irrelevant Alternatives (IIA) axiom (Luce 1959) – which states that preference ranking between two alternatives does not depend on the inclusion or exclusion of additional alternatives. In the previous example, IIA suggests that the preference rankings between coins A, B, and C should not be influenced by the presence of coin A'.

Real-world Evidence

The question of whether the dominance effect matters and can be observed in real-world settings has become a topic of debate (Lichters et al. 2015). Some have argued that one reason for the limited real-world evidence is that dominated alternatives are rare in the real-world, perhaps because they would be quickly eliminated from the market, or because firms may not believe it is profitable to offer them (Wu and Cosguner 2020). In this work, we test for the dominance effect in a real and consequential setting – fiverr.com, an online marketplace for freelance digital services – and find that dominance relationships can actually be common, underscoring the need to better understand its relevance in the real world.

Several challenges emerge when testing for the dominance effect in the real-world, which the Fiverr setting is uniquely able to overcome. First, customers must be able to observe the dominated-dominating relationship between two products in order for the effect to arise. While detectability of such relationships is likely greater in stylized lab experiments, the real-world is usually more complex with a greater number of attributes and alternatives, and greater information-gathering costs. This added complexity can mask dominance relationships in the real world, where they are likely to go unnoticed because greater effort is required for customers to detect them (Huber et al. 1982, Huber et al. 2014, Simonson 2014). On Fiverr, the assortment of alternatives, along with their key vertical attributes, are displayed side-by-side on the webpage, allowing customers to easily compare them, facilitating the identification of dominance relationships.

Another difficulty in studying the dominance effect in the real world is that products often have horizontal attributes (e.g., brand, design, size, etc.) among which consumers have heterogenous preferences. When this is the case, products that are deemed dominated by some may not be seen as dominated by others, which poses a challenge if dominance relationships at the individual level are not observable to the researcher. For this reason, prior literature has typically replied on vertical attributes where preferences are homogenous (e.g., quality ratings, since higher quality is always preferred, all else equal). On Fiverr, sellers are freelancers rather than well-known brands, reducing horizontal differentiation. Furthermore, to the degree that horizontal differentiation between sellers exists on Fiverr, it would work against finding an effect, since the ‘dominated’ alternative may not always be viewed as dominated by the consumer.

Lastly, not all settings may give rise to the dominance effect due to a lack of preference uncertainty. Since the dominance effect has been theorized to arise from the preference construction process (Bettman et al. 1998), situations in which consumers already have pre-existing preferences (e.g., for repeat-purchase products), may not generate a dominance effect. On Fiverr, repeat purchases are rare, and the assortment of services shown can vary on each visit, so customers must construct their preferences from the assortment of alternatives they are presented with. Importantly though, we show that the degree of preference uncertainty customers on Fiverr face varies by product category – in some product categories customers are able to sample prior work from the seller, while in others they cannot – allowing us to test whether variation in preference uncertainty moderates the dominance effect.

To the best of our knowledge, only three papers attempted demonstrating the dominance effect using real-world data (Doyle et al. 1999, Geyskens et al. 2010, Wu and Cosguner 2020)[1], and we provide important contributions that build on prior work. Geyskens et al. (2010) studied a natural experiment that occurred with the introduction of economy and premium private labels in the corn flakes and canned soup categories in the United Kingdom. However, as the authors note, “some of the context effects work in opposite directions and cannot be separately identified,” which was the case for the dominance effect. Wu and Cosguner (2020) modeled the dominance effect in diamond sales of a leading online jewelry retailer. However, since they could only observe aggregated market-level data – they were not able to determine which alternatives each customer was actually presented with, and therefore the specific dominance relationships in the

[1] One additional paper testing whether the dominance effect can improve hand hygiene in real-world food-processing factories was retracted due to anomalies in the data (Li et al. 2019).

assortment. Lastly, Doyle et al. (1999) tested for the dominance effect among baked beans with an experiment run at a grocery store. While technically a real-world demonstration, this experiment provides lab-in-the-field evidence where a product was artificially introduced in an isolated setting. Our work provides several improvements over prior research: we are able to isolate the dominance effect from other context effects, our detailed data allows us to observe the specific assortment of options presented to each individual customer rather than aggregated data, and the data are generated from consequential choices across many product categories from customers around the world. These features allow for a cleaner test of the following hypothesis:

H1: The dominance effect is significant in real-world settings.

The Role of Preference Uncertainty

Why might the dominance effect occur? Prior work has implicated the process by which consumers construct preferences from a given assortment (Bettman et al. 1998). For example, as proposed by Bettman et al. (1998), salient aspects of the choice task, such as dominance relationships, make choice heuristics based on relational properties of the options more accessible and therefore more likely to be used in decision-making. When dominance relationships are detected, their presence may activate the use of decision strategies that rely on such relationships between alternatives (e.g., componential context model proposed by Tversky and Simonson 1993, see also Kivetz et al. 2004). Using relationships between options as a basis for choice not only minimizes effort (since difficult tradeoffs between attributes are avoided (Shugan 1980)), but is also appealing since it provides a readily available justification for the choice – the fact that one option dominates another is a good reason to choose it (see Simonson (1989)). Several other non-cognitive theories have been proposed to explain why the dominating option appears better. For example, the dominating alternative can perceptually "seem better"

(Pocheptsova et al. 2009), or the similarity between the dominated and dominating alternatives (based on the Gestalt principle of grouping by similarity) can lead to a direct comparative process between them[2].

All these explanations share an important feature: They presume that the decision maker does not have clear pre-existing preferences over the alternatives, giving room for features of the context to shape their preferences[3]. For example, if individuals have pre-existing preferences between apples and oranges, then it would be unlikely to observe preference reversals when adding dominated options (e.g., a moldy orange or bruised apple) to the assortment. A similar experiment described in Frederick et al. (2014) found null results, and it is possible that null findings recently demonstrated in literature (e.g., Frederick et al. 2014, Yang and Lynn 2014) may be due to a lack of preference uncertainty, since it diminishes the ability of dominance relationships to influence preferences. This contention has also been put forth by the original authors of the dominance effect in response to the null findings in the literature (see discussion on "strong prior tradeoffs" in Huber et al. (2014)).

The Fiverr setting allows us to shed light on the proposed moderating role of preference uncertainty on the dominance effect. While each customer's degree of preference uncertainty cannot be observed, we can observe other variation associated with preference uncertainty: the degree to which customers can sample and experience prior work produced by each seller. Due to the standardized template that Fiverr uses for sellers in each product category, some are

[2] Other explanations for the dominance effect have also been proposed. For example, it can increase the decision weight consumers place on the attribute that differentiates similar products (Ariely and Wallsten 1995), it can shift the sensitivity associated with a given change along a particular attribute (i.e., range-based relative thinking; Bushong et al. 2021), or it can result from a neuroscientific process in which relative values are encoded (e.g., pairwise-normalization; Landry and Webb 2021).
[3] Two notable exceptions are Wernerfelt (1995) and Prelec et al. (1997), who argue that the effect can in part be explained by a rational model wherein consumers develop inferences about the location of their ideal points based on the stimuli and their relationships to one another.

experienceable, enabling sellers to showcase prior work they have done (e.g., logo design, voice over, whiteboard & animated explainers, etc.) while others are non-experienceable (e.g., translation, articles & blog posts, social media marketing, etc.). For example, when customers browse a product category, those in logo design can view past logos created by each seller (e.g., see Figure 1.2), in voice over they can listen to the voice of each seller by clicking a ‘play’ button, and in whiteboard & animated explainers can play a sample video from the seller. In a lab experiment mimicking the setup of Fiverr, we found that participants reported greater preference uncertainty, on average, in product categories where they were not able to experience prior work from the seller.

Why does the ability to experience the alternatives reduce preference uncertainty? Prior work has demonstrated the engaging nature of experience (Hoch 2002), which also amplifies differences among alternatives making the options more distinct and the choice easier (Hertwig and Pleskac 2010), especially when the experience is diagnostic. Other research in psychology (Fazio and Zanna 1978) and marketing (Smith and Swinyard 1983) finds that consumers place greater confidence and attitude-behavior consistency to judgments based on direct experience.

We propose that product experience, through its ability to reduce preference uncertainty, can attenuate the dominance effect. In fact, prior work has demonstrated a link between experience and the dominance effect in the context of lotteries: the ability to experience the payoff distribution of a lottery was associated with an attenuation of the dominance effect (Ert and Lejarraga 2018, Hadar et al. 2018). More formally, we test the following:

H2: The dominance effect is moderated by experienceability.

An alternative possibility is that rather than attenuating the dominance effect via its impact on preference uncertainty, experience can distract customers from noticing dominance relationships created by the other attributes, or simply eliminate the dominance relationships altogether. However, if this were the case, then a consequence is that dominated alternatives would not be avoided when products are experienceable – a possibility we investigate and rule out based on our results[4].

Underlying Mechanism

The vast amount of variation in the Fiverr data provides the ability to test key moderators that shed additional light on how the effect works. Specifically, we contrast two mechanisms that prior literature has argued underlie the dominance effect – one based on a Gestalt principle of grouping by similarity, where attention is drawn to similar groups of alternatives, and another that argues that the effect has a perceptual basis, where the dominating alternative appears to look better. Importantly, these mechanisms make opposite predictions for two of the moderators we explore: the count of dominated gigs in the assortment, and the magnitude of dominance.

The Gestalt account suggests that people group alternatives based on their similarity in attribute space, and that the similarity between dominated and dominating alternatives drives a direct pairwise comparison between them. But as a grouping becomes larger, for example through the addition of more dominated alternatives, it may make other, dissimilar, alternatives stand out more, receiving additional attention. In fact, prior work has found that adding additional dominated alternatives can make a dissimilar alternative a focal option, reducing the dominance effect (Hamilton et al. 2007) or even reversing it (Aaker 1991). Another implication of the Gestalt account is that reducing the similarity between the dominated and dominating options would reduce the likelihood that they are directly compared, attenuating the effect. For

[4] Another reason why it is unlikely that experience eliminates dominance relationships is that in our setting, the quality of a product's experience is likely correlated with the numeric attributes used to evaluate dominance (i.e., average ratings, number of ratings, and price).

example, as a dominated alternative becomes more inferior on one of its attributes, thereby decreasing its similarity with the dominating alternative and hence direct comparisons between them, the dominance effect should be attenuated. Some experimental work has found evidence consistent with this prediction (Wedell 1991, Wedell and Pettibone 1996), and computational models have also supported this possibility, though it depends on specific parameter values (Trueblood et al. 2014).

The perceptual account suggests that the dominance effect arises from simple, automatic processing where the dominating alternative just "seems better" in the presence of the dominated alternative. When Simonson (1989) asked participants to justify their choice of the dominating option, few participants used the relative advantage of the dominating option over the dominated option as a basis for their choice, and instead simply justified their choices by focusing on the option's attractiveness on one dimension. In other words, they were "typically unaware that the dominance relationship affects their perceptions of the options' attractiveness" and that the dominance relationship "creates the illusion" that the dominating option is preferred because of its attribute values (Dhar and Simonson 2003). By contrast, for the compromise effect, the choice of a middle option was justified on basis of its compromise position, indicating a more cognitive rather than perceptual mechanism for this context effect. Reinforcing these finding, Pocheptsova et al. (2009) found that when participants completed a resources depletion task impairing deliberate, careful processing, the dominance effect was magnified, though the compromise effect was attenuated. The perceptual mechanism suggests that as the number of dominated alternatives increases, it reinforces the perceptual superiority of the dominating alternative, strengthening the effect[5]. This account also suggests that a larger difference in magnitude

[5] A possible confound is that increasing the number of dominated alternatives increases the probability of dominance detection. To account for this possibility, we include and analysis that focuses on cases where the

between the dominated and dominating alternatives should be positively associated with the dominance effect size, since it enhances the perceptual contrast between the alternatives[6]. Recent experimental work (Padamwar et al. 2021) has found support for this prediction – range extensions amplified the dominance effect.

Formally, we test the following hypothesis:

H3: As the count of dominated alternatives increases, or the magnitude of dominance increases, the dominance effect is strengthened, consistent with a perceptual mechanism.

Current Research

We study the effect of dominance in a real and consequential setting – a marketplace for freelance digital services (fiverr.com). Our dataset allows us to observe, at the individual level, the assortment of sellers offering services (referred to as gigs) that customers are presented with when browsing the website. Critically, we also observe the numeric attributes of the gigs, which gigs customers clicked on and which were subsequently purchased, as well as various seller and customer characteristics.

This setting offers several advantages for studying the dominance effect in the real world. The Fiverr search page displays several key gig attributes (e.g., price, average rating, number of ratings, etc.; see Figure 1.2 for example) from which dominance relationships can be identified. The gigs comprising the assortment are updated at regular intervals based on a quasi-random process independent of a gig's dominance status, thereby creating variation in whether a given gig dominates another gig(s), is dominated by another gig(s), is both dominating and dominated, or is neither dominating nor dominated in every assortment of gigs displayed. In addition, the

dominating alternative is always directly adjacent to at least one dominated alternative, in an attempt to hold constant the probability of dominance detection.

[6] Relatedly, neuroeconomic range normalization models (Landry and Webb 2021) predict that the dominance effect should strengthen as the dominated alternative worsens.

ordering of the gigs on the first page of browse results are randomized; thus, the dominance status of a gig and its position are independent.

Another advantage of using the Fiverr dataset to study context effects is that the purchases are consequential. Since 80% of buyers using Fiverr at the time of this data are small or medium-sized businesses (SMBs), the purchases they make are of important financial relevance[7].

An additional key feature of the Fiverr platform is the exogenous variation in experiecneability, allowing us to test H2 using the Fiverr data. We further conducted a laboratory experiment to provide additional causal evidence for the effect of experienceability on the dominance effect.

Data

We obtained data for this study from Fiverr, one of the most popular online marketplaces for digital services (Similarweb 2022). The data contain the assortments of all search results that occurred on the platform, as well as which gigs were clicked on and which were purchased, from October 1, 2018 to October 2, 2019, as well as various buyer and seller characteristics captured by the platform. Specifically, the search results comprise of an assortment of freelance services offered (i.e., gigs) in a particular category, and the dataset captures which gigs appeared, their attributes at the time of the search, as well as the location on the screen where a gig is displayed vis-à-vis others (see Table A1.1 for a list of all variables and descriptions). The Fiverr marketplace offers services in a variety of categories (see Table A1.2 for a complete list), and prices range from $5 to $20,000, with over 400,000 unique gigs offered. Buyers and sellers are

[7] Since firms have been argued to act more rationally (see "Becker Conjecture" Thaler 2015) than consumers who may be uninformed or unincentivized, we view the Fiverr marketplace as a conservative setting to test for the dominance effect.

able to communicate through the platform, and Fiverr handles the financial transactions between them. Prices for the gigs are set by the seller, and Fiverr's revenue is generated from transaction and service fees (Cuofano 2019).

On the Fiverr website, users have two primary ways to initiate a search: they can either use the search bar where they can enter text to search ("search-bar-based searches"), or by clicking on a particular category of interest ("category-based searches", e.g., Website Design, Logo Design, Translation, etc.). We focused our analysis on category-based searches (which accounted for 37.2% of searches on the platform) without any search filters applied (32.5% of category-based searches) since they offer a key advantage – the assortment of gigs which appear on the first page are updated periodically in a quasi-random process, and the ordering of the assortment of 12 gigs that appear on the first page of results were randomized for every search (see Figure A1.1 for a randomization check). By contrast, the assortment composition and ordering of search results in search-bar-based searches was not randomized. Since gigs that appear in higher positions[8] in the assortment are more likely to be clicked on and purchased (see Figure A1.2), randomization in gig ordering is crucial to separate the effect of gig characteristics from position effects. We focused only on the assortment of gigs that appear on the first page due to the randomization since its composition. Furthermore, most clicks and purchases on the platform occur on the first page of results. Since mobile users of the website view a different layout and assortment size, we limited the analysis to searches that occurred on non-mobile platforms (89.2% of observations). An additional 12.9% of observations were removed due to incomplete data[9]. Lastly, we limited the analysis to assortments where 12 gigs were displayed

[8] Positions are numbered left to right, starting from top to bottom.
[9] Reasons for incomplete data were: missing an experienceability categorization (4.9% of observations), missing data for average rating or number of ratings (1.0%), or assortments with gaps created by missing positions (7.6% of

(76.9% of observations). Our resulting dataset, which was used for all analyses, consists of 51,609,096 observations, where each observation captures an impression of a gig in an assortment (4,300,758 assortments).

The distribution of average rating and number of ratings on the Fiverr platform are both highly skewed. 92% of gigs had an average rating between 4.8 and 5.0 (see Figure 1.3 for distribution of average rating). Number of ratings was top-coded at 1,000, and 17% of gigs had at least 1,000 ratings (see Figure 1.4 for distribution of number of ratings).

Gigs are priced in increments of $5, and the minimum gig price is $5. The modal price is also $5, the median price is $25, the mean price is $82, and the maximum price is $20,000. The standard deviation of prices is $328, and the interquartile range is $10-$55. In our dataset, we only observed the prices that were displayed to customers on the browse page (e.g., the price in Figure 1.2); it is possible that the final price customers paid may differ due to the specific packages they purchase. See Figure A1.3 for the distribution of prices.

Experienceability and Preference Uncertainty

Gigs on Fiverr are divided into product categories, and each category was exogenously classified by Fiverr according to experienceability (see Table A1.2 for a list of categories and their experienceability classification). 61 categories were classified as experienceable (comprising 74% of purchases), and 62 categories were classified as non-experienceable (comprising 26% of purchases). In experienceable categories buyers have the ability to sample a seller's work, while in a non-experienceable categories the buyer did not have this ability, even when clicking into a gig's page since the standardized template on Fiverr did not offer sellers in these categories a way to showcase their work.

---

observations). Based on discussions with Fiverr, missing gig data was due to rows missing from their database, which appeared to be at random.

To test whether customers face different degrees of preference uncertainty across experienceable and non-experienceable product categories, we conducted a pre-registered experiment mimicking the layout and design of Fiverr for the six most popular product categories (see Figure A1.4 for a screenshot and Appendix Study A1.1 for more details). After being presented with an assortment of twelve gigs from a randomly selected product category, participants were asked which one they would purchase. On average, participants reported greater preference certainty associated with their choice in the three experienceable categories (logo design, voice over, and whiteboard & animated explainers) than in the three non-experienceable categories (translation, articles & blog posts, and social media marketing), $t(344) = 2.10$, $p = .036$. This finding supports the theorized claim that experiencing samples from a seller reduces preferences uncertainty.

Methods

Dominance Definition

In our primary analysis, we defined a gig as dominated if there exists another gig in the assortment that has a higher average rating and greater number of ratings (or at least one of the two attributes is greater and the other equal in the case of weak dominance), offered at the same price. In other words, a given gig i is classified as dominated if there exists another gig j offered at the same price as gig i, and in the same assortment as gig i, where the following statement is true, and at least one of the inequalities is strict:

$$(number\_ratings_i \leq number\_ratings_j) \ \& \ (average\_rating_i \leq average\_rating_j)$$

Similarly, gig i would be classified as dominating if the following statement is true, and at least one of the inequalities is strict:

$$(number\_ratings_i \geq number\_ratings_j) \ \& \ (average\_ating_i \geq average\_rating_j)$$

Under this definition, it is possible for a gig to be both dominated and dominating – a situation that can happen in the real world. For example, consider three gigs, A, B, and C, all offered at the same price. Gig A has a 4.7 average rating, and 100 ratings, Gig B has a 4.8 average rating with 200 ratings, and Gig C has a 4.9 average rating with 300 ratings. In such an assortment, gig B would be classified as both dominating (since it dominates gig A) and dominated (since it is dominated by gig C). In our data, we classified all gigs in every assortment into one of four dominance categories: "dominated only", "dominating only", "both" (if they are both dominated and dominating), or "neither" (if they have no dominance relationship with another gig in the assortment). The distribution of dominance statuses is similar across experienceable and non-experienceable categories (see Figure A1.5). 93.1% of assortments had at least one dominance relationship – indicating that such relationships can be prevalent in the real world.

We excluded gigs that had no ratings (*number_ratings*=0; 4% of observations) when assessing dominance relationships and classified their dominance status as "neither" since it is possible that customers may not perceive low-rated gigs to dominate those with no ratings (gigs with no ratings may be new sellers).

Our primary methodology for assessing dominance focused on dominance relationships created by the average rating and number of ratings attributes, among equally-priced gigs in the same assortment. While this definition assumes that customers believe 'more is better' for both average rating and number of ratings, they need not place equal importance on them, and the importance can even interact (so long as more is always preferred)[10]. We adopted this narrow

[10] It is possible that among gigs with a low average rating, the 'more is better' assumption for number of ratings does not hold. Fortunately, gigs that have low average rating are rare – excluding gigs with 0 ratings (which are omitted when classifying dominance relationships) only about 1.9% of gigs have an average rating of less than 4.7 (out of 5).

definition, which holds price constant, because customers may infer a gig's quality from its price, and therefore view gigs with different prices as being categorically different, distorting comparisons between them. For example, if customers ratings are given conditional on the price at which a product was purchased, then a gig with an average rating of 4.9 for $5 may actually be *lower* quality than a gig with an average rating of 4.8 for $50, despite the higher rating. See Huber et al. (2014) for a discussion about how quality inferences from prices may distort the dominance effect. Another reason is that if customers are usually only interested in gigs within a relatively narrow price band, they may not make comparisons across gigs of sufficiently different prices. However, as a robustness check, we also repeat our analysis using an expanded definition of dominance that considered dominance within price bands.

While gigs have numerous horizontal attributes as well, we consider these to be random noise for the purposes of dominance classification, which would work against finding effects. We also note that our modeling approach holds constant all of a gig's attributes, so those cannot explain the observed differences. Additionally, while the search page shows a gig's starting price (which was used to classify dominance relationships), clicking on the gig's page may reveal a menu of prices for various options or packages. We consider this to be noise as well, though we also include a robustness check using clicks as an outcome, which would not be affected by this issue.

Modeling Approach

To estimate the causal impact of a gig's dominance status on purchases, we attempted to mimic an ideal experimental design. In such an experiment, one would hold a focal gig constant, but vary the assortment around it, so that in one condition it is dominating, in another condition it is dominated, and a control condition when it is neither dominating nor dominated. That way, any differences in the purchase probability of the focal gig would only stem from variation in its dominance status (see Figure 1.5 for conceptual examples). In our empirical approach, we accomplish this by including gig-level fixed effects, which allows us to identify how changes in a given gig's dominance status impact purchase probability. Specifically, to account for any changes in a gig's characteristics over time, we assigned a unique ID, *i*, to each combination of gig, price, average rating, number of ratings, type, seller level, and category (see Table A1.1 for variable descriptions). The following logistic regression equation was estimated using the "fixest" package version 0.9.0 (Bergé 2018) in R version 4.1.1 (R Core Team 2021):

$$\begin{aligned} purchase_{ia} = {} & \beta_d^E(dominance\_status_{ia} \text{ x } 1_{s \in Experienceable}) \\ & + \beta_d^{NE}(dominance\_status_{ia} \text{ x } 1_{s \in Non-Experienceable}) + \alpha_i + \alpha_t + \alpha_p + \alpha_c \\ & + \alpha_n + X_a\delta_a + \varepsilon_{ia} \end{aligned}$$

In the regression equation, $purchase_{ia}$ captures whether the gig with unique ID *i* in assortment *a* was purchased. The coefficients of interest, β_d^E and β_d^{NE}, capture the effect of changes in gig *i*'s dominance status (relative to the omitted 'neither dominating nor dominated' dominance status), when the assortment consisted of gigs in an experienceable or non-experienceable category, respectively. Fixed effects were included for gig (α_i), date (α_t), gig *i*'s position in the assortment (α_p), country of the buyer (α_c), and the count of gigs at the same price as gig *i* (α_n). The variables X_a are controls at the assortment level, to hold constant assortment and buyer characteristics (e.g., median price of the assortment, standard deviation of the prices in the assortment, etc.; see Table A1.3 for complete list)[11]. Thus, identification comes from

[11] To check for multicollinearity between the independent variables, we ran a correlation table including all the numeric variables (see Table A1.4). We found low correlations between the 'dominating only' variable – the key independent variable capturing the dominance effect – and the other variables. All were smaller than 0.1 in magnitude except its correlation with other dominance status variables (by construction) and the control variable capturing the number of gigs at the same price (which also occurs by construction, since dominance was defined within price).

variation in a gig's dominance status, since the fixed effects hold constant the gig's characteristics, and the control variables hold constant other assortment characteristics[12]. We also clustered the standard errors by assortment ID to account for violations of independence of purchases within a given assortment.

Main Results

Coefficients from the logistic regressions were converted to odds ratios. Odds ratios that are greater than one indicate an increase in the odds of purchasing a gig, while odds ratios less than one indicate a reduction in the odds of purchasing a gig, relative to the case when a gig is neither dominating nor dominated (the omitted reference category). While the results should strictly be interpreted as odds ratios, they can be interpreted approximately as risk ratios, since purchases are rare events in this setting (0.02% of observations, overall).

Baseline Results

In non-experienceable categories, the odds that a gig was purchased were 21% greater when it appeared in an assortment where it was dominating only, compared to when it was neither dominating nor dominated (OR = 1.21, $p < .001$). This finding provides evidence that the dominance effect can be found in the real world (H1). In experienceable categories, the effect was virtually eliminated – the odds that a gig was purchased was not statistically different when it appeared in an assortment where it was dominating only, compared to when it was neither dominating nor dominated (OR = 0.99, $p = 0.754$) [13]. The difference in the dominance effect across non-experienceable and experienceable categories ($p < .001$) provides evidence for

[12] To ensure that the random variation identifying the coefficients of interest is similar across experienceable and non-experienceable categories, Figure A1.6 plots the count of the unique dominance statuses for each gig by experienceability.

[13] Relative to being dominated only, the odds a gig was purchased when it was dominating only was 66% greater (OR = 1.66, $p < .001$) in non-experienceable categories, and 15% greater (OR = 1.15, $p < .001$) in experienceable categories.

moderation by experienceability (H2). As expected, when a gig is dominated it is significantly less likely to be purchased in both non-experienceable and experienceable categories. See Figure 1.6 with the dominance effect labeled, and Table A1.5 for regression results. An additional specification excluding all the control variables and fixed effects except those for gig revealed similar results (see Table A1.5 model 2).

Additional Moderators

To better understand the psychology underlying the dominance effect, we tested several additional moderators that help tease apart distinct mechanisms that give rise to the effect. One explanation suggests that the effect is perceptual, based on automatic or less vigilant processing (Malkoc et al. 2013), where the dominating alternative just "seems better" (Pocheptsova et al. 2009). Another mechanism we explore is based on the Gestalt principle of grouping by similarity, where attention is drawn to similar groups of alternatives. The similarity between the dominated and dominating alternatives leads to a pairwise comparative process between them, giving rise to the effect. The primary dominance definition we used makes the similarity particularly strong – the alternatives must be the same price, and we only enforce weak dominance on the average rating and number of ratings attributes, meaning one of them may be equal as well. These mechanisms make opposite predictions for two of the moderators we explore: the count of dominated gigs in the assortment, and the magnitude of dominance. In addition, utilizing the randomization in the ordinal position in which gigs appear on the screen, we test whether the distance between the dominated and dominating gigs moderates the effect.

Methods for Moderators

Count of dominated gigs. Our setting provides random variation in the count of dominated gigs there are relative to a given dominating gig. However, due to the concern that the count of dominated gigs in the assortment is correlated with the ordinal distance between a dominating gig and nearest dominated gig relative to it, we analyzed both of these moderators together in the same model in order to separate their effects.

Ordinal Distance. We exploit the randomization in gig position[14], which creates random variation in the visual proximity between the dominating gig and nearest dominated gig, to test how distance moderates the effect. We hypothesized that customers would be more likely to notice the dominance relationship when the gigs were positioned closer together, particularly when a dominating gig is associated with only a few dominated gigs in the assortment. To assess this relationship, we created two measures of distance between the dominating and nearest dominated alternative: (1) the position distance (the absolute value of the difference in position, numbered starting at the top left, moving from left to right and top to bottom; range = 1-11), (2) the Euclidean distance (range = 1-3.61). Euclidean distance was binned according to the following cutoffs: 1 (bin 1), ≤ 2 (bin 2), ≤ 3 (bin 3), or > 3 (bin 4). Note that the position distance measure would assign a distance of 1 between positions 4 and 5 despite the fact that position 5 appears on subsequent row, while the Euclidean distance between these positions is 3.16.

To test both the effect of ordinal distance and the count of dominated gigs, we added a three-way interaction between the dominating only variable, the count of gigs it dominates, and the distance measure, separately for experienceable and non-experienceable categories. We ran separate regressions for each of the two distance measures.

Magnitude of dominance. Prior research has studied how the placement of the dominated option in attribute space, relative to those of the other alternatives in the assortment moderates the dominance effect, and has found mixed results. Specifically, as a dominated alternative

[14] See Figure A1.1 for a randomization check.

becomes more inferior on one of its attributes, thereby increasing the range of that attribute in the assortment, it is unclear whether and how this change affects the strength of the dominance effect. The original studies on the dominance effect examined this question (Huber et al. 1982, Huber and Puto 1983), but did not find a significant relationship between range extension and the effect size[15].

To provide the cleanest test of magnitude effects, we focused only on cases where there was one dominated gig for each dominating one (56% of cases). The magnitude of dominance was calculated as the difference in the average rating and number of ratings between the dominating and dominated gigs. Since the distributions of the differences for both attributes were very skewed (see Figures 1.2 and 1.3 for the distributions of each attribute), they were logged. In our model, we included a three-way interaction term between the dominating only variable, the difference in average rating, and the difference in number of ratings, separately for experienceable and non-experienceable categories.

Results for Moderators

Count of dominated gigs and ordinal distance. Both the position distance and Euclidean distance models demonstrated similar pattern of results. Among non-experienceable categories, increasing the count of dominated gigs relative to a given dominating gig strengthened the dominance effect (position distance: $p < .001$; Euclidean distance: $p = .039$; see Figure 1.7 for a heatmap of the odds ratios using the position distance measure), and increasing the distance between the dominating and nearest dominated gig attenuated the effect, particularly for the position distance measure (position distance: $p = .007$; Euclidean distance: $p = .088$). In the

[15] In our real-world dataset, where a target product may have multiple competitors that span a relatively large range of the attribute space, increasing the attribute distance between a dominated and dominating gig may not always result in a range extension (see Figure 1.5).

position distance model, there was also a significant interaction indicating that the effect of distance attenuated as the count of dominated gigs increased. Among experienceable categories, there was a main effect of increasing the count of dominated gigs ($ps < .001$), and neither the main effect of distance nor their interaction was significant. The result that increasing the count of dominated gigs strengthens the dominance effect provides support for H3, suggesting that the dominance effect is consistent with a perceptual mechanism. See Figure 1.7 for results using the position distance measure and Table A1.6 for regression results[16].

Magnitude of dominance. Among non-experienceable categories, increasing the magnitude of dominance along only the average rating attribute or only the number of ratings was not significant ($ps > .15$). However, there was a significant and positive interaction, such that increases in both attributes was associated with a greater dominance effect ($p = .027$). These results – that increasing the magnitude of dominance accentuates the dominance effect – are also consistent with a perceptual account underlying dominance, providing additional support for H3. In experienceable categories, increasing the magnitude of dominance along only the number of ratings was significant ($p < .001$), however increases in only the average rating or the interaction between them was non-significant ($ps > .50$). See Figure 1.8 and Table A1.8 for regression results.

[16] A possible confound is that increasing the count of dominated gigs increases the probability of detecting a dominance relationship. One way to minimize this concern is to estimate the effect of the count of dominated gigs among cases where at least one dominated gig is always directly adjacent to (i.e., immediately to the left or right of) a dominating one, such that a dominance relationship is always likely to be noticed. See Table A1.7 for a regression implementing this approach. The results show that the effect persists – increasing the count of dominated gigs significantly increases the odds of purchasing the dominating gig, even among cases where at least one dominated gig is always adjacent to a dominating one.

Robustness Checks

As a robustness check, we re-ran the baseline model using an expanded definition of (weak) dominance that included price. Under this dominance definition, a given gig i was classified as dominating if there was another gig j offered within the price band[17] defined for gig i, and in the same assortment as gig i, where the following statement was true, and at least one of the inequalities was strict:

$$(number\_ratings_i \geq number\_ratings_j) \;\&\; (average\_rating_i \geq average\_rating_j) \;\&\; (price_i \leq price_j)$$

Dominated gigs were defined using the same definition with the direction of the inequalities reversed.

As an additional robustness check, we re-ran the baseline model using a linear probability model instead of a logistic model. This specification also serves to assuage concerns about the incidental parameters problem (Lancaster 2000), which can affect nonlinear panel data models with fixed effects. We also re-ran the baseline model using clicks instead of purchases as the outcome. Lastly, we included a simple alternative specification without gig fixed effects, that simply controls for price, average rating, and number of ratings.

Results for Robustness Checks

Dominance on average rating, number of ratings, and price. A similar pattern of results was observed using an expanded definition of dominance that included price. This definition reduced the proportion of neither dominating nor dominated gigs in the dataset from 49.1% to 30.3%, but increased the share of dominating only gigs (18.3% to 21.7%), dominated only gigs

[17] For gigs priced between $5-$30 the price band was $5, for gigs priced between $35-$60 the price band was $10, and for gigs priced greater than $60 the price band was $20.

(18.7% to 25.2%), and both dominating and dominated gigs (13.9% to 22.8%). In non-experienceable categories, the effect was observed (OR = 1.14, $p < .001$). However, in experienceable categories, the effect was attenuated (OR = 0.96, $p = 0.017$). See Table A1.9 for regression results.

Linear probability model. As a robustness check, we re-ran the baseline regression specification using a linear probability model. The results were similar to those from the logistic model. In non-experienceable categories, the probability that a gig was purchased was 48% greater when it was in an assortment where it was dominating only compared to an assortment where it was neither dominating nor dominated ($\beta = .0005$, $p < .001$). In experienceable categories, the effect was significantly attenuated – the probability that a gig was purchased was only 6% greater when it was dominating ($\beta = .000085$, $p = 0.026$; difference from non-experienceable: $p < .001$). See Table A1.10 for regression results.

Clicks. A similar pattern of results was observed when estimating the baseline model using clicks instead of purchases as the outcome, though the effect sizes were substantially smaller yet still significant. In non-experienceable categories, the odds that a gig was clicked was greater when it was in an assortment where it was dominating only compared to an assortment where it was neither dominating nor dominated (OR = 1.05, $p < .001$). A possible explanation for the smaller effect is that both curiosity and learning play a relatively larger role for clicks (compared to purchases), thereby reducing the relative share of clicks on similarly priced options (since a dominance relationship implies an equal price, as per our definition). This countervailing force present in clicks would serve to reduce the probability of a click on a given gig when it has a dominance relationship with another gig in the assortment. In experienceable categories, the effect was attenuated – the odds that a gig was purchased were lower when it was dominating only compared when it was neither dominating nor dominated (OR = 0.99, p = 0.019). See Table A1.11 for regression results.

