

Network Contexts and Social Identities Interact to Shape Beliefs and Behaviors

by

Joseph Michael Quinn

Department of Sociology  
Duke University

Date: \_\_\_\_\_

Approved

\_\_\_\_\_  
James W. Moody, Co-Supervisor

\_\_\_\_\_  
Lynn Smith-Lovin, Co-Supervisor

\_\_\_\_\_  
Ashley Harrell

\_\_\_\_\_  
Craig Rawlings

\_\_\_\_\_  
Sarah Gaither

Dissertation submitted in partial fulfillment of  
the requirements for the degree of Doctor  
of Philosophy in the Department of  
Sociology in the Graduate School  
of Duke University

2022

ABSTRACT

Network Contexts and Social Identities Interact to Shape Beliefs and Behaviors

by

Joseph Michael Quinn

Department of Sociology  
Duke University

Date: \_\_\_\_\_

Approved

\_\_\_\_\_  
James W. Moody, Co-Supervisor

\_\_\_\_\_  
Lynn Smith-Lovin, Co-Supervisor

\_\_\_\_\_  
Ashley Harrell

\_\_\_\_\_  
Craig Rawlings

\_\_\_\_\_  
Sarah Gaither

An abstract of a dissertation submitted in partial  
fulfillment of the requirements for the degree  
of Doctor of Philosophy in the Department of  
Sociology in the Graduate School of  
Duke University

2022

Copyright by  
Joseph Michael Quinn  
2022

## **Abstract**

This dissertation investigates the reciprocal relationship between micro-level beliefs and behaviors involving identity categories and macro-level features of social structure. Chapters 2 and 3 explore how social psychological processes intersect with persisting network exchange structures or environmental shifts to shape the beliefs or behaviors of embedded actors. Chapter 2 uses original survey data to investigate how beliefs about occupations shifted shortly after the Covid-19 pandemic began, and finds these belief changes persist into the following year for occupations made salient as the pandemic began.

Chapter 3 presents a novel experiment that assigned participants to exchange networks with different structures and identity compositions. The findings show that (a) persisting network arrangements effect pro-social behaviors in a similar way regardless of whether the networks contain actors with homogeneous or heterogeneous social identities, and that (b) interacting with dissimilar others over an extended period of time increases an actor's trust behavior toward unmet members of the out-group identity.

Chapter 4 extends insights from the previous chapters. The results of an agent-based computational experiment show that initial network arrangements can enable transitive tie formation between dissimilar others – shaping the macrostructure of the

network and emergence of homophily, not merely the beliefs and behaviors of the actors within them.

## **Dedication**

To the students of McKinley Technology High School. And to Rue.

# Contents

Abstract .....	iv
Dedication .....	vi
List of Tables .....	x
List of Figures .....	xi
Acknowledgements .....	xii
1. Introduction .....	1
2. Belief Stability and Change in Unsettled Times: Dynamics in Cultural Meanings about Occupations in a Pandemic .....	10
2.1 Introduction.....	10
2.2 Theoretical Background.....	13
2.2.1 Culture as Both Stable and Dynamic.....	13
2.2.2 Measuring Cultural Meanings: E, P, and A.....	14
2.2.3 Using Social Upheaval to Detect Changes in Personal Culture.....	16
2.3 Data and Methods .....	18
2.3.1 Participants and Design.....	18
2.3.2 Choosing and Classifying Occupations for Wave 2 and 3 .....	20
2.3.3 Cultural Meaning Measures .....	22
2.3.4 Analysis .....	22
2.4 Results .....	23
2.5 Conclusions and Discussion .....	35

3. How Identities and Exchange Structures Affect Present Cooperation and Future Behavior toward People from Different Groups .....	38
3.1. Introduction.....	38
3.2. Theoretical Background.....	40
3.2.1 Social Identity and Cooperation.....	40
3.2.2 Ongoing Cooperation and Emotion in Exchange Networks .....	42
3.2.3 Future Commitment Behavior as Trust and Trustworthiness .....	46
3.2.4 Bridging Identities and Exchange Structure.....	48
3.3. Data and Methods .....	52
3.3.1 Participants and Design.....	52
3.3.2 Procedure.....	55
3.3.3 The Exchange Structure Manipulation and Trust Game Outcome.....	57
3.3.3 The Identity Information Manipulation.....	62
3.3.4 Analytic Strategy .....	65
3.4. Results .....	67
3.4.1 Overall Giving Patterns in Exchange Networks .....	67
3.4.1 Affective Attachment.....	71
3.4.2 Future Trust and Trustworthiness by Social Identity .....	74
3.5. Discussion.....	76
4. Manipulating Initial Network Characteristics and Measuring Changes in the Emergence of Homophily: a Computational Experiment.....	79
4.1. Introduction.....	79
4.2. Background .....	80

4.3. Methods .....	87
4.3.1 The Organizational Setting .....	87
4.3.2 The Agents.....	89
4.3.2 Decisions Linking Agents to the Organization .....	90
4.4. Results .....	92
4.5. Discussion.....	96
5. Conclusion .....	100
References .....	105
Biography.....	117

## List of Tables

Table 1: Weighted Two-Sample t-Tests for Occupations in the Analysis Sample .....	25
Table 2: Linear Mixed Models Predicting EPA Mean Occupation Ratings .....	32
Table 3: Generalized linear models predicting giving to the focal beneficiary .....	70
Table 4: Multilevel linear models predicting affective attachment following a phase of the exchange task .....	73
Table 5: Linear Mixed Models Predicting Trust and Trustworthiness following Repeated Exchange Phases .....	76
Table 6: Parameterization of the Agents and Social Space .....	89

## List of Figures

Figure 1. Differences in Mean EPA Ratings in Survey Waves 2 and 3 vs. Wave.....	28
Figure 2. Evaluation and Potency Score Interactions.....	35
Figure 3: Key study manipulations.....	54
Figure 4: Giving to the focal beneficiary over time, by the key study manipulations.....	68
Figure 5. Initial Network Configurations for the Simulation Experiment.....	88
Figure 6. In-Group Bias (left) and Transitivity Bias (right) Plotted Against Change In Mean Assortativity in Simulated Networks.....	93
Figure 7. Transitivity Bias Levels Plotted Against Change in Mean Assortativity in Simulated Networks, Where In-Group Bias Levels = 0.9.....	94
Figure 8. Difference in Mean Assortativity Change Across In-Group and Transitivity Bias Conditions, for Networks With and Without Diverse Cliques.....	95

## **Acknowledgements**

To my committee for providing invaluable support and feedback on my dissertation, to Jim, Lynn, and Ashley for their mentorship during my time as a graduate student at Duke, and to the people who helped me refine the ideas in this document – especially Susan Jacobs, Bernard Coles, Liann Tucker, Tom Wolff, Em Maloney, Nico Restrepo, Kelly Quinn, and Sarah Haight. Thank you very much.

# 1. Introduction

The studies in this dissertation investigate how macro-level features of social structure may shape, and be shaped by, individual-level beliefs and behaviors involving social identities. Each study tests or extends ideas in social psychology, culture and cognition, and network theory to assess three interrelated lines of inquiry: how a person's beliefs about social identity categories can emerge or change when they are placed in particular cultural contexts or network structures for an extended period of time; whether changes to these belief changes impact the way people behave toward people from dissimilar social category groups outside of the contexts that caused belief changes; and the extent to which behaviors toward and beliefs about others may in turn influence the structure of social systems.

My dissertation chapters contribute to central topics in sociological research on the plausibility of changes to durable cultural beliefs, the problem of social order in diverse social systems, and the ways that social structure may reciprocally reinforce and potentially change individual-level behaviors and beliefs. At the same time, these basic questions are important for understanding the key determinants of (and opportunities for intervention) of patterns that are often seen as social problems in modern society. What people believe about a social category (such as an occupation, a racial group, a gender identity, or a political party affiliation) affects how they interact with the people who they see as members of different categorical groups, whether they choose to be

cooperative or generous when doing so, and the likelihood that they will form lasting ties with each other. These beliefs and behaviors, and the structural contexts that facilitate or inhibit their enactment, are the micro-level building blocks of contemporary social issues that emerge as macro-level patterns of racial segregation, gender inequality, and political polarization. As such, it is important to develop evidence about mechanisms that influence beliefs and behaviors enacted through social categorization in order for programs, policies, and social interventions that aspire to impact matters of cohesion and interaction between dissimilar groups of people can realize their desired outcomes.

My first study (Chapter 2) asks whether cultural beliefs about occupations changed during the Covid-19 pandemic – a time of major dislocation in American social and economic life. In particular, I evaluate whether dynamics in meaning are different for identities classified as “essential occupations” during the pandemic, and identities that were significantly more salient in cultural discourse (based on Google Trends data) in the first three months after the pandemic began in the U.S. I administered surveys to collect quantitative data on cultural beliefs drawn from Affect Control Theory (ACT) – evaluation, or goodness; potency, or powerfulness; and activity, or liveliness – about occupations measured shortly before, a few months after, then nearly one year after the outbreak of Covid-19 in the U.S.

Treating the pandemic as a persistent and transformative shock that displaced many people from their typical relationship to the institution of work, I use a linear mixed modeling strategy to estimate meaning changes across all three waves, and to explore whether meaning changes were different in magnitude or persistence for essential or salient occupations compared to other occupational identities. I find that evaluation and potency ratings of occupations decline at the beginning of the pandemic, but rebound over the longer term – a pattern that fits a “control model” of stable cultural meaning. At the same time, I find a different pattern for occupations that became more salient in the first three months after the pandemic began (based on Google Trends data) – a pattern indicative of an “updating model” of dynamic cultural meaning. These occupations were broadly protected from declines in evaluation that occurred shortly after the pandemic began. Longer-term differences for salient occupations persisted in both evaluation and potency meanings in the third survey wave, almost a year after the outbreak.

The results of this study augment contemporary discourse in cultural sociology about cultural beliefs and their stability in adult populations. While recent work emphasizes how uncommon it is to observe changes in the beliefs of adult individuals, my study of the Covid-19 pandemic as an unsettled time suggests that, to truly understand belief change, researchers should collect data around specific environmental

changes that make systematic retooling about salient cultural objects cognitively plausible.

My second study (Chapter 3) builds on this idea, looking to repeated network exchange structures as persistent contexts where people may develop different orientations toward interaction partners from different social identity groups. This study investigates how both the structure of a repeated exchange network and information about the identity group membership of their interaction partners (two well-known determinants of pro-social behavior) combine to shape cooperation and emotions toward others – both their ongoing social ties within their repeated exchange networks, and future unmet actors with whom they play a one-shot trust game. The study asks two fundamental questions: first, how do exchange structures and identity information combine to produce different levels of cooperation in repeated exchange? Do social structures that make individual contributions indistinguishable and facilitate a sense of shared responsibility produce higher levels of cooperation even when identity information allows for decisions to be patterned by in-group bias, or do certain identity information conditions undermine the “jointness” predictions linking exchange structure to cooperation frequency in social psychology? Second, do people in diverse repeated exchange networks behave more pro-socially toward out-group interaction partners outside of their network – or do people learn to find people from different social identity groups more trustworthy?

Leveraging a theoretical tradition known as the Affect Theory of Social Exchange (ATSE), the study in Chapter 3 involves a factorial experiment to assess the effect of different exchange structure features on cooperative and affective commitments that emerge in triads containing people with different social identities. Participants in the experiment were randomly assigned to salient social identity categories through a minimal group task. To manipulate the structural features of exchange networks, I assigned participants to play through three phases of repeated exchange in either a productive, reciprocal, or generalized network exchange structure. To manipulate each participant's exposure to different salient social identities, they were randomly assigned to a network that contained interaction partners with either homogeneous, mixed, or unknown social identities. Each of the three phases would then conclude with a one-shot trust game, where participants were paired with an unmet other of their the same or opposite identity group, and sent or returned points as measures of trust and trustworthiness.

I find that, when embedded in repeated exchange networks, people cooperate differently based on both the type of repeated exchange network structure in which they are embedded *and* the identity information about their interaction partners. Within each type of exchange structure, significant differences exist in the emergence of cooperation between actors embedded in homogeneous social identity and diverse social identity triads. These results also replicate prior scholarship in ATSE: the theory's classic

predictions about differences in the degree of cooperation between each network exchange structure type are robust to all social identity information conditions, and the inclusion of identity information does not moderate or suppress this powerful structural effect.

Results from the trust games that followed each phase or repeated exchange in the study show that people behave more pro-socially toward unmet out-group others after encountering people from their out-group in their previous repeated exchange network. While actors trust (i.e., send significantly more monetary units) unmet in-group others marginally more than out-group others, actors find unmet others who do not share their social identity more trustworthy (i.e., return significantly more monetary units) immediately after participating in a phase of repeated exchange with diverse identity composition, compared to participating in a prior homogeneous-identity phase. Moreover, these participants find unmet dissimilar alters significantly more trustworthy than unmet same-identity alters. The results show that the network structure of social exchange affects levels of cooperation and affective attachment even in the presence of diverse social categories, while shared social categories themselves directly and positively affect the degree of cooperation in repeated exchange networks. The results also show that prior exposure to repeated exchange in diverse groups can have a significant positive impact on the trustworthiness behavior between people from different social categories. In other words, a persistent structural context that exposes

people to people from a dissimilar social identity category can have a positive impact on cooperation and trustworthiness toward members of a relative out-group more broadly – even members of the outgroup that a person has not yet met.

The study in Chapter 4 builds on the first two studies with a more applied question: how might an organization manipulate its initial structure to minimize segregation along social identity categories? Prior research indicates in-group preferences and the principle of “triadic closure” – people preferring to become friends with the friends of their friends – intersecting with features of already-segregated networks in the (re)production of homophily in social systems. Top-down empirical studies of network ecology show that homophily in organizations is lower when the environment is structured in ways that encourage people to interact with dissimilar others, but these studies are constrained by limitations: without experimental manipulation, observed organizational networks are segregated from the outset, and both in-group preferences and closure preferences are estimated, not manipulated.

As such, the final study in my dissertation models the emergence of homophily in simulated networks of hypothetical organizations. Doing so allows me to computationally investigate about how homophily may emerge differently when (a) networks are not yet patterned by segregation along social categories in their initial state, and (b) when networks contain connectivity between agents from different social categories from the outset. By manipulating initial network structure in these two ways,

leaders of organizations would decrease the number of opportunities to form ties that satisfy an agent's closure preference while also satisfying their in-group preference, and increase the number of opportunities to form ties that satisfy an agent's closure preference that would result in tie formation with another agent from a different social identity group. In other words, preventing inheritance of segregated ties and creating explicit diverse foci should attenuate the emergence of segregation in the network of an organization.

While Chapters 2 and 3 explore the role of persistent environmental or network arrangements in shaping individual beliefs and behaviors toward different social identities, Chapter 4 varies individual behaviors of artificial agents (and their initial network arrangements) in a simulated network to investigate how preferences for similar others and the propensity for triadic and focal closure in networks can interact to produce different levels of homophilous sorting in organizational networks with different initial configurations. Using an agent-based model, I vary the in-group tie preferences and transitivity preferences of each agent, and manipulate the initial configuration of the network to either include or exclude explicit foci in the form of diverse fully-connected cliques, around which transitivity preferences can lead agents to be at greater risk for out-group tie formation. Preliminary results suggest that networks with focal diverse cliques yield less substantial increases in emergent homophily over time, even in the presence of strong and static in-group preferences.

A concluding section of the dissertation positions these findings as consequential to scholars attempting to understand and manipulate psychological and structural mechanisms that contribute to stratification and inequality.

## **2. Belief Stability and Change in Unsettled Times: Dynamics in Cultural Meanings about Occupations in a Pandemic**

### ***2.1 Introduction***

Social and economic life in the U.S. underwent rapid and dramatic change in March of 2020. The federal government declared the international Covid-19 pandemic a national emergency, and many American states began issuing stay-at-home orders and mandating the closure of non-essential businesses (AJMC, 2020). Employees that could work from home began to do so; people who worked in non-essential occupations and who could not work remotely were often unable to work at all. By April, over 23 million Americans were left unemployed as the unemployment rate increased by 10.3 percent in a single month to 14.7 percent, the highest rate and largest over-the-month increase since the Bureau of Labor Statistics began keeping records in 1948 (BLS, 2020). The precarity of work in form, function, and availability intertwined with a national dialogue and media coverage about the necessity of essential medical workers, as well as select blue-collar and service industry occupations.

Did this massive dislocation in the U.S. work force bring short- and long-term changes to the cultural meanings of occupational identities? A large body of empirical evidence tends to show that cultural meanings are usually very stable over time (see Heise, 2010 for a review), and that what cultural change does occur in society is gradual and driven mainly by population replacement, not individuals systematically retooling their beliefs (e.g., Kiley and Vaisey, 2020). Yet some theoretical accounts from cultural

sociology and anthropology – whose purveyors and critics alike clamor for more empirical clarification – argue that extended periods of social upheaval or substantial, persisting alteration of the material conditions of an environment may induce systematic cultural change within individuals embedded in such an unsettled context (DiMaggio, 1997; Swidler, 1986; Lizardo, 2021).

While such “unsettled times” are often ambiguously contextualized in cultural analysis (Quinn et al. 2018; Strauss 2018), the outbreak of the Covid-19 pandemic represents an unambiguous and remarkably large shift in Americans’ relationship to the work institution. Virtually all members of the culture were unable to work as they used to, more concerned about job stability and unemployment, and bombarded with a national discourse about labor considered vital to the most basic needs of citizens. Relationships to the work institution writ large became persistently salient. For some occupational identities, this persistent salience was not only extraordinary in magnitude but also accompanied by cultural narratives of unprecedented valorization, financial concern, or general appreciation. Grocery store workers were braving the “front lines” and risking exposure to keep the country from going hungry; hair-dressers were unable to go to work and in need of economic relief, and their former patrons realized that the value of a good haircut could not be taken for granted.

The outbreak of Covid-19 in the U.S. rapidly induced a set of pervasive cultural narratives about the precarity of work in general, and the newfound importance (or particular vulnerability) of certain jobs. This study shows how formal measures of

meaning for a set of essential and non-essential occupations changed during the social upheaval surrounding the Covid-19 pandemic, and explores whether differences in meaning shortly after the outbreak persisted into the following calendar year. Using survey strategy for quantifying cultural sentiments from a well-established theoretical and measurement framework known as Affect Control Theory (ACT), I estimate mean differences in the ways that members of the U.S. language culture score occupational identities on three psychologically important (Scholl, 2013) and conventionally stable (Heise, 2010) dimensions of affective meaning: ratings of evaluation (good to bad), potency (powerful to powerless), and activity (lively to inactive). I collected these “EPA ratings” data for 650 occupational identities in the fall and winter of 2019-2020, before the pandemic struck the U.S., and used quota sampling to match the U.S. population on sex, age, and race/ethnicity, and education. I then collected data for a subset of these identities – 41 occupations officially designated as essential by state governments and policy research centers, and 43<sup>1</sup> non-essential comparison occupations – at two later points in time after the outbreak began to prompt lockdown orders and novel employment concerns: first, in June and July of 2020; next, in January and February of 2021.

Using these data, I (1) show how EPA ratings for occupational identities changed in the short term, (2) assess whether overall EPA changes persist in the longer-term

---

<sup>1</sup> Analyses from a previous study using the same Wave 1 and Wave 2 data (Quinn et al., 2022) include 85 identities. The Wave 3 survey included the same 85 identities, but exploratory analyses led to the identification and removal of ratings for “telemarketer” as an influential outlier.

survey or revert to pre-pandemic levels, and (3) explore if and how ratings for essential occupations and occupations that achieved novel levels of salience immediately after the pandemic changed differently compared to non-essential and less salient occupational identities.

## ***2.2 Theoretical Background***

### **2.2.1 Culture as Both Stable and Dynamic**

Culture is often seen as a set of features or interpretive lenses shared by members of a group or society, and it tends to change slowly (Goodenough, 1961; Romney, 1991; Romney et al., 1987). Language, religious beliefs, types of cuisine, or behavioral tendencies are often considered part of a group's culture if they distinguish that group from some other groups and are shared by members of the group.

While cultures and the beliefs they carry change over time, population-wide changes are rarely considered the result of people changing their views. Ryder (1965) famously proposed that cohort replacement can serve as a key mechanism for society-level changes in attitudes, beliefs, and behavioral patterns. This perspective has found support in many empirical studies where researchers using longitudinal data show that features of culture can be so stable within individuals that societally held beliefs change primarily when older members die and younger cohorts with new socialization experiences replace them (Davis, 1975; Kiley & Vaisey, 2020; Schwadel, 2011).

Yet some theoretical work suggests that population-level cultural belief change may also occur because individuals in society change their beliefs in a collective way.

Proposed mechanisms for individually-driven systematic cultural change tend to involve substantial shifts in a community's environment or network structure (Swidler, 1986; Lizardo & Strand, 2011; Centola et al., 2018). While students of culture and cognition often debate whether actors process major environmental changes by either fundamentally retooling their beliefs or making sense of them in a more discursive way (Lizardo & Strand, 2011), they agree that substantial environmental change can cause poor fit between one's existing cultural beliefs and the conditions of their new environment.

Hypotheses about transient or durable forms of cultural change are difficult to test. Environmental changes that may induce either form of cultural change must be substantial and punctuated enough to lead to noticeable friction between a set of salient cultural beliefs held by actors within a culture and the realities of their new environmental conditions, and last for long enough that these moments of friction cannot be explained away as anomalous. The Covid-19 pandemic represents a substantial, punctuated, and persistent environmental change. As such, the resulting social dislocation represents ideal circumstances to assess the plausibility of cultural change driven by members of a culture collectively resituating some of their beliefs.

## **2.2.2 Measuring Cultural Meanings: E, P, and A**

There are many ways to measure shared cultural meanings. I use a well-known system developed by Osgood and colleagues (1957), which uses three discrete dimensions to measure the fundamental affective responses that members of a culture

share about actors and behaviors: evaluation (good-bad), potency (powerful-powerless) and activity (lively-quiet) (EPA). Osgood (1962, 1964) established that these three dimensions could be used to characterize a large array of concepts (from inanimate objects to social positions) and were useful in a wide range of cultures. Heise (1979, 2007, 2010) used these dimensions to develop a geometric space within which cultural meanings for identities, behaviors, emotions, personal characteristics, and social settings could all be simultaneously located. The positions of concepts within EPA space reflect fundamental sentiments about those concepts in a given language culture, which people act to maintain through social interaction.

Heise (2010) documented that EPA sentiments are both highly stable and widely shared within national language cultures. Much of the within-culture variation in EPA ratings of individual concepts is the result of measurement error or unique individual experience (e.g., rating the concept of “mother” immediately after having a negative conversation with your mother on the phone might be different from our general cultural understanding of what mothers should be and are usually). Recently, Freeland and Hoey (2018) demonstrated that shared fundamental EPA sentiments about occupational identities were a good indicator of the stable, widely shared status structure of occupations within the United States. The scale that they created using EPA profiles of occupations and an ACT-derived understanding of occupational status performed better than more traditional occupational prestige measures in predicting

traditional survey measures of job satisfaction, happiness, and feelings of being respected at work.

The ACT tradition views fundamental EPA sentiments as shared and stable reference signals that society maintains because people recreate their culture through interaction, but some researchers highlight that EPA values can indeed change over time. MacKinnon and Luke (2002) provide an excellent example of EPA measurements as an indicator of social change. While they found very high stability of EPA means in Canadian culture in surveys at two points in time, they also identified significant change in views about homosexuals, religious identities, and some political identities between 1981 and 1995. These findings are noteworthy because the sentiment changes coincided with social movements related to the identities for which views changed over a relatively short period of time. Though the study analyzed independent samples from two points in time, the findings suggest that the EPA sentiment differences could represent changes not fully explained by cohort effects. Similarly, Schneider and Schröder (2012) found that the meaning of manager had changed in the U.S. and Germany at the same time as macro-level changes in business relations and leadership styles. These studies echo Swidler's (1986) template for the different ways one might expect cultural meanings to evolve in "unsettled times."

### **2.2.3 Using Social Upheaval to Detect Changes in Personal Culture**

I use the tremendous changes that occurred in the occupational system in early 2020 to address the questions: do EPA values about occupations change in a collective,

consensual way when a natural disaster brings persistent material and structural changes to social life, and is there evidence that any changes persist for salient occupational identities as the prominence of these material and structural changes begins to subside a year later?

Beginning in mid-March 2020, much non-essential economic activity in the U.S. was stopped or dramatically curtailed as a result of the spread of Covid-19. Unemployment rose sharply as businesses and other activities that required face-to-face interaction closed to prevent viral spread. Medical facilities were filled to capacity as hospitals and associated services attempted to deal with the infected population. State and local governments declared some functions to be essential and mandated their continued operation, while others were deemed non-essential and ordered to shut down for varying lengths of time. Supply chains and delivery services struggled to help people stay fed and supplied with needed goods. First responders and other government service providers tried to continue their essential services while protecting themselves from infection. During the first several months of the pandemic – the timing of the Wave 1 survey – there was pervasive media coverage and social discourse about some work identities. Some occupations classified as essential work were embedded in dominant cultural narratives as “heroes” whose efforts provided citizens with food and medical care. Workers with other occupational identities were prominently discussed as “victims” whose services, perhaps formerly taken for granted, were rendered saliently impossible.

The pandemic continued to shape American social and economic life in the beginning of 2021 – the timing of the Wave 2 survey – but the economy had begun to recover, two vaccines were in early stages of a national distribution effort, and initial unemployment claims started to decline. While still salient, the decline in salience can help us tease apart whether rebounds began to occur or if changes observed at the beginning of the pandemic are reflected in the form of durably new meanings almost a year later.

As such, I examine the events surrounding the Covid-19 pandemic as a potential case of punctuated change. My paper asks three key questions: first, were the massive changes in work engagement and occupational activities surrounding the Covid-19 pandemic associated with changes in the affective meanings of work identities within the first few months of pandemic-induced social change? Second, what do mean EPA values look like almost a year after the beginning of this major social and economic dislocation compared to their baseline expected values? And third, did ratings for occupations that achieved novel levels of salience immediately after the pandemic exhibit a different pattern than ratings of less salient occupational identities?

## ***2.3 Data and Methods***

### **2.3.1 Participants and Design**

To answer these questions, I analyze the EPA meanings associated with 84 essential and non-essential occupational identities before and after the Covid-19 outbreak. Data for the study were collected by my study team through a set of three

online surveys programmed and run on the Qualtrics XM platform. Wave 1 pre-pandemic data were collected between May of 2019 and early March of 2020, as part of an ACT sentiment study about occupational prestige and social deference. The pre-pandemic survey included EPA ratings for 650 occupations, including all 570 occupations listed in the 2010 US Census.

After the Covid-19 outbreak and the enactment of lockdowns, mask orders, and social distancing policies, the study team replicated the previous survey measurement strategy for the 84 occupational identities from the Wave 1 survey that appear in the analysis sample of this paper (see Table 2 for a list of stimuli) at two different points in time. Wave 2, the proximal post-outbreak survey, was fielded from June 9 to July 16 of 2020. Wave 3, the distal post-outbreak survey, was fielded from January 8 to February 15 of 2020.

Acting as cultural informants, 3,614 unique respondents rated the 84 occupations analyzed in this paper. To mitigate respondent fatigue, each participant was assigned to a module that included questions on a subset of identities (between 31 and 59 identities per module<sup>3</sup>). This study includes responses from 16 modules in total – 12 from the pre-pandemic survey, two from the Wave 2 survey, and two from the Wave 3 survey. 31 occupations in the “core module” of our Wave 1 survey were rated by 1,131 respondents<sup>4</sup>, while all other Wave 1 modules included 99 to 132 valid responses. The Wave 2 modules were rated by 191 to 192 participants each, and the Wave 3 modules were rated by 296 participants each. In all cases, the number of raters per stimulus is

above the number usually used in ACT sentiment studies to achieve reliabilities over .90 for estimated means (Heise 2010).