Specification without gig fixed effects. A similar pattern of results was observed estimating a simple model including only gig dominance status interacted with experienceability, and controls for price, average rating, and number of ratings. In non-experienceable categories, the dominance effect was observed (OR = 1.14, p < .001). However, the effect in experienceable categories (OR = 1.04, p < .001) was attenuated (p < .001). See Table A1.12 for regression results.

Discussion

Our empirical analysis using real-world data suggests that dominance can indeed be observed in the real-world, and the effect was attenuated when consumers were able to sample or experience the product. This effect was robust to multiple definitions of dominance (how price was included), modeling approaches (logistic and linear models), regression specifications (inclusion of controls), and outcome measures (purchases and clicks).

We also identify important moderators for the dominance effect that shed light on how it operates. As the count of dominated gigs increased, the dominance effect strengthened – a phenomenon consistent with the perceptual account of dominance. Magnitude of dominance also moderated the effect in non-experienceable categories, specifically when the magnitude increased along both the average rating and number of ratings attributes, purchase likelihood increased – again consistent with the perceptual account. The effect of magnitude was attenuated in experienceable categories. The findings from both moderators are not consistent with the similarity-based mechanism for dominance, which makes predictions in the opposite direction. Our results also suggest that reducing the distance between a dominating gig and the nearest dominated gig strengthens the effect, particularly when there is just a single dominated gig, as this may increase the likelihood that the dominance relationship is noticed.

One caveat to our empirical approach is that the regression model assumes independence across observations. However, observations in the same assortment are likely not independent because purchasing one gig reduces the likelihood of purchasing another one in the assortment, though of assortments with purchases, 2.2% had multiple purchases, indicating it is still possible. We account for this non-independence by clustering the standard errors at the assortment level, and we also note that the results for clicks are less affected by this issue.

Laboratory Experiment

In the Fiverr data, experienceability was not randomized but instead naturally varied by category. To rule out the (unlikely) possibility that other systematic differences across categories besides experienceability drove the observed effects, we ran an experiment that directly manipulates the degree of experienceability of a given assortment of alternatives.

The experimental design and analysis were pre-registered and all data, analysis code, research materials, and pre-registrations for all studies are available at https://osf.io/8dasb/?view_only=973dbd5f9d204316ad0fcc7ffbe8e2ea. For all studies, we report all manipulations and measures, and recruited a minimum of 100 participants per cell. All sample sizes and exclusion criteria were determined in advance.

Methods

As outlined in our pre-registration (https://aspredicted.org/QZC_VWT), we aimed to recruit 400 participants who were undergraduate students at a large public university, and ended with a sample of 444 participants who completed the study. All participants were first asked to complete an audio check, which consisted of an audio clip asking, “What is seven plus five?”

intended to ensure that all participants could hear the audio from the experiment survey. As pre-registered, the survey was terminated for participants who responded to the audio check incorrectly, and we removed all cases of multiple responses from the same participant. We were left with a final sample of 395 participants (63% female, mean age = 20.73 years).

All participants were randomized into one of two conditions: non-experienceable or experienceable. In both conditions, participants were asked to imagine that they were interested in purchasing a song. They were given an assortment of three songs that had a similar melody but different arrangement, and were asked to select which song they would purchase, if they could only purchase one of them. All participants were shown a table with the price, average customer rating (out of 5 stars), and the number of ratings for all three songs. In the experienceable condition, participants were additionally given 33 second samples for each of the songs that they could listen to, and we randomized the pairings between the audio clips and the song attributes for each participant. We pre-tested whether the ability to sample the songs reduced participants' preference uncertainty, and found that it did ($t(188) = 4.56$, $p < .001$; see Appendix Study A1.2 for details).

The table of information that participants saw for the three songs always created a dominance relationship between the options. Every assortment contained two core songs – Songs A and B (see Table 1.1) – and the third song was randomized to be either Song A' or B'. Importantly, the attributes of Songs A' and B' were constructed such that the inclusion of Song A' made Song A dominating, while the inclusion of Song B' made Song B dominating. We adopted this approach to hold fixed the size of the assortment, to increase comparability to the Fiverr setting. A dominating song was defined as one that was superior to another one on all three attributes: lower price, higher average rating, and greater number of ratings. We randomized the order in which the songs appeared in the table, though we only allowed permutations where the dominated and dominating songs appeared side-by-side, to increase the detectability of the dominance relationship.

Finally, we asked participants an open-ended question about how they made their song choice and collected participants' gender and age.

Results

Song choice. When a song was dominating, 58.9% of participants chose it in the non-experienceable condition, while 40.4% of participants selected it in the experienceable condition ($\chi^2(1) = 12.76$, $p < .001$, $\varphi = 0.18$; see Figure 1.9; see also Table A1.13 model 1 for results from a logit model regressing choice of dominating song on condition). In the experienceable condition, 97.5% of participants sampled all 3 songs at least once.

Dominating Song. To test whether the choice share of the dominating song was different when song A or song B was dominating, we regressed the choice of the dominating song on condition (experienceable vs. non-experienceable), a variable capturing which song was dominating (A vs. B), and their interaction in a logistic regression. Both the variable capturing dominating song and interaction term were non-significant (ps > .40; see Table A1.13 model 2 for full regression results), indicating that whether song A or song B was dominating made little difference in both the experienceable and non-experienceable conditions.

Discussion

Our results suggest that the ability to sample or experience alternatives – which was associated with lower preference uncertainty – attenuates the dominance effect. In the non-experienceable condition, the dominating song attracted the majority of the choice share, regardless of which alternative was dominating. But when participants were able to experience the alternatives, the dominating song attracted significantly less of the choice share.

General Discussion

Our work is among the first to find evidence for the dominance effect in a real and consequential setting. While the prevalence of dominance relationships in the real world has been questioned, we find that in the real online setting we examine they can be very common, which underscores the need to better understand the effect in naturalistic settings. The results from the field and laboratory data found converging evidence for an important moderator of the effect – preference uncertainty. In settings where preference uncertainty was diminished, the dominance effect attenuated. Our results are also consistent with a perceptual mechanism based on simple, automatic processing underlying the phenomenon, where an alternative simply "seems better" when in the presence of an alternative it dominates.

The real-world Fiverr data we used to study dominance has many ideal features. Key attributes are neatly displayed to customers in a numeric form that facilitates comparisons, randomization in gig assortment composition and ordinal position allow for identification of the dominance effect, and customers do not typically have familiarity with offerings (since sellers are freelancers and repeat purchases are rare). Our analysis is also a conservative demonstration of dominance, since the dominance relationships were defined based on weak instead of strong dominance – only one attribute needed to be strictly dominated, while the other could be equal. Another advantage over prior work is that customers using the Fiverr platform have the option not to purchase any items in the assortment, unlike typical lab experiments that use a forced-choice design. Our results therefore provide evidence for the dominance effect under realistic conditions when customers have a no-choice option[18]. Our work also extends the prior literature by using a relatively large assortment – 12 items – instead of the typical 3-item assortment, capturing the complexity of choosing in a real marketplace.

While prior work has studied the dominance effect by introducing an additional dominated option to the assortment, our setting always used a fixed number of options – a 12-item assortment on Fiverr and a 3-item assortment in the laboratory experiment. While the laboratory experiment always had a single dominated option, the count of dominated options varied randomly in the Fiverr data. It is possible that replacing a non-dominated option with a dominated one mechanically generated the observed dominance effect. However, this need not necessarily happen since the Fiverr setting also includes a no-choice option, and also this alternative explanation is common to both experienceable and non-experienceable categories, yet we do not observe any overall effect in experienceable categories. In addition, this cannot explain our result demonstrating that the dominance effect strengthens as the magnitude of dominance increases (particularly in non-experienceable categories), since this analysis relied only on cases where there was always one dominated gig associated with a dominating gig.

Another possibility is that experience introduces a new set of attributes, which may remove dominance relationships that were defined solely based on the numeric attributes. On Fiverr, this could plausibly explain the elimination of the dominance effect in experienceable categories, yet we did not observe an analogous elimination of the effect of being dominated – it remained sizable and significant. In other words, dominated gigs were still less likely to be purchased in experienceable categories, suggesting customers still identified and responded to

[18] Dhar and Simonson (2003) argue that the inclusion of a no-choice option may actually strengthen the dominance effect by reducing the choice share of the target option in the absence of dominance, since the no-choice option allows customers to avoid decision difficulty.

the dominance relationship. While speculative, it is possible that the information gathered from experience would be correlated with the average rating and number of ratings, which would explain why the dominance relationship continues to hold even with the ability to experience the products.

Our results highlight the role of preference uncertainty – examined through the ability to sample or experience alternatives – as a key moderator of the dominance effect. In non-experienceable categories, where preference uncertainty was heightened due to the inability to experience alternatives, consumers may be more likely to construct preferences based on the numeric attributes used to establish dominance, thereby generating the observed dominance effect. In experienceable categories, preference may have been formed based on the sample of the product, circumventing the constructive process giving rise to the effect. In related work, Simonson (1989) found a stronger dominance effect when consumers expected to be evaluated by others, and suggested that this was due to heightened uncertainty of others' preferences compared to one's own preferences. Future work can investigate the specific ways in which preference uncertainty and product experience impact the dominance effect as well as other context effects.

Conclusion

Our work is among the first to show real-world evidence for the dominance effect, identifies preference uncertainty as a moderator for the effect, and finds support for a perceptual mechanism underlying it. Our unique year-long panel dataset allows us to observe the specific assortment of options and corresponding attributes that were presented to customers, as well as which products they clicked on and purchased. Our results show that in settings with heightened preference uncertainty, where customers could not sample or experience alternatives, a product was roughly 21% more likely to be purchased when it was dominating compared to when it is not. However, when preference uncertainty diminished and customers could experience alternatives, the effect was small and non-significant. These findings have broad implications for marketers, choice architects, user interface designers, and policymakers. When bringing products to customers, especially new ones that customers cannot experience, marketers can increase its purchase share by introducing a dominated option (or several) to the assortment.

Acknowledgements

We are grateful to Wendy Liu, Ayelet Gneezy, Uri Gneezy, Kenneth Wilbur, and Yuval Rottenstreich for their helpful input. We thank Louis Cohen for creating the songs used in the lab experiment.

Chapter 1, in full, has been submitted for publication of the material. Ariel Fridman, On Amir, and Karsten Hansen. The dissertation author was the primary investigator and author of this paper.

Figures

Coin A Coin B Coin C Coin A’

Figure 1.1. Collectible coins illustrating dominance relationships. Coin A’ is dominated relative to coin A (but not coins B or C) because it is the same product yet is more expensive, has a lower average rating, and was rated by fewer customers. Image source: Amazon.com.

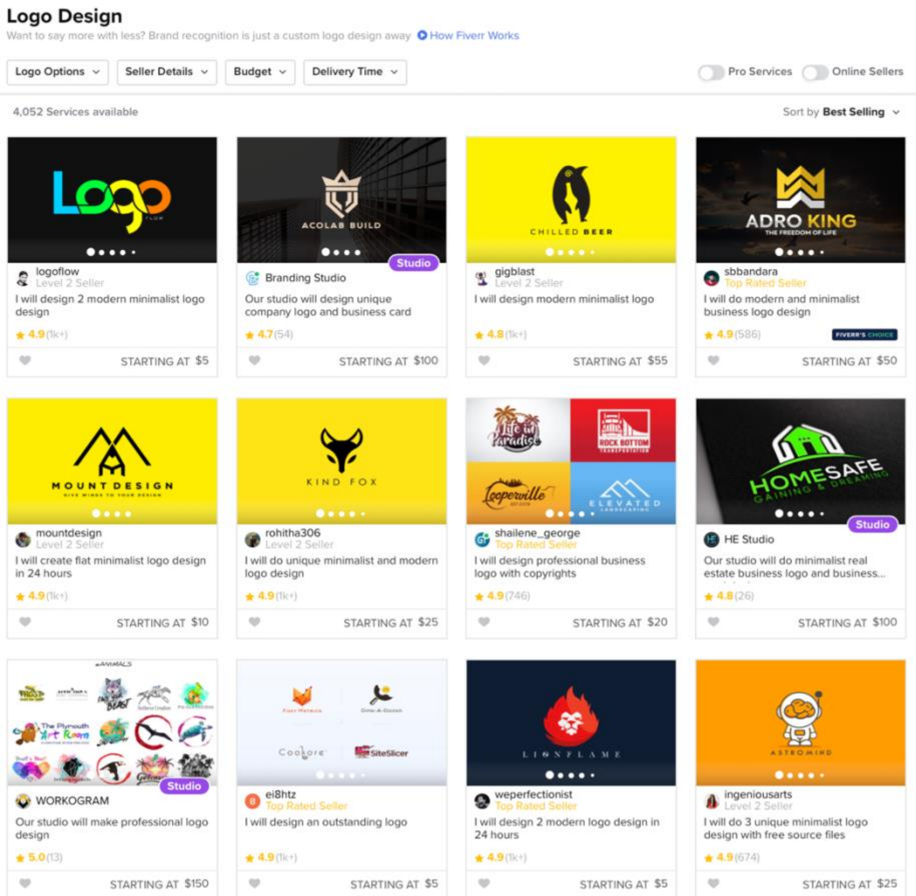

Figure 1.2. Screenshot from the logo design category on fiverr.com. The image shows an example of the assortment of 12 gigs that a customer would see when browsing in the Logo Design product category. Image source: Fiverr.com.

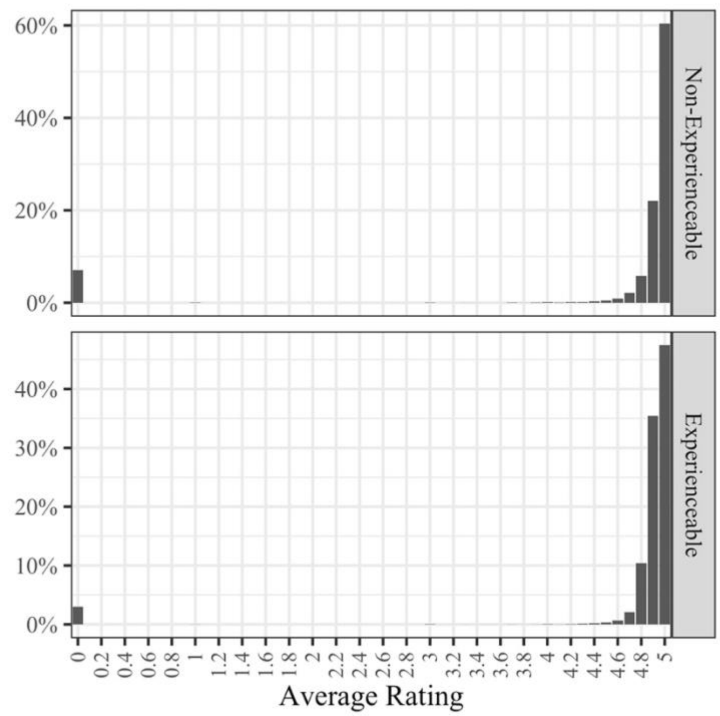

Figure 1.3. Distribution of average customer rating, by experienceability. The distribution covers all gigs included in the analysis. Ratings range from 0 to 5.

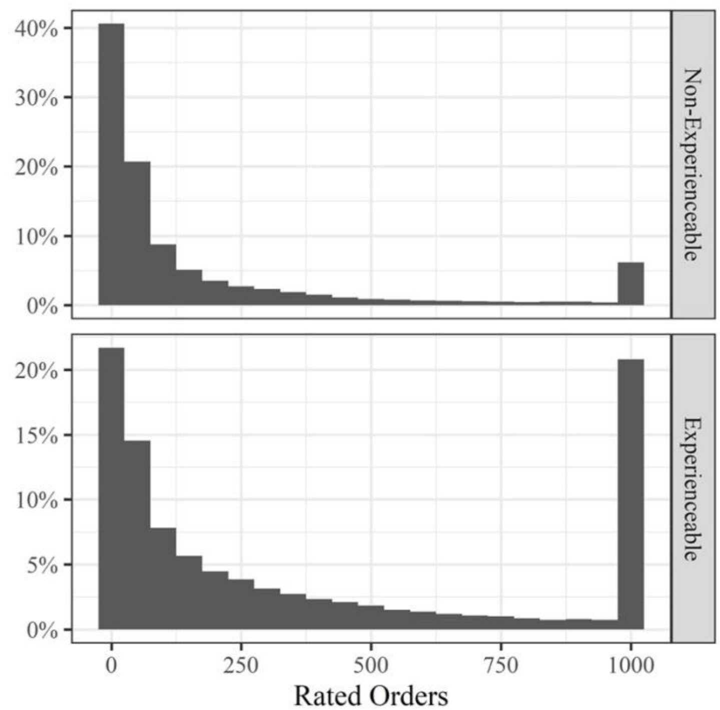

Figure 1.4. Distribution of number of ratings, by experienceability. The distribution covers all gigs included in the analysis. Width of the bins is 50. Number of ratings was top-coded at 1,000.

Figure 1.5A

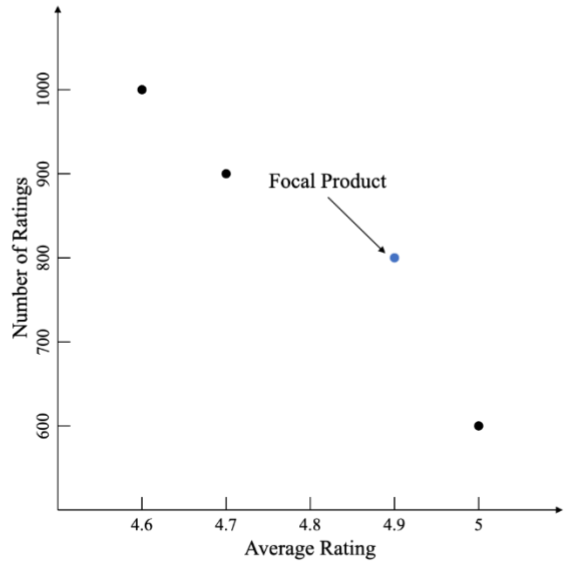

Figure 1.5B

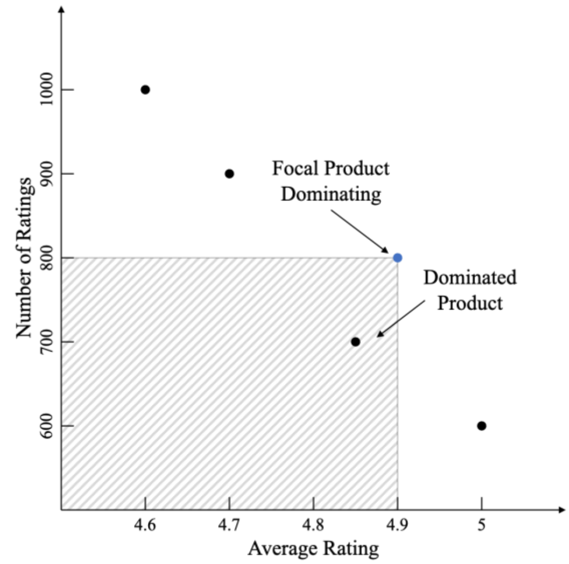

Figure 1.5C

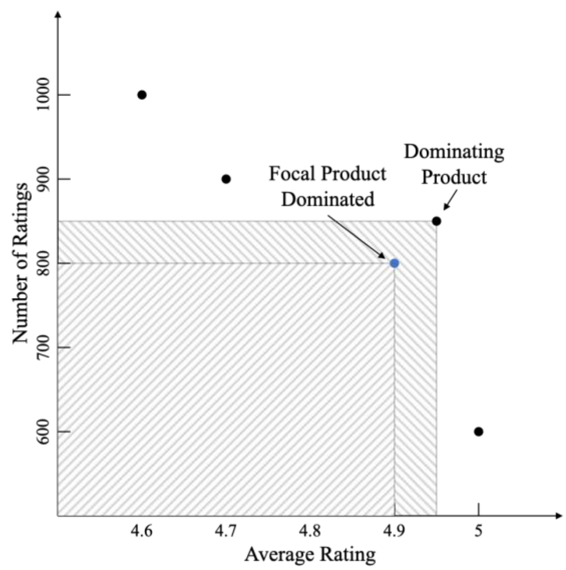

Figure 1.5D

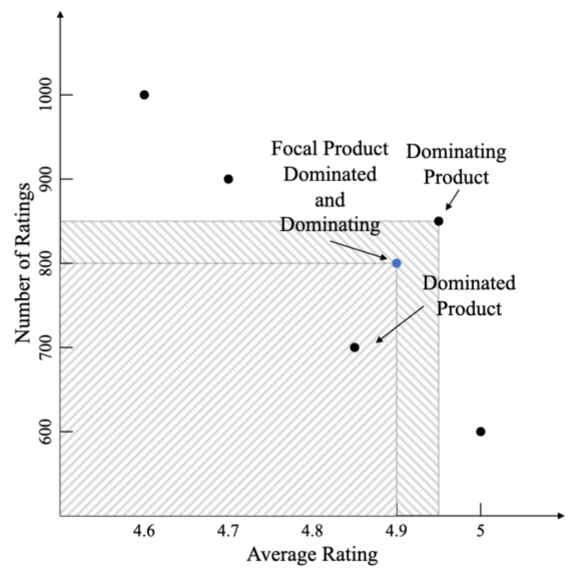

Figure 1.5. Conceptual Representation of Dominance Relationships. Products in the assortment are represented as points in the coordinate system, with the horizontal axis capturing the average rating and the vertical axis capturing the number of ratings. In figure 5A, no dominance relationships exist between the products. In figure 5B, the shaded box represents the combinations of attributes that would be dominated by the focal product, since the focal product would have a greater average rating, number of ratings, or both. The existence of a product in the shaded region makes the focal product dominating. In figure 5C, the focal product is dominated since it falls in the shaded box of another product. In figure 5D, the focal product is both dominated and dominating.

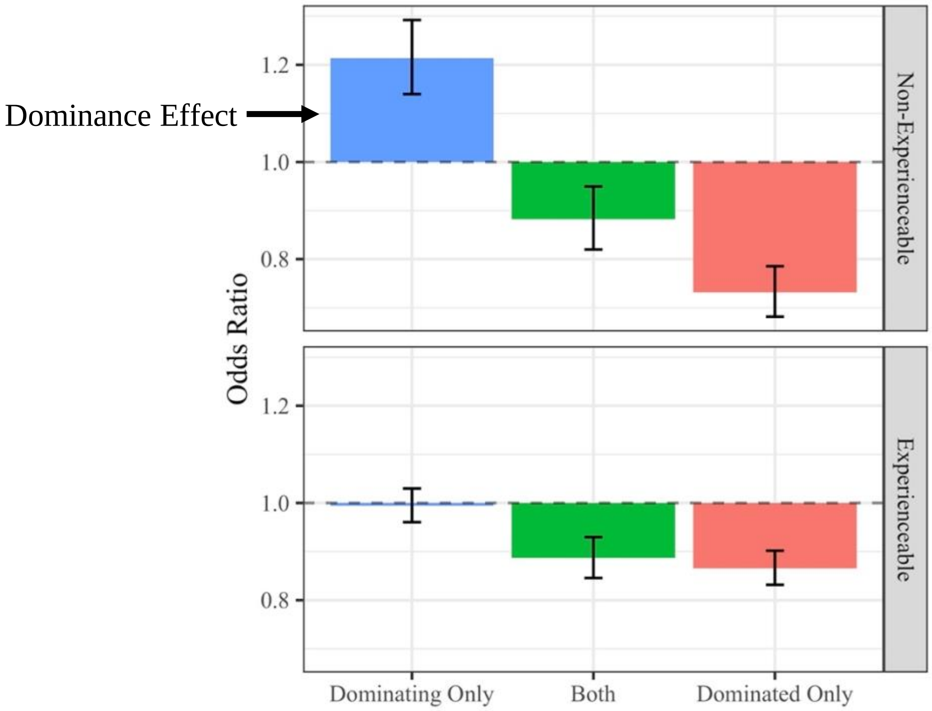

Figure 1.6. Baseline Model Results. The bars depict odds ratios of gig purchase relative to the control condition when a gig was neither dominating nor dominated. The blue bars depict the dominance effect: In non-experienceable categories, the odds that a gig was purchased were 21% greater when it was dominating (compared to when it was neither dominating nor dominated), though the dominance effect was eliminated in experienceable categories. Error bars represent 95% confidence intervals.

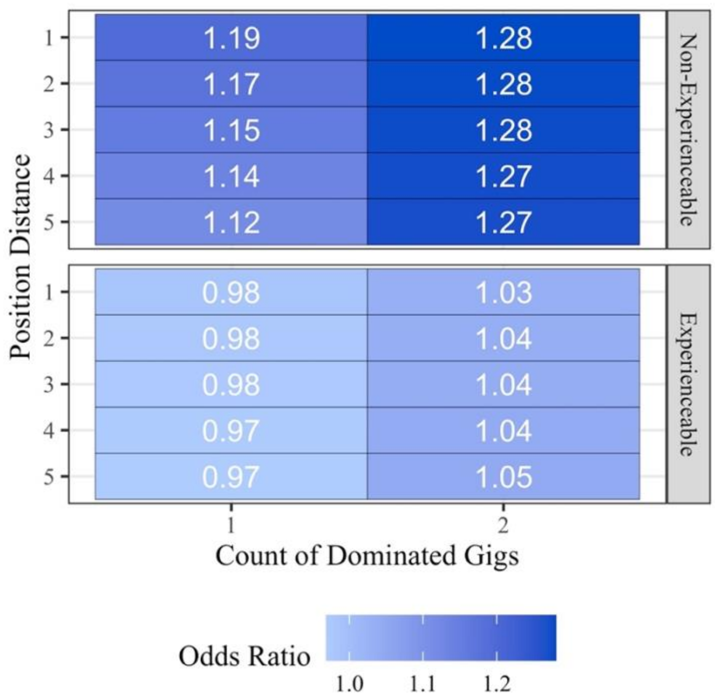

Figure 1.7. Odds Ratio Results for Count of Dominated Gigs by Position Distance. Values in each cell depict odds ratios of gig purchase. In non-experienceable categories, when there was just one dominated gig, decreasing the distance between the dominated and dominating gig strengthened the dominance effect. Increasing the count of dominated gigs also strengthened the effect. These effects were attenuated and non-significant in experienceable categories. Among gigs that were dominating, 56% had one dominated gig and 20% had two dominated gigs. Position distance was 5 or less in 80% of cases.

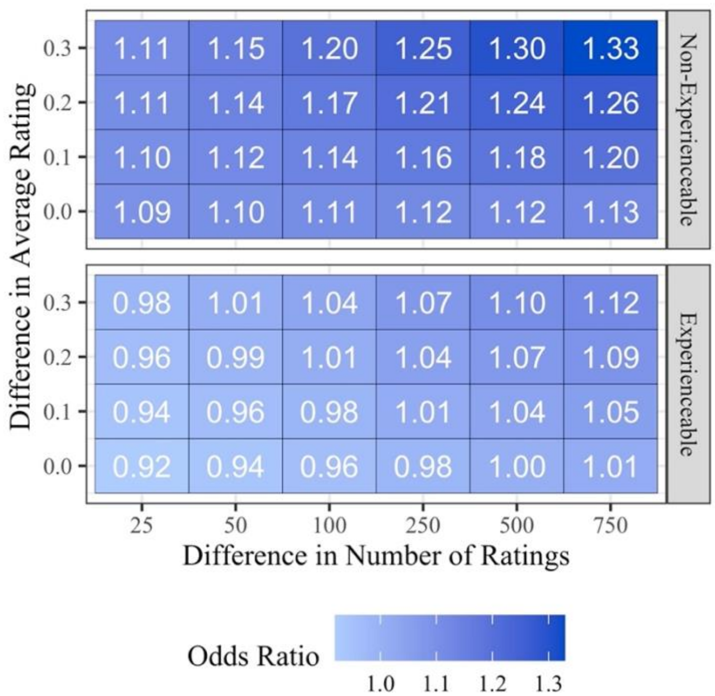

Figure 1.8. Odds Ratio Results by Magnitude of Dominance. Values in each cell depict odds ratios of gig purchase. In non-experienceable categories, a greater difference in both the number of ratings and average rating between the dominated and dominating gig was associated with a stronger dominance effect. This pattern was attenuated and non-significant in experienceable categories.

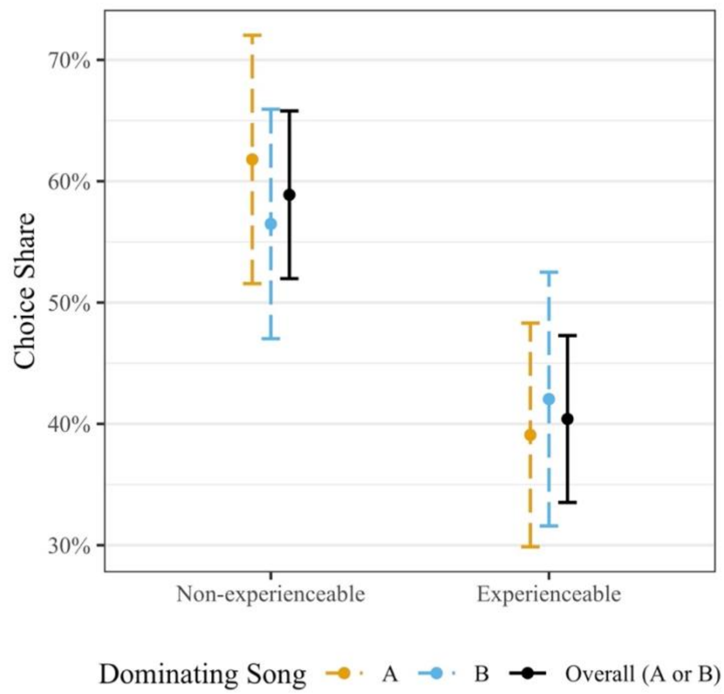

Figure 1.9. Results from Laboratory Study: Choice Share of Dominating Song by Condition. The dominance effect was stronger in the non-experienceable condition, where the dominating song was chosen 59% of the time; in the experienceable condition it was chosen only 40% of the time. Results were similar regardless of whether Song A or B was dominating. Error bars represent 95% confidence intervals.

Tables

Table 1.1. Song Attributes in Laboratory Experiment. Table shows attributes of the songs included in the experiment. Songs A and B were always included, and either Song A' or B' was randomly selected to be included as well. Song A' is dominated relative to Song A (but not Song B); similarly, Song B' is dominated relative to Song B (but not Song A).

	Song A'	**Song A**	**Song B**	**Song B'**
Price (cents)	95	90	70	80
Average Customer Rating (out of 5 stars)	4.8	4.9	4.7	4.6
Number of Ratings	96	119	135	82

Appendix

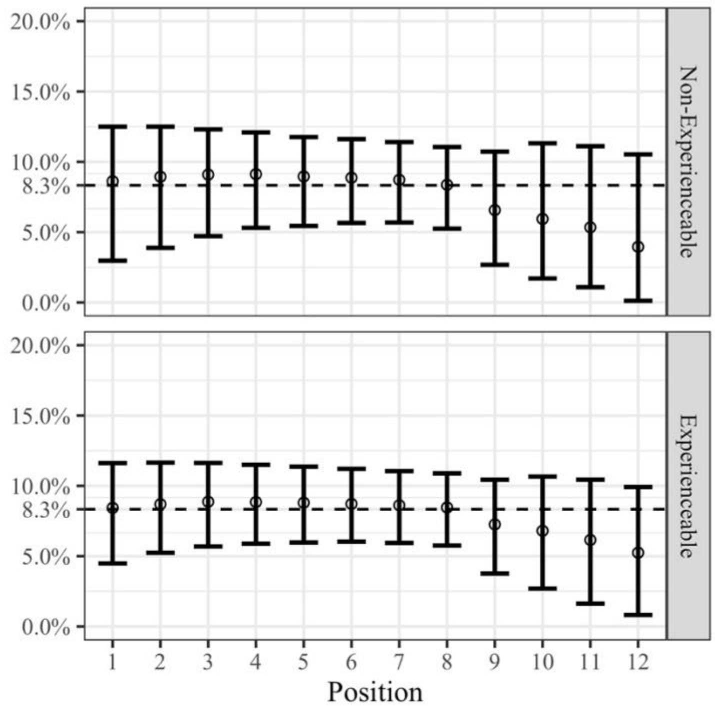

Figure A1.1. Gig position randomization check, by experienceability. For each gig that appeared in 100 or more assortments, the percentage of times it appeared in each position was computed. The points show the median percentage, and the error bars show the range between the first and third quartiles. The dashed horizontal line at 1/12 (8.3%) captures the expected percentage if gigs were sorted randomly.

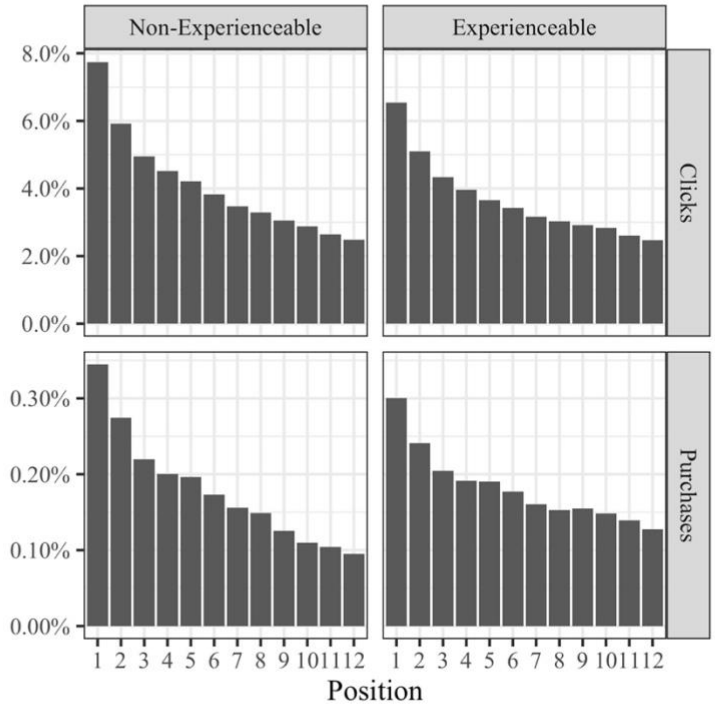

Figure A1.2. Share of clicks and purchases by position, by experienceability.

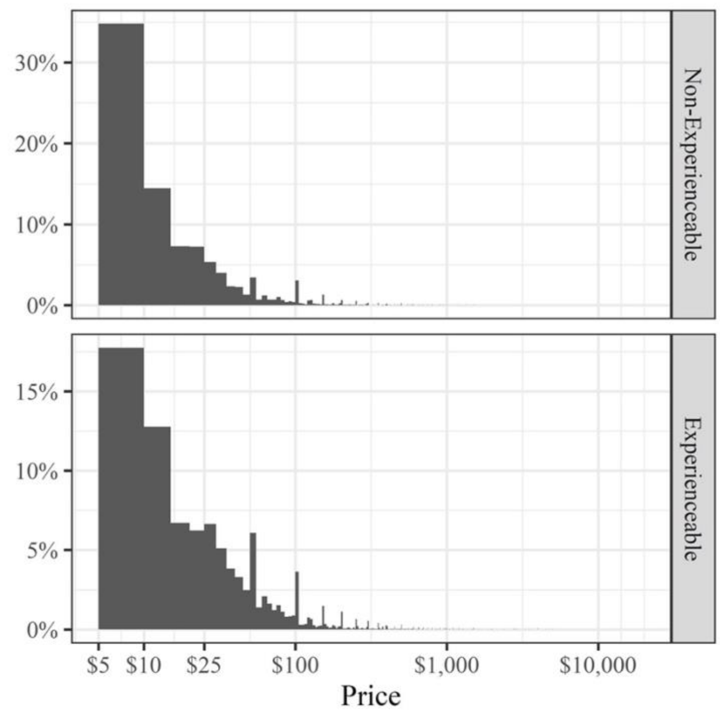

Figure A1.3. Distribution of gig prices, by experienceability. Horizontal axis was log-transformed.

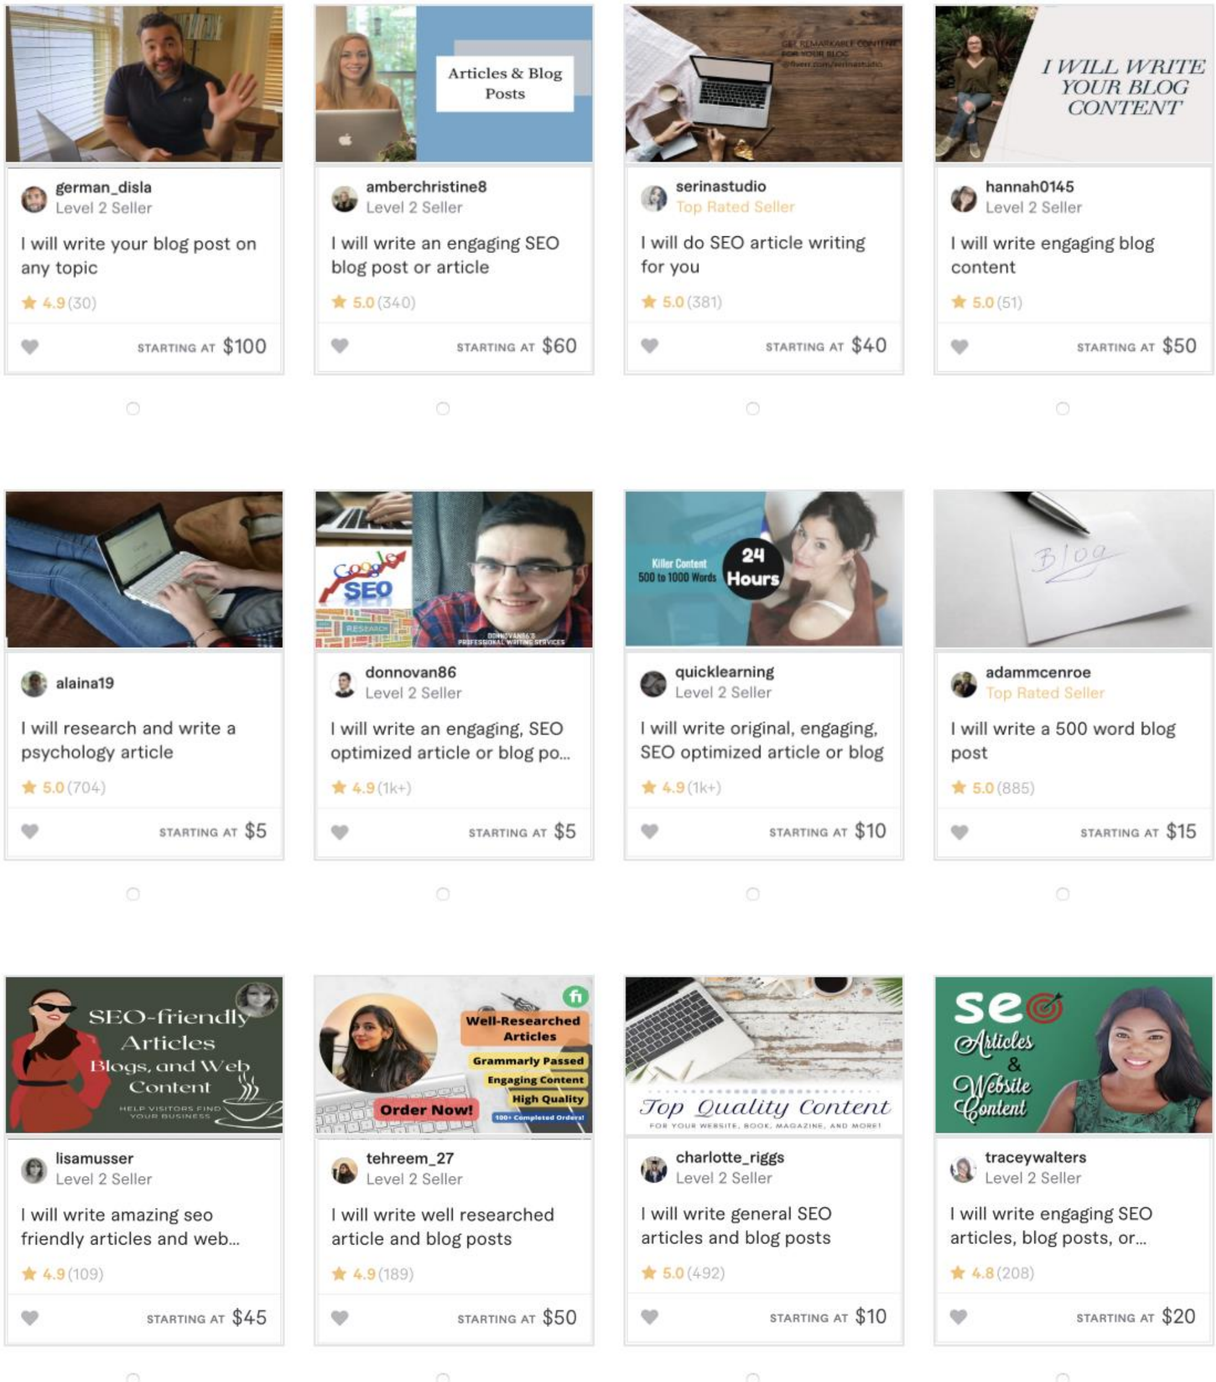

Figure A1.4. Screenshot from Study A1 for the Articles & Blog Posts Product Category.

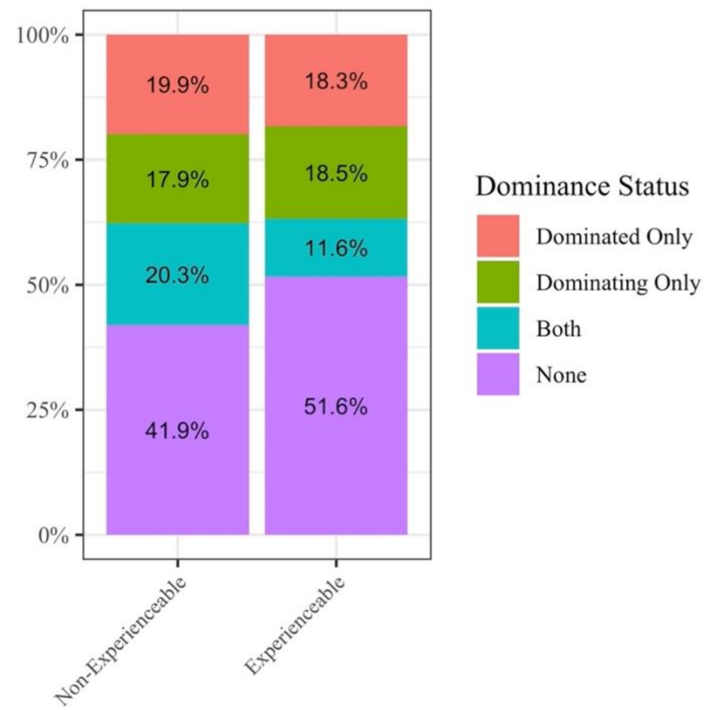

Figure A1.5. Distribution of dominance status, by experienceability.

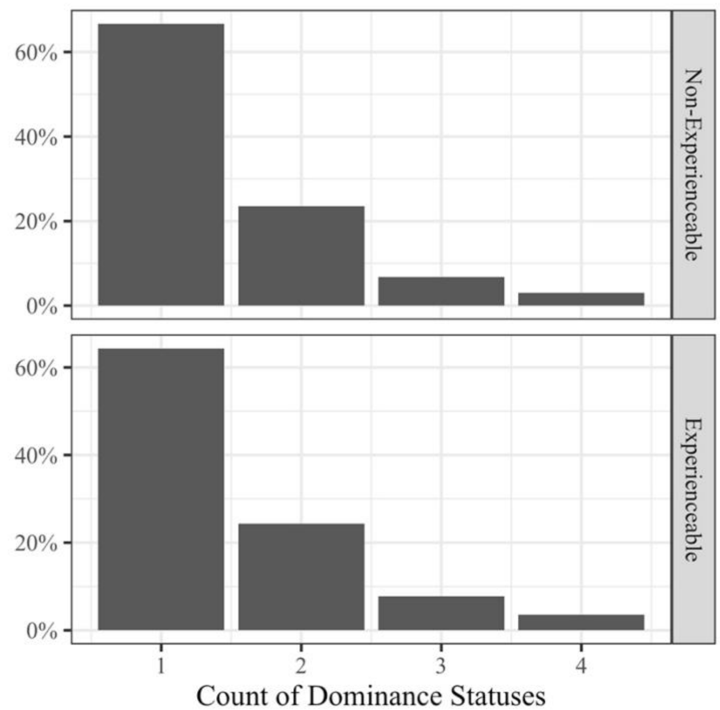

Figure A1.6. Distribution of count of dominance statuses for each gig, by experienceability.

Table A1.1. Fiverr Dataset Variables and Definitions.

Variable	**Definition**
id	Unique ID for each assortment of search results
gig_id	Unique ID for each gig
position	Position of a gig in the assortment of search results
price	Price of gig
pt	Date of search
event_time	Time of search
sub_category_id	Unique ID for each category
Experienceability	Classification of category by experienceability: experienceable, non-experienceable
average_rating	Average gig rating
rated_orders	Count of customer ratings for the gig
type	Gig type: default, pro, fiverr choice, studio, etc.
buyer_type	Type of buyer: guest, registered not converted (rnc), first time buyer (ftb), second time buyer (stb), repeat buyer, unknown
country	Buyer's country
seller_country	Seller's country
reg_platform	Buyer's platform: web, app, mobile_web
filter_tab	Filter applied: yes or no
rfm_segmentation	Recency, Frequency and Monetary segmentation of buyer: A, B, C, D, E, OTB, unknown
seller_level	Seller level: level one, level two, top rated, none
clean_os	Operating system of buyer: Windows, Mac, Linux, Chrome OS, iOS, Android, other, unknown
max_order_amount	Maximum dollar amount of the buyer's orders
avg_order_amount	Average dollar amount of the buyer's orders
total_orders_lifetime	Count of orders placed by the buyer
total_parent_orders_lifetime	Count of parent orders placed by the buyer
total_orders_cancellation	Count of orders canceled by the buyer
user_id	Unique ID for each buyer
view_id	Unique ID for each view of a gig's page
order_id	Unique ID for each order
parent_order_id	Unique ID for parent order

Table A1.2. List of Subcategories by Experienceability Classification.

Experienceable Subcategories		Non-Experienceable Subcategories	
3D Models & Product Design	Lyric & Music Videos	Articles & Blog Posts	Local Listings
3D Product Animation	Menu Design	Astrology & Readings	Market Research
Animated GIFs	Mixing & Mastering	Beta Reading	Marketing Strategy
Animation for Kids	Packaging Design	Book & eBook Writing	Music Promotion
Animation for Streamers	Photoshop Editing	Branding Services	Online Lessons
Arts & Crafts	Portraits & Caricatures	Business Names & Slogans	Press Releases
Banner Ads	Postcard Design	Business Plans	Product Descriptions
Book & Album Covers	Poster Design	Business Tips	Product Research
Brand Style Guides	Pranks & Stunts	Career Advice	Proofreading & Editing
Brochure Design	Presentation Design	Collectibles	QA
Business Cards & Stationery	Presentations	Content Marketing	Relationship Advice
Car Wraps	Producers & Composers	Convert Files	Research & Summaries
Cartoons & Comics	Product Photography	Creative Writing	Sales Copy
Catalog Design	Short Video Ads	Crowdfunding	Scriptwriting
Celebrity Impersonators	Social Media Design	Data Analysis & Reports	SEM
Character Animation	Sound Effects	Data Entry	SEO
Character Modeling	Spokespersons Videos	Databases	Social Media Marketing
Chatbots	Storyboards	Desktop Applications	Speechwriting
Flyer Design	T-Shirts & Merchandise	Domain Research	Spiritual & Healing
Game Design	Vector Tracing	E-Commerce Marketing	Support & IT
Game Development	Video Editing	Email Copy	Surveys
Game Trailers	Viral Videos	Email Marketing	Technical Writing
Graphics for Streamers	Visual Effects	Family & Genealogy	Transcripts
Greeting Cards & Videos	Vocal Tuning	Financial Consulting	Translation
Illustration	Voice Over	Flyer Distribution	User Testing
Infographic Design	Web & Mobile Design	Gaming	Video Marketing
Intros & Outros	Website Builders & CMS	Health, Nutrition & Fitness	Virtual Assistant
Jingles & Drops	Whiteboard & Animated Explainers	Influencer Marketing	Web Analytics
Live Action Explainers	WordPress	Lead Generation	Web Programming
Logo Animation	Your Message On...	Legal Consulting	Web Traffic
Logo Design		Legal Writing	Website Content

Table A1.3. List of control variables for assortment and buyer characteristics.

Variable Name	Description
Assortment Characteristics	
Price (median)	Median price of gigs in assortment
Price (SD)	Standard deviation of the prices of the gigs in assortment
Price (min)	Minimum price of gigs in assortment
Price (max)	Maximum price of gigs in assortment
Buyer Characteristics	
Total Orders Lifetime	Count of orders placed by the buyer
Total Orders Cancellation	Count of orders canceled by the buyer
Buyer Type	Type of buyer: guest, registered not converted (rnc), first time buyer (ftb), second time buyer (stb), repeat buyer, unknown
Operating System	Operating system of buyer: Windows, Mac, Linux, Chrome OS, iOS, Anrdoid, other, unknown
RFM Segmentation	Recency, Frequency and Monetary segmentation of buyer: A, B, C, D, E, OTB, unknown

Table A1.4. Correlation Table.

	Purchase	Experienceable	Dominated Only	Dominating Only	Both Dominated and Dominating	Price (median)	Price (SD)
Purchase	1.000	0.000	-0.002	0.008	0.008	-0.011	-0.015
Experienceable	0.000	1.000	-0.018	0.007	-0.112	0.268	0.168
Dominated Only	-0.002	-0.018	1.000	-0.227	-0.192	-0.101	-0.111
Dominating Only	0.008	0.007	-0.227	1.000	-0.190	-0.094	-0.094
Both Dominated and Dominating	0.008	-0.112	-0.192	-0.190	1.000	-0.362	-0.346
Price (median)	-0.011	0.268	-0.101	-0.094	-0.362	1.000	0.620
Price (SD)	-0.015	0.168	-0.111	-0.094	-0.346	0.620	1.000
Price (min)	-0.006	0.208	-0.056	-0.054	-0.169	0.717	0.278
Price (max)	-0.015	0.180	-0.111	-0.095	-0.345	0.647	0.995
Total Orders Lifetime	0.032	-0.034	0.000	-0.017	-0.008	0.131	-0.072
Total Orders Cancellation	0.012	-0.062	0.001	-0.011	0.009	0.042	-0.049
Position	-0.012	0.000	0.050	-0.073	-0.003	0.000	0.000
N gigs at same price	0.015	-0.151	0.283	0.273	0.512	-0.580	-0.554

	Price (min)	Price (max)	Total Orders Lifetime	Total Orders Cancellation	Position	N gigs at same price
Purchase	-0.006	-0.015	0.032	0.012	-0.012	0.015
Experienceable	0.208	0.180	-0.034	-0.062	0.000	-0.151
Dominated Only	-0.056	-0.111	0.000	0.001	0.050	0.283
Dominating Only	-0.054	-0.095	-0.017	-0.011	-0.073	0.273
Both Dominated and Dominating	-0.169	-0.345	-0.008	0.009	-0.003	0.512
Price (median)	0.717	0.647	0.131	0.042	0.000	-0.580
Price (SD)	0.278	0.995	-0.072	-0.049	0.000	-0.554
Price (min)	1.000	0.326	0.085	0.013	0.000	-0.282
Price (max)	0.326	1.000	-0.074	-0.052	0.000	-0.553
Total Orders Lifetime	0.085	-0.074	1.000	0.618	0.000	-0.028
Total Orders Cancellation	0.013	-0.052	0.618	1.000	0.000	0.002
Position	0.000	0.000	0.000	0.000	1.000	0.000
N gigs at same price	-0.282	-0.553	-0.028	0.002	0.000	1.000

Table A1.5. Baseline Model Results.

Dependent Variable:	Purchase	
Model:	(1)	(2)
Variables		
Non-Experienceable x Dominating Only	0.1936***	0.2521***
	(0.0320)	(0.0304)
Non-Experienceable x Dominated Only	-0.3125***	-0.4804***
	(0.0362)	(0.0353)
Non-Experienceable x Both	-0.1252***	-0.2333***
	(0.0375)	(0.0363)
Experienceable x Dominating Only	-0.0055	0.0350*
	(0.0177)	(0.0140)
Experienceable x Dominated Only	-0.1440***	-0.2630***
	(0.0207)	(0.0171)
Experienceable x Both	-0.1203***	-0.1112***
	(0.0242)	(0.0190)
Price (median)	0.0020***	
	(0.0003)	
Price (SD)	-0.0001	
	(0.0004)	
Price (min)	0.0077***	
	(0.0009)	
Price (max)	4.84e-6	
	(0.0001)	
Total Orders Lifetime	0.0002***	
	(4.84e-5)	
Total Orders Cancellation	-0.0016*	
	(0.0007)	
Fixed Effects		
Gig ID	Yes	Yes
Date	Yes	
Country	Yes	
Position	Yes	
N gigs at same price	Yes	
Buyer Type	Yes	
Operating System	Yes	
RFM Segmentation	Yes	
Fit statistics		
Observations	10,695,120	10,697,803
Within R^2	0.00199	0.00124
AIC	903,224.80	948,786.80

Clustered (assortment) standard-errors in parentheses
*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05*

Table A1.6. Results by Count of Dominated gigs and Distance Measures (Position and Euclidean Distance).

Dependent Variable:	Purchase	
Model:	(1) Position Distance	(2) Euclidean Distance
Variables		
Non-Experienceable x Dominating Only	0.1375**	0.1807*
	(0.0499)	(0.0714)
Non-Experienceable x Dominating Only x Distance	-0.0309**	-0.0733
	(0.0115)	(0.0430)
Non-Experienceable x Dominating Only x Dominated Count	0.0559***	0.0556*
	(0.0114)	(0.0269)
Non-Experienceable x Dominating Only x Distance x Dominated Count	0.0144*	0.0157
	(0.0063)	(0.0246)
Non-Experienceable x Dominated Only	-0.2632***	-0.2662***
	(0.0363)	(0.0363)
Non-Experienceable x Both	-0.0038	-0.0094
	(0.0385)	(0.0385)
Experienceable x Dominating Only	-0.0615*	-0.0709
	(0.0275)	(0.0392)
Experienceable x Dominating Only x Distance	-0.0093	0.0020
	(0.0063)	(0.0236)
Experienceable x Dominating Only x Dominated Count	0.0460***	0.0661***
	(0.0075)	(0.0169)
Experienceable x Dominating Only x Distance x Dominated Count	0.0061	-0.0129
	(0.0040)	(0.0153)
Experienceable x Dominated Only	-0.1125***	-0.1145***
	(0.0208)	(0.0208)
Experienceable x Both	-0.0228	-0.0271
	(0.0254)	(0.0254)
Price (median)	0.0020***	0.0020***
	(0.0003)	(0.0003)
Price (SD)	-7.92e-5	-8.17e-5
	(0.0004)	(0.0004)
Price (min)	0.0077***	0.0077***
	(0.0009)	(0.0009)
Price (max)	-4.96e-6	-4.28e-6
	(0.0001)	(0.0001)
Total Orders Lifetime	0.0002***	0.0002***
	(4.87e-5)	(4.87e-5)
Total Orders Cancellation	-0.0016*	-0.0016*
	(0.0007)	(0.0007)
Fixed Effects		
Gig ID	Yes	Yes
Date	Yes	Yes
Country	Yes	Yes
Position	Yes	Yes
N gigs at same price	Yes	Yes
Buyer Type	Yes	Yes
Operating System	Yes	Yes
RFM Segmentation	Yes	Yes

Table A1.6. Results by Count of Dominated gigs and Distance Measures (Position and Euclidean Distance). (Continued)

Dependent Variable:	Purchase	
Model:	(1) Position Distance	(2) Euclidean Distance
Fit statistics		
Observations	10,695,120	10,695,120
Within R^2	0.00223	0.00223
AIC	903,046.80	903,050.30

Clustered (assortment) standard-errors in parentheses
*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05*

Table A1.7. Results by Count of Dominated gigs with Adjacent Dominated Gig.

Dependent Variable:	Purchase
Model:	(1)
Variables	
Non-Experienceable x Dominating Only x Adjacent	0.1218*
	(0.0539)
Non-Experienceable x Dominating Only x Adjacent x Dominated Count	0.0693***
	(0.0093)
Non-Experienceable x Dominating Only x Non-Adjacent	0.0266
	(0.0407)
Non-Experienceable x Dominating Only x Non-Adjacent x Dominated Count	0.0921***
	(0.0135)
Non-Experienceable x Dominated Only	-0.2646***
	(0.0363)
Non-Experienceable x Both	-0.0064
	(0.0385)
Experienceable x Dominating Only x Adjacent	-0.0564
	(0.0308)
Experienceable x Dominating Only x Adjacent x Dominated Count	0.0494***
	(0.0063)
Experienceable x Dominating Only x Non-Adjacent +	-0.1024***
	(0.0226)
Experienceable x Dominating Only x Non-Adjacent + x Dominated Count	0.0681***
	(0.0090)
Experienceable x Dominated Only	-0.1128***
	(0.0208)
Experienceable x Both	-0.0230
	(0.0254)
Price (median)	0.0020***
	(0.0003)
Price (SD)	-7.89e-5
	(0.0004)
Price (min)	0.0077***
	(0.0009)
Price (max)	-5.03e-6
	(0.0001)
Total Orders Lifetime	0.0002***
	(4.87e-5)
Total Orders Cancellation	-0.0016*
	(0.0007)

Table A1.7. Results by Count of Dominated gigs with Adjacent Dominated Gig. (Continued)

Dependent Variable:	Purchase
Model:	(1)
Fixed Effects	
Gig ID	Yes
Date	Yes
Country	Yes
Position	Yes
N gigs at same price	Yes
Buyer Type	Yes
Operating System	Yes
RFM Segmentation	Yes
Fit statistics	
Observations	10,695,120
Within R^2	0.00223
AIC	903,050.00

Clustered (assortment) standard-errors in parentheses
*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05*
Note: 'Adjacent' refers to the case when at least one dominated gig appears immediately to the left or right of a dominating one.

Table A1.8. Results by Magnitude of Dominance.

Dependent Variable: Model:	Purchase (1)
Variables	
Non-Experienceable x Dominating Only	0.0617 (0.0740)
Non-Experienceable x Dominating Only x Dominated Count: 1 x NR Difference	0.0086 (0.0149)
Non-Experienceable x Dominating Only x Dominated Count: 1 x AR Difference	-0.4838 (0.3377)
Non-Experienceable x Dominating Only x Dominated Count: 1 x AR Difference x NR Difference	0.1686* (0.0765)
Non-Experienceable x Dominating Only x Dominated Count: 2+ x NR Difference	0.0503*** (0.0144)
Non-Experienceable x Dominating Only x Dominated Count: 2+ x AR Difference	0.7665 (0.4152)
Non-Experienceable x Dominating Only x Dominated Count: 2+ x AR Difference x NR Difference	-0.0276 (0.0941)
Non-Experienceable x Dominated Only	-0.2837*** (0.0365)
Non-Experienceable x Both	-0.0575 (0.0384)
Experienceable x Dominating Only	-0.1780*** (0.0375)
Experienceable x Dominating Only x Dominated Count: 1 x NR Difference	0.0289*** (0.0068)
Experienceable x Dominating Only x Dominated Count: 1 x AR Difference	0.1289 (0.2723)
Experienceable x Dominating Only x Dominated Count: 1 x AR Difference x NR Difference	0.0379 (0.0571)
Experienceable x Dominating Only x Dominated Count: 2+ x NR Difference	0.0511*** (0.0068)
Experienceable x Dominating Only x Dominated Count: 2+ x AR Difference	0.2813 (0.3599)
Experienceable x Dominating Only x Dominated Count: 2+ x AR Difference x NR Difference	-0.0262 (0.0756)
Experienceable x Dominated Only	-0.1217*** (0.0209)
Experienceable x Both	-0.0651** (0.0252)
Price (median)	0.0019*** (0.0003)
Price (SD)	-7.33e-5 (0.0004)
Price (min)	0.0077*** (0.0009)
Price (max)	-6.2e-6 (0.0001)
Total Orders Lifetime	0.0002*** (4.82e-5)
Total Orders Cancellation	-0.0016* (0.0007)

Table A1.8. Results by Magnitude of Dominance. (Continued)

Dependent Variable:	Purchase
Model:	(1)
Fixed Effects	
Gig ID	Yes
Date	Yes
Country	Yes
Position	Yes
N gigs at same price	Yes
Buyer Type	Yes
Operating System	Yes
RFM Segmentation	Yes
Fit statistics	
Observations	10,695,120
Within R^2	0.0022
AIC	903,088.70

Clustered (assortment) standard-errors in parentheses
*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05*
Note: NR and AR refer to number of ratings and average rating, respectively. The variables AR Difference and NR Difference were both log transformed.

Table A1.9. Baseline Model Results Including Price Dominance.

Dependent Variable:	Purchase
Model:	(1)
Variables	
Non-Experienceable x Dominating Only	0.1337***
	(0.0350)
Non-Experienceable x Dominated Only	-0.4046***
	(0.0375)
Non-Experienceable x Both	-0.2336***
	(0.0374)
Experienceable x Dominating Only	-0.04*
	(0.0167)
Experienceable x Dominated Only	-0.2100***
	(0.0179)
Experienceable x Both	-0.2169***
	(0.0186)
Price (median)	0.0017***
	(0.0003)
Price (SD)	-6.87e-5
	(0.0004)
Price (min)	0.0076***
	(0.0009)
Price (max)	-7.46e-6
	(0.0001)
Total Orders Lifetime	0.0002***
	(4.84e-5)
Total Orders Cancellation	-0.0016*
	(0.0007)
Fixed Effects	
Gig ID	Yes
Date	Yes
Country	Yes
Position	Yes
N gigs at same price	Yes
Buyer Type	Yes
Operating System	Yes
RFM Segmentation	Yes
Fit statistics	
Observations	10,695,120
Within R^2	0.00228
AIC	902,998.30

Clustered (assortment) standard-errors in parentheses
*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05*

Table A1.10. Linear Probability Model Results.

Dependent Variable:	Purchase
Model:	(1)
Variables	
Non-Experienceable x Dominating Only	0.0005***
	(6.08e-5)
Non-Experienceable x Dominated Only	-0.0005***
	(5.42e-5)
Non-Experienceable x Both	-0.0003***
	(7.45e-5)
Experienceable x Dominating Only	8.51e-5*
	(3.83e-5)
Experienceable x Dominated Only	-0.0003***
	(3.82e-5)
Experienceable x Both	-0.0002***
	(5.36e-5)
Price (median)	-1.22e-7
	(8.11e-8)
Price (SD)	2.67e-7
	(2.83e-7)
Price (min)	9.01e-7***
	(1.7e-7)
Price (max)	-1.61e-7
	(8.2e-8)
Total Orders Lifetime	1.37e-6***
	(1.49e-7)
Total Orders Cancellation	-3.78e-6
	(2.39e-6)
Fixed Effects	
Gig ID	Yes
Date	Yes
Country	Yes
Position	Yes
N gigs at same price	Yes
Buyer Type	Yes
Operating System	Yes
RFM Segmentation	Yes
Fit statistics	
Observations	51,609,096
Within R^2	1.94E-05
AIC	-177,014,841.40

Clustered (assortment) standard-errors in parentheses
*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05*

Table A1.11. Baseline Model Results for Clicks.

Dependent Variable:	Click
Model:	(1)
Variables	
Non-Experienceable x Dominating Only	0.0507***
	(0.0069)
Non-Experienceable x Dominated Only	-0.0934***
	(0.0074)
Non-Experienceable x Both	-0.0486***
	(0.0084)
Experienceable x Dominating Only	-0.0106*
	(0.0045)
Experienceable x Dominated Only	-0.0540***
	(0.0050)
Experienceable x Both	-0.0420***
	(0.0060)
Price (median)	-9.68e-5**
	(3.66e-5)
Price (SD)	0.0003***
	(6.46e-5)
Price (min)	0.0009***
	(6.44e-5)
Price (max)	-9.19e-5***
	(1.87e-5)
Total Orders Lifetime	4.54e-5***
	(1.38e-5)
Total Orders Cancellation	0.0008***
	(0.0002)
Fixed Effects	
Gig ID	Yes
Date	Yes
Country	Yes
Position	Yes
N gigs at same price	Yes
Buyer Type	Yes
Operating System	Yes
RFM Segmentation	Yes
Fit statistics	
Observations	35,075,090
Within R^2	0.00021
AIC	14,476,808.30

Clustered (assortment) standard-errors in parentheses
*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05*

Table A1.12. Alternative Specification Without Gig Fixed Effects.

Dependent Variable:	Purchase
Model:	(1)
Variables	
Non-Experienceable x Dominating Only	0.2915***
	(0.0188)
Non-Experienceable x Dominated Only	-0.2918***
	(0.0214)
Non-Experienceable x Both	0.0155
	(0.0200)
Experienceable x Dominating Only	0.0495***
	(0.0108)
Experienceable x Dominated Only	-0.2315***
	(0.0120)
Experienceable x Both	0.0079
	(0.0128)
Fixed Effects	
Price	Yes
Average Rating	Yes
Number of Ratings	Yes
Product Category	Yes
Fit statistics	
Observations	50,856,200
Within R^2	0.00124
AIC	1,319,075.60

Clustered (assortment) standard-errors in parentheses
*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05*

Table A1.13. Results from Laboratory Experiment.

Dependent Variable:	Choice of Dominating Song	
Model:	(1)	(2)
Variables		
(Intercept)	0.3591*	0.4810*
	(0.1450)	(0.2190)
Condition: Experienceable	-0.7478***	-0.9245**
	(0.2051)	(0.2940)
Dominating Song: B		-0.2202
		(0.2931)
Condition: Experienceable x Dominating Song: B		0.3428
		(0.4140)
Fit statistics		
Observations	395	395
R^2	0.02478	0.02614
AIC	538	541.25

IID standard-errors in parentheses
*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05*

Study A1.1. Fiverr Pre-test: Experienceability and Preference Uncertainty.

We pre-tested whether the ability to experience products was associated with greater preference certainty in an experimental setting mimicking Fiverr. We report all manipulations and measures, and recruited a minimum of 100 participants per cell. All sample sizes and exclusion criteria were determined in advance.

Methods

As outlined in our pre-registration (https://aspredicted.org/WJB_443), we aimed to recruit 375 U.S. participants from Amazon Mechanical Turk. All participants were first asked to complete the same audio check as the main experiment. As pre-registered, the survey was terminated for participants who responded to the audio check incorrectly, and we removed all cases of multiple responses from the same participant, as well as participants who responded to the survey from a mobile device. We were left with a final sample of 347 participants (46% female, mean age = 40.48 years).

All participants were presented with 12 products in a randomized order from one randomly chosen product category. The six most popular product categories on Fiverr were selected to be included in the experiment (3 were experienceable: Logo Design, Voice Over, Whiteboard & Animated Explainers; 3 were non-experienceable: Translation, Articles & Blog Posts, Social Media Marketing). The presentation of the products imitated the look and feel of browsing on Fiverr, and each product was presented in the same way as on Fiverr (see Figure A1.5 for an example from the Articles & Blog Posts product category). All assortments had an identical distribution of price, average rating, and number of ratings, and there were no dominance relationships in any assortment. In the voice over product category, participants could play a sample audio clip from the seller, and in the Whiteboard & Animated Explainers product category, participants could play a sample video animation from the seller.

All participants were asked to select which option they would purchase, if they could only purchase one of them. Subsequently, all participants were asked the dependent variable question: how certain they were that the choice they selected was the right one for them on a scale from 0 (not at all certain) to 10 (completely certain). Lastly, we collected participants' gender and age.

Results

We found that preference certainty was greater in the experienceable condition ($M = 6.77$, $SD = 2.53$) than in the non-experienceable condition ($M = 6.17$, $SD = 2.79$, $t(344) = 2.10$, $p = .036$). 45.9% of participants who were presented with products in either the Voice Over or Whiteboard & Animated Explainers product categories clicked to play a sample of at least one product.

Discussion

These findings indicate that in a setting mimicking Fiverr, participants reported greater preference certainty in product categories that were experienceable than in those that were non-experienceable.

Study A1.2. Laboratory Experiment Pre-test: Experienceability and Preference Uncertainty.

We pre-tested whether the ability to experience the song alternatives used in the laboratory experiment was associated with preference uncertainty. We report all manipulations and measures, and recruited a minimum of 100 participants per cell. All sample sizes and exclusion criteria were determined in advance.

Methods

As outlined in our pre-registration (https://aspredicted.org/K14_MVF), we aimed to recruit 200 U.S. participants from Amazon Mechanical Turk. All participants were first asked to complete the same audio check as the main experiment. As pre-registered, the survey was terminated for participants who responded to the audio check incorrectly, and we removed all cases of multiple responses from the same participant. We were left with a final sample of 198 participants (42% female, mean age = 40.94 years).

Participants were randomized into one of two conditions: non-experienceable or experienceable. In both conditions, participants were asked to imagine that they were interested in purchasing a song, and given a choice set of two songs. All participants were shown a table with songs A and B from Table 1 (in a randomized order). In the experienceable condition, participants were additionally given 33 second samples for each of the songs that they could listen to. For each participant in the experienceable condition, we randomly selected two of the three songs samples used in the main experiment, and we randomized the pairings between the audio clips and the song attributes.