Respondents to all three surveys were U.S. citizens drawn from Qualtrics panels. A quota sampling method was used to approximately match the sample of respondents who completed each survey module to the population proportions reported by the 2010 Census in its categorical measures for age, race and ethnicity, and gender. The marginal distributions from the Census for the matched variables, as well as level of education, were then used to construct sample weights via iterative proportional fitting for each of the surveys. Respondents who failed attention checks were not included in the analyses or Ns reported above.

### **2.3.2 Choosing and Classifying Occupations for Wave 2 and 3**

Persistent salience of a social object like an identity is a key ingredient of cognitively plausible cultural belief change. To develop the stimulus list for the Wave 2 survey that sampled heavily on occupational identities made salient during the transition into the pandemic, the study team compiled a list of essential occupations identified in two sources: state COVID-19 executive orders and the Center for Economic and Policy Research's demographic study of workers in frontline industries (Rho et al., 2020). Because these sources often listed industries rather than specific occupations, we built a list of Census occupations within these industries, and flagged all occupational identities in the Wave 1 survey that matched this list of essential workers.

All five authors of Quinn et al. (2022) then used a blind voting system to select candidates from the 650 occupations in the pre-pandemic survey who were considered essential and non-essential, to ensure that occupations included in our study were both officially identified as essential and would likely be seen as important to core societal function by the general public. Occupations that received three or more votes as “essential” in addition to being flagged as essential occupations in the previous step became the essential occupations in our analysis sample. Non-essential occupations included all identities for which none of the five authors voted in favor of classification as an “essential” occupational identity. The Wave 2 and Wave 3 surveys included ratings of these occupations.

To obtain a more direct measure of cultural salience for each occupational identity, I obtained weekly U.S.-wide popularity rankings of each stimulus as a search term from Google Trends. Data from Google Trends has been successfully used as a topical salience or collective attention within a culture by researchers across the social sciences (Burrows and Savage 2014; Chykina and Crabtree 2018; Jun et al. 2018;). Occupations with mean weekly Google search popularity ratings that were 20 percentage points more salient in the 12 weeks that followed the outbreak of Covid-19 in the U.S., compared to mean popularity in the 12 weeks preceding the pandemic’s social dislocation, were then flagged as culturally salient. This grouping excludes essential occupations that were not central to public discourse (e.g., welders), and includes non-essential occupations that were prominently profiled in news media for being uniquely

impacted as victims of the pandemic (e.g., hairdressers). Some occupations lacked sufficient search term frequency to be identified as salient or not salient. These analyses are excluded from models that use this cultural salience indicator.

### **2.3.3 Cultural Meaning Measures**

For each of the 84 occupations, we calculate two means and variances: one for the ratings before the Covid-19 outbreak, and one for the post-outbreak data. Ratings were recorded as continuous positions (to the 10<sup>th</sup> decimal place) along the slider bar from 4.0 to -4.0 (with 0 in the center). Each observation was assigned a sample weight based on marginal distributions of age, gender, race/ethnicity, and education level from the 2010 Census through iterative proportional fitting in R. Weights were used when computing all mean EPA scores presented in our analyses to both account for minimal differences in the samples and to capture national cultural sentiments to the best of our ability.

### **2.3.4 Analysis**

The analyses presented in this paper proceed in two parts that build from the findings of a study comparing Wave 1 and Wave 2 data (Quinn et al. 2022). First, I conduct an array of two-sample unpaired t-tests between EPA ratings taken in Wave 1 and Wave 3, applying a Bonferroni correction to account for the large number (3x84) tests conducted, and use the adjusted result to determine whether there is compelling evidence that the pandemic influenced longer-term cultural meanings of occupations for any of the three dimensions. I plot these Wave 3 mean differences from Wave1 next to

Wave 2 differences from Wave 1 to describe broader trends in mean rating values over time.

In the second set of analyses, I use a linear mixed modeling strategy to estimate mean differences from pre-pandemic Wave 1 ratings to both the proximal Wave 2 survey ratings and the more distal Wave 3 survey ratings. In each model, I include interaction terms that allow ratings for an occupation in a given post-outbreak wave to interact with both essential worker classification and the Google Trends indicator of cultural salience. These interactions act as a test of the “unsettled times” hypothesis, and allow me to estimate the marginal effects of rating changes for essential or salient occupational identities and assess persistence in changes for each subgroup over time.

## **2.4 Results**

A previous study (Quinn et al. 2022) analyzed the first two waves of data involved in this paper, and found that mean evaluation and potency ratings for occupational identities declined shortly after the Covid-19 outbreak in the Wave 2 survey compared to ratings collected immediately before the pandemic in the Wave 1 survey. Wave 3 survey data allow us to assess change between pre-pandemic cultural meanings and more distal meanings almost a year after the pandemic began, and as the social upheaval associated with the spring of 2020 began to subside.

Replicating part of the analytic strategy from this previous work, I first report the results of 84 Welch-adjusted two-sample t-tests within each of the three EPA dimensions for all occupational identities in the study (Table 1) to assess broad patterns of change

and stability in cultural meanings of occupations in the study. The first sample for each t-test in this paper is the set of ratings for a given identity the pre-outbreak Wave 1 survey. The second sample is the set of ratings for the same identity in the Wave 3 survey, which was fielded almost a year after the pandemic began in the U.S. I use a Bonferroni correction within each dimension to account for the higher family-wise error rate expected when conducting multiple hypothesis tests. The presence of *any* statistically significant p-values in a given dimension that persist after the correction provides evidence in favor of rejecting the universal null hypothesis for the family of tests – that there are no significant differences in the Wave 1 and wave 3 means for an EPA dimension of interest (Lee & Lee, 2018).

Prior analyses of Wave 2 data showed significant negative shifts in evaluation and potency (see Quinn et al. 2022), but overall patterns in the EPA data from Wave 3 tell a different story. The columns for evaluation and potency ratings in Table 1 clearly show that mean estimates in both dimensions are not significantly different between the Wave 1 and Wave 3 surveys for the majority of individual occupational identities in the sample. Regardless of significance, the majority of evaluation and potency means for Wave 3 are marginally more positive (51 and 63 occupations, respectively) than the Wave 1 means for most of the individual tests. Of the five uncorrected mean differences in evaluation meeting criteria for significance at an alpha level of 0.05 or lower, all were “essential” identities, and four changed in the positive direction. More of the individual tests in the potency dimension met criteria for significance across essential and non-

**Table 1: Weighted Two-Sample t-Tests for Occupations in the Analysis Sample**

N for Wave 1 means ranged from 99 to 1,132 depending on the module of the survey; the N for Wave 3 means is 296 for all occupations. All analyses that include Wave 3 data exclude “tele-marker” as an influential outlier. The left-hand column of the table divides identities in the analysis sample into “non-essential” and “essential” categories. The eight occupations that were classified as “salient” based on Google Trends data are Blackjack Dealer, Childcare Worker, Fast Food Worker, Grocery Bagger, Hairdresser, Healthcare Manager, Mailman, and Restaurant Hostess. Five of these salient occupations are also essential occupations; three are non-essential occupations.

Occupation	Evaluation			Potency			Activity		
	Weighted Mean		BFC	Weighted Mean		BFC	Weighted Mean		BFC
	BL	Wave 3		BL	Wave 3		BL	Wave 3	
Accountant	1.50	1.53		1.22	1.29		-0.01	0.16	
Actor	1.24	1.35		1.37	1.78		1.58	1.72	
Aerobics Instructor	1.78	1.76		0.65	1.13 *		2.32	2.14	
Artist	1.68	1.78		0.28	0.88 *		-0.52	0.02 *	
Barber	2.11	2.04		0.54	0.38		0.80	0.46 *	
Bartender	1.83	1.89		0.59	0.73		2.23	2.02	
Blackjack Dealer	0.82	1.08		0.82	0.91		1.08	1.42	
Bridal Consultant	1.77	1.75		0.59	0.84		0.51	0.95	
Cafeteria Server	1.72	1.86		-0.38	-0.48		1.26	1.05	
Casino Cashier	1.22	1.36		0.07	0.66 *		1.32	1.02	
Casino Manager	0.91	0.90		2.20	2.11		1.80	1.72	
Cinematographer	0.98	1.36		0.74	1.47 **		0.51	1.24 **	
Cook	1.95	1.87		0.97	1.11		2.00	1.85	
Cosmetologist	1.80	1.81		0.48	0.74		0.20	0.84 **	
Designer	1.69	1.43		1.10	1.48 *		0.41	0.80	
Dietitian	1.68	1.68		0.39	0.90 *		-0.38	0.33 **	
Door-to-Door Salesman	0.05	0.52		-0.40	-0.32		0.51	1.26 **	
Event Planner	2.02	1.85		1.04	1.58 **		1.40	1.51	
Gift Wrapper in a Department Store	1.59	1.64		-0.61	-0.49		0.39	0.77	
Hairdresser	1.69	1.94		0.50	0.71		1.11	1.09	
Housekeeper in a Private Home	1.95	1.89		-0.38	-0.29		0.82	0.78	
Jewelry Maker	1.78	1.64		0.42	0.84		-0.44	-0.22	
Lawyer	0.63	0.77		2.42	2.49		1.39	1.20	
Librarian	2.48	2.44		0.22	0.31		-1.44	-1.28	
Library Assistant	1.98	1.97		-0.48	0.17 **		-1.91	-0.93 ***	**
Manager of a Movie Theater	1.40	1.43		0.59	1.00		0.33	0.55	
Motion Picture Projectionist	1.38	1.26		0.64	0.72		0.01	0.41	
Museum Curator	1.80	1.78		0.66	0.99		-0.35	-0.17	
Musician	2.07	1.95		1.39	1.49		2.32	2.08	
Parking Lot Attendant	0.84	1.17		-0.30	-0.13		-0.06	0.54 *	
Piano Tuner	1.49	1.65		0.49	0.48		-0.15	0.35	
Professional Athlete	1.49	1.28		2.52	2.39		2.82	2.58	
Real Estate Agent	1.33	1.49		1.07	1.52 *		0.81	1.20	
Restaurant Hostess	1.89	2.05		-0.21	0.52 **		1.65	1.47	
Restaurant Operator	1.67	1.79		1.96	1.85		1.67	1.38	
Security Guard	1.38	1.52		1.23	1.42		-0.08	0.59 *	
Sign Painter	1.32	1.25		-0.06	0.26		0.13	0.32	
Singer	1.61	1.67		1.60	1.94		1.91	2.02	
Ticket Taker	1.32	1.34		-0.33	-0.06		0.09	0.19	
Tour Guide	1.95	2.10		0.31	0.90 *		1.31	1.77 *	
Travel Agent	1.91	1.77		0.73	0.65		0.39	0.60	
Waitress	1.79	1.97		-0.60	0.07 *		2.02	1.71	
Writer	1.94	2.07		0.71	1.17 *		-1.17	-0.63 *	

Non-Essential

(continued)

Bank Manager	1.44	1.43	2.07	1.73 *	0.33	0.07
Bulldozer Mechanic	1.14	1.34	1.38	1.71	1.74	1.90
Bus Driver	1.72	1.74	0.48	0.19	0.76	0.66
Cashier	1.60	1.75	-0.39	-0.12	1.22	0.86 +
Childcare Worker	1.67	2.30 **	0.57	1.35 **	1.54	1.77
Diesel Motor Mechanic	1.27	1.33	1.22	1.45	1.61	1.68
Elder Care Aide	2.14	2.12	0.58	1.05 +	-0.21	0.35 +
Elementary School Teacher	2.57	2.56	1.30	1.40	1.46	1.11 +
Farm Owner and Operator	2.37	1.91 **	1.74	1.56	2.06	1.74 +
Fast Food Worker	0.84	1.46 +	-0.72	-0.23	1.67	1.93
Firefighter	2.98	2.77	2.78	2.73	2.90	2.77
Forklift Driver	1.42	1.14	1.13	1.11	1.73	1.77
Grocery Bagger	1.43	1.63	-0.75	-0.41	0.89	1.21
Healthcare Manager	1.43	1.67	1.61	1.83	0.66	0.93
Home Health Aide	1.84	2.16	0.44	1.04 **	0.25	0.90 +
Janitor	1.88	1.68	-0.51	-0.33	0.75	0.83
Legislator	-0.13	0.41 **	2.07	1.98	0.69	0.39
Licensed Practical Nurse	2.32	2.27	1.26	1.80 +	0.98	1.50 +
Locomotive Engineer	1.47	1.71	1.66	1.94	1.84	1.74
Mail Sorter	1.20	1.48	-0.26	0.35 +	0.31	1.23 ***
Mailman	1.92	2.10	0.42	0.85	0.82	1.02
Manager	1.23	1.42	1.97	1.88	1.34	1.41
Medical Assistant	2.19	2.01	0.78	1.21	0.89	1.22
Medical Technician	2.15	2.29	1.43	1.64	1.21	1.19
Paramedic	2.85	2.84	2.27	2.30	2.59	2.41
Pharmacist	1.77	2.21 +	1.60	1.85	-0.02	0.69 **
Pharmacy Aides	1.63	1.90	0.58	0.55	0.42	0.63
Physician	2.39	2.32	2.47	2.44	1.19	1.37
Post Office Clerk	1.36	1.42	0.01	0.43	0.10	0.41
Railroad Conductor	1.69	1.52	1.92	1.62	1.39	1.65
Registered Nurse	2.65	2.63	1.78	1.82	1.91	1.68
School Counselor	1.91	2.05	1.24	1.30	0.12	0.33
School Principal	1.44	1.68	2.35	2.41	0.93	1.11
Shelf-Stocker in a Grocery Store	1.12	1.46	-1.03	-0.17 **	0.52	1.34 **
Social Worker	1.30	1.70	1.10	1.55	0.30	0.58
Subway Operator	1.20	1.43	1.34	1.13	1.43	1.44
Supervisor of a Truck Delivery Service	0.94	1.29	1.54	1.63	1.35	1.30
Surgeon	2.49	2.54	2.86	2.82	1.35	1.27
Truck Driver	1.74	1.33	1.13	1.39	1.36	1.51
Truck Driver's Helper	0.89	1.11	0.36	0.15	1.11	1.03
Welder	1.52	1.47	1.29	1.23	1.82	1.67

N Occupations = 84

\*  $p < .05$ ; \*\*  $p < .01$ ; \*\*\*  $p < .001$  (two-tailed tests)


essential categories (20 in total), but no persisting significance in either the evaluation and potency t-test families persists after applying a Bonferroni correction to each<sup>2</sup>. This result – both in the magnitude of p-values in the tests, and the direction of even non-

<sup>2</sup> Changes to telemarketer were significant, positive, and robust to the Bonferroni correction, but the estimate was well-outside the interquartile range and represented a highly influential outlier in the context of analyses in this paper, so it was excluded from the tests reported in this paper.

significant change – stands in stark contrast to comparisons of means from the pre-pandemic Wave 1 survey and the proximal post-pandemic Wave 2 survey which show substantial short-term declines in evaluation and potency.

Results in the activity dimension follow a different pattern. Almost all uncorrected significant differences in the family of pairwise tests for mean activity ratings are also positive, but the magnitude of some increases in mean activity are more substantial in comparison to the other two dimensions. Because of this, persisting significance after applying a Bonferroni correction to the activity t-test family leads to a rejection of the null hypothesis and provides strong evidence that mean activity ratings increased between the Wave 1 survey and the more distal post-outbreak Wave 3 survey.

While the analytic strategy does not allow for *post hoc* inferences about specific occupations, it is clear that some occupations were viewed as significantly more active almost a year after the Covid-19 outbreak began. The family of unpaired t-tests capturing mean differences in the evaluation and potency dimensions, however, appear more stable in comparison. Both the direction and magnitude of these differences between Wave 1 and Wave 3 – in every dimension – are the opposite of the differences between Wave 1 and Wave 2. Not only do the declines shown in the more proximal Wave 2 survey appear to have rebounded in the more distal Wave 3 survey; in some instances, they appear to have superseded pre-pandemic estimates.



**Figure 1. Differences in Mean EPA Ratings in Survey Waves 2 and 3 vs. Wave 1**

Note: Occupations are sorted from left to right by magnitude the change for an occupational identity from Wave 1 to Wave 2.

Figure 1 decomposes the mean EPA differences from Wave 1 across both post-outbreak waves for all 84 occupations in the analysis sample. Each point on the x-axis represents a mean evaluation rating for an occupational identity in the Wave 1 survey. Black circles show the difference between Wave 1 and Wave 2 evaluation means for each identity, and gray diamonds show the difference between the Wave 1 and Wave 3 evaluation means for the same identities. The rated occupations are plotted in order of least to greatest mean difference between the Wave 1 survey and the Wave 2 survey. Grey diamonds can be interpreted as residuals in terms of their position relative to either the line of Wave 2 black circles *or* the Wave 1 survey axis.

As panels A and B of Figure 1 clearly show, changes in evaluation and potency that occurred shortly after the U.S. Covid-19 outbreak (i.e., in Wave 2) are mostly negative across the sample of occupations. But by early 2021 (i.e., in Wave 3), the post-outbreak evaluation and potency changes that were observed in Wave 2 broadly disappear, and even trend in the positive direction relative to pre-outbreak Wave 1 measurements. Unlike the estimates from the Wave 2 survey, the Wave 3 means for evaluation and potency are distributed more evenly around the axis indicating no change from Wave 1 estimates. In the activity dimension (Figure 2C), prior analyses did not find evidence that mean ratings in Wave 2 were substantively different from Wave 1 mean ratings of the occupations in the study (see Quinn et al. 2022), while activity

ratings appear to be significantly more positive in Wave 3 compared to Wave 2 and Wave 1 surveys.

While there is evidence for dynamics in cultural beliefs over the course of the study period, it is important to note that the system of cultural meanings in the work institution remained largely stable during the socially tumultuous period of inquiry. The correlation between pre-outbreak and post-outbreak means in each dimension for the 84 work identities is exceptionally high. Correlations of Wave 1 values with both Wave 2 and Wave 3 values are above .91 in all three EPA dimensions. There is a clear, strong and positive association between pre-outbreak mean ratings with Waves 2 and 3, for both essential and non-essential identities.

The final analysis formally estimates the direction and magnitude of changes in mean EPA ratings within occupations and between survey waves, and explores whether short-term or long-term changes are moderated by their classification as essential (from state and federal reports) or salient (from Google Trends). I specify a set of linear mixed models to estimate EPA values as a function of survey wave and essential or salience status (Table 2). The mixed effects approach allows me to capture differences in mean ratings for each EPA dimension across the three survey waves within both essential and non-essential occupational groups and cultural salience flags (the “fixed” components), while accounting for variation in overall mean rating differences between occupations (the “random” component). Each occupation appears in the analysis dataset for the

mixed model of a given EPA dimension three times, once for each point in time. The intercept of each model is a random parameter representing the value of the average Wave 1 mean value of the relevant EPA dimension for non-essential and non-salient occupations.

The coefficient associated with the Wave 2 and Wave 3 variables in Table 2 are the estimated effect of the short-term (i.e., by June of 2020) or longer-term (i.e., by February of 2021) post-outbreak transition on the mean value of a given dimension rating when both essential and salience status is zero, or for non-essential and non-salient occupations. The parameters “Wave 2” and “Wave 3” are binary indicators that identify whether an observation contains a mean rating from one of the post-outbreak waves, as opposed to the Wave 1 reference category.

Essential and salience statuses are two separate binary fixed parameters<sup>3</sup> that identify whether an occupation was classified as essential or salient (0 = no; 1 = yes). Coefficients associated with each of these indicators in Table 2 represent the estimated effect of essential status on the difference in evaluation, potency, or activity scores during Wave 1, before the pandemic. Note that any significance in these coefficients would be driven by pre-pandemic differences in the EPA scores of the occupations included in our analysis sample (a feature of our study design), rather than inherent

---

<sup>3</sup> The essential and salience indicators classify a broadly orthogonal set of occupational identities. The correlation of these binary fixed effects within the models reported in Table 3 are -0.045.

differences in how U.S. culture evaluates the occupations classified as essential or salient versus not.

The analysis sample for the mixed models includes 63 occupational identities, as 21 occupations contained data that were either too sparse (e.g., “supervisor of a truck delivery service”) or too oversaturated, often because of the term captured a broader set of meanings than an occupational identity (e.g., artist), to include in the analysis sample<sup>4</sup>.

**Table 2: Linear Mixed Models Predicting EPA Mean Occupation Ratings**

Coefficient	Model 1 (Evaluation)		Model 2 (Potency)		Model 3 (Activity)	
	Estimate	P-Value	Estimate	P-Value	Estimate	P-Value
(Intercept)	1.660 ***	<0.001	0.790 ***	<0.001	0.745 ***	<0.001
Wave 2	-0.282 ***	<0.001	-0.142 **	0.002	-0.092	0.061
Wave 3	0.027	0.478	0.239 ***	<0.001	0.179 ***	<0.001
Essential Status	0.166	0.203	0.554 **	0.008	0.367	0.084
Saliency Status	-0.286	0.167	-1.015 **	0.003	0.296	0.376
Wave 2 x Essential	0.072	0.169	-0.129 *	0.038	-0.078	0.249
Wave 3 x Essential	0.015	0.780	-0.121	0.052	-0.940	0.163
Wave 2 x Saliency	0.323 ***	<0.001	0.156	0.112	0.007	0.951
Wave 3 x Saliency	0.292 ***	<0.001	0.267 **	0.007	0.039	0.715
N (63 x 3)	189		189		189	
Marginal Pseudo R <sup>2</sup>	0.091		0.192		0.0647	
Conditional Pseudo R <sup>2</sup>	0.925		0.963		0.952	

\* p < .05; \*\* p < .01; \*\*\* p < .001

<sup>4</sup> The findings reported in Table 2 are robust to the inclusion of the excluded identities with imputed saliency indicators of 0, and to analyses that randomly assigned values for the saliency indicator to missing cells.

The model results show that evaluation and potency decline significantly in the short-term Wave 2 survey, but both rebound substantially by the time the Wave 3 survey is administered. In fact, potency ratings are significantly more positive in Wave 3 compared to the Wave 1 survey. Activity ratings in Wave 2 did not appear to be significantly different from Wave 1, but significantly increased in Wave 3.

The models in Table 2 indicate that salient and non-salient occupational identities had different potency ratings in Wave 1 (-1.015,  $p < 0.01$ ), as did Wave 1 mean potency ratings for essential and non-essential identities (0.554,  $p < 0.01$ ). Again, it is important to note that these differences are a feature of the occupations in our analysis sample, which contains purposively selected occupational identities categorized as essential and non-essential. We preserve the coefficients for readers to understand this design aspect, but do not interpret them as real differences between the population of essential and non-essential workers from before the pandemic.

For occupations that spiked in cultural salience during the first three months after the U.S. outbreak, Covid-19 had a different effect on mean evaluation ratings. Short-term mean evaluation ratings in the Wave 2 survey for salient occupations are significantly more positive (0.323) than non-salient occupations ( $p < 0.001$ ). This change in evaluation for salient occupations compared to non-salient occupations persisted instead of rebounding: the longer-term mean evaluation ratings in the Wave 3 survey remain significantly more positive for salient occupations (0.291) compared to non-

salient occupations ( $p < 0.001$ ). Broadly speaking, it seems that salient occupations were largely protected from the decline in evaluation that accompanied U.S. culture's transition to the pandemic, and this emergent difference in meaning has persisted beyond the initial social discontinuity of the pandemic's outbreak. In the presence of the salience indicator, essential occupation status did not appear to significantly moderate short term or long term changes in mean evaluation ratings<sup>5</sup>.

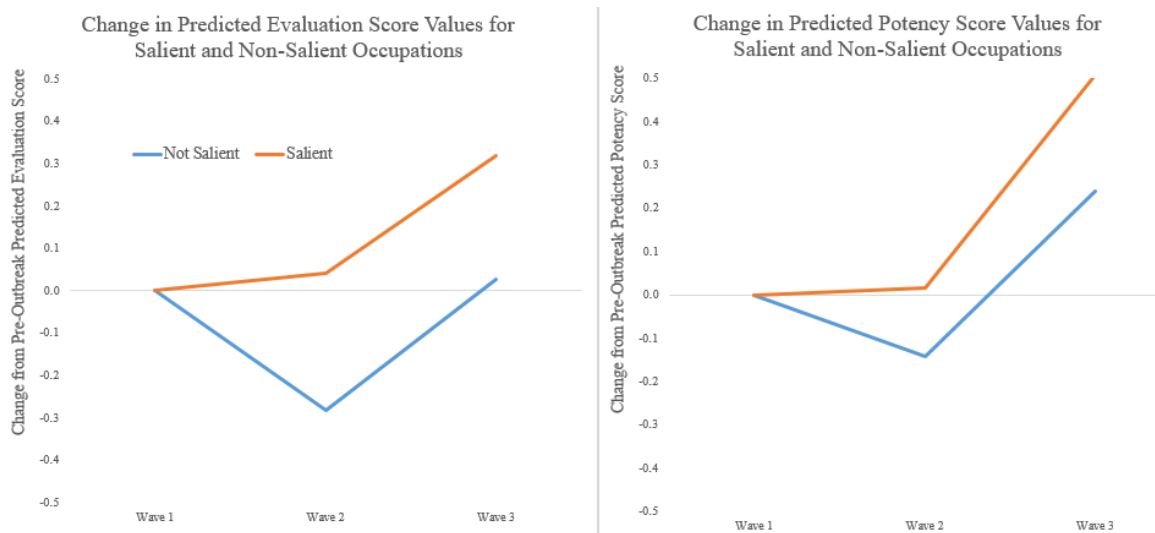
By contrast, Model 2 shows that short-term changes in mean potency ratings *are* moderated by the essential status of an occupational identity, as the difference in Wave 1 and Wave 2 mean potency values is somewhat more negative (-0.129) for essential occupations compared to non-essential occupations ( $p < 0.05$ ). Yet the moderating effect of salience in long-term mean potency rating changes from Wave 3 are far more pronounced (0.267,  $p < 0.01$ ). Changes in mean activity ratings occur solely in Wave 3 relative to Wave 1 (0.179,  $p < 0.001$ ), and post-outbreak wave ratings for activity do not interact with the salience or essential occupation indicators.

Figure 2 portrays the interactions for evaluation and potency, using the equations derived from Model 1 and Model 2 to estimate predicted values for the difference in post-outbreak mean scores for both salient and non-salient occupations from Wave 1 to Wave 3. These plots show that beliefs for salient occupational identities

---

<sup>5</sup> Note that, while Quinn et al. (2022) found essential occupational identities did moderate the relationship between Wave 1 and Wave 2 evaluation ratings, models that replace the essential status indicator with a cultural salience indicator render this interaction non-significant, and fit the data more effectively. I include both terms and their interactions with each Wave in the models above for the sake of comparison.

appear to have emerged as the pandemic began, and persisted beyond the initial months in which the social dislocation made these identities momentarily salient.



**Figure 2. Evaluation and Potency Score Interactions**

## **2.5 Conclusions and Discussion**

The results of the study show two important patterns. First, while some individual occupations declined significantly in evaluation and potency shortly after the pandemic began to impact social and economic life in the U.S., and short-term mean ratings were significantly more negative, these values broadly rebounded nearly a year later. As a population of cultural objects, cultural meanings of occupational identities in this study appear to be generally stable in the face of a major social dislocation like the pandemic. At least some individual occupations in Wave 3 have significantly higher mean activity ratings, which could be explained by the material reality of the improved

relationship that Americans had with the work institution and social life more broadly during the timing of the survey compared to the timing of the Wave 2 survey.

In the context of ACT's theoretical framework, the overall short-term negative effect of the pandemic outbreak could be viewed as victim derogation. This interpretation is consistent with the notion that workplace dislocation has made occupational identities victims of the pandemic. Essential occupations have been especially salient as socially necessary and worthy of public praise during the pandemic, which may explain the lower decline in mean evaluation scores for these identities compared to non-essential occupations. While essential workers are clearly impacted by the pandemic, public discourse about their value and importance to our societal functions could attenuate victim derogation.