All participants were asked to select which song they would purchase, if they could only purchase one of them. Subsequently, we asked participants how certain they were that the choice they selected was the right one for them on a scale from 0 (not at all certain) to 10 (completely certain). Lastly, we collected participants' gender and age.

Results

We found that the song choices participants made in the experienceable and non-experienceable conditions were not significantly different. The choice share of Song A was 53.5% in the non-experienceable condition and 45.5% in the experienceable condition ($\chi 2(1) = 0.99$, $p = .32$, $\varphi = 0.08$). This likely occurred because the audio clips were randomly paired with song attributes for each participant, and because the experience was not very differentiating – all songs had a similar melody. However, preference certainty was greater in the experienceable condition ($M = 7.81$, $SD = 1.90$) than in the non-experienceable condition ($M = 6.42$, $SD = 2.35$, $t(188) = 4.56$, $p < .001$). 98.0% of participants in the experienceable condition sampled both songs at least once.

Discussion

These findings indicate that the ability to sample and experience songs reduced participants' reported preference uncertainty associated with their choice, even though experience did not systematically alter choices.

References

Aaker J (1991) The Negative Attraction Effect? A Study of the Attraction Effect Under Judgment and Choice. *ACR North American Advances* NA-18.

Ariely D, Wallsten TS (1995) Seeking Subjective Dominance in Multidimensional Space: An Explanation of the Asymmetric Dominance Effect. *Organizational Behavior and Human Decision Processes* 63(3):223-232.

Bateson M, Healy SD, Hurly TA (2002) Irrational choices in hummingbird foraging behaviour. *Animal Behaviour* 63(3):587-596.

Bergé L (2018) Efficient estimation of maximum likelihood models with multiple fixed-effects: the R package FENmlm. *CREA Discussion Papers* (13).

Bettman JR, Luce MF, Payne JW (1998) Constructive Consumer Choice Processes. *Journal of Consumer Research* 25(3):187-217.

Bushong B, Rabin M, Schwartzstein J (2021) A Model of Relative Thinking. *The Review of Economic Studies* 88(1):162-191.

Cuofano G. (2019) How Does Fiverr Work And Make Money? Fiverr Business Model In A Nutshell. *FourWeekMBA.*

Dhar R, Simonson I (2003) The Effect of Forced Choice on Choice. *Journal of Marketing Research* 40(2):146-160.

Doyle JR, O'Connor DJ, Reynolds GM, Bottomley PA (1999) The robustness of the asymmetrically dominated effect: Buying frames, phantom alternatives, and in-store purchases. *Psychology & Marketing* 16(3):225-243.

Ert E, Lejarraga T (2018) The Effect of Experience on Context-dependent Decisions. *Journal of Behavioral Decision Making* 31(4):535-546.

Fazio RH, Zanna MP (1978) Attitudinal qualities relating to the strength of the attitude-behavior relationship. *Journal of Experimental Social Psychology* 14(4):398-408.

Frederick S, Lee L, Baskin E (2014) The Limits of Attraction. *Journal of Marketing Research* 51(4):487-507.

Geyskens I, Gielens K, Gijsbrechts E (2010) Proliferating Private-Label Portfolios: How Introducing Economy and Premium Private Labels Influences Brand Choice. *Journal of Marketing Research* 47(5):791-807.

Griffin D, Liu W, Khan U (2005) A New Look at Constructed Choice Processes. *Marketing Letters* 16(3):321-333.

Hadar L, Danziger S, Hertwig R (2018) The Attraction Effect in Experience-based Decisions. *Journal of Behavioral Decision Making* 31(3):461-468.

Hamilton R, Hong J, Chernev A (2007) Perceptual Focus Effects in Choice. *Journal of Consumer Research* 34(2):187-199.

Heath TB, Chatterjee S (1995) Asymmetric Decoy Effects on Lower-Quality versus Higher-Quality Brands: Meta-analytic and Experimental Evidence. *Journal of Consumer Research* 22(3):268-284.

Hertwig R, Pleskac TJ (2010) Decisions from experience: Why small samples? *Cognition* 115(2):225-237.

Hoch SJ (2002) Product Experience Is Seductive. *Journal of Consumer Research* 29(3):448-454.

Huber J, Puto C (1983) Market Boundaries and Product Choice: Illustrating Attraction and Substitution Effects. *Journal of Consumer Research* 10(1):31-44.

Huber J, Payne JW, Puto C (1982) Adding Asymmetrically Dominated Alternatives: Violations of Regularity and the Similarity Hypothesis. *Journal of Consumer Research* 9(1):90-98.

Huber J, Payne JW, Puto CP (2014) Let's be Honest about the Attraction Effect. *Journal of Marketing Research* 51(4):520-525.

Kahneman D, Tversky A (1984) Choices, values, and frames. *American Psychologist* 39(4):341-350.

Kelman M, Rottenstreich Y, Tversky A (1996) Context-Dependence in Legal Decision Making. *The Journal of Legal Studies* 25(2):287-318.

Kivetz R, Netzer O, Srinivasan V (2004) Alternative Models for Capturing the Compromise Effect. *Journal of Marketing Research* 41(3):237-257.

Lancaster T (2000) The incidental parameter problem since 1948. *Journal of Econometrics* 95(2):391-413.

Landry P, Webb R (2021) Pairwise normalization: A neuroeconomic theory of multi-attribute choice. *Journal of Economic Theory* 193:105221.

Lea AM, Ryan MJ (2015) Irrationality in mate choice revealed by túngara frogs. *Science*.

Li M, Sun Y, Chen H (2019) RETRACTED:The Decoy Effect as a Nudge: Boosting Hand Hygiene With a Worse Option. *Psychological Science* 30(1):139-149.

Lichters M, Sarstedt M, Vogt B (2015) On the practical relevance of the attraction effect: A cautionary note and guidelines for context effect experiments. *AMS Review* 5(1):1-19.

Luce RD (1959) *Individual choice behavior, a theoretical analysis* (Wiley, New York).

Mailath GJ, Postlewaite A (2003) The Social Context of Economic Decisions. *Journal of the European Economic Association* 1(2-3):354-362.

Malkoc SA, Hedgcock W, Hoeffler S (2013) Between a rock and a hard place: The failure of the attraction effect among unattractive alternatives. *Journal of Consumer Psychology* 23(3):317-329.

Padamwar PK, Dawra J, Kalakbandi VK (2021) The impact of range extension on the attraction effect. *Journal of Business Research* 126:565-577.

Pocheptsova A, Amir O, Dhar R, Baumeister RF (2009) Deciding without Resources: Resource Depletion and Choice in Context. *Journal of Marketing Research* 46(3):344-355.

Prelec D, Wernerfelt B, Zettelmeyer F (1997) The Role of Inference in Context Effects: Inferring What You Want from What Is Available. *Journal of Consumer Research* 24(1):118-125.

R Core Team (2021) *R: A Language and Environment for Statistical Computing* (R Foundation for Statistical Computing, Vienna, Austria).

Schuck-Paim C, Pompilio L, Kacelnik A (2004) State-Dependent Decisions Cause Apparent Violations of Rationality in Animal Choice. *PLOS Biology* 2(12):e402.

Schwartz JA, Chapman GB (1999) Are More Options Always Better?: The Attraction Effect in Physicians' Decisions about Medications. *Medical Decision Making* 19(3):315-323.

Shafir S, Waite TA, Smith BH (2002) Context-dependent violations of rational choice in honeybees (Apis mellifera) and gray jays (Perisoreus canadensis). *Behavioral Ecology and Sociobiology* 51(2):180-187.

Shugan SM (1980) The Cost Of Thinking. *Journal of Consumer Research* 7(2):99-111.

Similarweb. (2022). fiverr.com Traffic Analytics & Market Share.

Simonson I (1989) Choice Based on Reasons: The Case of Attraction and Compromise Effects. *Journal of Consumer Research* 16(2):158-174.

--- (2014) Vices and Virtues of Misguided Replications: The Case of Asymmetric Dominance. *Journal of Marketing Research* 51(4):514-519.

Smith RE, Swinyard WR (1983) Attitude-Behavior Consistency: The Impact of Product Trial versus Advertising. *Journal of Marketing Research* 20(3):257-267.

Tetlock PE (1985) Accountability: The neglected social context of judgment and choice. *Research in Organizational Behavior* 7:297-332.

Thaler RH (1999) Mental accounting matters. *Journal of Behavioral Decision Making* 12(3):183-206.

--- (2015) *Misbehaving: The making of behavioral economics.* (W W Norton & Co, New York, NY, US).

Thomadsen R, Rooderkerk RP, Amir O, Arora N, Bollinger B, Hansen K, John L, Liu W, Sela A, Singh V, Sudhir K, Wood W (2018) How Context Affects Choice. *Customer Needs and Solutions* 5(1):3-14.

Trueblood JS, Brown SD, Heathcote A (2014) The multiattribute linear ballistic accumulator model of context effects in multialternative choice. *Psychological Review* 121:179-205.

Trueblood JS, Brown SD, Heathcote A, Busemeyer JR (2013) Not Just for Consumers: Context Effects Are Fundamental to Decision Making. *Psychological Science* 24(6):901-908.

Tversky A, Simonson I (1993) Context-Dependent Preferences. *Management Science* 39(10):1179-1189.

Wedell DH (1991) Distinguishing among models of contextually induced preference reversals. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 17:767-778.

Wedell DH, Pettibone JC (1996) Using Judgments to Understand Decoy Effects in Choice. *Organizational Behavior and Human Decision Processes* 67(3):326-344.

Wernerfelt B (1995) A Rational Reconstruction of the Compromise Effect: Using Market Data to Infer Utilities. *Journal of Consumer Research* 21(4):627-633.

Wu C, Cosguner K (2020) Profiting from the Decoy Effect: A Case Study of an Online Diamond Retailer. *Marketing Science* 39(5):974-995.

Yang S, Lynn M (2014) More Evidence Challenging the Robustness and Usefulness of the Attraction Effect. *Journal of Marketing Research* 51(4):508-513.

CHAPTER 2

INDIVIDUALS PREFER TO HARM THEIR OWN GROUP RATHER THAN HELP AN OPPOSING GROUP

Rachel Gershon* and Ariel Fridman*

*denotes equal contribution

Rady School of Management, University of California, San Diego, La Jolla, CA, 92093, USA

Abstract

Group-based conflict enacts a severe toll on society, yet the psychological factors governing behavior in group conflicts remain unclear. Past work finds that group members seek to maximize relative differences between their in-group and out-group ("in-group favoritism") and are driven by a desire to benefit in-groups rather than harm out-groups (the "in-group love" hypothesis). This prior research studies how decision makers approach tradeoffs between two net-positive outcomes for their in-group. However, in the real world, group members often face tradeoffs between net-negative options, entailing either losses to their group or gains for the opposition. Anecdotally, under such conditions, individuals may avoid supporting their opponents even if this harms their own group, seemingly inconsistent with "in-group love" or a harm minimizing strategy. Yet, to the best of our knowledge, these circumstances have not been investigated. In six pre-registered studies, we find consistent evidence that individuals prefer to harm their own group rather than provide even minimal support to an opposing group across polarized issues (abortion access, political party, gun rights). Strikingly, in an incentive-compatible experiment, individuals preferred to subtract more than three times as much from their own group rather than support an opposing group, despite believing that their in-group is more effective with funds. We find that identity concerns drive preferences in group decision making, and individuals believe that supporting an opposing group is less value-compatible than harming their own group. Our results hold valuable insights for the psychology of decision making in intergroup conflict as well as potential interventions for conflict resolution.

Introduction

Group conflicts are a pervasive feature of society. Yet, despite the extensive literature on the topic, a unified understanding of the psychology underlying decision making in group conflicts remains elusive. Prior work has documented two broad principles governing group-based decision making. First, group members exhibit in-group favoritism (Tajfel et al. 1971, Turner 1975). That is, individuals prefer to create a favorable comparison between their in-group and the out-group, even leading to choices that prioritize relative gains compared to the out-group over greater absolute gains for their in-group (see Social Identity Theory; Abrams and Hogg 1990, Hogg 2016). Second, past work theorizes that group members are driven by a cooperative motive to help the in-group ("in-group love") rather than an aggressive motive to hurt the out-group ("out-group hate"; Halevy et al. 2008, Halevy et al. 2012). Critically, these principles are derived from studies in which participants chose between outcomes that are all ultimately favorable to the in-group. However, real-world decision making often entails making choices where harm is unavoidable (Volz et al. 2017, Berman and Kupor 2020). Groups may have to choose between in-group losses and out-group gains, a circumstance that, to the best of our knowledge, has not been previously studied, and reveals that individuals' decisions cannot be explained by existing theories.

Consider the example of Montgomery, Alabama, where only white residents were allowed to use the publicly funded Oak Park Pool until, in 1959, a federal court deemed the segregated pool unconstitutional. The white town council then faced two options that they considered unfavorable: give Black families access to the pool or close one of the town's favorite gathering spots. Previous work on in-group favoritism and the dominant role of "in-group love" would predict that white Montgomery residents would avoid harming their own group even at the cost of extending pool use to Black citizens. Yet famously, the white citizens of Montgomery closed the pool. Other public resources, such as parks and zoos, were closed to all across the country to defy similar rulings (McGhee 2021). The Oak Park Pool is an exemplar of the fact that, anecdotally, when faced with two counter-attitudinal choices – aid the out-group or harm the in-group – group members may avoid showing support for the out-group, even at the apparent expense of their own side (Morewedge et al. 2018, Van Boven et al. 2018, Benen 2021, Friedersdorf 2021). However, there has never been a rigorous investigation of how individuals navigate unfavorable choices in inter-group conflicts, and whether there is a broad preference to harm one's own group rather than support the opposition.

Here, we develop a novel paradigm in which individuals must either deduct funds from their in-group or add funds to an opposing group to examine how group members make tradeoffs in lose-lose[19] intergroup conflicts. Our experiments were conducted across multiple countries (United States and United Kingdom), several polarized issues (abortion access, political party, and gun control), and various experimental measures (financial donations and incentive-compatible multiple price lists). Taken together, our results offer the first unambiguous evidence that individuals are so averse to showing support for an opposing out-group that they even prefer to do greater harm to their own group instead. Our finding was symmetrically exhibited by individuals on both sides of each issue we studied, and even among participants who identified only weakly with their side. However, the degree to which a participant identified with their side of the focal issue does play a moderating role – those with stronger attitudes in favor of their side (e.g., more strongly pro-choice) were more likely to choose to harm their own side (and willing to cause greater overall harm to their side) rather than help the opposing side.

[19] While the term "lose-lose" is often used to express a loss for each side, here we define lose-lose (win-win) to describe situations where both options are unfavorable (favorable) to the decision maker.

Our results reveal the central role of identity in decision making in polarized contexts. Identity often plays an important role in decision making (Akerlof and Kranton 2000), such as when and to whom we offer support (Swann Jr et al. 2010, Kessler and Milkman 2018). We propose that individuals aim to protect their group-based identity when facing intergroup conflict, and therefore behave in ways that best express their values, especially those that are central to their identity (Hitlin 2003). Previous work finds that individuals prefer expressions of support (e.g., "I support Democrats") to expressions of opposition (e.g., "I oppose Republicans"), because support is considered more "value expressive" (Zhong et al. 2008, Catapano and Tormala 2021). We would therefore expect group members to choose actions of support over actions of opposition when both options convey their values (i.e., choosing to help their own side rather than harm their opponent). However, what would individuals choose when the options are unfavorable, inconsistent with their values (i.e., a lose-lose choice between harming their own group or helping their opponent)? In such situations, since both options express values that are counter-attitudinal, our Identity-Support model suggests that individuals will choose the least value expressive option, thereby best protecting their identity. Our results support these predictions – individuals believe that helping an opposing group is more harmful to their identity than inflicting equivalent harm to their in-group, even when this leads to a worse relative standing for their in-group. Critically, we find that by ameliorating identity concerns through shifting perceived in-group norms (Borsari and Carey 2003, Cialdini 2003, Hogg and Reid 2006, Gerber and Rogers 2009, Allcott 2011), individuals become more likely to support the opposing group, providing a practical way to achieve more constructive outcomes in group conflict.

While there are numerous existing models of group decision making, the Identity-Support model can uniquely explain individual behavior for both favorable (win-win) and unfavorable (lose-lose) choices with opposing groups. We address other frameworks below in Table 2.1. First, it is possible that individuals view these decisions as a zero-sum tradeoff (Wilkins et al. 2015, Davidai and Ongis 2019), wherein a gain for one group is perceived as an equivalent loss for the other group. In this case, group members should be indifferent between the two options. It is also possible that group members consider which side is more effective at using funds to pursue their mission and then choose based on harm minimization for their cause. Across our studies, we find that group members typically believe their own group is more effective with funds, so those motivated by harm minimization should prefer to add to their own (more effective) side in win-win scenarios and avoid subtracting from their own side in lose-lose scenarios.

The social value orientation literature (e.g., Bornstein et al. 1983), which arbitrarily assigns in-groups and out-groups (i.e., a minimal group paradigm) to study group-member decision making, describes additional motives that may also guide decision making for real-world opposing groups. For example, individuals may simply prefer allocations that maximize the payoff for the in-group (analogous to being motivated by in-group love (Halevy et al. 2008, Halevy et al. 2012)), maximize the relative difference in payoff between their in-group and opposing group (consistent with findings of in-group favoritism (Tajfel et al. 1971, Turner 1975)), or, as demonstrated in Bornstein et al. (1983), either minimize the difference in payoffs to each side or maximize joint profit in favor of the in-group. Additionally, recent work on negative partisanship and increased out-group animosity (Abramowitz and Webster 2016, Abramowitz and Webster 2018, Finkel et al. 2020), suggests that individuals may be primarily motivated to minimize the payoff to the opposing side. Table 2.1 summarizes how individuals guided by each of these motives would choose when faced with win-win and lose-lose scenarios. Notably, the Identity-Support model is the only model that can explain both findings across scenarios.

Finally, previous work on prospect theory (Tversky and Kahneman 1991, 1992), finding that people experience losses more strongly than equivalent gains, also does not make clear predictions in this context. For example, when faced with a lose-lose scenario, while participants technically choose between a loss or gain (of funding), they may encode both adding to the opposing group and subtracting from the in-group as losses for their side. That is, any gains for the opposing group can feel like losses for the in-group and vice versa. However, if this is not the case and individuals are indeed more affected by losses to their side than gains to the opposing side, then our findings would be inconsistent with predictions based on prospect theory as well.

Open science. All study designs and analyses were pre-registered and all data, analysis code, research materials, and pre-registrations are available at https://osf.io/gzxke/. Data was analyzed using R, version 4.1.2, and the package ggplot2, version 3.3.5 (Wickham et al. 2016). For all studies, we reported all manipulations and measures, and recruited a minimum of 100 participants per condition. All sample sizes and exclusion criteria were determined in advance.

Results

Study 1: Individuals Prefer to Harm their own Group rather than Support the Opposition

In Study 1, we develop a novel paradigm to investigate how individuals behave in group conflicts in which they must choose between supporting an opposing group or harming their own group. In our paradigm, we asked participants to indicate their position on several polarized issues in two sub-studies conducted with participants from the United States and United

Kingdom. Given the similarities in their designs and hypotheses, we report these studies together, noting only where they differ. In the U.S. study (Study 1A; N = 797, matched to U.S. census data on age, sex, and ethnicity), issues included abortion access, gun control, and political party affiliation (Democratic or Republican party). In the UK study (Study 1B; N = 393, matched to UK census data on age, sex, and ethnicity), participants were asked about political party affiliation (Labour or Conservative Party). After indicating their attitude toward the relevant issues (on a 6-point Likert scale[20], which we subdivide into weak, medium, and strong attitudes; see Methods section), participants learned that real donations would be made to organizations supporting each side of the partisan divide (e.g., both a pro-choice and pro-life organization). Participants were then asked how they would choose to alter the donation (in Study 1A they rated each of the three issues separately, randomly ordered), and were informed that for 10 randomly chosen participants, their choices would actually be implemented. Participants were randomly assigned to one of two between-subjects experimental conditions: (a) a win-win condition or (b) a lose-lose condition. In the win-win condition, which serves as an experimental control to conceptually replicate prior findings (Halevy et al. 2008, Halevy et al. 2012), both options altered the donation in ways that were favorable given the participant's stated attitude: either add $1 to the donation going to the organization on their side or subtract $1 from the donation going to the organization on the opposing side. In the lose-lose condition, both options altered the donation in ways that were unfavorable given the participant's stated attitude: either add $1 to the donation going to the organization on the opposing side or subtract $1 from the donation going to the organization on their side. Finally, participants reported which side of each

[20] We note that for this and all subsequent studies, we asked all participants to indicate their position on each issue without allowing them to express indifference. While this introduces noise by forcing truly indifferent participants to choose a side, it would not bias the results in a particular direction and works against finding an overall effect.

cause was more effective at using funds to pursue their mission. Specifically, they were asked "Do you believe that [Pro-life/Pro-gun/Republican/Conservative] or [Pro-choice/Anti-gun/Democratic/Labour] organizations are more effective at pursuing their mission? In other words, which one is able to do more with each dollar they receive?"

In the win-win condition, our results conceptually replicate and extend prior studies. We find that participants are more likely to choose to support their side by adding $1 to the organization on their side of each cause (72.5%) rather than harm the opposition by subtracting $1, $t(594) = 13.87$, $p < .001$, 95% CI = [69.7%, 75.2%][21] (see Figure 2.1). Thus, these results reproduce the findings of the in-group love model in the context of a win-win choice in our sample, extending previous results to natural groups (as opposed to minimal group paradigms).

The results of our lose-lose condition, however, contradict the predictions of the in-group love model. Whereas in-group love predicts that individuals will support their opposition to avoid harming their in-group, participants in the lose-lose condition predominantly chose not to help the opposition (helping the opposing group was chosen by only 25.8% of participants), preferring instead to harm their own group almost three-quarters of the time, $t(594) = -13.85$, $p < .001$, 95% CI = [23.0%, 28.7%] (see Figure 2.1).

The preference to harm the in-group rather than support the opposition is robust and consistent across a variety of subsamples. First, we find similar effects for both the nationally representative sample in the United States (24.6%; $t(398) = -12.92$, $p < .001$, 95% CI = [21.6%, 27.9%]) and the nationally representative sample in the United Kingdom (32.7%; $t(195) = -4.74$, $p < .001$, 95% CI = [26.4%, 39.5%]). Second, the same pattern holds across every issue we tested (abortion: (27.8%; $t(398) = -8.52$, $p < .001$, 95% CI = [23.6%, 32.4%]); gun control:

[21] Statistical tests are from logistic regressions, and account for repeated observations from each participant across issues by clustering the standard errors by participant.

(25.3%; t(398) = -9.39, p < .001, 95% CI = [21.3%, 29.8%]); party support: (24.7%; t(594) = -11.71, p < .001, 95% CI = [21.4%, 28.3%]). Finally, the preference to harm the in-group rather than support the opposition is present on both sides of the ideological spectrum[22]: (liberals: 26.4%; t(477) = -11.46, p < .001, 95% CI = [23.1%, 29.9%]; conservatives: 24.1%; t(235) = -8.59, p < .001, 95% CI = [19.6%, 29.2%]). The results from the win-win condition were also consistent and significant across these robustness checks (all ps < .001). See Table 2.2 for a summary of results across all studies.

Although individuals prefer to deduct funding from their in-group rather than support the opposing group in this study with real consequences, one possible explanation is that participants view this as the better outcome for their cause. Indeed, if individuals believe that the opposing group is more effective at advancing their interests per dollar spent than their own group, then reducing equivalent funding to the in-group maximizes the relative difference in outcomes between groups in a lose-lose decision. However, our results show the reverse: participants indicated that they view organizations supporting their side of an issue to be *more* effective in spending donation money to achieve their goal (see Appendix Results for Study 1). To verify that our results are not driven primarily by individuals who believe the opposition is more effective with funds, we conducted an ancillary analysis in which we studied the subset of participants who believe organizations on their side of an issue are strictly more effective with funds (43.8% of observations). These participants in the lose-lose condition *still* preferred to "harm" their side (supporting the opposing side was chosen by only 22.8% of these participants; t(402) = -10.82, p < .001, 95% CI = [19.2%, 27.0%].

[22] 69% of participants in the U.S. sample held either liberal (pro-choice, anti-gun ownership, and Democrat) or conservative (pro-life, pro-gun ownership, and Republican) positions on all three issues.

Finally, we conducted an initial test of our central theory that individuals make decisions in group conflicts on the basis of protecting their identity (developed in-depth in Studies 3 and 4). As previously outlined, individuals believe that acts of support are more value-expressive than acts of opposition. Individuals may choose the least value expressive option when offered two unfavorable choices, thereby opposing their own group rather than supporting the opposing group. We predicted that this desire to protect their identity (and therefore the choice not to support the opposition) should be stronger for participants with stronger group identities, as assessed by stronger attitudes about the underlying issue. Indeed, in the lose-lose condition, we find that those with strong attitudes were even less likely to choose to add $1 to the opposing side (vs. subtract $1 from their side), compared to those with medium ($\beta = .89$, $SE = .17$, $z = 5.21$, $p < .001$, $OR = 2.43$, 95% CI = [1.74, 3.40]) or weak ($\beta = 1.05$, $SE = .18$, $z = 6.01$, $p < .001$, $OR = 2.87$, 95% CI = [2.04, 4.05]) attitudes. In the win-win condition, we find that those with strong attitudes were more likely to choose to add $1 to their own side (vs. subtract $1 from the opposing side), compared to those with medium ($\beta = -.36$, $SE = .17$, $z = -2.13$, $p = .03$, $OR = .70$, 95% CI = [.50, .97]) or weak ($\beta = -.83$, $SE = .16$, $z = -5.10$, $p < .001$, $OR = .44$, 95% CI = [.32, .60]) attitudes.

Taken together, these results establish that not only do group members prefer to deduct funds from their in-group rather than contribute an equivalent amount to their opposition, but they make this choice despite explicitly believing that this leaves their group worse off than the alternative. While motives such as harm minimization, in-group favoritism, and in-group love cannot explain the findings from this study, the results are consistent with predictions from the Identity-Support model. Having established that our findings in win-win scenarios are consistent with previous work, we focus on decision making in our novel lose-lose paradigm in the remaining studies.

Study 2: Quantifying the Aversion to Supporting an Opposing Group

Whereas Study 1 establishes that individuals prefer to deduct a given amount of funds from their in-group rather than add the same amount to the opposition, Study 2 quantifies the degree to which individuals prefer to harm their own group rather than support an opposing group. We approach this quantitative analysis by eliciting participants' indifference amount between harming their own group and supporting the opposition using an incentive-compatible choice titration procedure (Becker et al. 1964; see Methods section for more detail). For this choice titration analysis, we recruited 300 U.S. participants from Amazon's Mechanical Turk with a final sample of 268 following our pre-registered exclusions. Given that we found consistent evidence across issues in Study 1, here we focused on a single polarizing issue: abortion access.

We informed participants that the researchers would make two $10 donations: one to a pro-life organization and another to a pro-choice organization. Participants were asked to choose how to alter the donation amount in a series of 14 choices. For each choice they could select to either add $1 to the opposing organization's donation or subtract an amount (sequentially from $0.10 to $10, order counterbalanced) from their own side's organization (similar to a price list; Andersen et al. 2006). To incentivize responses, we informed participants that for 1 in 10 participants, chosen at random, we would actually make donations to both organizations and randomly select one of their 14 choices to alter the donation amount (for similar elicitation and bonus procedures, see (Lerner et al. 2013, DeSteno et al. 2014, Dorison et al. 2020)). As in Study

1, participants also reported their beliefs about the relative effectiveness of pro-life and pro-choice organizations, as well as their attitude strength toward the issue.

Our results show that participants' indifference amounts had a mean value significantly greater than $1 (M = $3.85, median = $1.50, max = $10, SD = $4.10, $t(267) = 11.39$, $p < .001$, $d = .70$, 95% CI = [.56, .83]), and the majority of participants (65%) had an indifference amount greater than $1 ($\chi^2(1, N = 268) = 23.29$, $p < .001$; see Figure 2.2). In other words, on average participants required almost $4 to be subtracted from the donation going to their organization, to be indifferent toward adding $1 to the opposing organization. Strikingly, 28% of all participants chose to entirely forgo the $10 donation to their side rather than add $1 to the opposing organization. Lastly, we found that those with greater attitude strength (i.e., more strongly pro-choice or pro-life) required more funds to be subtracted from their side to be indifferent toward adding $1 to the opposing side, and we also replicated the finding that participants believed that organization on their side are more effective at spending their donation money to achieve their goal (for details on both results, see Appendix Results for Study 2).

In sum, Study 2 demonstrates the strength of participants' preferences when facing lose-lose choices. Despite believing that their own side is more effective with funds, group members preferred to subtract, on average, more than three times as much from their own group rather than give a small amount of support to the opposing group.

Study 3: Identity Concerns Trump Effectiveness Considerations for Lose-Lose Decisions

Studies 1-2 provide evidence across contexts that individuals are so averse to supporting an opposing group that they prefer to harm their own group instead. Moreover, the finding that individuals with stronger attitudes toward the focal issue (and therefore stronger identity-relevance; Hogg and Smith 2007, Smith and Hogg 2008) are more prone to harming their in-group rather than helping the opposing group offers preliminary evidence that identity considerations govern decision making in group conflicts. In Study 3, we directly test the hypothesis that identity concerns, as opposed to effectiveness considerations, underlie the psychology of decision making in group conflicts involving lose-lose choices.

To test the relative contributions of identity concerns and effectiveness considerations, we recruited 400 U.S. participants from Amazon's Mechanical Turk with a final sample of 393 following our pre-registered exclusions. After indicating whether they identified more strongly as Republican or Democrat (binary choice), participants were asked to make a lose-lose choice, identical to the U.S. political party choice from Study 1. Subsequently, participants responded to an effectiveness and an identity concern question, both on a 7-point Likert scale, asked in a randomized order. The effectiveness question was similar to the one used in the previous studies, and the identity concern question asked participants whether adding $1 to the opposing side or subtracting $1 from their side "would make you feel like a worse [Democrat/Republican]? In other words, which option most undermines your identity as a [Democrat/Republican]?".

As in our previous studies, less than half of participants chose to add $1 to the organization supporting the opposing political party (36.39%), $\chi^2(1, N = 393) = 38.59$, $p < .001$, for both Democrats and Republicans (both Ps < .001). Moreover, as in the previous studies, the overall effectiveness measure was positive ($M = 0.61$, $SD = 1.35$), $t(392) = 8.95$, $p < .001$, $d = .45$, 95% CI = [.35, .55]), indicating that, on average, participants view the organization on their side to be more effective at using donated funds to pursue their mission.

Critically, the identity concern measure revealed that participants believed adding $1 to the opposition undermined their partisan identity more than deducting $1 from their own party ($M = -0.22$, $SD = 1.99$, $t(392) = -2.16$, $p = .03$, $d = -0.11$, 95% CI = [-0.21, -.01]). The fact that individuals perceive their group-based identities to be at greater risk when supporting an opposing group (vs. harming their own) offers a clear rationale for why they choose in-group harm over supporting the opposing group. However, to explicitly test the contributions of identity concerns and effectiveness considerations on choice, we regressed participant choice (0 = subtract \$1, 1 = add \$1) on both our identity and effectiveness measures using a logistic regression. The identity concern measure was positively associated with choice, such that participants were more likely to choose to add to the opposing side when they believed that subtracting \$1 from their side undermined their identity more ($\beta = .65$, SE = .07, z = 8.77, $p < .001$, OR = 1.91, 95% CI = [1.66, 2.22]). The effectives measure was not significantly associated with choice ($\beta = -.09$, SE = .09, z = -0.97, p = .33, OR = .91, 95% CI = [.76, 1.09]). A similar pattern of results was observed when choice was regressed on each measure in separate regressions. See Appendix Results for Study 3 for regression results, and Figure 2.3 for plots of both measures.

In summary, Study 3 explicitly establishes that individuals feel that their group identity suffers more when supporting the opposition than when harming the in-group. Moreover, in a direct test of the relative contributions of identity concerns and effectiveness considerations, we find that whereas the identity measure is significantly correlated with the choice to harm the in-group, there is no significant association between the effectiveness measure and individual decision making in lose-lose group conflicts.

These findings align with our Identity-Support model of group decision making, pointing to the important role identity plays in decisions involving inter-group conflict. In a pre-registered supplemental study (Appendix Study A2.1), we test another key element of the model, building on relevant research (Catapano and Tormala 2021) – that acts of support are more value-expressive than acts of opposition. Participants read about and evaluated another participant from their in-group (Republican or Democrat) who had to make either a lose-lose or win-win allocation (as in the U.S. political party choice from Study 1) and were told which option the participant chose. Based on this decision, they were asked to assess how strongly they believed this target identified with their political party. In the win-win scenario, those who opted to support their in-group were perceived as identifying more strongly with their party than those who subtracted money from the opposing group. In the lose-lose scenario, those who opted to support the opposing side were perceived as more weakly identifying with their party than those who subtracted from their in-group. This suggests that participants do in fact believe that acts of support are more value-expressive than acts of opposition using our paradigm. In line with the Identity-Support model, participants should prefer the choice that best promotes or protects their identity, and therefore choose the most value-expressive option (supporting their own side) when offered pro-attitudinal choices and avoid the most value-expressive option (supporting the other side) when offered counter-attitudinal choices.

Study 4: Increasing the Salience of a Group-Based Identity Decreases the Probability of Supporting the Opposing Group

Whereas Study 3 offers correlational evidence that identity considerations govern the preference to harm one's in-group to avoid supporting the opposition, in Study 4 we causally test the relationship between identity concerns and choice in lose-lose group conflicts. Building on previous work showing that identities are malleable and making certain identities more salient can affect preferences and behaviors (LeBoeuf et al. 2010, Van der Werff et al. 2014, Kessler and Milkman 2018), we hypothesized that strengthening group identity salience would lead to an increase in the probability of harming one's own group to avoid supporting the opposing group.

To test our hypothesis, participants were randomized into one of two conditions: in the identity-strengthened condition, participants were asked to write about an event, story, or personal experience where they strongly identified with their political party. In the control condition, participants wrote about what they do on a typical Monday evening. Participants then made the same choice as in Study 3 – add $1 to the donation going to the organization supporting the opposing political party or subtract $1 from the donation going to the organization supporting their own political party.

We recruited 500 participants from Amazon's Mechanical Turk, with a final sample of 497 following our pre-registered exclusions. Consistent with our previous studies, in the control condition, the proportion of participants choosing to add $1 to the organization supporting the opposing political party was less than 50% (41.11%), $\chi^2(1, N = 270) = 8.18$, $p = .004$. In the identity-strengthened condition, the proportion of participants choosing to add $1 was significantly lower (30.4%) than in the control group, $\chi^2(1, N = 497) = 5.67$, $p = .017$, $\varphi = 0.11$ (see Figure 2.4). This result held among the subset of participants who believed organizations on their side of a cause are more effective with funds (74% of participants; $\chi^2(1, N = 369) = 5.36$, $p = .021$, $\varphi = 0.13$), Democratic participants (62% of participants; $\chi^2(1, N = 308) = 4.24$, $p = .039$, $\varphi = 0.12$), and was directional but non-significant among the (relatively smaller) subset of Republican participants (38% of participants; $\chi^2(1, N = 189) = 1.20$, $p = .274$, $\varphi = 0.09$). As in prior studies, we also found that participants believed that organizations on their side are more effective at spending their donation money to achieve their goal (see Appendix Results for Study 4). In sum, our results provide additional evidence for the role of identity on decision making in group conflict, demonstrating a causal effect of identity salience on the decision to harm one's in-group rather than support the opposition.