The second important pattern the study shows is that occupations that were made salient by the transition to the pandemic developed substantial changes in mean evaluation early on in the pandemic, and these mean differences persist almost a year later in spite of receding conditions and a decline in their salience since. This pattern corresponds to theoretical research on socialization in the culture and cognition literature, which posits that systematic individual-level belief change in a culture can be brought about by "unsettled times" (Swidler 1986). The results of this study suggest that measurable changes in cultural meaning about occupations may emerge and persist in

the presence of a large exogenous shock that places cultural actors in a context where their prior meanings fail to map appropriately to particularly salient cultural objects.

### **3. How Identities and Exchange Structures Affect Present Cooperation and Future Behavior toward People from Different Groups**

#### **3.1. Introduction**

Sociologists have long argued that micro-level social interactions between individuals shape, and are shaped by, the meso- and macro-level structures in which they are embedded (Stryker 1980; Fine 2012; Turner 1988; Giddens 1984; Lawler, Thye, and Yoon 2009; Blau 1964; Simmel 1950). One middle-range theoretical attempt to reconcile the problem of social order, the *Affect Theory of Social Exchange* (ATSE) (Lawler 2001), develops and tests predictions about what structural forms of exchange are more likely to facilitate cooperative exchange and social cohesion. ATSE researchers have since argued and shown (Lawler, Thye, and Yoon 2008) that the structure of exchange networks in which people are embedded affects their levels of cooperation and affective attachment. Highly joint exchange task structures (i.e., those that make contributions non-separable and induce a sense of shared responsibility amongst people in the network) cause people to behave more cooperatively with and develop emotional commitments toward their ongoing exchange partners.

At the same time, social psychological research has clearly shown that people frequently engage in preferential treatment toward those with whom they share social category memberships. Ample interdisciplinary research shows that people are more likely to trust (Romano et al. 2017; Foddy, Platow, and Yamagishi 2009; Tanis and

Postmes 2005), connect with (Melamed et al. 2020; McPherson, Smith-Lovin, and Cook 2001; Marsden 1987; Schelling 1971), and cooperate with (Yamagishi, Jin, and Kiyonari 1999; Yamagishi and Mifune 2009; Melamed et al. 2020; Balliet, Wu, and De Dreu 2014; Romano et al. 2021; Whitham 2018) those with whom they share a salient social identity. This micro-level behavior can act as a barrier to macro-level social order in diverse communities. As same-identity bias is enacted and reinforced through repeated dyadic interactions, norms of reciprocity trust, solidarity, and cooperation tend to emerge at higher levels in homogeneous communities compared to more diverse collectives (Coleman 1988; Baldassarri and Abascal 2020; Melamed et al. 2020; Habyarimana et al. 2009).

Scholars studying how social structures shape cooperation and affective attachment within ATSE have theorized about (Lawler, Thye, and Yoon 2016; 2021), and called for research to empirically evaluate (Lawler 2019), the impact of social identity information on the emergence of cooperation and affective commitment. An ideal test of this proposition would include exchange between people who do not share a social category. Moreover, as theoretical accounts of ATSE argue that highly joint repeated exchange results in emotional commitments to distal social objects like group identities, the theory implies that the structure of one's *prior* repeated exchange network can impact the way that they interact in the *future* toward unmet people from different groups. A test of this proposition represents a contribution to the ATSE, and has clear

implications for more applied intervention-based social policy research on creating cohesion and trust between people from different groups.

This study extends these lines of inquiry, and reports the results of an experiment designed to ask two previously unanswered questions at the intersection of social structure and social identity. First, I explore whether – and if so, how – network exchange structure interacts with same-identity bias in the production of ongoing cooperation with existing ties. I leverage insights from ATSE to make competing predictions about ongoing cooperation in productive, reciprocal, and generalized exchange when interaction partners are embedded in networks where identities are known and homogeneous, compared to networks where identities are known and heterogeneous, or unknown. Second, I test how the structure and identity composition of prior repeated exchange network impacts a person’s future behavior toward unmet interaction partners. Specifically, I ask whether engaging in prior repeated exchange with dissimilar others affects a person’s *future* trust and trustworthiness behavior toward unmet others, particularly when those others possess an out-group social identity category.

## **3.2. Theoretical Background**

### **3.2.1 Social Identity and Cooperation**

Social actions—including decisions to exchange valuable benefits with others—commonly involve tension between self-interest and what is best for the collective. This

is because cooperating with other people frequently comes with a short-term individual cost in return for long-term collective benefits. This conflict, and its implications for understanding social order (Hobbes 1651; Fehr and Gintis 2007), motivates a large body of research on the patterns and determinants of cooperative behavior (Simpson and Willer 2015). One important area in this literature is how social identities, aspects of the self that stem from membership in a group or category (Billig and Tajfel 1973; Tajfel 1981; Tajfel et al. 1971; Hogg 2020), impact cooperation.

Shared social identities facilitate cooperation and prosocial behavior across a variety of contexts and identities. These include identities ranging from political affiliations (Fowler and Kam 2007; Iyengar and Westwood 2015) and racial or ethnic groups (Abascal 2015), to group assignments in summer camps (Sherif et al. 1961) or in the lab (Whitham 2018), university affiliations (Foddy, Platow, and Yamagishi 2009; Melamed et al. 2020), and even relatively meaningless social categories, like dot over- vs. under-estimators or painting preferences (Tajfel et al. 1971; Simpson 2006; Aksoy 2015, 2019). The benefits of shared social identity on cooperation are highly robust (Balliet, Wu, and De Dreu 2014) and cross-cultural (Romano et al. 2017; Romano et al. 2021).

At the same time, scholars have increasingly recognized that people *do* engage in cooperative endeavors with out-group others (Aksoy 2015; Romano, Balliet, and Wu 2017; Bettencourt et al. 1992). Members of heterogeneous groups and communities often manage to get along (Baldassarri and Abascal 2020). Studies under the umbrella of

intergroup contact theory (see Pettigrew and Tropp 2006) show that structural (Blau 1977) and interpersonal (Allport 1954) features of group dynamics can concurrently generate cooperation and cohesiveness between people with different social identities. And the Common Ingroup Identity Model, backed by decades of empirical work (Dovidio, Gaertner, and Kafati 2000; Gaertner and Dovidio 2014; Gaertner, Dovidio, and Bachman 1996; Gaertner et al. 1999; Gaertner et al. 1993; Riek et al. 2010; Vezzali et al. 2015; Cunningham 2005; Lemay Jr and Ryan 2021), proposes that key environmental and social-structural factors increase the salience of shared identities (e.g., “American”) and reduce preferential treatment between actors along the lines of their more narrow social identities (e.g., “Democrat” or “Republican”).

These findings complicate the clear evidence that people typically deploy perceptions and actions that favor others who share their social identity. Instead, same-identity bias may operate differently based on the structure in which people interact. In the section to follow, we review a key structural and sociological determinant of cooperation: the structure of exchange, which shapes the patterns and processes by which benefits flow between interaction partners.

### **3.2.2 Ongoing Cooperation and Emotion in Exchange Networks**

As noted above, sociologists in the social exchange tradition have long argued that micro-level social interactions are shaped by the structures in which those interactions occur. The *affect theory of social exchange* (Lawler 2001) argues that the

structure of exchange shapes the degree to which valuable benefits flow through networks, producing stable patterns of cooperation, positive emotion, and social order.

According to the theory, these micro social orders are shaped by the *jointness* of the exchange task. Jointness is high when the exchange activity (1) makes peoples' contributions to task success non-separable, or indistinguishable from the contributions of others, and (2) generates perceptions of shared responsibility for task success or failure. When repeated exchange entails high task jointness – as in jointly-authored papers or team work projects where the product is a result of joint efforts, attributed to the group rather than to any given member – the resulting cooperation rates are higher, as are emotional bonds toward broader and more distal social units than immediate relations, such as the network itself (Lawler 2001; Lawler, Thye, and Yoon 2000, 2008).

By contrast, exchange activities where the contributions of interactants are separable and perceptually individualistic, like providing comments on a colleague's paper, have relatively lower task jointness. ATSE predicts that these forms of exchange, in turn, yield lower levels of cooperative exchange, and weaker commitments in general – especially, according to the theory, toward more distal social units.

This study focuses on three specific exchange structures explored by ATSE scholars.<sup>6</sup> *Productive exchange* involves jointly produced benefits: interactants jointly

---

<sup>6</sup>A fourth is *negotiated exchange* (Barrera 2007; Kuwabara 2011; Lawler et al. 2008; Molm 1994; Molm et al. 2006; Molm et al. 2009), which entails zero-sum bargaining over a set pool of resources (as in agreeing on a price at which a seller will sell to a buyer). Exchange is bilateral and made on agreement between both parties. We did not include

produce a shared good by unilaterally providing benefits to produce the good, and receiving benefits from it (Lawler, Thye, and Yoon 2000; Emerson 1976; Molm 1994). When people collaborate on a team work project, scholars co-author a paper together, or multiple people commit to a business partnership, they are all seeking to accomplish something they cannot do alone and potentially receive significant rewards from mutual cooperation. Interdependence is high in productive exchange: the exchange partners must coordinate their efforts to produce the collective good and share in the collective benefits; the largest payoffs are from mutual cooperation (Lawler, Thye, and Yoon 2008; Kollock 1998). Thus in productive exchange, nonseparability and joint responsibility for the shared outcome are high, and ATSE predicts it to yield particularly high levels of exchange frequency and positive emotions to the social unit (Lawler, Thye, and Yoon 2008).

*Reciprocal exchange* is a form of direct exchange: two or more actors can give directly to each other over time. As in productive exchange, giving is a unilateral act, done without knowing whether others will give. But unlike productive exchange, reciprocal exchange involves giving directly to one person with the opportunity to receive directly from that person at some later point in time. Examples including paying for the check when out for drinks with a friend, doing a favor for a neighbor, or

---

negotiated exchange in the experiment described in the next section. Instead, our focus here is on how *unilateral* decisions to provide benefits to others facilitate exchange when social identities are shared or not.

providing comments on a colleague's work. Because it involves direct and distinct acts of giving from one party to another, jointness is not as high as in productive exchange (which involves the joint production of a collective benefit shared by the group, rather than distinct decisions to give directly); as a result, nonseparability and sense of shared responsibility is lower compared to productive exchange (Lawler, Thye, and Yoon 2008).

*Generalized exchange* also involves unilateral acts of giving to another party that entail risk and uncertainty about receiving benefits in return, but it is indirect. This is because in generalized exchange, the flow of benefits does not return directly to those who have provided them – instead, A's giving to B may mean instead that A will be helped by some third party, C, in the future, rather than directly from B. Examples include acts of mentorship, or stopping to help a stranded motorist. As with reciprocal exchange, in generalized exchange giving behaviors are highly separable, rather than joint, and thus are less likely to generate a sense of shared responsibility and unleash high levels of cooperation and affective attachment (Lawler, Thye, and Yoon 2008). These acts further entail a high risk of uncertainty and the likelihood that one's giving will not be reciprocated by some third party in the future (Molm 1994; Molm, Collett, and Schaefer 2007; Lawler 2001); this risk is reflected in relatively lower levels of giving in generalized exchange systems (Molm, Collett, and Schaefer 2007; Whitham 2021).

### 3.2.3 Future Commitment Behavior as Trust and Trustworthiness

The ATSE also argues that the ongoing structure of repeated exchange affects future commitment behaviors toward others based on the extent to which a structure leads to cooperation and causes emotional commitments to form between interaction partners, their relations, and broader social units like the network itself or the social categories of other actors within the network (Lawler 2001). In other words, exchange structures with greater task jointness are *also* expected to induce higher levels of commitment toward the larger groups to which actors within an exchange network are affiliated – such as the exchange network itself, and the social identity categories or organizational affiliations of actors within the exchange network. This “nested commitments” proposition has become a recent topic of theoretical clarification efforts (Lawler, Thye, and Yoon 2016; 2021), but remains broadly untested in an empirical context (Lawler 2018).

One consequential set of commitment behaviors is trust and trustworthiness (Yamagishi and Yamagishi 1994). Researchers measure these two distinct outcomes by placing two interactants in a trust game where each makes a choice about sending and returning points without knowledge of their partner’s decision. Participants assigned to the “sender” role make a decision about how much they trust their partner by determining how many monetary units they will keep to themselves versus sending to that partner. Points that they send to their partner are multiplied. Participants assigned

to the “returner” role make a decision about how trustworthy they find their interaction partner to be by determining what proportion of the multiplied points shared by the sender they will return versus keep for themselves. Trust behavior involves higher levels of risk, as senders cannot be certain that they will not lose everything they share with the returner. Though trustworthiness behavior still involves a decision on the amount of valuable benefits a returner will share with their sender, the returner can only distribute benefits that have been allocated to them by the sender to begin with.

Prior studies of social exchange distinguishing between and measuring the emergence of trust and trustworthiness have shown that one or the other can emerge in dyadic relations as a result of repeated participation in prior reciprocal (Molm et al. 2000) or negotiated exchange (Barrera 2007), and argue that trust cannot develop without repeated demonstrations of one’s trustworthiness (Molm et al. 2000).

If, as ATSE proposes, the structure or identity composition of an exchange network impacts on the affective commitments that a person cultivates toward out-group others as distal social objects, that person would be more likely to demonstrate greater levels of trustworthiness behavior toward unmet interaction partners in one-shot trust games. The same is true of trust behavior, though the higher-risk nature of trust decisions would make such an outcome less likely than trustworthiness.

### 3.2.4 Bridging Identities and Exchange Structure

While empirical work has found support for the ATSE framework, it has evolved outside of the interdisciplinary literature on same-identity bias. Likewise, research on the effects of same-identity bias on cooperative exchange in networks has not yet incorporated potentially clarifying insights from ATSE about the influence of exchange structure. A handful of studies have demonstrated the benefits of shared social categories within generalized (Whitham 2018; Buchan, Croson, and Dawes 2002) or reciprocal (Yamagishi and Kiyonari 2000) exchange. But it is unclear how same-identity bias operates across the different exchange forms. As such, here we ask how social categories intersect with variation in the features of different exchange structures to alter ATSE's predicted levels and patterns of giving. Our study thus integrates the findings on same-identity bias in cooperation with ATSE's insights on the structuring of exchange relations in triadic networks, and its proposition of distal social unit commitments from repeated exchange.