Finally, we note that Study 4 was designed to also address the potential confound that participants may disproportionately focus on the most negative ways in which the opposing group would use donated funds, but do not similarly consider the most positive ways in which the in-group would use donated funds (Rozin and Royzman 2001). Such a difference could explain the strong aversion to supporting the opposing side that we find across lose-lose choices (although would not explain the win-win preference to help one's own side rather than harm the opposition). Consequently, in this study, we also specified that all donations would go to "administrative costs (e.g., maintaining the organization's website)" to hold constant the use of donations, ensuring that participants would imagine similar donation uses for each organization. We therefore conclude that the preference to harm one's own group rather than support an opposing group is not explained by different imagined uses of the funds by the in-group and opposing group.

Study 5: Modulating Group Norms Alters Decision Making in Group Conflicts

In Study 5, we test a practical method for shifting behavior in lose-lose group conflict, specifically testing whether shifting perceived in-group norms alters individual decision making. Since our results suggest that individuals make choices to protect and promote their group identity, we hypothesized that decision making will be sensitive to group norm information. In the absence of clear norm information, we consistently find that individuals avoid supporting the opposing group. In Study 5, we test whether providing participants with alternate group norm information (i.e., others in your in-group chose to support the opposing group) will increase the choice share supporting the opposing group over harming their own side, as the norm serves as a powerful guideline for making choices that maintain an identity consistent with the in-group.

To test whether modulating group norms alters decision making, participants in Study 5 were asked to report their position on abortion access ("very much against abortion access" to "very much in favor of abortion access") and were then randomized into one of three conditions: control, norm-add, or norm-subtract. Participants in all conditions chose between adding \$1 to a donation going to an organization supporting the opposing side or subtracting \$1 from a donation going to their side of the issue. In the norm-add condition, participants were also informed that in a previous study, 70% of participants who shared their views on abortion access chose to add to the opposing side rather than subtract from their own and that one of those participants had said the following: "I care way too much about my cause to take money away from it". In the norm-subtract condition, participants were instead told that 70% of previous participants on their side of the cause had chosen to subtract from their in-group rather than add to the opposing group. The statement from the previous participant was changed to: "I dislike the other side way too much to give them money". Finally, as in our prior studies, each participant indicated which of the two sides of the cause they believe is more effective at pursuing its mission. We recruited 653 participants from Amazon's Mechanical Turk, with a final sample of 635 following our pre-registered exclusions.

In the control condition, we replicated our fundamental finding: the proportion of participants choosing to add \$1 to the opposing side's donation was significantly less than 50% (39.2%), $\chi^2(1, N = 212) = 9.55$, $p = .002$ (see Figure 2.5).

In the norm-add condition, the key test of the power of group norms, we found a significant increase in the proportion of participants choosing to add \$1 in the norm-add condition (57.7%), compared to participants in the control condition $\chi^2(1, N = 420) = 13.72$, $p < .001$, $\varphi = 0.19$. In fact, the proportion of participants in the norm-add condition choosing to add

\$1 to the opposing side was significantly greater than 50% ($\chi^2(1, N = 208) = 4.62$, $p = .032$). The results of the norm-add condition show that norms-based interventions about group identities can powerfully shift decision making away from harming the in-group and toward supporting the opposing group.

In the norm-subtract condition, the proportion of participants choosing to add \$1 to the opposing side (36.7%) was not statistically different from the control group, $\chi^2(1, N = 427) = .17$, $p = .680$, $\varphi = 0.02$. The similarity between the norm-subtract and control conditions implies that in the absence of an explicit group norm (as in the control condition), the default norm is not to support the opposition even at the expense of harming one's own side. We further verified that the norm-add and norm-subtract conditions were statistically different, $\chi^2(1, N = 423) = 17.79$, $p < .001$, $\varphi = 0.21$. In a supplementary analysis (see Appendix Results for Study 5), we found that the norms manipulation can shift behavior even for those with strong attitudes toward an issue. We also replicated the result that participants believed the organization on their side is more effective at spending donation money to achieve their goal.

Discussion

In the present work, we investigate how individuals prefer to adjudicate a lose-lose choice in intergroup conflict: harm their in-group or support their opposition. We operationalize this choice by giving study participants the option to either deduct funds from organizations within their in-group or add the same amount of funding to an opposing organization. Such choices help to separate various motives that could be driving decision making, and remarkably, we find that even though individuals report that organizations in their in-group (vs. opposing group) are more effective with funds, they choose to deduct from their (more effective) in-group rather than add an equivalent amount of funds to the opposition. Indeed, individuals are so averse to providing any support to the opposing group that they, on average, accepted triple the amount of financial loss to their in-group to avoid any gains for the other side (Study 2). We reproduce our main findings across both sides of an array of group conflicts (abortion, gun control, political party) and in multiple countries (United States and United Kingdom; Study 1) to illustrate that the preference to harm one's in-group to avoid supporting the opposing group is a robust, fundamental feature of individual decision making in group conflicts.

Moreover, we explored the role of identity concerns to understand the psychology underlying the preference to harm one's in-group rather than support the opposition. We found that whereas the strength of an individuals' group identity strongly correlates with the decision to harm the in-group rather than support the opposition, individual assessments of group efficacy were uncorrelated with choice (Study 3). Manipulation of identity salience modulated the choice to harm the in-group versus support the opposition (Study 4), further illustrating the central role of identity considerations in decision making within group conflicts. Finally, we demonstrated a practical method to alter preferences in inter-group conflicts: shifting perceptions of in-group norms lead to corresponding changes in behavior – individuals who were told that other in-group members were willing to support the opposing group became more likely to do the same (Study 5).

Identity concerns as the central driver of decision making in group conflicts

Previous models of individual psychology in groups, such as in-group favoritism (Tajfel et al. 1971, Brewer 1979, Abrams and Hogg 1990, Hewstone et al. 2002, Hogg 2016) and in-group love (Halevy et al. 2008, Halevy et al. 2012) examined decision making using win-win scenarios, which cannot explain our findings in lose-lose scenarios. In a win-win context, in which individuals choose between various gains for the in-group and/or losses for the out-group, past work has found that individuals will seek the best relative outcome for their in-group ("in-group favoritism") while avoiding unnecessarily harming the out-group ("in-group love" rather than "out-group hate"). While this literature used "minimal" groups where trivial differences created in-group and out-group distinctions, we replicate the preference to help one's in-group rather than harm the out-group using natural groups for pre-existing polarizing issues. However, we find that in a lose-lose context, individuals choose to financially harm their in-group rather than support an opposing out-group. This is a violation of both in-group love and in-group favoritism, as the alternative choice – supporting the opposition – maximizes the relative position of the in-group, because organizations on one's side are generally viewed as more effective with funds than opposing organizations. In fact, participants even chose to accept triple the amount of financial losses to their own group to avoid supporting the opposition, illustrating that group members were not acting to establish the most favorable comparison between their in-group and the opposing group. Rather than in-group love, the results from lose-lose scenarios appears to be evidence for the opposite – out-group hate – in line with recent work on negative partisanship, finding that partisans are demonstrating increasingly negative affect towards the opposing party (Abramowitz and Webster 2016, Abramowitz and Webster 2018). Among political partisans in the United States, "out-party hate" was recently found to be stronger than "in-party love" (Finkel et al. 2020).

We synthesize prior work on support-framing (Zhong et al. 2008, Catapano and Tormala 2021) and propose the Identity-Support model, which can parsimoniously explain our findings across win-win and lose-lose scenarios. The model suggests that individuals act in group conflicts to promote their identity, and they do so primarily by providing support to causes they believe in (and avoid supporting causes they oppose; see also Appendix Study A2.1). Simply put, in win-win contexts, supporting the in-group is more expressive of one's identity as a group member than harming the opposing group, thereby leading to a preference for in-group support. In lose-lose contexts, supporting the opposing group is more negatively expressive of one's identity as a group member than harming the in-group, resulting in a preference for in-group harm. Therefore, the principle that individuals make decisions in group conflicts to promote and protect their identity, primarily by allocating their support in ways that most align with their values, offers a single framework that predicts individual behavior in group conflicts in both win-win and lose-lose contexts.

Alternative Explanations and Related Literature

Although our findings offer strong support for the role of identity considerations in group conflict, our results do not address whether these identity concerns are driven by a motivation to maintain or boost one's self-image (Bodner and Prelec 2003, Gneezy et al. 2012, Tonin and Vlassopoulos 2013) or their reputation (i.e., for social approval; Fehr and Fischbacher 2002, Ariely et al. 2009). As an initial test of whether the aversion to supporting the opposing side is driven primarily by reputation concerns, we ran a pre-registered supplemental study (see Appendix Study A2.2). As in previous studies, we offered participants a lose-lose tradeoff, but here we also manipulated whether the choice was explicitly anonymous or would be made public. If participants' decision is based on public compliance or a desire for social approval (Deutsch and Gerard 1955, Andreoni and Petrie 2004), we would expect the effect to be stronger when making their choice publicly (vs. privately). However, we found that individuals' preferences did not differ when their choices were public versus private and that they preferred to harm their group rather than help the opposing group in both conditions. Our results suggest that the motive to express one's values by avoiding out-group support is internalized. However, some work suggests that social influence may still be at play as individuals sometimes act as though they are being observed by a third party even when they are not (as in the anonymous condition; Sackeim and Gur 1979, Baldwin and Holmes 1987, Bodner and Prelec 2003, Chaudhry and Loewenstein 2019). Nevertheless, this supplemental study provides further evidence for the robustness of the aversion to helping the opposing side regardless of whether others would learn about their decision.

In the same supplemental study, we examined a possible alternative explanation for the pattern of preferences we observe. Participants in our experiments may have chosen to subtract funds from their side (rather than add to the opposing side) because it feels easier to undo or reverse (e.g., by making an additional donation to their side later). By contrast, participants may believe it is more difficult to "undo" the addition of funding to the opposing side. We therefore asked study participants to explain why they chose the option they selected. Of the 497 participants in the study, only four mentioned reversibility as an explanation for their choice to subtract funding, suggesting that this is not the primary driver for the preference to harm one's own side in our studies.

This work studies decision making for polarizing issues where individuals may have deeply held beliefs. We chose polarizing contexts because of the importance of improved decision making around these contentious issues and our specific interest in intergroup conflict. While previous research finds that deeply held beliefs or sacred values leads to behavior that frequently departs from normative theory (Tetlock et al. 2000, Tetlock 2003), this prior work does not make clear predictions for which choice individuals will make in the lose-lose scenarios in our studies. That is, individuals with deeply held beliefs tying themselves to their in-group would likely have a strong aversion to both harming the in-group and helping the out-group, though it is unclear which would prevail based on this literature. We also note that the preference to harm one's in-group persists even among individuals for whom these are not strongly held beliefs (i.e., those who report weaker attitudes towards their side of the issue). We would not expect the preference to harm one's own group rather than help an out-group to emerge for all out-groups, but rather for out-groups with which individuals do not want to align themselves (White and Dahl 2006, 2007) or for groups that directly oppose the decision maker's beliefs.

Implications for Better Outcomes in Group Conflicts

One striking facet of our work is that individuals resolve lose-lose decisions in group conflicts in ways that leave their own in-group in a worse relative position than if they had simply supported the opposition. When generalized across both sides of several issues, our work points to the possibility that identity concerns may act as a barrier to better outcomes for both groups. Therefore, groups engaged in conflict may realize mutual gains if individuals are less averse to supporting the opposition. Building on an extensive literature on in-group norms (Cialdini et al. 1990, Cialdini et al. 2006, Goldstein et al. 2008, Miller and Prentice 2016), we demonstrated that shifting group norms can modulate individuals' aversion toward showing support for the opposing group. While many accounts suggest that the United States is becoming more affectively and ideologically polarized (Pew 2019, Klein 2020, Dimock and Wike 2021), an emerging literature on "false polarization" suggests that intergroup conflict is exacerbated by misperceptions about the magnitude and consistency of out-group members' beliefs (Chambers et al. 2006, Fernbach and Van Boven 2021). In fact, recent work finds that Americans often tolerate and even show admiration for in-group members who seek to understand the out-group (Heltzel 2019), indicating both sides may have a desire for cooperation. Our findings offer a practical approach that has the potential to increase cooperation: providing information about in-group norms may reduce group members' identity concerns, thereby allowing for behaviors that support the out-group when advantageous. Future research might further examine the nature of this norm belief and test realistic and effective methods for increasing the likelihood to work with the opposition, such as modeling cooperative behavior by high-status in-group members.

Our findings add to a literature on how psychological barriers impede the advancement of important causes (Van Boven et al. 2018). In contexts in which accommodating two groups' desires is crucial for progress, how do we compromise when both sides would rather harm their own cause than make concessions in which the opposition gains any benefits? For example, a congressperson wishing to cross the aisle to support legislation may be hindered by the assumption that it would be a sign of disloyalty to her constituents. In an era of high perceived polarization, understanding how identity concerns and beliefs about group norms shape these decisions is critical. Otherwise, these psychological barriers are likely to impede progress, not only for the causes we oppose, but also for those we most strongly support.

Materials and Methods

Overview

All experiments were approved by the UC San Diego IRB, and all participants gave their informed consent to participate.

Studies 1A and 1B

Study 1A was conducted in January 2022. As outlined in our pre-registration (https://aspredicted.org/cz62t.pdf), we aimed to recruit a nationally representative sample of 800 U.S. participants through Prolific and ended up with a sample of 801 participants who completed the study (50.6% female, mean age = 45.17 years). We excluded 4 participants who failed the reading check, leaving us with a final sample of 797 participants.

Participants were randomized into one of two conditions (lose-lose or win-win) to make a hypothetical choice. All participants first reported their position on three issues, presented in a randomized order: abortion access (“very much against abortion access” to “very much in favor of abortion access”), gun control (“very much against gun control” to “very much in favor of gun control”), and political party (“strongly Republican” to “strongly Democratic”). Responses were captured using a 6-point Likert scale to prevent participants from expressing indifference, as alignment to a side of each issue was required for the scenario assignment. We used these responses to classify participants as having either weak (3, 4 on the scale), medium (2, 5 on the scale), or strong (1, 6 on the scale) attitudes.

For each of the three issues, participants were told that, as part of the study, donations would be made to organizations supporting each side, and that they would need to make a choice about how to alter the donation amount. We informed participants that we would randomly select 10 of them and adjust one of the donation amounts based on their choice and actually make the donations on their behalf. For each issue, brief descriptions of each organization’s mission were provided. For example, for the abortion access issue, participants read: “The mission of the Pro-life organization is to reduce access to abortions. The mission of the Pro-choice organization is to increase access to abortions.” No organizations were referred to by name to avoid any associations a participant may have with a particular organization. All the scenarios and corresponding binary choices were presented in a randomized order.

For each cause, participants were asked to select one of two options. In the win-win condition, both options altered the donation in ways that were favorable given the participant’s stated attitude: either add $1 to the donation going to the organization on their side or subtract $1 from the donation going to the organization on the opposing side. In the lose-lose condition, both options altered the donation in ways that were unfavorable given the participant's stated attitude: either add $1 to the donation going to the organization on the opposing side or subtract $1 from the donation going to the organization on their side.

After responding to all three scenarios, participants reported which side of each issue had organizations that they believed to be more effective at pursuing their mission. Participants were specifically asked "which one is able to do more with each dollar they receive?". Responses were collected on a 7-point Likert scale for all three issues, in a randomized order: abortion access ("pro-life organization are more effective" to "pro-choice organizations are more effective"), gun control ("pro-gun organizations are more effective" to "anti-gun organizations are more effective"), and political party ("Republicans are more effective" to "Democrats are more effective").

Study 1B was conducted in February 2022 and was identical, except we collected participants from the UK, and only focused on a single issue – political partisanship. As outlined in our pre-registration (https://aspredicted.org/m4nq5.pdf), we recruited a nationally representative sample of 400 UK participants through Prolific (50.3% female, mean age = 44.68 years). We excluded 7 participants who failed the reading check, leaving us with a final sample of 393 participants. All participants reported their political position on the following 6-point Likert scale: "Strongly Conservative Party" to "Strongly Labour Party". As with Study 1A, we used these responses to classify participants as having either weak (3, 4 on the scale), medium (2, 5 on the scale), or strong (1, 6 on the scale) attitudes.

For analyses across issues, we combined the datasets collected from Studies 1A and 1B. To test whether attitude strength moderated the participants' choices, we regressed their choice

(0 = subtract \$1, 1 = add \$1) on their condition, and the interaction between condition and attitude strength (as a categorical variable), using a logistic regression.

Study 2

Study 2 was conducted in December 2020. We recruited 300 U.S. participants from Amazon's Mechanical Turk (MTurk; 53% female, mean age = 36.71 years). As outlined in our pre-registration (https://aspredicted.org/cn7ry.pdf), we excluded participants who switched more than once between the left and right-hand choices (10.7% of participants). This fraction of exclusion is within the typical range observed in prior studies involving multiple price lists (Bruner 2011). All remaining participants passed the pre-registered reading check and there were no duplicated MTurk IDs, so there were no additional exclusions, resulting in a final sample of 268 participants which was used for all analyses.

All participants reported the extent to which they are against or in favor of abortion access on a 12-point Likert scale ("very much against abortion access" to "very much in favor of abortion access"). We informed participants that the researchers would make two \$10 donations, one to a pro-life organization and another to a pro-choice organization. Participants were then asked to choose how to alter the donation amount in a series of 14 choices, where for each choice they could select either a right-hand side or left-hand side option (similar to a price list; Andersen et al. 2006). The right-hand side option was always to add \$1 to the donation going to the opposing organization. The left-hand side option was to subtract \$X from the donation going to the organization on the participant's side of the cause, where X took the values 0.10, 0.25, 0.50, 0.75, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10 (see Figure 2.6). We randomized whether participants viewed X in ascending or descending order. To incentivize the responses, we informed participants that for

1 in 10 participants, chosen at random, we would actually make donations to both organizations and randomly select one of their 14 choices to alter the donation amount.

The outcome of interest was each participant's indifference amount. We assumed that the indifference amount is at the midpoint of the subtract amounts on either side of the switch. For example (see Figure 2.6), if a participant switches from preferring to subtract \$3 from their side (instead of adding \$1 to the other side), to preferring to add \$1 to the other side (instead of subtracting \$4 from their side), the indifference amount must be in the interval between \$3 and \$4 and was coded as the midpoint (\$3.50). However, our results are robust to coding the indifference amount as the lower bound of each interval instead of the midpoint. Using the lower bound is a highly conservative measure since it might underestimate each participant's true indifference amount, which lies in-between the endpoints of each interval. If a participant selected the left-hand option for every choice, we coded their indifference amount as \$10. If a participant selected the right-hand option for every choice, we interpreted their indifference amount to be between \$0 and \$0.10 and coded their indifference amount as the midpoint (\$0.05). There was no significant difference in indifference amounts by price list order ($p = .40$), so we collapsed across the ascending and descending conditions.

After this series of choices, we assessed beliefs about the relative effectiveness of the organizations as in Study 1. Participants were asked "Do you believe that the Pro-Life or Pro-Choice organization is more effective at pursuing their mission? In other words, which one spends their donation money more effectively?" Responses were collected on a 7-point Likert scale, allowing participants to report equal effectiveness.

Study 3

Study 3 was conducted in February 2022. As outlined in our pre-registration (https://aspredicted.org/ja6un.pdf), we aimed to recruit 400 U.S. MTurk participants and ended up with a sample of 401 participants who completed the study (51% female, mean age = 39.57 years). We excluded 2.0% of participants who failed the reading check or had duplicate MTurk IDs, leaving us with a final sample of 393 participants.

Participants first indicated whether they identified more strongly as Republican or Democrat (binary choice) and were subsequently asked to make a lose-lose choice, identical to the U.S. political party choice from Study 1. We told participants we would pick 10 of them at random and make the donations according to their adjusted donation amounts. Subsequently, participants responded to an effectiveness question, and an identity concern question, in a randomized order. The effectiveness questions asked participants, "Do you believe that the Republican or Democratic Party is more effective at using donated funds to pursue their mission? In other words, which one is able to do more with each dollar they receive?" (1 = Republicans are more effective, 7 = Democrats are more effective). The identity question asked, "Which of these two options would make you feel like a worse [Democrat/Republican]? In other words, which option most undermines your identity as a [Democrat/Republican]?" (1 = Definitely adding $1 to the [Republican/Democratic] organization, 4 = Both choices equally undermine my identity as a [Democrat/Republican], 7 = Definitely subtracting $1 from the [Democratic/Republican] organization). The first option in square brackets was selected for participants who identified as Democrats, and the second was selected for Republicans.

Study 4

Study 4 was conducted in June 2021. As outlined in our pre-registration (https://aspredicted.org/gu45s.pdf), we recruited 500 U.S. MTurk participants (50% female, mean age = 39.51 years). We excluded 0.6% of participants who failed the reading check or had duplicate MTurk IDs, leaving us with a final sample of 497 participants. All participants reported whether they identified more strongly as Republican or Democrat (binary choice). Participants were then randomized into one of two conditions: control or identity-strengthened. Participants in the identity-strengthened condition were then asked to write about an event, story, or personal experience where they strongly identified with their political party. In the control condition, participants were asked to write about what they do on a typical Monday evening.

Subsequently, all participants were asked to make a lose-lose choice, similar to Study 1 – participants had to choose between adding $1 to the donation going to the organization supporting the opposing political party or subtracting $1 from the donation going to the organization supporting their political party. We also specified that all donations would go to administrative costs (e.g., maintaining the organization's website).

As with our previous studies, we also asked participants whether they believe Republican or Democratic organizations are more effective at pursuing their mission, using the same scale as Studies 1 and 2, except with 6 points.

Study 5

Study 5 was conducted in March 2020. As outlined in our pre-registration (https://aspredicted.org/cz2kf.pdf), we aimed to recruit 650 U.S. MTurk participants and ended up with a sample of 653 participants who completed the study (55% female, mean age = 36.18

years). We excluded 2.8% of participants who failed the reading check or had duplicate MTurk IDs, leaving us with a final sample of 635 participants.

All participants reported the extent to which they are against or in favor of abortion access on a 12-point Likert scale ("very much against abortion access" to "very much in favor of abortion access"). Participants were then randomized into one of three conditions: control, norm-add, or norm-subtract.

In the control condition, participants were informed that the experimenter would be making donations to a pro-life and a pro-choice organization, and that they would have to choose how to alter the amount – add $1 to the donation going to the opposing side or subtract $1 from the donation going to their side. The norm-add condition was identical, except we also informed participants that in a previous study, 70% of MTurkers who shared their views on abortion access chose to add to the opposing side rather than subtract from their own and that one of those participants had said the following: "I care way too much about my cause to take money away from it". In the norm-subtract condition, participants were instead told that 70% of previous participants on their side of the cause had chosen to subtract from their in-group rather than add to the opposing group. The statement from the previous participant was changed to: "I dislike the other side way too much to give them money".

As with our previous studies, we also asked participants to indicate which of the two sides of the cause they believe is more effective at pursuing its mission, using a 6-point scale.

Acknowledgements

We would like to thank Charles Dorison, Eric VanEpps, David Tannenbaum, Ayelet Gneezy, On Amir, Wendy Liu, and Lisa Ordóñez for providing feedback on this work.

Chapter 2, in full, is a reprint of previously published material as it appears in the Proceedings of the National Academy of Science, 119(49), e2215633119, Rachel Gershon and Ariel Fridman. The dissertation author was one of the primary investigators and authors of this paper.

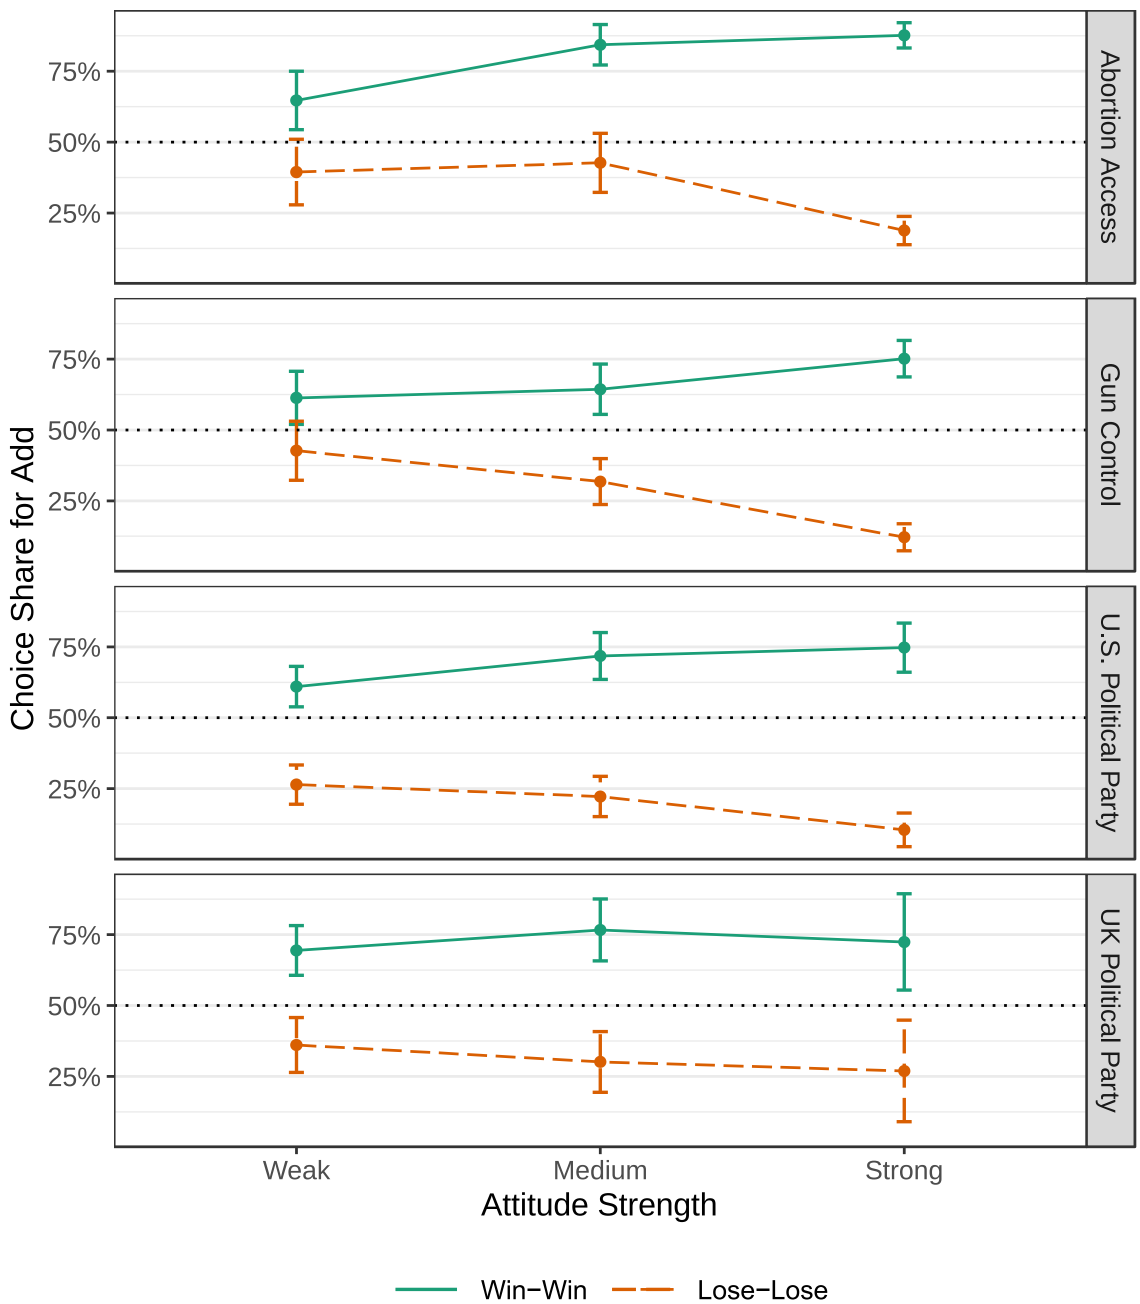

Figure 2.1. Choice share in each condition for all four issues by attitude strength in Studies 1A (N = 797) and 1B (N = 393). The vertical axis shows the proportion of participants choosing to add funds (in the lose-lose condition: add funds to opposing group vs. subtract from their in-group; in the win-win condition: add funds to the in-group vs. subtract from the opposing group). Error bars represent 95% confidence intervals.

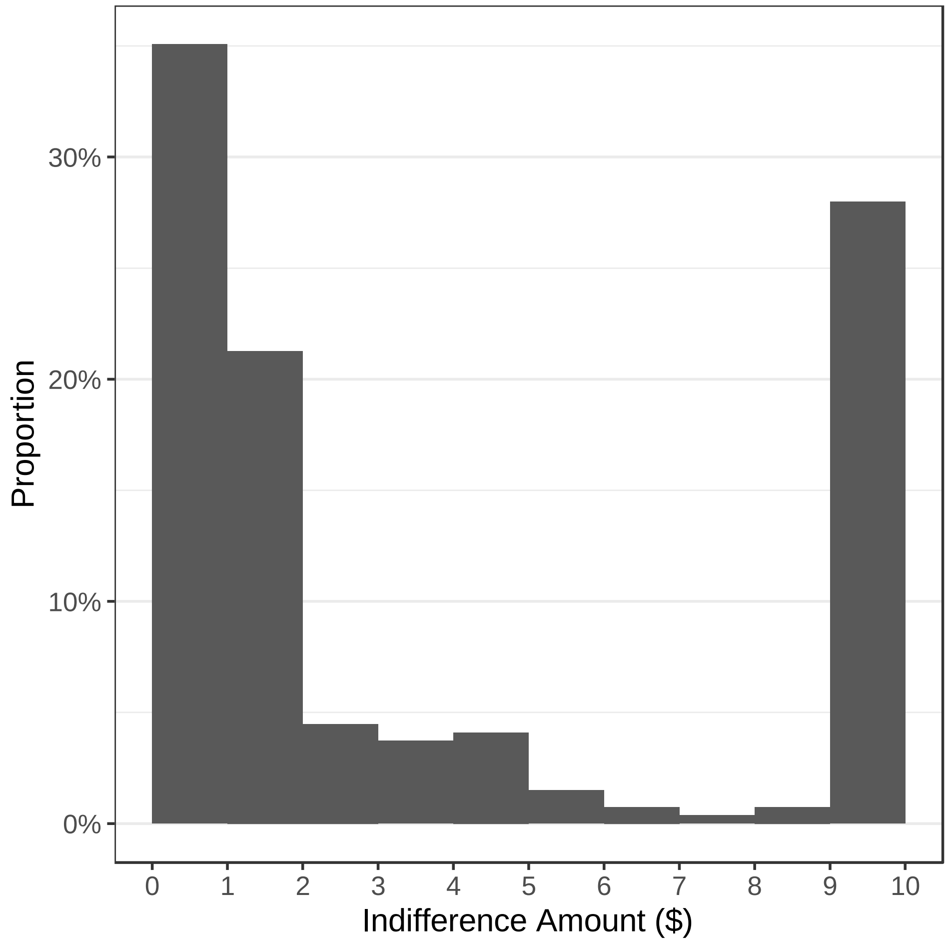

Figure 2.2. Results from Study 2 (N = 268): distribution of participants' indifference amounts.

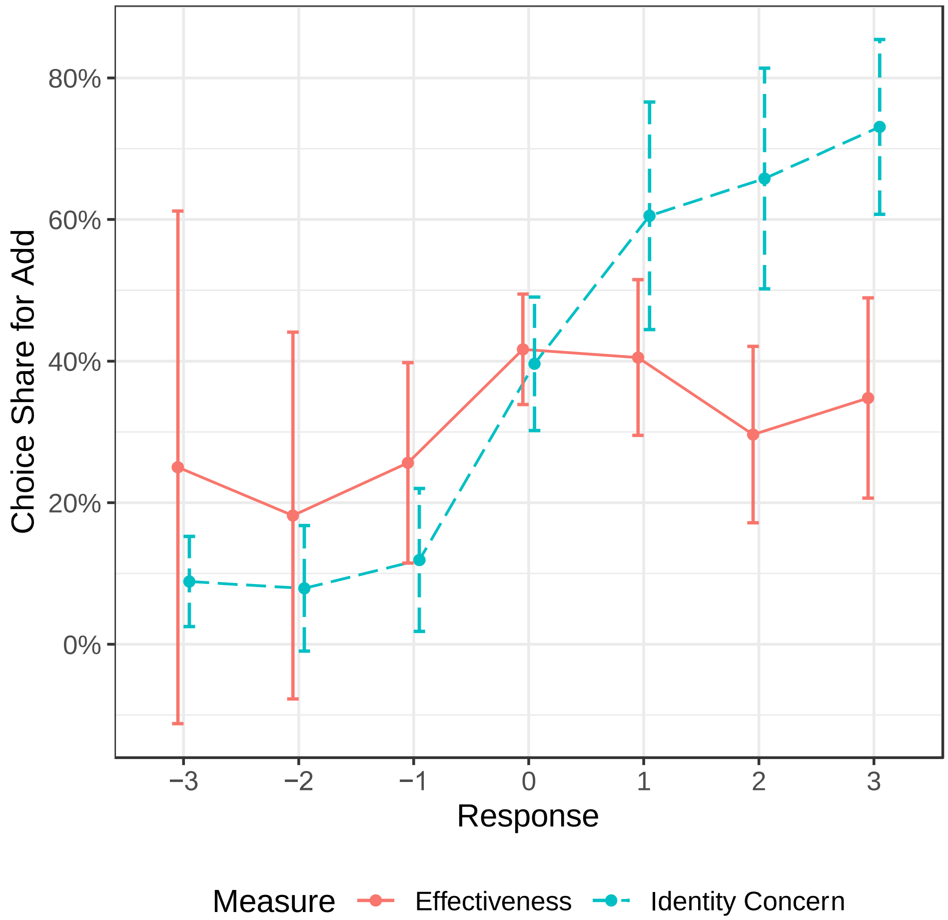

Figure 2.3. Results from Study 3 (N = 393): choice share by measure. The vertical axis shows the proportion of participants who chose to add funds to the donation for the opposing organization (vs. subtract from their side). The horizontal axis captures participants' responses on the effectiveness and identity concern measures (which were both centered at 0). For the effectiveness measure, more positive values indicate the in-group is more effective at using donated funds (vs. the out-group; 0 indicates equal effectiveness). For the identity concern measure, more positive values indicate subtracting $1 from the in-group undermines identity more (vs. adding $1 to the opposing group; 0 indicates both equally undermine identity). Error bars represent 95% confidence intervals.