My first set of hypotheses is about ongoing cooperation in exchange networks – whether and how the effects of exchange structure and social identity intersect. Will the main effects of sharing a social identity and being embedded in triadic exchange networks with different degrees of task jointness appear when we manipulate both of these features simultaneously? Beginning with an attempt to corroborate the predictions (Lawler 2001) and findings (Lawler, Thye, and Yoon 2008) from ATSE, we expect that

the three exchange structures included in our study will induce distinct cooperative exchange rates as follows:

*Hypothesis 1: All else equal, productive exchange structures will facilitate high levels of cooperative exchange and affective attachment, followed by reciprocal exchange, followed by generalized exchange.*

Likewise, the large body of research on same-identity bias in cooperation implies a clear expectation about the effects of shared social identity on cooperation within any particular exchange structure:

*Hypothesis 2: All else equal, exchange structures where interactants share a known social identity will produce higher levels of cooperative exchange and affective attachment, compared to exchange structures where they do not.*

The third hypothesis about cooperation and affect in ongoing exchange focuses explicitly on the presence or absence of an interaction between types of exchange structure and the presence of social identity information in a network. As sharing social categories make people feel more assured in higher-risk situations (Yamagishi and Kiyonari 2000; Kiyonari 2002; Yamagishi, Jin, and Kiyonari 1999; Kiyonari and Yamagishi 2004), and not sharing or not knowing identities can generate a greater sense of uncertainty (Molm 2010), I propose that different minimal group identity information conditions will interact with the exchange structure conditions:

*Hypothesis 3: Shared social identities will interact with exchange structure: when interactants share a known social identity, the difference between productive, reciprocal, and generalized exchange will be reduced, compared to when they do not.*

The last two hypotheses I test in this study depart from the first three by focusing on the effects of past repeated exchange structure and identity information on *future* trust and trustworthiness behavior toward unmet others. The ATSE's nested commitments proposition – which states that exchange tasks with greater jointness are more likely to lead to emotional commitments toward distal social units compared to exchange tasks with lower levels of jointness – implies the first of these two hypotheses:

*Hypothesis 4: productive exchange structures will facilitate highest levels of trust and trustworthiness behaviors toward unmet alters, followed by reciprocal exchange, followed by generalized exchange.*

The final hypothesis makes a prediction about the difference between being in a diverse exchange network with one dissimilar other versus two dissimilar others. As many studies show that cooperation emerges from repeated network exchange, cooperation that emerges in diverse exchange structures provides interactants with the opportunity to develop generalized expectations about out-group others as a social object rather than idiosyncratic expectations about individual actors from an out-group (i.e., "people from my out-group tend to behave cooperatively because I have seen more than one person from that group being cooperative in a group with me," versus "this single person in my

network from my out-group is being cooperative with me, but I do not know what to expect from the next actor from my out-group"). Building on findings that repeated exchange in can shape trust and trustworthiness behavior (Molm 2000), and evidence on the role of intergroup contact in positive assessments of out-group others (see Pettigrew and Tropp 2015), I propose:

*Hypothesis 5: Actors who were previously embedded in a social exchange network containing multiple interaction partners from their relative out-group will engage in more trust and trustworthiness behavior when they encounter unmet out-group alters in a follow-up one-shot trust game.*

In the section to follow, I describe an experiment designed to test the first three hypotheses about ongoing exchange and the last two about future trust and trustworthiness. Following the exchange task procedures of Lawler and colleagues (2008), I embedded participants in triads where interactants shared a known social identity, compared to two variations on a control condition: one where social identities were known but not shared, and another where social identities were unknown. Then, at the end of repeated exchange phases, I place participants in two rounds of a one-shot trust game where they play the role of sender and receiver.

### **3.3. Data and Methods**

#### **3.3.1 Participants and Design**

Study participants were student and non-student adults belonging to a university-affiliated research participant pool. They were required to have access to a laptop or desktop computer with a stable internet connection in order to complete the study for payment. Consistent with past experimental work in this area, participants completed the exchange task in triads (Lawler, Thye, and Yoon 2000, 2008; Molm, Collett, and Schaefer 2006; Molm, Schaefer and Collett 2007). The study involved three phases of the exchange task; participants were embedded in a new triad with new alters in each phase. In order to pair participants with different alters in each of the three triads, we required at least nine participants to be enrolled in each session; otherwise, it was cancelled. Our sample includes data from 30 study sessions in total; 29 sessions contained nine participants each, while one session included 12 participants.

Each study session took place entirely online. If a participant was unable to progress through the session due to a technical issue or went idle for more than two minutes during the real-time decision-making task, the study program filled in for them for the current round of decision-making and all subsequent rounds, as described in more detail below. This approach prevented the study from being cancelled mid-session because we could not form, or continue interactions in, a triad. These programmed decisions (9% of all 14,166 rounds of interaction we observed) were dropped from all

analyses reported below. The result was a final sample containing 251 participants, who completed 12,871 exchange task decisions. The Online Appendix demonstrates that our results do not differ whether the programmed decisions are included or, as they are in sections to follow, excluded from analyses.

The study contained two manipulations, shown in Figure 3: *exchange structure* and *information about identities*. *Exchange structure* was manipulated between sessions; all nine triads (three triads in each of three phases) within a session were randomly assigned to the same structure. Eleven of our 30 sessions were randomly assigned to the *generalized exchange* condition, 9 sessions to the *reciprocal exchange* condition, and 10 sessions to the *productive exchange* condition. *Information about identities* was manipulated at the phase level. In the *unknown identities* phase, participants did not receive information about the others' identities. There were two phases with *known identities*: in the *homogeneous identities* phase, participants made exchange decisions with two others with whom they shared a social identity (described in more detail below). In the *diverse identities* phase, the triad contained either one or two alters who did *not* share ego's identity. For example, in the *diverse identities* triads shown in Figure 3, from A's perspective the triad contains two alters who do not share A's identity (i.e., B and C). From B's (and C's) perspective, the triad contains one alter who does not share ego's identity (i.e., A).




Figure 3: Key study manipulations

Each row shows diagrams of the three levels of the *exchange structure* manipulation; each column shows one of the three phases in each study session, where we manipulated *information about identities*. Participants completed each triad in one row, depending on their session’s randomly assigned exchange structure, in random order. For simplicity all labels are identical (ABC); in the actual rounds of the task, the identifying letters varied. The example *known identities* triads are shown from the perspective of a participant whose own social category was colored blue, and who was paired with two others who did not share their identity in the *diverse identities* phase. Other participants saw different colors depending on whether their social category was colored blue or green and whether they were paired with one or two others who did not share their identity in the *diverse identities* phase. See the next section for more information about the identities manipulation.

Participants completed three phases, each time in a new triad with new others; thus they were exposed to all three levels of the information about identities manipulation. The order in which participants completed the three phases was randomized at the session level. Following each phase of repeated exchange, participants would then play two rounds of a trust game to make one trust decision and one trustworthiness decision about a previously-unmet interaction partner. The following section provides more details about the manipulations and study procedures.

### **3.3.2 Procedure**

The study took place online, and was programmed using zTree (Fischbacher 2007) and zTree Unleashed (Duch, Grossmann, and Lauer 2020). Prior to the main study task, participants read a consent form detailing the study procedures and their expected payment for participating. They proceeded to answer basic demographic questions (age, gender, and whether they were a university student or not) and completed a Klee and Kandinsky task to manipulate identities. The task is a standard method of establishing social identities in the laboratory (Aksoy 2015, 2019; Billig and Tajfel 1973; Ridgeway et al. 1998; Ridgeway and Erickson 2000; Simpson 2006; Yamagishi and Kiyonari 2000) on the basis of a relatively neutral trait: preference for abstract paintings. Participants were told that they would view pairs of paintings by two artists—abstract works by the artists Paul Klee and Wassily Kandinsky—and choose the painting from the pair that they

preferred the most. Then, they and the others completing the study would be classified as either Klees or Kandinskys, based on which artist they preferred relative to the other.

In reality, whether participants were told that they were a Klee or a Kandinsky was randomly assigned by the computer program, and was not actually dependent on their answers in the painting task. We used a minimal group identity task to randomly assign identities, rather than relying on a more “natural” social identity like a demographic category (e.g., gender or race), for two reasons. First, assigning identities rather than exploiting the demographic categories or existing group affiliations of participants ensured that participants were able to behave without bringing prior biases and expectations about “real” social categories into their interactions. Second, random assignment allowed us to distribute participants within each session so that each participant could successfully be paired in a homogeneous identities triad (i.e., paired with two alters who had been categorized as a Klee if ego had been categorized as a Klee) in one phase, and in a triad with diverse identities (i.e., paired with at least one alter who had been categorized as a Kandinsky if ego had been categorized as a Klee) in another phase. Meta-analyses have demonstrated that artificially established identities like those used in the Klee and Kandinsky task are as likely as more “naturally” based categories to generate favoritism toward one’s own social category in the lab (Balliet, Wu, and De Dreu 2014).

After answering five painting preference questions, participants were told that based on their answers they were either a Klee or a Kandinsky, and that in the upcoming portion of the study they would complete a decision-making task with other participants who had also been identified as Klees or Kandinskys.

### **3.3.3 The Exchange Structure Manipulation and Trust Game Outcome**

The computer program randomly assigned the session to either the generalized, reciprocal, or productive exchange structure. An initial set of instructions, held constant across conditions, informed participants that the session included at least eight other people. In the upcoming task they would be connected at random to exactly two of the others, and those others might be Klees, Kandinskys, or both (more specific information was provided about identities later, as described in the next section). The decision-making task, which they would complete with the two others, consisted of multiple rounds. In each round, they would make a binary decision to either transfer or keep an endowment of points. Points had value: participants' total points earned throughout the study would be translated into a dollar amount that would determine their study payment.

Subsequent instructions varied based on the exchange structure. Before turning to the specific details of the task by exchange structure condition, first note that we followed prior work manipulating exchange structure by holding constant the costs and benefits of exchange across the different structures (Cheshire, Gerbasi, and Cook 2010;

Kuwabara 2011; Lawler, Thye, and Yoon 2008; Molm, Collett, and Schaefer 2007; Molm, Takahashi, and Peterson 2000; Yamagishi and Cook 1993). Specifically, we followed the exchange structure payoffs used by Lawler and colleagues (2008) exactly. Regardless of the exchange structure to which the session was assigned—that is, across all exchange conditions—choosing not to exchange yielded a default profit of 300 points for the participant. On the other hand, a viable exchange behavior yielded a benefit of 800 points (Lawler, Thye, and Yoon 2008).

In the generalized exchange condition, participants were told that they could send points to a specific other and receive points from a different other, as shown in a black-and-white example on their screen (reproduced in the first row of the first column in Figure 3). In this example, it was explained, A could decide whether to send points to B, B could decide whether to send points to C, and C could decide whether to send points to A. Those who decided not to send points in a round would keep a default payoff of 300 points. If a participant decided instead to send their points (i.e., engage in exchange) in a round, then the recipient would receive 800 points.

In the reciprocal exchange condition, participants could send points to one or both of the others, as shown in the black-and-white example shown in the second row of Figure 3. Thus, A could send points to B and/or to C, and could receive points from B and/or C. Again, those who decided not to send points to anyone kept a default payoff of 300 points. But if they decided to engage in an exchange with one or both of the

others, the other(s) with whom the participant chose to exchange received a benefit of 800 points. As a result, each actor could either keep 300 points by not giving to either alter, or could opt to make one or two exchanges that entailed the alter receiving 800 points.

In productive exchange, participants made decisions to send points to a “joint fund” shared by the triad, rather than sending points to specific others in the triad. Likewise, participants received points from the joint fund rather than from specific alters (see the third row of the first column in Figure 3). Giving to the joint fund produced benefits dependent on the number of others who also gave to the joint fund, as is standard in productive exchange and in experimental research on productive exchange structures (Lawler 2001; Lawler, Thye, and Yoon 2000). If one person gave to the joint fund, that person produced 800 points for the joint fund. These points in the joint fund would then be distributed equally to all actors in the triad (i.e., producing 267 profit points for each person in the triad; the two who did not give also received their 300-point default payoff). If two people sent their resources to the joint fund, they produced 2400 points for the joint fund. Thus all actors in the triad received 800 profit points ( $2400/3$  actors) from the joint fund (the one who did not make an exchange also received their 300-point default payoff). Finally, giving points to the joint fund when all others gave to the joint fund resulted in a total profit of 4800 points for the joint fund, yielding 1600 points ( $4800/3$  actors) for each actor—identical to the possible profits in the

reciprocal exchange structure, where all exchanging with all others resulted in a total profit of 1600 points for all (Lawler, Thye, and Yoon 2008).

Note that in the generalized exchange condition, each person could only benefit from one viable exchange (i.e., B could only benefit from A, and so on). As a result, the maximum benefit any actor could receive from an exchange partner in a round of generalized exchange was 800 points (plus the 300-point default profit if they chose not to exchange). On the other hand, those in productive and reciprocal exchange could benefit from both alters' exchange behaviors, and thus could receive benefits of up to 1600 points from others' exchange behaviors in a round. Put differently, because productive and reciprocal exchange entailed the possibility of exchanging with up to two alters in the triad, the maximum possible profit was by necessity lower in the generalized exchange condition. Lawler and colleagues noted this issue as well (2008, pg. 530). To resolve it (as Lawler and colleagues did), within the generalized exchange condition, we ran some sessions (N = 4 of the 11 generalized exchange sessions) where giving entailed either keeping 300 profit points or giving 1600 (rather than 800) profit points to the alter to whom one could give. That is, in this ancillary generalized exchange condition it was possible to earn 1600 points from successful exchange, just as it was in the reciprocal or productive exchange conditions.

All participants were informed that decisions would be simultaneous, made without anyone knowing what others had given before making the decision to give.