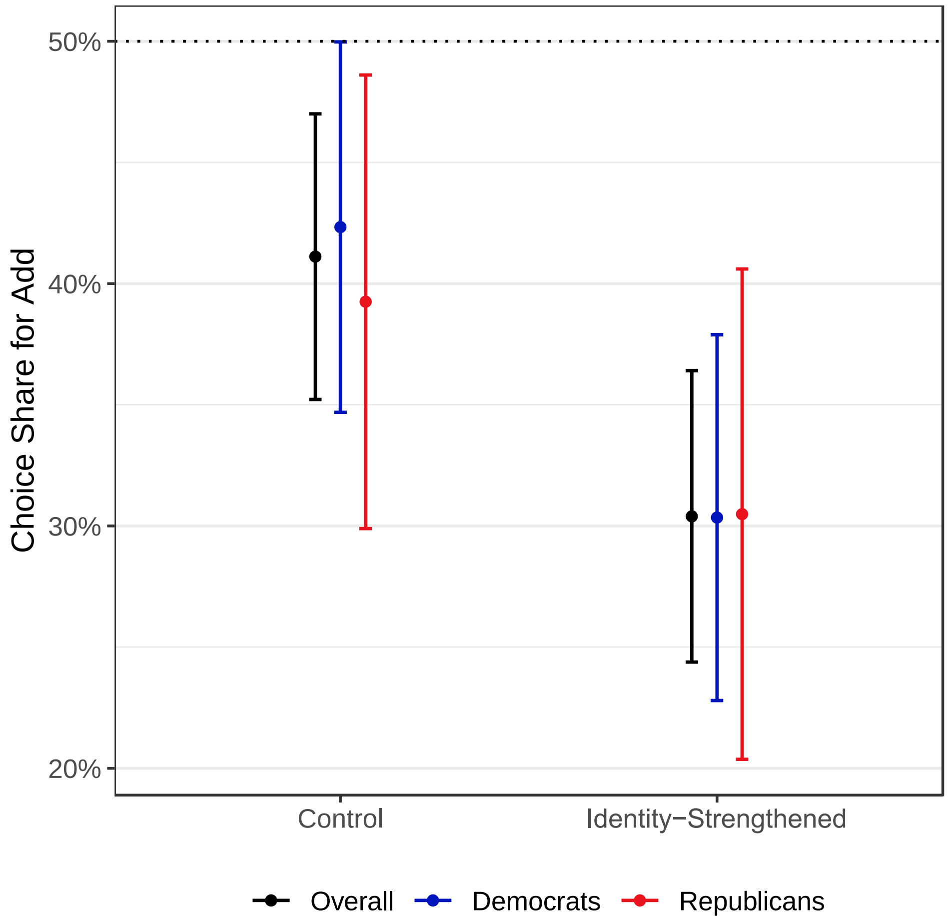

Figure 2.4. Results from Study 4 (N = 497): choice share by condition. The vertical axis shows the proportion of participants who chose to add funds to the donation to the opposing organization (vs. subtract from their side). Error bars represent 95% confidence intervals.

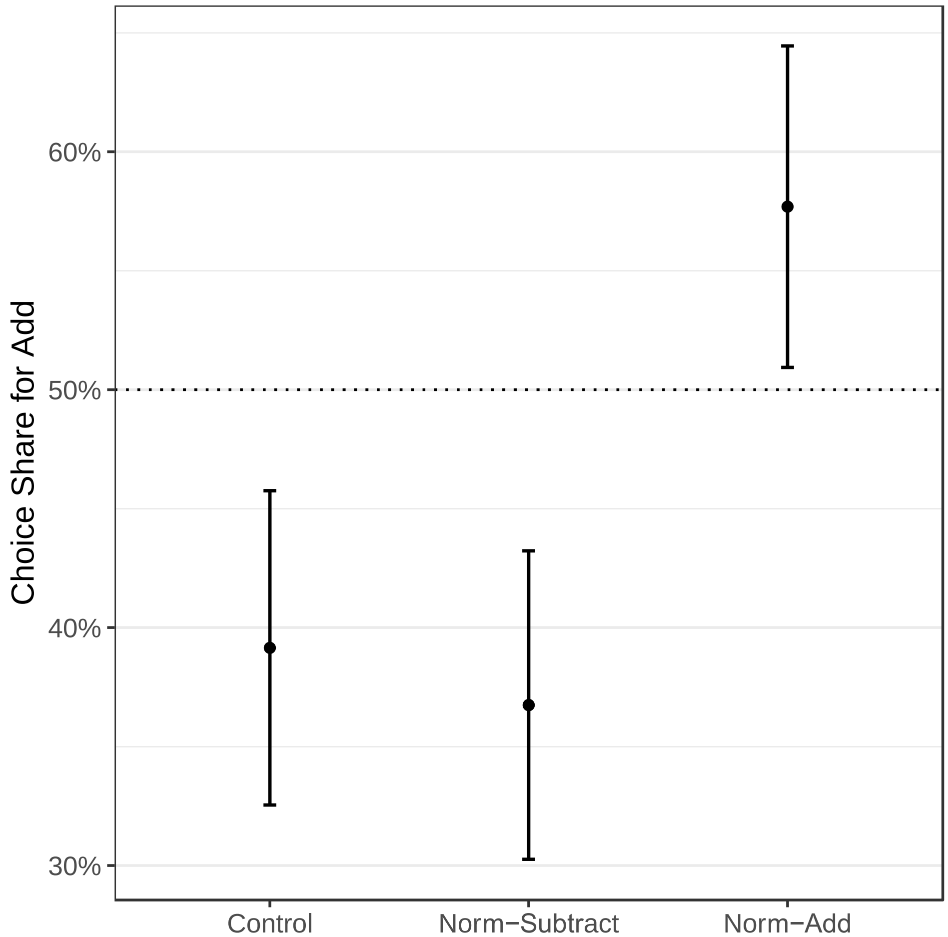

Figure 2.5. Results from Study 5 (N = 635): choice share by condition. The vertical axis shows the proportion of participants who chose to add funds to the donation to the opposing organization (vs. subtract from their side). Error bars represent 95% confidence intervals.

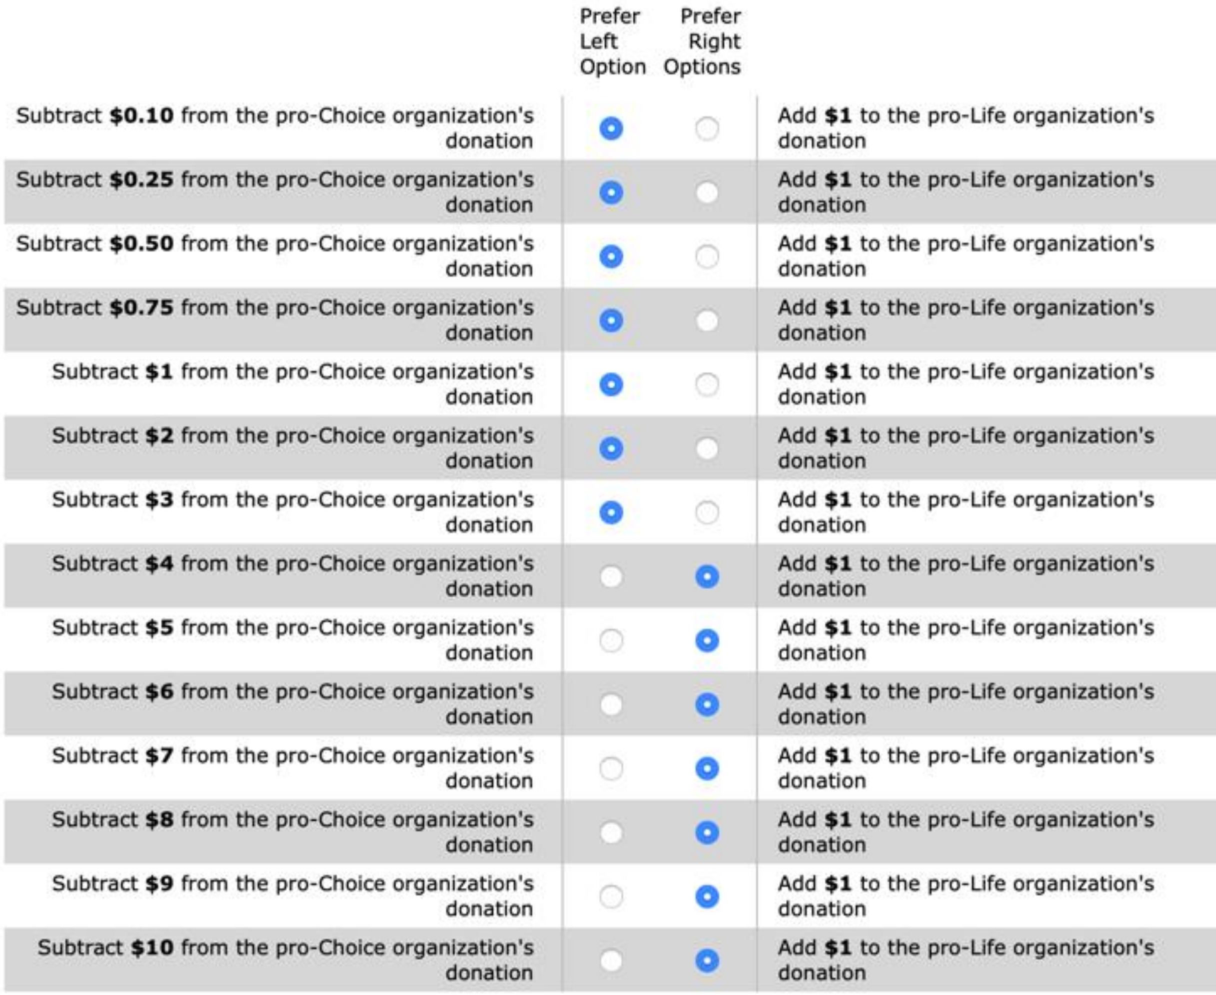

	Prefer Left Option	Prefer Right Options	
Subtract **$0.10** from the pro-Choice organization's donation	●	○	Add **$1** to the pro-Life organization's donation
Subtract **$0.25** from the pro-Choice organization's donation	●	○	Add **$1** to the pro-Life organization's donation
Subtract **$0.50** from the pro-Choice organization's donation	●	○	Add **$1** to the pro-Life organization's donation
Subtract **$0.75** from the pro-Choice organization's donation	●	○	Add **$1** to the pro-Life organization's donation
Subtract **$1** from the pro-Choice organization's donation	●	○	Add **$1** to the pro-Life organization's donation
Subtract **$2** from the pro-Choice organization's donation	●	○	Add **$1** to the pro-Life organization's donation
Subtract **$3** from the pro-Choice organization's donation	●	○	Add **$1** to the pro-Life organization's donation
Subtract **$4** from the pro-Choice organization's donation	○	●	Add **$1** to the pro-Life organization's donation
Subtract **$5** from the pro-Choice organization's donation	○	●	Add **$1** to the pro-Life organization's donation
Subtract **$6** from the pro-Choice organization's donation	○	●	Add **$1** to the pro-Life organization's donation
Subtract **$7** from the pro-Choice organization's donation	○	●	Add **$1** to the pro-Life organization's donation
Subtract **$8** from the pro-Choice organization's donation	○	●	Add **$1** to the pro-Life organization's donation
Subtract **$9** from the pro-Choice organization's donation	○	●	Add **$1** to the pro-Life organization's donation
Subtract **$10** from the pro-Choice organization's donation	○	●	Add **$1** to the pro-Life organization's donation

Figure 2.6. Screenshot of the series of choices made by participants in Study 2. This example is for a pro-choice participant, with the choices listed in ascending order.

Table 2.1. Group decision-making theories. The ✓ indicates a preference for the given choice, whereas 50% indicates an indifference between the choices. For each scenario (win-win and lose-lose), the column shaded in gray is the option selected by most participants in our studies.

	Win-Win		Lose-Lose	
	Add $1 to in-group	**Subtract** $1 from opposing group	**Add** $1 to opposing group	**Subtract** $1 from in-group
Identity-Support Model	✓			✓
Zero-sum beliefs	50%	50%	50%	50%
Harm minimization (based on effectiveness considerations)	✓		✓	
Maximize in-group payoff	✓		✓	
Maximize relative payoff for in-group	50%	50%	50%	50%
Minimize payoff difference	50%	50%	50%	50%
Maximize joint payoff, in favor of in-group	✓		50%	50%
Minimize opposing group payoff		✓		✓

Table 2.2. Summary table for Studies 1-5 of the percentage of participants preferring to support the opposing group (for lose-lose choices) and support the in-group (for win-win choices; Study 1 only). For Study 2, the percentage captures the proportions of participants with an indifference amount of less than $1 (i.e., would rather add $1 to the opposing group than subtract $1 (or less) from the in-group).

	Lose-Lose (% choosing to support opposing group)	Win-Win (% choosing to support in-group)	Participant Attitudes
Study 1			
Abortion Access	27.8%	81.9%	Pro-Choice: 76%; Pro-Life: 24%
Gun Control	25.3%	68.3%	Anti-Gun Ownership: 75%; Pro-Gun Ownership: 25%
U.S. Politics	20.8%	67.6%	Democrats: 70%; Republicans: 30%
UK Politics	32.7%	72.1%	Labour Party: 64%; Conservative Party: 36%
Study 2 Abortion Access	35.1%		Pro-Choice: 74%; Pro-Life: 26%
Study 3 U.S. Politics	36.4%		Democrats: 65%; Republicans: 35%
Study 4 U.S. Politics			
Control	41.1%		Democrats: 60%; Republicans: 40%
Identity-Strengthened	30.4%		Democrats: 64%; Republicans: 36%
Study 5 Abortion Access			
Control	39.2%		Pro-Choice: 67%; Pro-Life: 33%
Norm-Subtract	36.7%		Pro-Choice: 71%; Pro-Life: 29%
Norm-Add	57.7%		Pro-Choice: 72%; Pro-Life: 28%

Appendix

Supplemental Results for Study 1

To test whether participants believed that organizations supporting their side of an issue are more effective at spending donation money to achieve their goal, we centered and re-oriented the effectiveness scale, such that 0 indicates equal effectiveness, positive values indicate organizations on the participant's side are more effective, and negative values indicate organizations on the opposing side are more effective. The overall effectiveness rating was positive ($M = 0.35$, $SD = 1.58$), $t(1{,}189) = 9.52$, $p < .001$, 95% CI = [.28, .43].

The belief that organizations within one's in-group are more effective with funds than those on the opposing side is consistent across each of the four issues and for participants on each side of the four issues ($Ps < .001$; Figure A2.1), however there are two notable exceptions. First, for pro-life Americans, the result does not attain statistical significance ($p = .43$) but is directionally consistent. The other exception is among "pro gun control" Americans, who indicate that they view organizations that are "anti gun control" as more effective ($M = -0.40$, $SD = 1.64$), $t(593) = -5.95$, $p < .001$, $d = -.24$, 95% CI = [-.33, -.16], such that individuals in this group are technically making a harm minimizing choice by not supporting the opposition. This deviation by "pro gun control" participants likely reflects the strong influence of the National Rifle Association and can serve as evidence that participant responses are not simply a result of motivated beliefs or partisan cheerleading (Bullock et al. 2013), as they do appear to consider the real-world impact of such organizations, even when it is not in their favor. Despite these two exceptions, the overall pattern of results makes it clear that individuals generally perceive organizations on their side to be more effective than those on the opposing side.

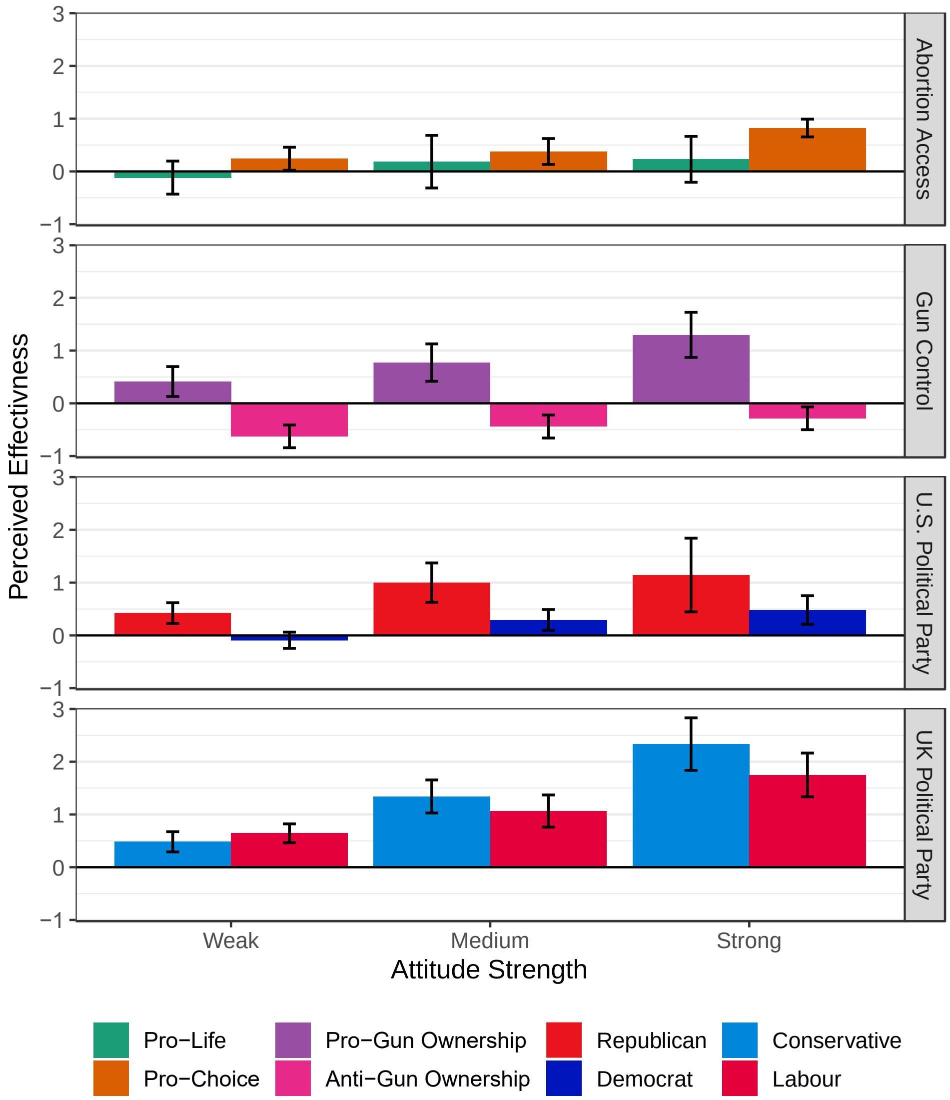

Figure A2.1. Effectiveness measure from Studies 1A and 1B for all four issues by attitude strength. Positive (negative) values indicate that organizations on the participant's (opposing) side of a cause are more effective at managing their funds. Error bars represent 95% confidence intervals.

Supplemental Results for Study 2

We tested whether attitude strength moderates the reported indifference amounts. As in Study 1, we first centered and re-oriented the 12-point attitude strength scale, such that the middle numbers (6, 7) indicate the weakest attitude strength and values closer to the extremes indicate stronger attitudes (i.e., more strongly pro-choice or pro-life). In a linear regression model, we regressed participants' indifference amount on this adjusted attitude strength measure (as a continuous variable[23]). We found that greater attitude strengths were positively associated with indifference amounts (β = .50, SE = .15, t(266) = 3.33, p < .001, 95% CI = [.20, .79]; see Figure A2.2). The correlation between attitude strength and indifference amount also addresses concerns of information leakage in the stimuli (Prelec et al. 1997), wherein participants may make inferences from the fact that there were fewer choices less than (vs. greater than) $1 in the choice list. While prior work finds that those with weaker attitudes tend to be more susceptible to contextual effects (Lavine et al. 1998), we find here that those with weaker (vs. stronger) attitudes subtracted less from their own group.

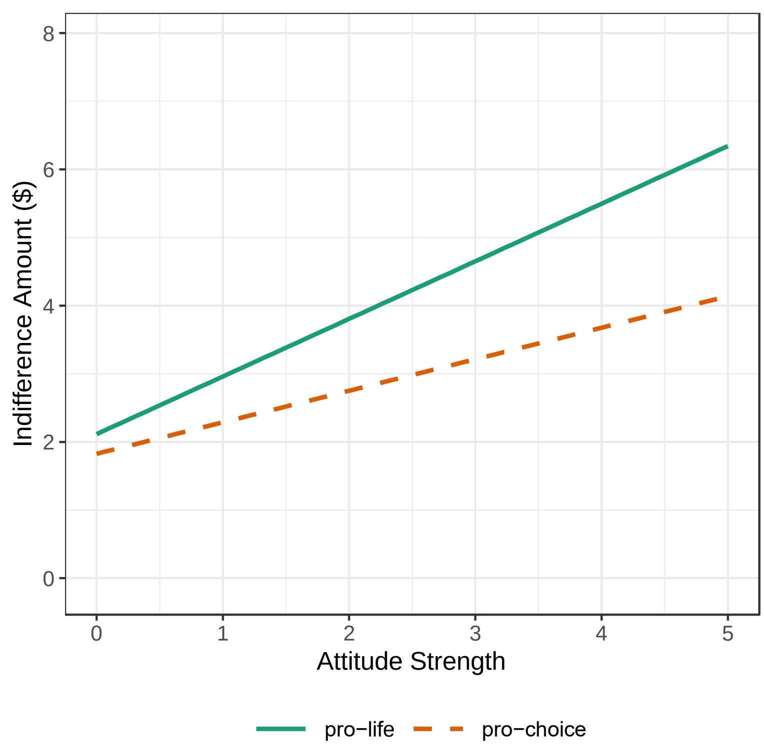

Figure A2.2. Lines depict linear regression of indifference amount on attitude strength. Shaded regions represent 95% confidence intervals.

---

[23] While we coded attitude strength as continuous, similar results were obtained when we created three buckets (weak, moderate, strong) and coded it as categorical as in Study 1, with the difference between weak and strong attitudes attaining statistical significance (p < .001).

As in Study 1, we also tested whether participants believed that organizations supporting their side of an issue are more effective at spending donation money to achieve their goal. We replicated the result from Study 1, finding a positive overall effectiveness rating (M = 1.26, SD = 1.65), t(267) = 12.49, p < .001, d = .76, 95% CI = [.63, .90]). This pattern held significantly across pro-choice and pro-life participants (both Ps < .03). Indeed, most participants (64%) indicated that organizations on their side are more effective in spending their funds (χ^2(1, N = 268) = 20.99, p < .001), and 23% indicated equal effectiveness.

Supplemental Results for Study 3

See Table A2.1 for regression results for Study 3.

Table A2.1. Regression results for Study 3.

	Dependent variable: choice to add $1		
Model:	(1)	(2)	(3)
intercept	-0.556***	-0.607***	-0.576***
	(0.133)	(0.123)	(0.116)
identity concern measure	0.647***	0.637***	
	(0.074)	(0.073)	
effectiveness measure	-0.09		0.028
	(0.092)		(0.078)
Observations	393	393	393
Log Likelihood	-204.154	-204.629	-257.589
Akaike Inf. Crit.	414.308	413.258	519.179
Note:	*p<0.05; **p<0.01; ***p<0.001		

Supplemental Results for Study 4

As in prior studies, we verified that the overall effectiveness rating was positive (M = 0.64, SD = 1.27), t(496) = 11.28, p < .001, d = .51, 95% CI = [.41, .60], indicating that participants believe that organizations supporting their political party are more effective than those on the opposing side. This pattern held across Democratic and Republican participants, and across both conditions (all Ps < .001).

Supplemental Results for Study 5

We investigated the efficacy of group norms for shifting decision making as a function of reported attitude strength, which was measured prior to the norms intervention. We conducted this analysis using a logistic regression, where we regressed participants' choice (0 = subtract $1, 1 = add $1) on the participants' condition (control, norm-add, norm-subtract), attitude strength (as a continuous variable), and their interaction. The results showed no significant interaction between condition and attitude strength (Ps > 0.40). However, there was a main effect of attitude strength, such that those with stronger attitudes toward the cause were more likely to harm their side than support the opposing side ($\beta = -.28$, SE = .08, $z = -3.33$, $p < .001$, OR = .76, 95% CI = [.64, .89]; see Figure A2.3). Nevertheless, these results imply that norms have the potential to shift behavior even for those with strong attitudes toward an issue.

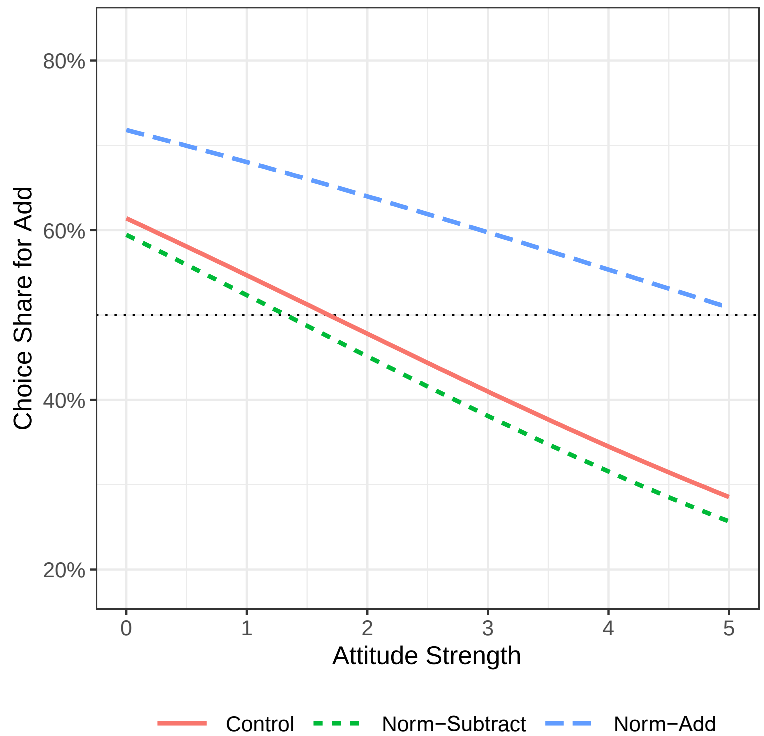

Figure A2.3. Lines depict logistic regression of choice share on attitude strength by condition. Shaded regions represent 95% confidence intervals.

As a robustness check, we verified that all of our choice results replicated when limiting the sample to the 74% of participants who believed that organizations on their side of a cause are more effective with funds, and the only difference was that the proportion of participants in the norm-add condition choosing to add $1 was not statistically different from 50% at conventional levels (57.6%, $\chi^2(1, N = 165) = 3.06$, $p = .080$). Moreover, as in the previous studies, participants indicated that they viewed organizations supporting their side of a cause to be more effective in spending their donation money to achieve their goals. The overall effectiveness rating was positive ($M = 0.82$, $SD = 1.40$), $t(634) = 14.83$, $p < .001$, $d = .59$, 95% CI = [.50, .67], indicating that participants believe organizations on their side are more effective than those on the opposing side. This pattern held across pro-choice and pro-life participants, and across all three conditions (all $Ps < .001$).

Study A2.1

The study was designed to test how people perceive others who have made either a win-win or lose-lose decision. Specifically, we tested whether evaluations of their identity as an in-group member differed based on how they chose, to shed light on whether people typically choose in ways that best protect others' perception of their identity.

Methods

Study A2.1 was conducted in August 2022. As outlined in our pre-registration (https://aspredicted.org/3wz2v.pdf), we recruited and obtained a sample of 400 U.S. MTurk participants who completed the experiment (49.5% female, mean age = 39.66 years). We excluded 1% of participants who failed the reading check or had duplicate MTurk IDs, leaving us with a final sample of 396 participants.

The study had a 2 (choice type: win-win vs. lose-lose) x 2 (decision: add vs. subtract) between-subjects design. All participants first reported whether they identified more strongly as Republican or Democrat (binary measure). Next, in the win-win (lose-lose) condition, participants were told about another participant in a previous study who shared their political preference and faced a win-win (lose-lose) choice, and were informed about how they chose. In the win-win condition, the choice described was between adding $5 to a donation to support the in-group (which always matched the participant's political preference) or subtract $5 from a donation supporting the opposing group. In the lose-lose condition, the choices were reversed: add $5 to a donation to support the opposing group or subtract $5 from a donation supporting the in-group. Participants were also informed that this choice was incentivized.

Subsequently, all participants responded to the same question that formed the key dependent variable. Participants reported how strongly they thought the participant they were evaluating identified with their in-group (Democrat/Republican party) on a 7-point Likert scale (1 = identifies very weakly; 7 = identifies very strongly).

Results

As pre-registered, we regressed the dependent variable on condition (choice type and decision), and their interaction. We found that all regressors were statistically significant, and revealed the same pattern for Democrat and Republican participants. In the lose-lose condition, participants thought that those who had chosen to subtract \$5 from their side identified more strongly with their in-group compared to those who added \$5 to the opposing side ($\beta = 1.02$, SE = .24, $t(392) = 4.19$, $p < .001$). In the win-win condition, participants thought that those who had chosen to subtract \$5 from the opposing side identified less strongly with their in-group than those who added \$5 to their in-group ($\beta = -1.95$, SE = .24, $t(392) = -8.05$, $p < .001$).

Discussion

This study provides support for the identity-support model by demonstrating that actions of support are perceived as more powerful expressions of one's identity as an in-group member than actions of opposition. When faced with win-win choices, those who opted to support their side (vs. harm the opposing side) were perceived to be stronger in-group members. Similarly, when faced with a lose-lose choice (where both options are counter-attitudinal), the action of support to the opposing side (vs. harming the in-group) was seen as a stronger expression against the in-group.

In addition, consistent with findings from Study 1 that most participants prefer to harm their side when faced with lose-lose choices and help their side when faced with win-win choices, this study suggests that people's choices are consistent with those that best protect their identity as an in-group member, highlighting the role of identity in these decisions.

Study A2.2

The goal of this study was to examine potential moderators of our main finding. Specifically, we tested whether varying the degree of anonymity of the choice (i.e., whether it is made in private or public) and the size of the stakes ($1 or $100) influenced choices.

Additionally, we asked participants to describe how they made their choice. In an exploratory analysis, we examined these free-responses to check for a possible alternative explanation for the pattern of preferences we observed: participants may have chosen to subtract donations because it was the easier option to undo or reverse (e.g., by making an additional donation later).

Methods

Study A2.2 was conducted in January 2022. As outlined in our pre-registration (https://aspredicted.org/tj6qc.pdf), we recruited 500 U.S. MTurk participants and ended up with a final sample of 501 participants who completed the experiment (47% female, mean age = 39.79 years). We excluded 0.8% of participants who failed the reading check or had duplicate MTurk IDs, leaving us with a final sample of 497 participants.

The study had a 2 (anonymity: public vs. private) × 2 (stakes: $1, $100) between-subjects design. All participants reported whether they identified more strongly as Republican or Democrat on a 6-point scale ("strongly Republican" to "strongly Democratic").

In all conditions, participants were informed that the experimenter would be making a donation to a Democratic and Republican organization, and that they would have to choose how to alter the amount. In the public condition, participants were asked to imagine that their decision was not anonymous and that their friends would be informed of what they chose. In the private condition, participants were asked to imagine that the decision would be completely anonymous and nobody would ever find out what they chose. Next, all participants were asked to make a lose-lose choice, similar to Study 1. In the $1 condition, participants chose between adding $1 to the donation going to the opposing political party or subtracting $1 from the donation going their political party. In the $100 condition, the add and subtract amounts were both increased from $1 to $100.

Afterwards, all participants were asked to write a few sentences explaining why they chose the option they selected. Lastly, we asked participants whether they believe Republican or Democratic organizations are more effective at pursuing their mission, using the same scale as Study 1.

Results

Across all conditions, the proportion of participants who chose to add money to the opposing side was 38.0%, which was significantly different from 50%, $\chi^2(1, N = 497) = 28.02$, $p < .001$. There was no significant difference between Democratic or Republican participants ($p = .85$).

Anonymity. The proportion of participants who chose to support the opposing side in the public condition (35.5%) and private condition (40.1%) was not statistically different, $\chi^2(1, N = 497) = 1.15$, $p = .28$, $\varphi = .05$.

Stakes. The proportion of participants who chose to support the opposing side in the $1 condition (33.5%) was lower than in the $100 condition (42.9%; $\chi^2(1, N = 497) = 4.31$, $p = .04$, $\varphi = .10$.

To test for interactions between conditions, we regressed participants' choice on anonymity condition, stakes condition, and their interaction, using a logistic regression model. No significant interaction was found ($p = .71$).

Reversibility. The free responses were examined to determine whether participants subtracted funds from their side because they considered this action to be easier to reverse (e.g., through extra-experimental donations or fundraising actions) than adding to the opposing side. Four responses, amounting to less than 1% of participants, alluded to this motive in their explanation for their decision[24]. An additional participant, who chose to add funds to the opposing side, also mentioned they would "make up for" their choice by adding funds to their side too at a later time[25], demonstrating that extra-experimental donations could be used to justify both adding and subtracting funds.

Effectiveness measure. As in the previous studies, participants indicated that they viewed organizations on their political side to be more effective at spending their donation money to achieve their goals. The overall effectiveness rating was positive ($M = 0.71$, $SD = 1.43$), $t(496) = 11.13$, $p < .001$, $d = .50$, 95% CI = [.41, .59], indicating that participants believe organizations supporting their political party are more effective than those on the opposing side. This pattern held across Democratic and Republican participants, and across all conditions (all $Ps < .001$).

Discussion

This study found that participants' preference to harm their own group over supporting the opposing group is not influenced by degree of anonymity, suggesting that the motive to express ones' values is internalized. We also found that the effect was attenuated with higher stakes ($100 vs. $1), though was not eliminated. This may suggest that harm minimization

[24] The 4 responses were: "… I felt like I could always donate money myself later to make up for the subtraction, whereas the reverse would not be possible. That made the choice to subtract more preferable."; "I could always donate another dollar to the republicans later to even it back out. I would never be able to have an opportunity to take away the dollar that I gave to the Democrats, though."; "I would take away from my own party because I could help the party find new ways to gain ground and make up for the loss. Giving to the other party would simply give them more resources."; and "…After withdrawing $100 from the Democratic Party, I plan to donate $200 to the Democratic Party by adding my donation.".

[25] The response stated: "…I would be able to donate additional money to the Democratic organization at a later time to make up for this prior action."

motives play a larger role when consequences are larger. We also did not find strong evidence that participants opted to subtract funds because it was more easily reversible than adding funding to the opposing side.

References

Abramowitz AI, Webster S (2016) The rise of negative partisanship and the nationalization of US elections in the 21st century. *Electoral Studies* 41:12-22.

Abramowitz AI, Webster SW (2018) Negative partisanship: Why Americans dislike parties but behave like rabid partisans. *Political Psychology* 39:119-135.

Abrams DE, Hogg MA (1990) *Social identity theory: Constructive and critical advances* (Springer-Verlag Publishing).

Akerlof GA, Kranton RE (2000) Economics and identity. *The quarterly journal of economics* 115(3):715-753.

Allcott H (2011) Social norms and energy conservation. *Journal of public Economics* 95(9-10):1082-1095.

Andersen S, Harrison GW, Lau MI, Rutström EE (2006) Elicitation using multiple price list formats. *Experimental Economics* 9(4):383-405.

Andreoni J, Petrie R (2004) Public goods experiments without confidentiality: a glimpse into fund-raising. *Journal of public Economics* 88(7-8):1605-1623.

Ariely D, Bracha A, Meier S (2009) Doing good or doing well? Image motivation and monetary incentives in behaving prosocially. *American Economic Review* 99(1):544-55.

Baldwin MW, Holmes JG (1987) Salient private audiences and awareness of the self. *Journal of personality and social psychology* 52(6):1087.

Becker GM, DeGroot MH, Marschak J (1964) Measuring utility by a single-response sequential method. *Behavioral science* 9(3):226-232.

Benen S (2021) GOP starts to pay a price for opposing, then promoting, relief bill. May 7, 2021, https://www.msnbc.com/rachel-maddow-show/maddowblog/gop-starts-pay-price-opposing-then-promoting-relief-bill-n1266674.

Berman JZ, Kupor D (2020) Moral choice when harming is unavoidable. *Psychological science* 31(10):1294-1301.

Bodner R, Prelec D (2003) Self-signaling and diagnostic utility in everyday decision making. *The psychology of economic decisions* 1(105):26.

Bornstein G, Crum L, Wittenbraker J, Harring K, Insko CA, Thibaut J (1983) On the measurement of social orientations in the minimal group paradigm. *European Journal of Social Psychology* 13(4):321-350.

Borsari B, Carey KB (2003) Descriptive and injunctive norms in college drinking: a meta-analytic integration. *Journal of studies on alcohol* 64(3):331-341.

Brewer MB (1979) In-group bias in the minimal intergroup situation: A cognitive-motivational analysis. *Psychological bulletin* 86(2):307.

Bruner DM (2011) Multiple switching behaviour in multiple price lists. *Applied Economics Letters* 18(5):417-420.

Bullock JG, Gerber AS, Hill SJ, Huber GA (2013) Partisan Bias in Factual Beliefs about Politics.

Catapano R, Tormala ZL (2021) Do I support that it's good or oppose that it's bad? The effect of support-oppose framing on attitude sharing. *Journal of Personality and Social Psychology*.

Chambers JR, Baron RS, Inman ML (2006) Misperceptions in intergroup conflict: Disagreeing about what we disagree about. *Psychological Science* 17(1):38-45.

Chaudhry SJ, Loewenstein G (2019) Thanking, apologizing, bragging, and blaming: Responsibility exchange theory and the currency of communication. *Psychological review* 126(3):313.

Cialdini RB (2003) Crafting normative messages to protect the environment. *Current directions in psychological science* 12(4):105-109.

Cialdini RB, Reno RR, Kallgren CA (1990) A focus theory of normative conduct: recycling the concept of norms to reduce littering in public places. *Journal of personality and social psychology* 58(6):1015.

Cialdini RB, Demaine LJ, Sagarin BJ, Barrett DW, Rhoads K, Winter PL (2006) Managing social norms for persuasive impact. *Social influence* 1(1):3-15.

Davidai S, Ongis M (2019) The politics of zero-sum thinking: The relationship between political ideology and the belief that life is a zero-sum game. *Science Advances* 5(12):eaay3761.

DeSteno D, Li Y, Dickens L, Lerner JS (2014) Gratitude: A tool for reducing economic impatience. *Psychological science* 25(6):1262-1267.

Deutsch M, Gerard HB (1955) A study of normative and informational social influences upon individual judgment. *The journal of abnormal and social psychology* 51(3):629.

Dimock M, Wike R (2021) America Is Exceptional in Its Political Divide. *Pew: Trust Magazine*, March 29, 2021, https://www.pewtrusts.org/en/trust/archive/winter-2021/america-is-exceptional-in-its-political-divide.

Dorison CA, Wang K, Rees VW, Kawachi I, Ericson KM, Lerner JS (2020) Sadness, but not all negative emotions, heightens addictive substance use. *Proceedings of the National Academy of Sciences* 117(2):943-949.

Fehr E, Fischbacher U (2002) Why social preferences matter–the impact of non-selfish motives on competition, cooperation and incentives. *The economic journal* 112(478):C1-C33.

Fernbach PM, Van Boven L (2021) False Polarization: Cognitive Mechanisms and Potential Solutions. *Current Opinion in Psychology*.

Finkel EJ, Bail CA, Cikara M, Ditto PH, Iyengar S, Klar S, Mason L, McGrath MC, Nyhan B, Rand DG (2020) Political sectarianism in America. *Science* 370(6516):533-536.

Friedersdorf C (2021) The Conservatives Who'd Rather Die Than Not Own the Libs. Sep 24, 2021, https://www.theatlantic.com/ideas/archive/2021/09/breitbart-conservatives-john-nolte-vaccine/620189/.

Gerber AS, Rogers T (2009) Descriptive social norms and motivation to vote: Everybody's voting and so should you. *The Journal of Politics* 71(1):178-191.

Gneezy A, Gneezy U, Riener G, Nelson LD (2012) Pay-what-you-want, identity, and self-signaling in markets. *Proceedings of the National Academy of Sciences* 109(19):7236-7240.

Goldstein NJ, Cialdini RB, Griskevicius V (2008) A room with a viewpoint: Using social norms to motivate environmental conservation in hotels. *Journal of consumer Research* 35(3):472-482.

Halevy N, Bornstein G, Sagiv L (2008) “In-group love” and “out-group hate” as motives for individual participation in intergroup conflict: A new game paradigm. *Psychological science* 19(4):405-411.

Halevy N, Weisel O, Bornstein G (2012) “In-group love” and “out-group hate” in repeated interaction between groups. *Journal of Behavioral Decision Making* 25(2):188-195.

Heltzel G (2019) Seek and ye shall be fine: attitudes towards political perspective-seekers. University of British Columbia.

Hewstone M, Rubin M, Willis H (2002) Intergroup bias. *Annual review of psychology* 53(1):575-604.

Hitlin S (2003) Values as the core of personal identity: Drawing links between two theories of self. *Social psychology quarterly*:118-137.

Hogg MA (2016) Social identity theory. *Understanding peace and conflict through social identity theory* (Springer), 3-17.

Hogg MA, Reid SA (2006) Social identity, self-categorization, and the communication of group norms. *Communication theory* 16(1):7-30.

Hogg MA, Smith JR (2007) Attitudes in social context: A social identity perspective. *European Review of Social Psychology* 18(1):89-131.

Kessler JB, Milkman KL (2018) Identity in charitable giving. *Management Science* 64(2):845-859.

Klein E (2020) *Why we're polarized* (Simon and Schuster).

Lavine H, Huff JW, Wagner SH, Sweeney D (1998) The moderating influence of attitude strength on the susceptibility to context effects in attitude surveys. *Journal of Personality and Social Psychology* 75(2):359-373.

LeBoeuf RA, Shafir E, Bayuk JB (2010) The conflicting choices of alternating selves. *Organizational behavior and human decision processes* 111(1):48-61.

Lerner JS, Li Y, Weber EU (2013) The financial costs of sadness. *Psychological science* 24(1):72-79.

McGhee HC (2021) *The Sum of Us: What Racism Costs Everyone and How We Can Prosper Together.Racism has a cost for everyone* (One World/Ballantine).

Miller DT, Prentice DA (2016) Changing norms to change behavior. *Annual review of psychology* 67:339-361.

Morewedge CK, Tang S, Larrick RP (2018) Betting your favorite to win: Costly reluctance to hedge desired outcomes. *Management Science* 64(3):997-1014.

Pew (2019) Partisan Antipathy: More Intense, More Personal. https://www.people-press.org/2019/10/10/the-partisan-landscape-and-views-of-the-parties/.

Prelec D, Wernerfelt B, Zettelmeyer F (1997) The Role of Inference in Context Effects: Inferring What You Want from What Is Available. *Journal of Consumer Research* 24(1):118-125.

Rozin P, Royzman EB (2001) Negativity bias, negativity dominance, and contagion. *Personality and social psychology review* 5(4):296-320.

Sackeim HA, Gur RC (1979) Self-deception, other-deception, and self-reported psychopathology. *Journal of consulting and clinical psychology* 47(1):213.

Smith JR, Hogg MA (2008) Social identity and attitudes.

Swann Jr WB, Gómez Á, Dovidio JF, Hart S, Jetten J (2010) Dying and killing for one’s group: Identity fusion moderates responses to intergroup versions of the trolley problem. *Psychological Science* 21(8):1176-1183.

Tajfel H, Billig MG, Bundy RP, Flament C (1971) Social categorization and intergroup behaviour. *European journal of social psychology* 1(2):149-178.

Tetlock PE (2003) Thinking the unthinkable: Sacred values and taboo cognitions. *Trends in cognitive sciences* 7(7):320-324.

Tetlock PE, Kristel OV, Elson SB, Green MC, Lerner JS (2000) The psychology of the unthinkable: taboo trade-offs, forbidden base rates, and heretical counterfactuals. *Journal of personality and social psychology* 78(5):853.

Tonin M, Vlassopoulos M (2013) Experimental evidence of self-image concerns as motivation for giving. *Journal of Economic Behavior & Organization* 90:19-27.

Turner JC (1975) Social comparison and social identity: Some prospects for intergroup behaviour. *European journal of social psychology* 5(1):1-34.

Tversky A, Kahneman D (1991) Loss aversion in riskless choice: A reference-dependent model. *The quarterly journal of economics* 106(4):1039-1061.

--- (1992) Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty* 5(4):297-323.

Van Boven L, Ehret PJ, Sherman DK (2018) Psychological barriers to bipartisan public support for climate policy. *Perspectives on Psychological Science* 13(4):492-507.

Van der Werff E, Steg L, Keizer K (2014) Follow the signal: when past pro-environmental actions signal who you are. *Journal of Environmental Psychology* 40:273-282.

Volz LJ, Welborn BL, Gobel MS, Gazzaniga MS, Grafton ST (2017) Harm to self outweighs benefit to others in moral decision making. *Proc Natl Acad Sci U S A* 114(30):7963-7968.

White K, Dahl DW (2006) To be or not be? The influence of dissociative reference groups on consumer preferences. *Journal of Consumer Psychology* 16(4):404-414.

--- (2007) Are all out-groups created equal? Consumer identity and dissociative influence. *Journal of Consumer Research* 34(4):525-536.

Wickham H, Chang W, Wickham MH (2016) Package ‘ggplot2’. *Create Elegant Data Visualisations Using the Grammar of Graphics. Version* 2(1):1-189.

Wilkins CL, Wellman JD, Babbitt LG, Toosi NR, Schad KD (2015) You can win but I can't lose: Bias against high-status groups increases their zero-sum beliefs about discrimination. *Journal of Experimental Social Psychology* 57:1-14.

Zhong C-B, Phillips KW, Leonardelli GJ, Galinsky AD (2008) Negational categorization and intergroup behavior. *Personality and Social Psychology Bulletin* 34(6):793-806.

CHAPTER 3

INCREASED GENEROSITY UNDER COVID-19 THREAT

Ariel Fridman, Rachel Gershon, and Ayelet Gneezy

Rady School of Management, University of California, San Diego, La Jolla, CA, 92093, USA

Abstract

In the face of crises – wars, pandemics, and natural disasters – both increased selfishness and increased generosity may emerge. In this paper, we study the relationship between the presence of COVID-19 threat and generosity using a four-year longitudinal dataset (N = 696,942) capturing real donations made before and during the pandemic, as well as allocations from a six-month dictator game study (N = 1,003 participants) during the early months of the pandemic. Consistent with the notion of "catastrophe compassion" and contrary to some prior research showing a tendency toward self-interested behavior under threat, individuals across both datasets exhibited greater financial generosity when their county experienced COVID-19 threat. While we find that the presence of threat impacted individual giving, behavior was not sensitive to threat level. Our findings have significant societal implications and advance our understanding of economic and psychological theories of social preferences under threat.

Introduction

During major crises, such as natural disasters, wars, and now the COVID-19 pandemic, two conflicting behaviors may emerge: increased selfishness or increased generosity. Selfishness is innate to our survival instincts (Dawkins 1976, Rachlin 2002). Evidence suggests that when facing adverse circumstances, individuals' may shift away from other-regarding practices (Erikson 1976), arguably because fear and uncertainty resulting from increased risk perceptions (Slovic and Weber 2013) drive self-preservation (Rodrigues et al. 2009). Along these lines, research finds that the presence of threat can decrease individuals' willingness to engage in charitable activities and civic duties (e.g., paying taxes and reporting a crime (Poulin et al. 2012)) and that generosity toward in-group and out-group members may decrease following a natural disaster (Vardy and Atkinson 2019). Considering these findings, one might expect that individuals experiencing COVID-19 threat would, on average, behave more selfishly than those not experiencing threat. Indeed, at an early stage of the COVID-19 pandemic, survey data representing 35 countries showed that, despite government appeals, those who felt more threatened were more likely to engage in selfish stockpiling (Garbe et al. 2020), putting the health and well-being of others at risk (U.S. Attorney urges public to report potential hoarding of supplies needed to fight COVID-19 2020). Across the world, many of the widespread product shortages during the early months of the COVID-19 pandemic were triggered by consumers purchasing resources in excess of their actual need (e.g., hoarding toilet paper (Westbrook 2020) and masks (McNeil 2020)).

On the other hand, there is evidence suggesting that groups facing a common threat often demonstrate stronger social cohesion (Fritz 1961, Gilligan et al. 2014) and more cooperative, communal behaviors (Turkel 2002, Bowles and Gintis 2003, A. Gneezy and Fessler 2012, Bauer et al. 2016). As proposed by Jamil Zaki's "catastrophe compassion" theory (Zaki 2020), disasters may promote an increased sense of community and altruism (Glynn et al. 2003). Indeed, experiencing high (vs. low) stress can increase trust and sharing behaviors (Dovidio and Morris 1975, von Dawans et al. 2012). Perceived threat may also promote the expansion of social connections, as observed in monkeys in response to the environmental instability caused by Hurricane Maria (Testard et al. 2021). Considering these findings, one might predict that experiencing and bearing witness to the devastating effects of the COVID-19 pandemic would promote generosity. Consistent with this proposition, a survey conducted by the Lilly Family School of Philanthropy (Mesch et al. 2020) during the COVID-19 pandemic found that nearly half of respondents supported their communities in a variety of ways, for example by continuing to pay individuals and businesses for services that could not be rendered (Adamczyk 2020).

Interestingly, research suggests that the impact of threat could go either way – increasing or decreasing generosity. For example, when primed with resource scarcity, individuals became more competitive, causing some to become more selfish while others to exhibit greater generosity, depending on the context (e.g., behavior observability) (Roux et al. 2015). Additional research finds evidence that communities experiencing disasters could simultaneously undergo positive and negative behavioral change (Nurmi et al. 2012).

We examine the relationship between COVID-19 threat and generosity using two independent longitudinal datasets. The first dataset, provided by Charity Navigator (CN), the world's largest independent charity evaluator, consisted of actual charitable-giving data spanning July 2016 through December 2020 (N = 696,942 donations). For each donation, the data included the donation amount, the charities benefited, each charity's assigned category (e.g., environment and human services), and the donor's location. In addition, CN assigned each donor a unique identifier, which allowed us to observe within-person differences in donation behavior in both the presence and absence of COVID-19 threat.

The second dataset, which sheds light on the relationship between COVID-19 threat and generosity in a more controlled setting, consisted of individuals' (N = 1,003 U.S. participants) allocations from an incentivized dictator game. In the dictator game, one player (the *dictator*) receives $10 and makes a unilateral decision on how to divide it between themselves and another, typically unknown, individual (Kahneman et al. 1986, Charness and Gneezy 2008). Rather than maximizing their own financial payoff (i.e., allocating $0), experimental evidence shows that "dictators" often choose to give some of their money to recipients (Henrich et al. 2004, List 2007, Engel 2011). In our study, participants played an incentive-compatible dictator game monthly, from March to August 2020. Importantly, at the start of this period, COVID-19 threat was only present in 10% of participants' counties (see Table A3.1), allowing us to observe their behavior when threat was first introduced in most counties. Notably, while the dictator game has previously been used to capture generosity at a single point in time, there are few cases of its use in a longitudinal setting (Vardy and Atkinson 2019, Arechar and Rand 2022). Finally, our research is unique in the use of longitudinal dictator game data jointly with longitudinal archival data to show convergent evidence of changes in giving behavior.

While our observational datasets do not lend themselves to causal claims, it is reasonable to infer that the presence of threat would increase generosity (Dovidio and Morris 1975, Bowles and Gintis 2003, Glynn et al. 2003, A. Gneezy and Fessler 2012, von Dawans et al. 2012, Gilligan et al. 2014, Bauer et al. 2016, Adamczyk 2020, Mesch et al. 2020, Zaki 2020, Testard et al. 2021), while reverse causality is highly unlikely. See Discussion for a more detailed explanation.

Our large-scale longitudinal datasets provide real-world evidence that people exhibited greater generosity during a time where some theories and experts predicted the opposite due to the economic downturn associated with the pandemic. While our analyses consider various levels of threat, we found that only the presence (vs. absence) of threat was associated with greater generosity. Our findings contribute to economic and psychological theories of social preferences, suggesting that people come together in the presence of a shared threat and demonstrate a willingness to support others, despite the uncertainty surrounding their own health and financial well-being.

Results

Across both datasets, we observe increased generosity in the presence of COVID-19 threat in participants' geographic location.

Our analyses controlled for potential confounds at the national level (e.g., stimulus payouts and the Black Lives Matter movement) by including date fixed effects in all regressions. See Methods section for details.

Charity Navigator

We first analyzed the data with a relatively simplified approach by treating COVID-19 threat as a binary variable—whether any COVID-19-related deaths (vs. no COVID-19 deaths) occurred in each calendar month. In this analysis, we compared the proportion of counties that increased their overall donation amount as a function of whether the county had experienced COVID-19 threat. Compared with March 2019, 78% of counties that experienced threat increased the total amount donated in March 2020. Of the counties that did not face threat, 55% increased giving ($\chi^2(1, N = 440) = 24.75$, $P < .001$; see Figure 3.1). We found similar results for the comparison between April 2019 and April 2020 (see Figure A3.1). After April 2020, only 33 counties or fewer did not face threat, though results were directionally the same for all months except August and October. While we treat counties with no COVID-19 deaths as "no threat", it is possible that some people in these counties still experienced COVID-19 threat, suggesting that our analysis is conservative. These results present initial model-free evidence for our main finding – that a greater proportion of counties experiencing threat increased giving, compared with those facing no threat.

To examine the relationship between COVID-19 threat level and county-level donation amount, we ran a regression analysis using county-level data aggregated by month-year. We captured the COVID-19 threat level using a seven-day average of daily new deaths per million in each county, which we averaged over all days in each month and binned into four categories by quantiles (no, low, medium, or high threat; log transformation generated similar results, see Tables A3.2 and A3.3). We binned the threat variable to allow for non-linear relationships between threat and charitable giving. We regressed log-transformed aggregated giving amounts on threat level and included county and month-year fixed effects. This analysis showed that, overall, giving through CN's platform increased across all threat levels compared with no threat. On average, county-level giving increased 31.6% under low threat (SE = .06, t = 4.94, P < .001), 28.5% under medium threat (SE = .07, t = 3.86, P < .001), and 32.9% under high threat (SE = .05, t = 6.10, P < .001), relative to periods of no threat in the county. All pairwise comparisons across low, medium, and high threat levels were non-significant (P > .32), suggesting insensitivity to threat magnitude.

To examine whether our findings held within individuals (i.e., among repeat donors), we analyzed individual-level data, which allows us to rule out potential selection bias (i.e., changes in donor characteristics before and during the COVID-19 pandemic). Of those who donated in

2020, 32% were repeat donors, meaning they made donations through the platform more than once. We captured the level of COVID-19 threat using the same seven-day lagged moving average of daily new deaths per million (without month-level averaging), considering "no threat" as our baseline. A regression of log-transformed giving amounts on threat level, including individual and date fixed effects, revealed that repeat donors' giving increased significantly across all charity categories by 3.4% under high threat (SE = .02, t = 2.11, P = .040), and non-significantly by 1.3% under medium threat (P = .348), and 2.8% under low threat (P = .063). All pairwise comparisons across low, medium, or high threat levels were non-significant (P > .07).

Although this effect is smaller than the effect we observe with the county-level model, including interaction terms with the charity category revealed a highly significant increase in donations to human services charities – organizations that provide direct services to those in need (see Table A3.4). While one might expect an increase in donations to health charities, this category includes organizations such as Planned Parenthood and Cure Alzheimer's Fund, which are not directly related to the COVID-19 pandemic. In contrast, human services charities include food banks and homeless services. Among repeat donors, donations to human services charities increased by 8.4% under low threat (SE = .02, t = 4.48, P < .001), 6.7% under medium threat (SE = .02, t = 4.47, P < .001), and 8.0% under high threat (SE = .01, t = 5.39, P < .001), relative to periods of no threat in the donor's county (Figure 3.2). All pairwise comparisons across low, medium, or high threat levels were non-significant (P > .29). Collapsing all charity categories except human services revealed no significant difference in donations in response to the presence of COVID-19 threat (Ps > .55 for low, medium, and high threat). Notably, our analysis revealed a significant interaction of COVID-19 threat with human services charities (Ps < .001), indicating that under COVID-19 threat in one's county, repeat donors were significantly more likely to donate to human services charities than to other types of charities (see Table A3.4). Together, these findings suggest giving to human services charities increased under threat but not at the expense of donations to other charity categories.

In an additional analysis, we found that the observed county-level increases in giving did not vary by county-level median household income (see Table A3.5 for results).

Dictator Game

An analysis of our six-wave longitudinal dictator-game data also showed that participants gave significantly more under the presence of COVID-19 threat. We used the same threat measure and individual-level regression model as the CN analysis. Our outcome measure was participants' allocation decisions, and we used wave fixed effects instead of date fixed effects. We found that within-person giving increased by approximately \$0.25 (8.6%) under low threat (SE = .10, $t = 2.58$, $P = .013$), \$0.38 (13.1%) under medium threat (SE = .11, $t = 3.45$, $P = .001$), and \$0.24 (8.3%) under high threat (SE = .09, $t = 2.56$, $P = .014$), relative to periods of no threat in the participant's county (Figure 3.3 and Table A3.6). Percentages were calculated relative to a mean allocation of \$2.92. All pairwise comparisons across low, medium, and high threat levels were non-significant ($P > .10$). Our analysis further revealed a significant interaction between allocation amount and gender, indicating women gave more than men under threat (collapsed across low, medium, and high threat levels, $P = .033$). We found no interaction between allocation amount and age. Unlike some other COVID-19-related behaviors (Calvillo et al. 2020, van Holm et al. 2020, Fridman et al. 2021), we found no difference in giving patterns based on political affiliation.

Discussion

Researchers have long argued that experiencing threat influences social preferences (Zaki 2020), but offer different predictions for how. While some propose that those facing threat will become more generous, others predict increased selfishness. Leveraging data from a naturally-occurring state of emergency – the COVID-19 pandemic – we investigated the relationship between local COVID-19 threat and generosity.

The present work offers several important findings. First, analyses of both datasets show that individuals exhibited greater financial generosity under COVID-19 threat. The CN data indicates contributions were directed primarily toward charities in the human services category – organizations that help mitigate the effects of COVID-19. Second, although we examined local COVID-19 threat, increased generosity often emerged in support of non-local organizations (CN data) or unidentified individuals (dictator game data). CN's data also show that the increase in donations was exhibited by both repeat and new donors, suggesting an overall increase in giving, as opposed to a mere allocation shift. We note that although both datasets demonstrate an association between the presence of threat and generosity, we found insensitivity to the level of threat. Although merely speculative, this pattern is consistent with research demonstrating scope insensitivity in emotionally charged settings (Hsee and Rottenstreich 2004, Weber and Johnson 2009, Urminsky and Kivetz 2011). Finally, the increased generosity observed across both datasets is particularly intriguing in light of expert predictions, based on historical data, that the economic downturn caused by the pandemic would lead to reduced giving (Beer 2020), and the fact that a record-high majority of Americans reported a worsening financial situation during the same period (Jones 2020). Prior work suggests that when people experience such financial scarcity, they may engage in extreme, even immoral, behaviors to acquire financial wealth

(Prediger et al. 2014, Sharma et al. 2014). Yet analyses of both our datasets clearly shows that in this particular circumstance, individuals were, on average, more willing to part with their financial resources.

From a methodological perspective, our results lend credibility to the dictator game as a reliable measure of real-world generosity (Levitt and List 2007, U. Gneezy and Imas 2017), because our dictator-game findings are consistent with CN's field data. Thus, this work adds to a growing discussion in the literature regarding the validity of lab findings (Charness et al. 2013, Kessler and Vesterlund 2015). The present research is also unique with respect to the longitudinal nature of our data, which, as noted in a recent call for the integration of such data (Chintagunta and Labroo 2020), is largely absent from behavioral research, and particularly rare in the context of major crises (Testard et al. 2021).

Although both datasets show that the presence of local COVID-19 threat is associated with increased generosity, our observational data does not necessarily lend itself to causal claims. With that in mind, we believe it is reasonable to infer that the presence of threat would influence generosity (Dovidio and Morris 1975, Bowles and Gintis 2003, Glynn et al. 2003, A. Gneezy and Fessler 2012, von Dawans et al. 2012, Gilligan et al. 2014, Bauer et al. 2016, Adamczyk 2020, Mesch et al. 2020, Zaki 2020, Testard et al. 2021); it is improbable, however, that the increase in generosity would trigger an increase in local COVID-19 threat. While we cannot rule out the possibility that there could be another variable that influences both factors to produce the observed correlation, our analyses and contemplation do not point to any likely candidates. However, it is possible that the presence of COVID-19 threat influences other variables (e.g., local lockdowns, media attention, etc.), which in turn lead to an increase in generosity. Nonetheless, we cautiously conclude that the presence of local COVID-19 threat led, either directly or indirectly, to an increase in generosity.

We are unable to discern what mechanism underlies the observed increase in generosity. Individuals may have been motivated to give more when experiencing threat as a result of increased feelings of sympathy (Loewenstein and Small 2007), a desire to regain a sense of agency and self-efficacy (Cryder et al. 2013, U. Gneezy et al. 2014), mortality salience (Jonas et al. 2002, Dunn et al. 2020), or simply to experience positive emotions (e.g., warm glow (Andreoni 1990, Park et al. 2017)) during a stressful period. It is also possible our findings reflect self-interested behavior if people increased their donations to benefit themselves, thinking that doing so would help to mitigate the negative effects of COVID-19 in general. However, it is unlikely that this motive explains the dictator game findings, where giving benefits a randomly chosen anonymous individual. Based on our CN data, it also seems likely individuals donated to honor those who passed away during the pandemic or were otherwise affected (e.g., sick, lost a loved one, etc.). Indeed, a closer look at the CN data shows that the proportion of donations made "in memory of" someone was significantly greater in 2020 than every prior year, although this motivation is likely less salient in our longitudinal dictator-game findings.

Since our CN data represents individuals with financial means to give, we cannot rule out the possibility that our findings apply to a subset of the population. However, we found similar results using the dictator-game data, which likely represents a less affluent demographic (Goodman et al. 2013), and an analysis comparing changes in giving on CN by county-level median household income also found no significant differences.

While our work focuses on the early period of the COVID-19 pandemic, additional research is needed to understand the dynamics of the relationship between threat and generosity in the longer term, as well as once the crisis has ended (e.g., see (Vardy and Atkinson 2019) for a dictator-game study before and after Cyclone Pam in 2015).

Finally, while our results show an increase in financial generosity under COVID-19 threat, it is possible that the pandemic also resulted in selfish behavior unobserved in our data (e.g., hoarding resources). Future research can investigate when, why, and for whom these divergent reactions may occur.

This work adds to our understanding of human behavior during times of crisis. Amidst the uncertainty, fear, and tragedy of the pandemic, we find a silver lining: people became more financially generous toward others in the presence of COVID-19 threat.

Materials and Methods

Human Subject Protections

The longitudinal dictator game received ethical approval from the UC San Diego Institutional Review Board (Project #191273XX) and was carried out in accordance with the guidelines and regulations for a human research study. Informed consent was obtained for all participants. No personal identifying information was collected to ensure participants' privacy.

COVID-19 Threat

To capture COVID-19 threat, we used the Center for Systems Science and Engineering at Johns Hopkins University (Dong et al. 2020) time-series data of daily new deaths at the county level, which also includes U.S. Census Bureau population data by county. Our measure represented a seven-day lagged moving average of daily new deaths per million, which we binned into four categories based on population-weighted quantiles, taken over the entire timespan of the COVID-19 data in 2020: no deaths (0), bottom third (1), middle third (2), and top third (3). Consistent with prior work (Pagliaro et al. 2021), we used deaths rather than confirmed cases for the threat measure, because it is more consistent over time, as changes in confirmed cases may capture changes in testing availability, particularly at the beginning of the pandemic (Cumulative Cases , Wu et al. 2020).

To merge participant data from the dictator game and donor data from CN with our COVID-19 threat measure, we mapped participants' zip codes to the corresponding county using the U.S. Department of Housing and Urban Development data (HUD USPS ZIP Code Crosswalk Files | HUD USER).

County-level median household income data was obtained from the U.S. Department of Agriculture's Economic Research Service (USDA ERS - County-level Data Sets 2021).

Charity Navigator

To measure changes in generosity over time, we analyzed a dataset of donations made through CN. For each donation, the data specified the following: date, donation amount, charity name, a unique donor identifier (anonymized email address), and donor zip code. After excluding incomplete observations (0.9%; see Appendix Text A3.1), we were left with 696,942 donations.

Dictator Game

We recruited a representative panel of U.S. residents (see Appendix Text A3.2) on Amazon's Mechanical Turk platform to respond to monthly survey waves from March to August 2020. The sample consisted of 1,003 unique participants. The first wave included 998 participants, ranging between 605 to 755 thereafter (5 participants did not respond to the dictator game in the first wave, but did so on subsequent waves; for attrition details, see Table A3.1). After excluding incomplete observations due to missing county information (0.01%), we were left with 4,272 observations.

Participants played the dictator game (Forsythe et al. 1994, Engel 2011) on all six survey waves and were always assigned to the dictator role. Participants allocated an amount between $0 and $10 to a randomly selected participant. To incentivize responses, we informed participants that, in each wave, one randomly selected participant would receive a $10 bonus, which would be split between them and another randomly selected participant according to their decision. Our survey also included demographic questions capturing age, gender, and political party affiliation.

Regression Models

For our primary analysis, we examined the relationship between COVID-19 threat and giving using multiple regression.

Charity Navigator. In our county-level specification, we aggregated the data to the month-year level by county and estimated the following model:

$$log(y_{cmy}) = \sum_{s=1}^{3}(\beta_s 1_{threat_{cmy}=s}) + \alpha_c + \alpha_{my} + \varepsilon_{cmy},$$

where y_{cmy} is the sum of the amount donated over all individuals in county *c* in month *m* in year *y*, α_c are county-level fixed effects, and α_{my} are month-year fixed effects. $threat_{cmy}$ is the average threat level in county *c* in month *m* in year *y*, where *s* indexes the threat level and coefficients $\beta_s$measure the effect of COVID-19 threat on giving, relative to no threat. We log-transformed the amount donated because it was right-skewed. To test whether the effect of threat varied by median household income, we ran a similar specification, including interaction terms for median household income and $threat_{cmy}$. Standard errors were clustered by state and the regression was weighted by county population.

In our individual-level specification, we exploited within-person variation in giving over time, estimating the following model:

$$log(y_{id}) = \sum_{s=1}^{3}(\beta_s 1_{threat_{c_i d}=s}) + \alpha_i + \alpha_d + \varepsilon_{id},$$

where y_{id} is the amount donated by individual *i* on date d, α_i are individual-level fixed effects, and α_d are date fixed effects. $threat_{c_i d}$is the threat level in county *c* of individual *i* on date *d*. To estimate the effect of COVID-19 threat on donations for each category of charities, we ran a similar specification, including charity-category dummy variables and interaction terms for $threat_{c_i d}$and each charity category. Standard errors for both regressions were clustered by individual and state.

Dictator Game. For the dictator-game analysis, we used a similar approach and estimated the following model:

$$y_{iw} = \sum_{s=1}^{3}(\beta_s 1_{threat_{c_i w}=s}) + \alpha_i + \alpha_w + \varepsilon_{iw},$$

where y_{iw} is the amount given by individual *i* on survey wave *w*, $\alpha_i$are individual-level fixed effects, and α_w are survey-wave fixed effects. $threat_{c_i w}$ is the threat level in county c of individual *i* on wave *w*, where *s* indexes the threat level and coefficients β_s measure the effect of COVID-19 threat on giving, relative to no threat. Standard errors were clustered by individual and state.

See Appendix Text A3.3 for additional methodological information.

Data Availability

Our materials, anonymized behavioral data, and additional analyses, including robustness checks, are available at https://osf.io/2rykb/.

Acknowledgements

We thank UC San Diego Global Health Initiative (GHI) for providing funding, and Charity Navigator for sharing their anonymized data with us. We also thank Kristen Duke, Uri Gneezy, Minah Jung, Elizabeth Keenan, Katherine Milkman, Leif Nelson, Sally Sadoff, Marta Serra-Garcia, and Kenneth Wilbur for their insightful feedback on this work.

Chapter 3, in full, is a reprint of previously published material as it appears in Scientific Reports 12, 4886, Ariel Fridman, Rachel Gershon, and Ayelet Gneezy. The dissertation author was the primary investigator and author of this paper.