After everyone made their decision, they could see their profits. Round-level profits consisted of both the points they earned from the others they could receive points from, if they received them, plus the points they kept if they opted not to exchange. Then, they would go on to the next of several repeated rounds of exchange in the same triad. After the instructions, participants completed a set of quiz questions to check their understanding. Missed questions were followed by a detailed description of the correct answer.

Following each phase of repeated exchange – in which interactants were “treated” through their exposure to a level of cooperation and affective attachment produced through the structure and identity composition of their exchange network – participants then played two trust games – one as a sender making a trust decision, the next as a returner making a trustworthiness decision. Both decisions were made simultaneously, and could not be informed by the behavior of the partner. In all trust game decisions, the identity information of the interactants was made salient. Half of the trust games, participants possessed a shared social identity; in the other half, participants possessed different social identities.

As senders making trust decisions, respondents were given 300 points and instructed to decide how many points (from 0 to 300) they would share with their sender. Any points sent to the returner would triple in value. Senders were also told that returners would simultaneously decide what proportion of the tripled points they will

send back to the sender. Returners were provided the same instructions when making trustworthiness decisions and deciding what proportion of the shared points they would return to the sender.

### **3.3.3 The Identity Information Manipulation**

While reading the instructions for the basic exchange structure task, participants could see an example matching the exchange structure condition of their session, with nodes labeled A, B, and C (see the first column of Figure 3). They were informed that while completing the task with the two others, they would see an image very similar to the example in the study-wide instructions, except the image would contain the actual random letter identifiers assigned to themselves and others in the study with whom they were paired.

Additionally, in the two known identities conditions (homogeneous identities and diverse identities), participants were told that when the actual task began, they and the others would be identified with both their letter identifiers and colored circles: blue if the participant was a Klee and green if they were a Kandinsky. By contrast, those in the unknown identities condition were told that they would see a figure similar to the one in the example from the instructions during the phase: it would have black nodes containing their and the others' letter identifiers, but the Klee or Kandinsky identities of each participant in their structure would not be shown.

When the actual task began, in the unknown identities condition participants were paired without regard to identities and nodes were shown in black and white (first column of Figure 3). In the homogeneous identities condition, they were linked to two others in the session with whom they shared a Klee or Kandinsky identity. Thus, from their perspective they were embedded in a triad that showed either all-blue (see the second column of Figure 3) or all-green nodes to indicate that they had been paired with two other Klees (or Kandinkys). In the diverse identities condition, Klee participants were paired with either one or two Kandinsky alters, and vice versa (third column of Figure 3).

Though they were not told the number of rounds in advance, participants proceeded to complete 18 rounds of the decision-making task (as in Lawler, Thye, and Yoon 2000). We have already noted that occasionally, participants had technical issues that prevented them from completing the online study, or went idle for more than two minutes and were removed from the study. When that happened, the program filled in for the participant for all remaining rounds. Programmed responses entailed giving points if it was the first round of the phase, and, in all subsequent rounds, mimicking the others' exchange behaviors directed at the programmed actor in the previous round. In generalized exchange, this meant giving if the programmed actor had received points in the previous round. In reciprocal exchange, this meant giving to the other if the programmed actor had received points from that specific other in the previous round.

For productive exchange, the programmed actor gave if both others gave, did not give if both others did not give, and gave at a 50% chance if one other gave and one did not in the previous round. Put differently, pre-programmed actors engaged in tit-for-tat behavior. Participants were never aware if or when an alter had dropped out and was replaced with pre-programmed responses.

After 18 rounds of the task, participants answered a series of questions about their “experience with the two others to whom [they] were connected in the task” designed to measure affective attachment toward the triad (Lawler, Thye, and Yoon 2008).

Participants rated their experience working with the others in five items, each measured on a nine-point scale: bad-good, detached-attached, disloyal-loyal, disconnected-connected, and negative-positive. Participants would then participate in two rounds of the trust game described in the previous section with unmet interaction partners.

After these questions and the trust games, one phase was complete. Because information about identities was manipulated within the session, participants completed all three conditions, each in one of three phases. Participants did not know ahead of time that there would be multiple phases. They were informed only after the first phase that they would be assigned to a new triad. We emphasized (accurately) that they would be paired with new others, and not the two people with whom they had just interacted. This was followed by an image of the new triad to which they had been assigned—homogeneous identities, diverse identities, or unknown identities (Figure 3) and a new

set of decisions and questionnaire; this process repeated one more time for the third information about identities condition.

After three phases of the task—each consisting of 18 exchange rounds, the post-phase questionnaire, and two rounds of a trust game with an unmet interaction partner—the study was complete. Participants entered responses to open-ended questions about their perceptions of the study, including suspicion about deception, and read a debriefing screen. The study took about 75 minutes. After completing the session participants received an Amazon e-gift card, ranging from \$15 to \$20, depending on their actual earnings in the task.

### **3.3.4 Analytic Strategy**

In the repeated exchange phases, decisions to engage in exchange were binary choices for participants, but differences in how each exchange structure is arranged yield unique exchange opportunities for each type of structure. Those in the reciprocal exchange condition made two distinct decisions per round—whether to exchange with each of two alters. By contrast, participants in the productive and generalized exchange conditions made only one decision per round: whether to transfer points to the joint fund, or to the one other in the triad to whom they were able to transfer, respectively.

To compare giving behavior across these different settings (i.e., where ego made one decision or two decisions per round), we relied on one dichotomous outcome: giving to what we call here the focal beneficiary. In generalized exchange, the focal

beneficiary is the one other to whom the participant could give. In productive exchange, the focal beneficiary is the joint fund. In reciprocal exchange, the focal beneficiary corresponds to the structurally equivalent focal beneficiary in generalized exchange: the other that is counterclockwise in the “chain” (i.e., if ego is A, the focal beneficiary is B for both generalized and reciprocal exchange, see Figure 3). Thus half of the decisions were dropped from the reciprocal exchange condition to analyze one binary outcome in each round, as in productive and generalized exchange.

In exchange game analyses, data are nested in four levels: (a) experimental sessions that included (b) individual participants, each of whom completed (c) three phases; within these phases, participants made decisions in (d) 18 rounds of exchange. Because of this, we estimated our binary outcome—whether the actor engaged in exchange with the focal beneficiary or not in a given round—using generalized linear mixed models with random intercepts for the session, participant, and phase to account for the dependencies between the four levels of observation.

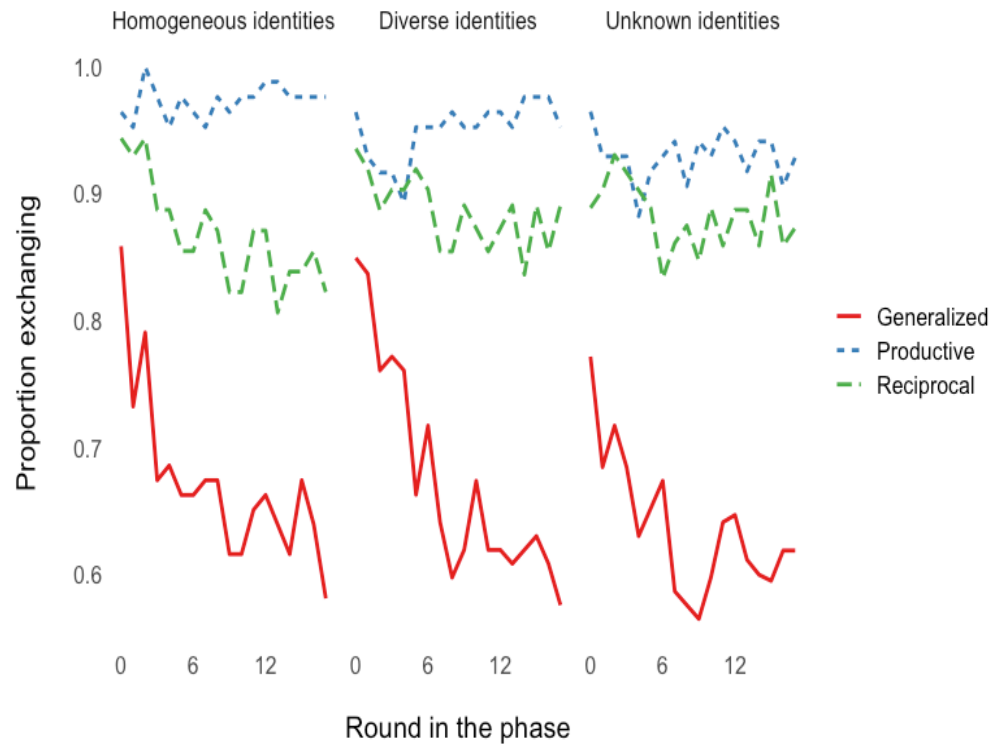
Trust game analyses conceive of the structure and composition of a network from a prior phase of exchange as treatment. As such, indicators of the structure of the previous exchange network and the number of out-group interactants the exchange network contained are constructed and used to predict trust and trustworthiness behavior. These analyses focus exclusively on network exchange structures where identity information was salient – the homogeneous and mixed conditions.

### **3.4. Results**

#### **3.4.1 Overall Giving Patterns in Exchange Networks**

First, I ask if differences in the amount of cooperation between productive, reciprocal, and generalized exchange structures are consistent with previous findings. In other words, is ATSE's empirical prediction robust to the inclusion of information about exchange partners' social identities? Figure 4 displays the proportion of participants who decided to exchange with the focal beneficiary across the rounds of the phase, based on our two key manipulations: exchange structure and information about identities. This figure reflects a pattern of giving that is remarkably consistent with previous work in ATSE across all three levels of our information about identities manipulation: giving was highest in productive exchange and lowest in generalized exchange. Figure 2 also suggests that differences in the patterns of giving behavior for each exchange structure emerged over time. We therefore controlled for round in all analyses to follow.

Our estimates underscore the pattern shown in the descriptive results. Table 1, Model 1 demonstrates that exchange structure was a significant predictor of giving to the focal beneficiary. In line with past work (Lawler, Thye, and Yoon 2008) and Hypothesis 1, giving rates were significantly higher in productive ( $B = 4.96, p < .001$ ) and reciprocal exchange ( $B = 3.44, p < .001$ ), compared to generalized exchange. Exchange patterns in the former two conditions also differed; the likelihood of giving to the focal



**Figure 4: Giving to the focal beneficiary over time, by the key study manipulations.**

beneficiary was higher in productive than in reciprocal exchange ( $B = 1.51, p < .05$ ).

These patterns, along with the others described further below, held after controlling for whether generalized exchange yielded a benefit of 800 (standard generalized exchange condition) or 1600 (ancillary generalized exchange condition) profit points to the other (see the Online Appendix, Table S1, Model 1). Additional ancillary analyses controlling for the demographic variables we collected (gender, age, and whether the participant was a University student or not), which did not impact results.

Our second key manipulation was whether own and others' social categories were known or not and, if so, whether the triad contained members whose identities were

homogeneous or diverse. Results show that when identities were known and those identities were homogenous (i.e., all exchange partners shared the same social category), cooperation was promoted in the exchange task. Exchange rates were higher in triads with homogeneous identities, compared to those in our first control condition—where identities were unknown ( $B = .57, p < .05$ , Table 3, Model 1). This finding is consistent with much past work that has found evidence for favoritism toward those sharing one's social category in cooperation. Exchange rates were also in the expected direction for the difference between homogeneous and diverse identities (our alternative control condition) triads, though the difference did not reach significance, possibly because, so that there was no deception in the real-time interaction task, our diverse identities triads by necessity included interactions with same-identity others.

Exchange rates toward focal alters in triads with diverse identity information (i.e., that included a mix of exchange partners with observable variation in identity categories) did not significantly differ from the unknown identities condition. Recall that each participant in a given diverse identities triad could either be a member of the “majority” (i.e., one other in the triad shared ego's identity, and the other did not) or “minority” (both others did not share ego's identity). We controlled for minority or majority status in these analyses to test whether participants experienced the manipulation differently because of it. But we find that exchange behaviors in diverse

identities triads were similar to those in unknown identities triads regardless of whether participants were in the relative majority or minority in their triad ( $p > .39$ ). Likewise,

**Table 3: Generalized linear models predicting giving to the focal beneficiary.**

All tests are two-tailed. \*  $p < .05$ , \*\*  $p < .01$ , \*\*\*  $p < .001$ . The reference categories are *Generalized exchange* and *Unknown identities*. *Round in phase* is coded such that 0 = the first round and 17 = the final (18<sup>th</sup>) round of the phase. Four-level generalized linear model with random intercepts for phases, participants, and sessions;  $N = 12,871$  decisions nested in 735 phases nested in 251 participants nested in 30 sessions.

	<b>Model 1</b>	<b>Model 2</b>
	<i>B (SE)</i>	<i>B (SE)</i>
<b><u>Fixed effects</u></b>		
<b>Productive exchange (P)</b>	4.96 (.64) <sup>***</sup>	4.78 (.74) <sup>***</sup>
<b>Reciprocal exchange (R)</b>	3.44 (.66) <sup>***</sup>	3.59 (.76) <sup>***</sup>
<b>Homogeneous identities (H)</b>	.57 (.28) <sup>*</sup>	.55 (.39)
<b>Diverse identities, ego is minority</b>	.37 (.43)	
<b>Diverse identities, ego is majority</b>	.16 (.49)	
<b>Diverse identities (D)</b>		.50 (.38)
<b>Round in the phase</b>	-.07 (.01) <sup>***</sup>	-.07 (.01) <sup>***</sup>
<b>P x H</b>		.85 (.71)
<b>P x D</b>		-.22 (.66)
<b>R x H</b>		-.57 (.65)
<b>R x D</b>		.15 (.69)
<b>Intercept</b>	2.18 (.45) <sup>***</sup>	2.16 (.47) <sup>***</sup>
<b><u>Random effects</u></b>		
<b>Phase intercept variance</b>	4.74	4.60
<b>Participant intercept var.</b>	6.23	6.24
<b>Session intercept var.</b>	.82	.85

those in diverse identities triads did not differ in their giving regardless of their minority or majority status ( $B = .16$ ,  $p = .74$ ). These results taken on the whole suggest that it is the

particular combination of having access to information about identities and those identities are homogeneous that facilitates cooperative exchange, supporting Hypothesis 2.