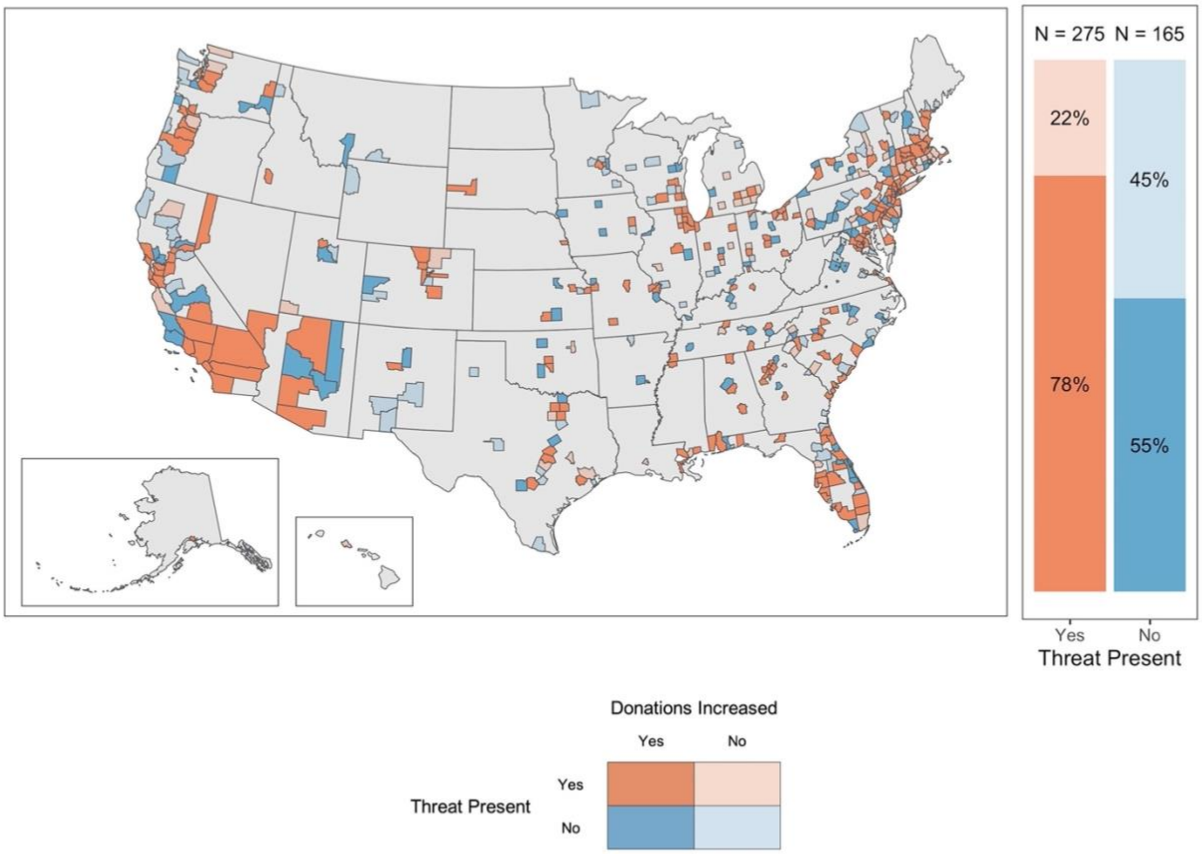

Figure 3.1. Counties by COVID-19 Threat and Donation Changes – March 2019 vs. March 2020. Orange [blue] represents the presence [absence] of threat in March 2020. Darker [lighter] shades indicate an increase [decrease] in giving across all charity categories relative to March 2019. The map shows U.S. counties with inset maps for counties in Alaska and Hawaii. The chart on the right shows the proportion of counties in each "threat present" and "donations increased" group.

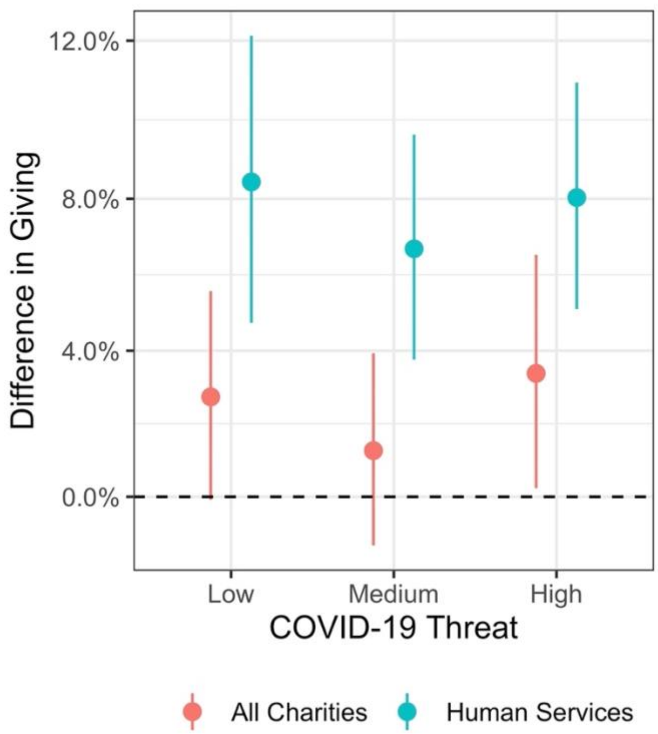

Figure 3.2. Charity Navigator Donations. The vertical axis captures the difference between giving under threat relative to no threat (dashed line). The horizontal axis indicates threat levels in each participant's county at the time of the donation. Points and error bars represent regression coefficient estimates and 95% confidence intervals, respectively. Note the "All charities" category includes human services charities.

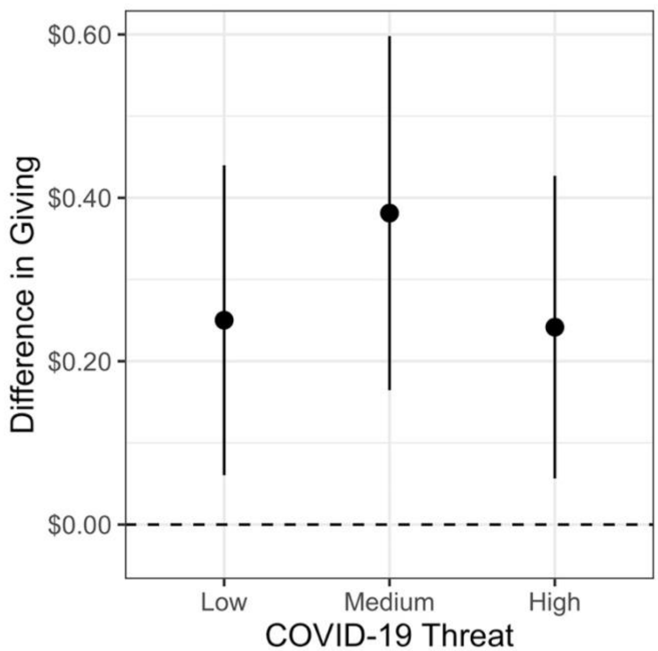

Figure 3.3. Dictator Game Allocations. The vertical axis captures the difference between giving under threat relative to no threat (dashed line). The horizontal axis indicates threat levels in each participant's county at the time of the survey. Points and error bars represent regression coefficient estimates and 95% confidence intervals, respectively.

Appendix

Text A3.1. Additional Information for Charity Navigator.

The Charity Navigator's dataset used in this analysis spans from July 22, 2016, through December 10, 2020; the dictator-game was administered monthly from March 16 to August 16, 2020.

We excluded incomplete entries (i.e., missing email addresses, zip codes not matched to counties, or counties missing COVID-19 data) from our CN dataset, accounting for 0.9% of observations. We included donations that were refunded in our analysis (about 0.14% of donations), because almost all donations from September 21, 2020, onward were not yet determined to be refunded or not.

Text A3.2. Additional Information for Dictator Game.

To incentivize completion of all survey waves, we informed participants in Wave 1 that their payment would increase for subsequent surveys and that those completing the first three waves would enter a $100 raffle. Participants were paid $0.30 for wave 1, $0.40 for wave 2, $0.60 for waves 3 and 4, $1.00 for wave 5, and $1.20 for wave 6. The median survey completion time was 5.5 minutes.

Our panel represented the broad and diverse population of the U.S. The first-wave sample included participants from all 50 states (except Wyoming) and Washington D.C., with an age range of 18 to 82 years old (mean = 38.48, median = 35). Approximately half of our participants (53%) identified as male, 46% as female, and .6% as other. The racial makeup in our sample was 80% White, 9% Asian, 6% Black or African American, 4% multiple racial or ethnic identities, and 1% other. Relative to the U.S. Census (2019) (U.S. Census Bureau QuickFacts: United

States 2019) estimates, our sample over-represents White and Asian individuals and under-represents Black or African American individuals and other racial groups.

Text A3.3. Additional Methodological Information.

All analyses were conducted using R (version 4.0.2), and regressions were run using the package "fixest" (version 0.6.0).

For the measure of COVID-19 threat, we determined category cutoffs using county-level daily data spanning January 22, 2020 (the earliest available) to December 31, 2020, for every county in the U.S. We classified observations with a value of 0 on the COVID-19-threat measure as no threat, and classified the remaining observations as low, medium, or high threat based on whether they fell in the bottom, middle, or top third (population-weighted), respectively.

In both the dictator game and CN dataset, we mapped each participant's zip code to the corresponding county, allowing us to merge each dataset with the COVID-19-threat data. If a zip code matched multiple counties, we matched it to the county with the larger population.

Counties by COVID−19 Threat and Donation Changes − April 2019 vs. April 2020

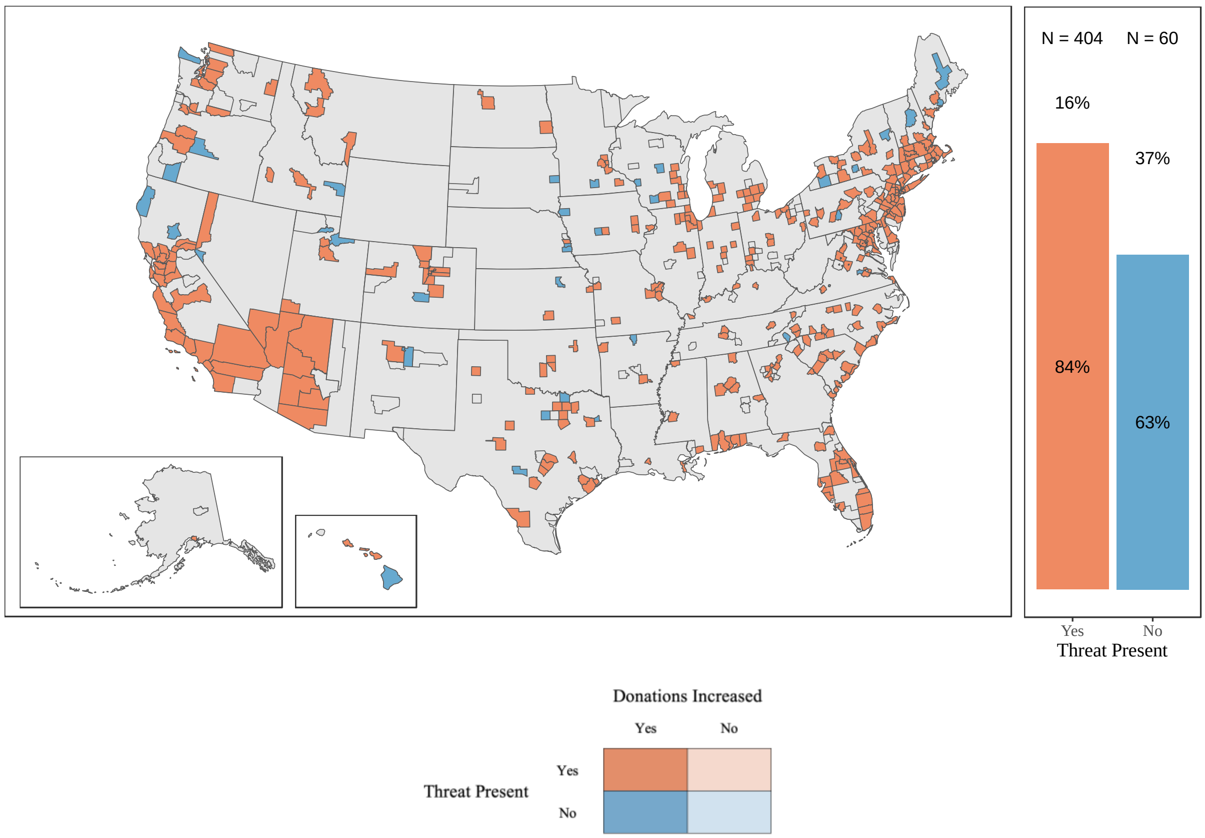

Figure A3.1. April 2019 vs. 2020 County Map. Orange [blue] represents the presence [absence] of threat in April 2020. Darker [lighter] shades indicate an increase [decrease] in giving across all charity categories relative to April 2019. The main map shows U.S. counties with inset maps for counties in Alaska and Hawaii. The chart on the right shows the proportion of counties in each "threat present" and "donations increased" group. Similar to the results shown for March in Figure 3.1, a greater proportion of counties facing threat increased giving, compared with those facing no threat. Specifically, compared with April 2019, 84% of counties experiencing threat increased the total amount donated in April 2020. Of the counties that did not face a threat, 63% increased giving ($\chi2(1, N = 464) = 14.13$, $P < .001$).

Table A3.1. Dictator-Game Attrition.

To rule out differential attrition, we tested whether the composition of our sample (e.g., age, gender, and political party) changed over time. Specifically, we tested whether participants who responded to waves 2–6 were significantly different at baseline (wave 1) from the entire sample in wave 1. The only significant change detected (Ps < .05) was with respect to participants' age, though the differences were small—the average age was 38.5 in the first wave, and among participants who responded to subsequent waves, the baseline average age ranged from 39.9 and 40.8. We found no other systematic pattern of attrition among participants.

Table shows the baseline (wave 1) characteristics of respondents to each survey wave. For gender, only the proportion of females is shown; participants identifying as "another gender" constituted around 0.65% of the sample in each wave. For all variables, we tested whether participants who responded to waves 2-6 were significantly different at baseline (wave 1) from the full sample at baseline (chi-squared test for gender and threat, t-test for all others). The number of observations across variables vary somewhat because respondents were allowed to skip questions. Significance codes: *** P < .001, ** P < .01, * P < .05.

Variable	Wave Participants	Wave 1 Mean (SD)	N
Age	1	38.48 (12.21)	997
	2	39.93 (12.54)*	762
	3	40.59 (12.49)***	654
	4	40.80 (12.43)***	608
	5	40.14 (12.47)**	651
	6	40.33 (12.39)**	666
Gender (female)	1	46.19%	998
	2	46.19%	762
	3	46.02%	654
	4	45.72%	608
	5	45.86%	652
	6	47.08%	667

Table A3.1. Dictator-Game Attrition. (Continued)

Variable	**Wave Participants**	**Wave 1 Mean (SD)**	**N**
Political Party (1 = strongly Republican; 6 = strongly Democratic)	1	3.98 (1.49)	998
	2	3.95 (1.50)	762
	3	4.00 (1.50)	654
	4	3.95 (1.52)	608
	5	3.97 (1.52)	652
	6	4.00 (1.49)	667
Dictator Game Allocations ($0-$10)	1	2.97 (2.54)	998
	2	2.97 (2.44)	757
	3	2.90 (2.49)	649
	4	2.91 (2.47)	604
	5	2.91 (2.47)	649
	6	2.95 (2.46)	665
No COVID-19 Threat	1	89.05%	986
	2	89.52%	754
	3	89.01%	646
	4	90.37%	602
	5	90.08%	645
	6	90.47%	661

Table A3.2. Charity Navigator Regression Table – Logged Threat Level. Table shows full regression results of CN models using the log transformed threat measure. The dependent variable for all models was log-transformed donation amounts. Including the category interactions in the individual-level model reduced the number of observations due to missing category labels. Significance codes: *** p < .001, ** p < .01, * p < .05, † p < .1.

	County-Month Aggregation	**Individual-Level**	**Individual-Level**
Threat Level (logged)	0.0396* (0.0196)	0.0096 (0.0058)	
Category: Arts, Culture, Humanities			-0.1438*** (0.0127)
Category: Community Development			0.1892*** (0.0076)
Category: Education			0.0182* (0.0082)
Category: Environment			0.0500*** (0.0047)
Category: Health			0.1159*** (0.0047)
Category: Human and Civil Rights			0.0694*** (0.0067)
Category: Human Services			0.1378*** (0.0051)
Category: International			0.1927*** (0.0053)
Category: Religion			0.0707*** (0.0203)
Category: Research and Public Policy			-0.6576*** (0.0218)
Threat Level (logged) x Category: Animals			-0.0085 (0.0077)
Threat Level (logged) x Category: Arts, Culture, Humanities			0.0146 (0.0115)
Threat Level (logged) x Category: Community Development			0.0050 (0.0104)
Threat Level (logged) x Category: Education			0.0152 (0.0109)
Threat Level (logged) x Category: Environment			-0.0107 (0.0093)
Threat Level (logged) x Category: Health			-0.0142† (0.0083)
Threat Level (logged) x Category: Human and Civil Rights			0.0120 (0.0078)
Threat Level (logged) x Category: Human Services			0.0339*** (0.0068)
Threat Level (logged) x Category: International			-0.0071 (0.0064)
Threat Level (logged) x Category: Religion			0.0249 (0.0258)
Threat Level (logged) x Category: Research and Public Policy			-0.0177 (0.0204)
Fixed-Effects:			
County	Yes	No	No
Month-Year	Yes	No	No
Individual	No	Yes	Yes
Date	No	Yes	Yes

Table A3.2. Charity Navigator Regression Table – Logged Threat Level. (Continued)

	County-Month Aggregation	**Individual-Level**	**Individual-Level**
S.E. Clustered by:	State	Individual & State	Individual & State
Weights:	County Population	None	None
Observations	116,480	696,942	617,657
R^2	0.82726	0.83448	0.85416
Within R^2	0.0000568	2.31E-05	0.06418

Table A3.3. Dictator-Game Regression Table – Logged Threat Level. Table shows full regression results of the dictator-game model using the log transformed threat measure. The dependent variable was the allocation amount. Significance codes: *** p < .001, ** p < .01, * p < .05.

Threat Level (logged)	0.0705 (0.0429)
Fixed-Effects:	
Individual	Yes
Wave	Yes
S.E. Clustered by:	Individual & State
Weights:	None
Observations	4,272
R^2	0.72841
Within R^2	0.00114

Table A3.4. Charity Navigator Regression Table. Table shows full regression results of CN models described in the text. The dependent variable for all models was log-transformed donation amounts. Including the category interactions in the individual-level model reduced the number of observations due to missing category labels. Significance codes: *** p < .001, ** p < .01, * p < .05, † p < .1.

	County-Month Aggregation	**Individual-Level**	**Individual-Level**
Threat Level: Low	0.3163*** (0.0641)	0.0275† (0.0145)	
Threat Level: Medium	0.2849*** (0.0738)	0.0128 (0.0135)	
Threat Level: High	0.3294*** (0.0540)	0.0339* (0.0161)	
Category: Arts, Culture, Humanities			-0.1433*** (0.0124)
Category: Community Development			0.1882*** (0.0083)
Category: Education			0.0168* (0.0079)
Category: Environment			0.0510*** (0.0045)
Category: Health			0.1159*** (0.0049)
Category: Human and Civil Rights			0.0694*** (0.0064)
Category: Human Services			0.1334*** (0.0050)
Category: International			0.1924*** (0.0053)
Category: Religion			0.0695** (0.0215)
Category: Research and Public Policy			-0.6568*** (0.0218)
Threat Level: Low x Category: Animals			-0.0087 (0.0199)
Threat Level: Medium x Category: Animals			-0.0117 (0.0178)
Threat Level: High x Category: Animals			0.0088 (0.0238)
Threat Level: Low x Category: Arts, Culture, Humanities			0.0015 (0.0257)
Threat Level: Medium x Category: Arts, Culture, Humanities			-0.0159 (0.0234)
Threat Level: High x Category: Arts, Culture, Humanities			0.1020** (0.0326)
Threat Level: Low x Category: Community Development			0.0346 (0.0309)
Threat Level: Medium x Category: Community Development			0.0165 (0.0240)
Threat Level: High x Category: Community Development			0.0149 (0.0295)
Threat Level: Low x Category: Education			0.0554† (0.0330)
Threat Level: Medium x Category: Education			0.0168 (0.0320)
Threat Level: High x Category: Education			0.0433 (0.0304)

Table A3.4. Charity Navigator Regression Table. (Continued)

	County-Month Aggregation	**Individual-Level**	**Individual-Level**
Threat Level: Low x Category: Environment			-0.0318 (0.0236)
Threat Level: Medium x Category: Environment			-0.0066 (0.0222)
Threat Level: Low x Category: Environment			-0.0318 (0.0236)
Threat Level: Medium x Category: Environment			-0.0066 (0.0222)
Threat Level: High x Category: Environment			0.0011 (0.0241)
Threat Level: Low x Category: Health			0.0005 (0.0221)
Threat Level: Medium x Category: Health			-0.0262† (0.0154)
Threat Level: High x Category: Health			-0.0038 (0.0231)
Threat Level: Low x Category: Human and Civil Rights			0.0261 (0.0211)
Threat Level: Medium x Category: Human and Civil Rights			0.0020 (0.0154)
Threat Level: High x Category: Human and Civil Rights			0.0529* (0.0234)
Threat Level: Low x Category: Human Services			0.0843*** (0.0188)
Threat Level: Medium x Category: Human Services			0.0670*** (0.0150)
Threat Level: High x Category: Human Services			0.0803*** (0.0149)
Threat Level: Low x Category: International			0.0067 (0.0163)
Threat Level: Medium x Category: International			0.0053 (0.0126)
Threat Level: High x Category: International			-0.0055 (0.0182)
Threat Level: Low x Category: Religion			0.0442 (0.0529)
Threat Level: Medium x Category: Religion			0.0286 (0.0582)
Threat Level: High x Category: Religion			0.0797 (0.0636)
Threat Level: Low x Category: Research and Public Policy			-0.0194 (0.0435)
Threat Level: Medium x Category: Research and Public Policy			0.0017 (0.0444)
Threat Level: High x Category: Research and Public Policy			-0.0289 (0.0503)
Fixed-Effects:			
County	Yes	No	No
Month-Year	Yes	No	No
Individual	No	Yes	Yes
Date	No	Yes	Yes
S.E. Clustered by:	State	Individual & State	Individual & State

Table A3.4. Charity Navigator Regression Table. (Continued)

	County-Month Aggregation	**Individual-Level**	**Individual-Level**
Weights:	County Population	None	None
Observations	116,480	696,942	617,657
R^2	0.82735	0.83448	0.8542
Within R^2	0.00063	0.0000554	0.06446

Table A3.5. Charity Navigator County-Level Models, Median Household Income (MHI) Interactions. Table shows regression results using county-level CN data, which include interactions between threat level and county-level median household income (MHI). Threat level is categorical in the left two columns, and log-transformed in the right two. MHI was included either as a continuous variable (labeled "MHI"), or median split (1 = county MHI is greater than or equal to the U.S.-wide MHI, 0 = otherwise, labeled "Above Median MHI"). Interactions with log-transformed MHI obtained similar results. Significance codes: *** p < .001, ** p < .01, * p < .05, † p < .1.

	Continuous MHI	**Median Split MHI**	**Continuous MHI**	**Median Split MHI**
Threat Level: Low	0.1840† (0.0939)	0.2649*** (0.0619)		
Threat Level: Medium	-0.1063 (0.1641)	0.2221** (0.0799)		
Threat Level: High	0.2021 (0.1552)	0.3180*** (0.0578)		
Threat Level: Low x MHI	1.77e-6 (1.58e-6)			
Threat Level: Medium x MHI	5.72e-6* (2.18e-6)			
Threat Level: High x MHI	1.92e-6 (2.43e-6)			
Threat Level: Low x Above Median MHI		0.0847 (0.0762)		
Threat Level: Medium x Above Median MHI		0.1242 (0.0788)		
Threat Level: High x Above Median MHI		0.0239 (0.0714)		
Threat Level (logged)			-0.0803 (0.0667)	0.0134 (0.0261)
Threat Level (logged) x MHI			1.78e-6† (9.37e-7)	
Threat Level (logged) x Above Median MHI				0.0500 (0.0308)
Fixed-Effects:				
County	Yes	Yes	Yes	Yes
Month-Year	Yes	Yes	Yes	Yes
S.E. Clustered by:	State	State	State	State
Weights:	County Population	County Population	County Population	County Population
Observations	115,492	115,492	115,492	115,492
R^2	0.82653	0.82652	0.82642	0.82641
Within R^2	0.00082	0.00072	0.00016	0.00012

Table A3.6. Dictator-Game Regression Table. Table shows full regression results of the dictator-game model described in the text. The dependent variable was the allocation amount. Significance codes: *** p < .001, ** p < .01, * p < .05.

Threat Level: Low	0.2502* (0.0968)
Threat Level: Medium	0.3813** (0.1106)
Threat Level: High	0.2417* (0.0946)
Fixed-Effects:	
Individual	Yes
Wave	Yes
S.E. Clustered by:	Individual & State
Weights:	None
Observations	4,272
R^2	0.72942
Within R^2	0.00482

References

Adamczyk A. (2020) Americans are generous in their charitable giving during Covid-19, but how they donate changed. *CNBC*, https://www.cnbc.com/2020/09/15/how-individuals-in-the-us-helped-their-neighbors-through-covid-19.html.

Andreoni J (1990) Impure altruism and donations to public goods: A theory of warm-glow giving. *The economic journal* 100(401):464-477.

Arechar AA, Rand DG (2022) Learning to be selfish? A large-scale longitudinal analysis of Dictator games played on Amazon Mechanical Turk. *Journal of Economic Psychology* 90:102490.

Bauer M, Blattman C, Chytilová J, Henrich J, Miguel E, Mitts T (2016) Can War Foster Cooperation? *Journal of Economic Perspectives* 30(3):249-274.

Beer J. (2020) How much will charitable giving decline? A new survey provides us with a starting point. *Philanthropy Daily*, https://www.philanthropydaily.com/how-much-will-charitable-giving-decline/.

Bowles S, Gintis H (2003) *The Origins of Human Cooperation. The Genetic and Cultural Evolution of Cooperation, P. Hammerstein ed* (MIT Press, Cambridge MA).

Calvillo DP, Ross BJ, Garcia RJ, Smelter TJ, Rutchick AM (2020) Political ideology predicts perceptions of the threat of covid-19 (and susceptibility to fake news about it). *Social Psychological and Personality Science* 11(8):1119-1128.

Charness G, Gneezy U (2008) What's in a name? Anonymity and social distance in dictator and ultimatum games. *Journal of Economic Behavior & Organization* 68(1):29-35.

Charness G, Gneezy U, Kuhn MA (2013) Experimental methods: Extra-laboratory experiments-extending the reach of experimental economics. *Journal of Economic Behavior & Organization* 91:93-100.

Chintagunta P, Labroo AA (2020) It's About Time: A Call for More Longitudinal Consumer Research Insights. *Journal of the Association for Consumer Research* 5(3):240-247.

Cryder CE, Loewenstein G, Seltman H (2013) Goal gradient in helping behavior. *Journal of Experimental Social Psychology* 49(6):1078-1083.

Cumulative Cases. *Johns Hopkins Coronavirus Resource Center*, https://coronavirus.jhu.edu/data/cumulative-cases.

Dawkins R (1976) *The Selfish Gene* (Oxford University Press).

Dong E, Du H, Gardner L (2020) An interactive web-based dashboard to track COVID-19 in real time. *The Lancet Infectious Diseases* 20(5):533-534.

Dovidio JF, Morris WN (1975) Effects of stress and commonality of fate on helping behavior. *Journal of Personality and Social Psychology* 31(1):145-149.

Dunn L, White K, Dahl DW (2020) A Little Piece of Me: When Mortality Reminders Lead to Giving to Others. *Journal of Consumer Research* 47(3):431-453.

Engel C (2011) Dictator games: a meta study. *Experimental Economics* 14(4):583-610.

Erikson KT (1976) Loss of communality at Buffalo Creek. *The American Journal of Psychiatry* 133(3):302-305.

Forsythe R, Horowitz JL, Savin NE, Sefton M (1994) Fairness in Simple Bargaining Experiments. *Games and Economic Behavior* 6(3):347-369.

Fridman A, Gershon R, Gneezy A (2021) COVID-19 and vaccine hesitancy: A longitudinal study. *PloS one* 16(4):e0250123.

Fritz CE (1961) *Disaster* (Institute for Defense Analyses, Weapons Systems Evaluation Division).

Garbe L, Rau R, Toppe T (2020) Influence of perceived threat of Covid-19 and HEXACO personality traits on toilet paper stockpiling. *PLOS ONE* 15(6):e0234232.

Gilligan MJ, Pasquale BJ, Samii C (2014) Civil War and Social Cohesion: Lab-in-the-Field Evidence from Nepal. *American Journal of Political Science* 58(3):604-619.

Glynn SA, Busch MP, Schreiber GB, Murphy EL, Wright DJ, Tu Y, Kleinman SH, Group ftNRS (2003) Effect of a National Disaster on Blood Supply and Safety: The September 11 Experience. *JAMA* 289(17):2246-2253.

Gneezy A, Fessler DMT (2012) Conflict, sticks and carrots: war increases prosocial punishments and rewards. *Proceedings of the Royal Society B: Biological Sciences* 279(1727):219-223.

Gneezy U, Imas A (2017) Lab in the field: Measuring preferences in the wild. *Handbook of economic field experiments*, vol. 1 (Elsevier), 439-464.

Gneezy U, Keenan EA, Gneezy A (2014) Avoiding overhead aversion in charity. *Science* 346(6209):632-635.

Goodman JK, Cryder CE, Cheema A (2013) Data collection in a flat world: The strengths and weaknesses of Mechanical Turk samples. *Journal of Behavioral Decision Making* 26(3):213-224.

Henrich JP, Boyd R, Bowles S, Fehr E, Camerer C, Gintis H (2004) *Foundations of human sociality: Economic experiments and ethnographic evidence from fifteen small-scale societies* (Oxford University Press on Demand).

Hsee CK, Rottenstreich Y (2004) Music, pandas, and muggers: on the affective psychology of value. *Journal of Experimental Psychology: General* 133(1):23.

HUD USPS ZIP Code Crosswalk Files | HUD USER.

Jonas E, Schimel J, Greenberg J, Pyszczynski T (2002) The Scrooge effect: Evidence that mortality salience increases prosocial attitudes and behavior. *Personality and Social Psychology Bulletin* 28(10):1342-1353.

Jones JM. (2020) U.S. Personal Finances: Future More Concerning Than Present. *Gallup.com*, https://news.gallup.com/poll/308936/personal-finances-future-concerning-present.aspx.

Kahneman D, Knetsch JL, Thaler R (1986) Fairness as a constraint on profit seeking: Entitlements in the market. *The American economic review*:728-741.

Kessler JB, Vesterlund L (2015) The external validity of laboratory experiments: Qualitative rather than quantitative effects.

Levitt SD, List JA (2007) What do laboratory experiments measuring social preferences reveal about the real world? *Journal of Economic perspectives* 21(2):153-174.

List JA (2007) On the interpretation of giving in dictator games. *Journal of Political economy* 115(3):482-493.

Loewenstein G, Small DA (2007) The scarecrow and the tin man: The vicissitudes of human sympathy and caring. *Review of general psychology* 11(2):112-126.

McNeil DG, Jr. (2020) Mask Hoarders May Raise Risk of a Coronavirus Outbreak in the U.S. *The New York Times* (2020-01-29), Health.

Mesch D, Osili U, Skidmore T, Bergdoll J, Ackerman J, Sager J (2020) COVID-19, Generosity, and Gender: How Giving Changed During the Early Months of a Global Pandemic.

Nurmi J, Räsänen P, Oksanen A (2012) The norm of solidarity: Experiencing negative aspects of community life after a school shooting tragedy. *Journal of Social Work* 12(3):300-319.

Pagliaro S, Sacchi S, Pacilli MG, Brambilla M, Lionetti F, Bettache K, Bianchi M, Biella M, Bonnot V, Boza M (2021) Trust predicts COVID-19 prescribed and discretionary behavioral intentions in 23 countries. *PloS one* 16(3):e0248334.

Park SQ, Kahnt T, Dogan A, Strang S, Fehr E, Tobler PN (2017) A neural link between generosity and happiness. *Nature Communications* 8(1):1-10.

Poulin MJ, Holman EA, Buffone A (2012) The Neurogenetics of Nice: Receptor Genes for Oxytocin and Vasopressin Interact With Threat to Predict Prosocial Behavior. *Psychological Science* 23(5):446-452.

Prediger S, Vollan B, Herrmann B (2014) Resource scarcity and antisocial behavior. *Journal of Public Economics* 119:1-9.

Rachlin H (2002) Altruism and selfishness. *The Behavioral and brain sciences* 25:239-50; discussion 251.

Rodrigues SM, LeDoux JE, Sapolsky RM (2009) The influence of stress hormones on fear circuitry. *Annual review of neuroscience* 32:289-313.

Roux C, Goldsmith K, Bonezzi A (2015) On the Psychology of Scarcity: When Reminders of Resource Scarcity Promote Selfish (and Generous) Behavior. *Journal of Consumer Research* 42(4):615-631.

Sharma E, Mazar N, Alter AL, Ariely D (2014) Financial deprivation selectively shifts moral standards and compromises moral decisions. *Organizational Behavior and Human Decision Processes* 123(2):90-100.

Slovic P, Weber EU (2013) Perception of risk posed by extreme events. *Regulation of Toxic Substances and Hazardous Waste (2nd edition)(Applegate, Laitos, Gabba, and Sachs, Editors), Foundation Press, November 2010*.

Testard C, Larson SM, Watowich MM, Kaplinsky CH, Bernau A, Faulder M, Marshall HH, Lehmann J, Ruiz-Lambides A, Higham JP (2021) Rhesus macaques build new social connections after a natural disaster. *Current biology*.

Turkel G (2002) Sudden Solidarity and the Rush to Normalization: Toward an Alternative Approach. *Sociological Focus* 35(1):73-79.

U.S. Attorney urges public to report potential hoarding of supplies needed to fight COVID-19. (2020). The United States Department of Justice, https://www.justice.gov/usao-mt/pr/us-attorney-urges-public-report-potential-hoarding-supplies-needed-fight-covid-19.

U.S. Census Bureau QuickFacts: United States. (2019).

Urminsky O, Kivetz R (2011) Scope insensitivity and the "mere token" effect. *Journal of Marketing Research* 48(2):282-295.

USDA ERS - County-level Data Sets. (2021), https://www.ers.usda.gov/data-products/county-level-data-sets/.

van Holm E, Monaghan J, Shahar DC, Messina J, Surprenant C (2020) The impact of political ideology on concern and behavior during COVID-19. *Available at SSRN 3573224*.

Vardy T, Atkinson QD (2019) Property Damage and Exposure to Other People in Distress Differentially Predict Prosocial Behavior After a Natural Disaster. *Psychological Science* 30(4):563-575.

von Dawans B, Fischbacher U, Kirschbaum C, Fehr E, Heinrichs M (2012) The Social Dimension of Stress Reactivity: Acute Stress Increases Prosocial Behavior in Humans. *Psychological Science* 23(6):651-660.

Weber EU, Johnson EJ (2009) Mindful judgment and decision making. *Annual review of psychology* 60:53-85.

Westbrook A (2020) Opinion | People Around the World Are Panic-Buying … Toilet Paper? *The New York Times* (2020-03-12), Opinion.

Wu SL, Mertens AN, Crider YS, Nguyen A, Pokpongkiat NN, Djajadi S, Seth A, Hsiang MS, Colford JM, Reingold A, Arnold BF, Hubbard A, Benjamin-Chung J (2020) Substantial underestimation of SARS-CoV-2 infection in the United States. *Nature Communications* 11(1):4507.

Zaki J (2020) Catastrophe Compassion: Understanding and Extending Prosociality Under Crisis. *Trends in Cognitive Sciences* 24(8):587-589.