Model 2 in Table 3 includes interactions between our key manipulations. Because relative majority or minority position within the diverse identities condition did not differentially predict giving in Model 1, we collapsed across majority and minority participants in Model 2. We find no evidence that exchange structure or information about identities interact to differentially impact cooperation ( $ps > .23$ ). Instead, we retain Model 1 in Table 3. Our results show that both the form of exchange and shared social identities among exchange partners appear to be distinct, orthogonal factors affecting exchange behavior, failing to find support for Hypothesis 3.

### **3.4.1 Affective Attachment**

After the 18 rounds of each phase of the exchange task, participants completed a measure of affective attachment. We averaged across the participant's responses on the five items to produce a single measure of affective attachment for each participant in each phase (Cronbach's alpha = .96). Just as we controlled for round in the phase in the exchange task models, here we control for which measurement (i.e., phase) this was, of the three instances (once at the end of each phase) where we measured affective attachment.

Table 4, Model 1 reveals that overall—that is, regardless of our other key manipulation, information about identities— affective attachment was stronger in productive and reciprocal exchange, compared to generalized exchange ( $B = 1.91$  and  $B = 1.77$  respectively,  $ps < .001$ ). The difference between the former two exchange structures was not significant ( $B = .14$ ,  $p = .71$ ), suggesting that productive exchange and reciprocal exchange allowed for the development of similar levels of affective attachment. Likewise, regardless of exchange structure, identities also impacted affective attachment. Those in triads with homogeneous identities reported greater affective attachment compared to those where identities were unknown. They also reported marginally more affective attachment than those in diverse identities triads ( $B = .31$ ,  $p = .06$ ). Again, there was no difference in affective attachment between diverse identities and unknown identities triads ( $B = .01$ ,  $p = .96$ ).

Additionally, Model 2 in Table 4 demonstrates that differences in affective attachment in productive, reciprocal, and generalized exchange structures—and in homogeneous identities triads, compared to those with diverse identities or unknown identities— were mediated by the network-level proportion of successful exchanges made in the exchange task. We have already demonstrated that both exchange structure and homogeneous social identities facilitated cooperative exchange. Here we show that after controlling for the group-level proportion of cooperative exchange, the differences

**Table 4: Multilevel Linear Models Predicting Affective Attachment Following a Phase of the Exchange Task.**

All tests are two-tailed. \*  $p < .05$ , \*\*  $p < .01$ , \*\*\*  $p < .001$ . The reference categories are *Generalized exchange* and *Unknown identities*. Three-level multilevel linear model (maximum likelihood estimation) with random intercepts for participants and sessions;  $N = 702$  measurements nested in 249 participants nested in 30 sessions for each model.

	<b>Model 1</b>	<b>Model 2</b>
	<i>B (SE)</i>	<i>B (SE)</i>
<b><u>Fixed effects</u></b>		
<b>Productive exchange</b>	1.91 (.35)***	.05 (.20)
<b>Reciprocal exchange</b>	1.77 (.37)***	.38 (.20)
<b>Homogeneous identities</b>	.32 (.16)*	.10 (.11)
<b>Diverse identities</b>	.01 (.16)	-.10 (.11)
<b>Proportion giving during the exchange task</b>		6.44 (.25)***
<b>Phase 2</b>	.18 (.16)	.02 (.11)
<b>Phase 3</b>	.44 (.16)**	.23 (.11)*
<b>Intercept</b>	5.02 (.27)***	.96 (.22)***
<b><u>Random effects</u></b>		
<b>Participant intercept var.</b>	.98	1.05
<b>Session intercept var.</b>	.41	.00

in levels of affective attachment following the exchange task become non-significant.

Both structure and same-identity bias affected cooperative exchange rates in our experiment; in turn, these cooperative exchange rates promoted higher levels of affective attachment within the triad. These findings have implications for future work on cooperative exchange in homogeneous and non-homogeneous group settings, to which we turn in the concluding section.

### 3.4.2 Future Trust and Trustworthiness by Social Identity

Following each phase of repeated exchange and affective attachment questions, respondents then played two rounds of a one-shot trust game – one as a sender to measure trust, one as a receiver to measure trustworthiness – with an unmet interaction partner who was not a participant in their previous phases of repeated exchange. Table 5 reports the results from trust decisions (Model 1) and trustworthiness decisions (Model 2) made toward unmet alters.

Hypothesis 4, implied by ATSE, proposes that productive exchange structures will lead to the highest levels of trust and trustworthiness toward unmet alters, followed by reciprocal and then generalized exchange. Model 1 shows that the reality is more complicated. Trust behavior is significantly higher after phases of repeated exchange in productive (0.122,  $p < 0.01$ ) and reciprocal (0.187,  $p < 0.001$ ) exchange compared to the generalized exchange conditions, but productive exchange does not clearly supercede reciprocal – in fact, while differences between the estimates mentioned between productive and reciprocal exchange are not significant, reciprocal exchange has a larger estimated main effect on future trust behavior. Trustworthiness (Table 5, Model 2) – which was hypothesized to be the more sensitive measure due to the lower risk associated with the decision – does *not* differ significantly by prior exposure to different structures of repeated exchange.

Hypothesis 5 proposes that actors who were previously embedded in social exchange networks that contained multiple interaction partners from one's relative out-group social identity will engage in more trust and trustworthiness behavior when they encounter unmet out-group alters in the one-shot trust game, compared to those from homogeneous exchange network phases. Model 2 in Table 5 provides support for this hypothesis in terms of trustworthiness. While the model reports a negative main effect of having an out-group partner as a returner making a trustworthiness decision, participants who were embedded in diverse prior exchange networks that contained two out-group alters were significantly more likely to give more to an out-group trust game partner compared to those who were in diverse triads with only one out-group alter or homogeneous information exchange structure conditions (0.093,  $p < 0.05$ ).

The results show that being involved in repeated exchange with more than one dissimilar other leads to greater generalized beliefs about trustworthiness that are enacted behaviorally toward members of that out-group social category, consistent with hypothesis 5. I do not find the same relationship in the trust model, which is a higher risk decision. This indicates that, after the brief phase of repeated exchange, in-group preferences continue to shape consequential giving behaviors toward unmet others.

**Table 5: Linear Mixed Models Predicting Trust and Trustworthiness following Repeated Exchange Phases**

All tests are two-tailed. \*  $p < .05$ , \*\*  $p < .01$ , \*\*\*  $p < .001$ . The reference categories are *Generalized exchange* and *Unknown identities*.

	<b>Model 1</b>	<b>Model 2</b>
	<b>Trust</b>	<b>Trustworthy</b>
	<i>B (SE)</i>	<i>B (SE)</i>
<b><u>Fixed effects</u></b>		
<b>Productive Exchange</b>	.122 (.04)**	.003 (.03)
<b>Reciprocal Exchange</b>	.187 (.04)***	-.036 (.04)
<b>Trust Game: Out-Grp Sender</b>	.044 (.03)	-.052 (.02)*
<b>Trust Game: Prop Sent as Truster</b>	.486 (.06)***	.270 (.03)***
<b>Prior Phase: One Out-Grp Alter</b>	.003 (.03)	-.030 (.02)
<b>Prior Phase: Two Out-Grp Alters</b>	.009 (.04)	-.032 (.03)
<b>Phase 2</b>	-.015 (.02)	-.046 (.02)**
<b>Phase 3</b>	-.020 (.03)	-.043 (.02)*
<b>Prior-One x Trust-Out</b>	-.010 (.05)	.038 (.03)
<b>Prior-Two x Trust-Out</b>	-.016 (.06)	.093 (.05)*
<b>Intercept</b>	.363 (.04)***	.348 (.03)***
<b><u>Random effects</u></b>		
<b>Participant intercept variance</b>	.053 (.23)	.038 (.19)
<b>Session intercept variance</b>	.033 (.18)	.017 (.13)

### **3.5. Discussion**

The overarching goal of this study is to understand the role of same-identity bias in the context of different repeated network exchange structures, as well as future commitment behaviors toward unmet out-group others. Linking these two known determinants of cooperation and trust by manipulating both, I demonstrate that same-identity bias and the jointness of repeated exchange independently affect the emergence of social order, even in the presence of each other. Moreover, I show that features of

one's prior exchange network context can shape future trust and trustworthiness behaviors, in ways that attenuate in-group preferences in decision-making about the trustworthiness of an unmet person from an out-group.

My results clearly demonstrate that ATSE's structural proposition – that the jointness of an exchange task affects the frequency of cooperation and the nature of affective commitments toward one's exchange partners – is robust to the addition of social identity information. Productive exchange structures yielded the highest levels of cooperation, followed by reciprocal exchange, with generalized exchange leading to the lowest levels of cooperation (replicating Lawler, Thye, and Yoon 2008). These patterns occurred regardless of whether the interactants had access to social identity information or not, and, when it was available, whether identities were homogeneous or diverse. Our competing hypothesis, that identity information might suppress the expected differences in the frequency of exchange emerging between the exchange structure types, was guided by prior work suggesting that same-identity information may reduce uncertainty and risk in those exchange structures that are particularly risky. But our results clearly show that the structural arrangement of repeated exchange remains powerfully influential, independent of identities.

Our findings also show that exchange structure should be a consequential consideration for interdisciplinary research on same-identity bias, especially in studies that explore levers that mediate or moderate such biases in enduring and diverse

groups. Studies of intergroup contact theory (see Pettigrew and Tropp 2006), for example, tend to show that groups of people with different social identities can be cooperative and cohesive when members of diverse groups are made to feel equal in their ability to contribute cooperatively toward a shared goal that a presiding authority encourages. But it is unclear whether some or all of these conditions need to be present for the anticipated effects to emerge, or exactly how these behavioral scope conditions induce psychological and emotional orientations toward cooperating with dissimilar others.

Research from theories of social exchange such as ATSE develop a formal social-psychological mechanism to explain the impact of well-identified exchange structures on social dilemmas that are analogous to the nature of tasks found in studies of intergroup contact or diverse teams in organizations. By integrating insights on the relationship between exchange structure and affective commitments, we propose that ATSE provides researchers who study diverse group dynamics with a key explanation of the structurally-induced psychological processes that help generate the sometimes-mixed empirical outcomes of these research traditions.

## **4. Manipulating Initial Network Characteristics and Measuring Changes in the Emergence of Homophily: a Computational Experiment**

### ***4.1. Introduction***

Social ties tend to form between people with matching traits. This phenomenon, known as homophily, is among the most remarkable and robust empirical patterns observed in sociology (Lazarsfeld and Merton 1954; McPherson, Smith-Lovin, and Cook 2001). Thousands of analyses of observational data across many disciplines and in a variety of organizational settings (McPherson, Smith-Lovin, and Rawlings 2021) have shown that people are almost universally share observable characteristics more than chance would otherwise predict.

This paper uses a computational experiment to ask a simple question: how might an organization intervene to attenuate the emergence of homophilous clustering within its constituents? Many studies have used empirical data from longitudinal networks to provide compelling answers to this problem. Moody (2001) shows, for example, that increasing heterogeneity at the organizational level can actually increase segregation, as people can encounter enough similar others to satisfy their social needs by forming ties primarily within their own groups – unless opportunities for social interaction and tie formation are constrained to smaller (yet still heterogeneous) interaction foci of social interaction. Yet explanations of homophily in organizational settings are constrained by

the nature of available empirical data. Longitudinal network data in schools and organizations as they occur organically are inherently segregated.

By contrast, leaders of organizations often have the ability to manipulate the structure and connectivity of people within them by assigning people to smaller task groups and teams, or partitioning newcomers into discrete cohorts. Such decisions would not only decrease the degree of homophily in an organization's network at the outset: because patterns of homophily are induced in part by actors' inherent preferences for transitivity (i.e., to form ties with the friends of their friends), such manipulations would increase the probability that the friend of one's friend either does not yet exist or is a person who possesses an out-group social identity category.

In this paper, I use an agent-based simulation to investigate how initial network configurations can shape the emergence of segregation in a hypothetical organizational setting. In the simulation, I emphasize two key manipulations that exploit transitivity preferences to maximize the probability of cross-group tie formation (or minimize the probability of in-group tie formation): whether an organizational network contains focal cliques with heterogeneous actors – diverse teams – at the outset, and whether that initial network is already segregated or fully disconnected outside of the cliques.

## **4.2. Background**

Though many studies document its ubiquity, research on homophily often treats the pervasiveness of the descriptive pattern as its own explanation (Fiel 2021). The

causes of homophily are broadly assumed to arise from an interaction between the preferences of individuals for similar people and the social structures in which these preferences unfold. The former mechanism – often called *choice homophily* or in-group preference, comes from classical findings in social psychology that sharing attributes reduces perceptions of uncertainty about and challenges of interacting with others (Allport 1954; Festinger 1957; Werner and Parmelee 1979; Hamm 2000). As a result, trust and solidarity tends to form more easily between similar people as opposed to people who inhabit different observable social categories (e.g., Portes and Sensenbrenner 1993). Findings about in-group preferences are well-documented in dyadic interactions, and hold compelling face validity when they are assumed to operate identically in networks by generating homophily from the ground up. After all, birds of a feather flock together (Turner 1545), and “everybody knows that we like what resembles us, those who think and feel as we do” (Durkheim 2014:45).

The latter cause – often called *induced homophily*, indicts the structural features of a social system itself in the emergence and reproduction of homophilous ties. People can only form ties with other people with whom they are at risk to come into contact with (Blau 1977; Feld 1981). The risk set that one finds themselves in – the proportion of similar and dissimilar actors in their network (Schelling 1971), the distribution of these actors within the broader social structure (Blau 1977), and their connections to implicit and explicit foci such as teams, classrooms, and shared activities or interests (Feld 1981;

1982) – constrains their opportunities for new tie formation. The key mechanism through which social structure is believed to generate induced homophily is known as transitivity, the phenomenon where two individuals connected by a third party tend to become connected themselves (Rapoport 1953; Holland and Leinhard 1971).

Longitudinal network analyses show that an individuals' inclination toward transitivity lead them to form ties with the friends of their friends, who are already more likely to be situated in the social space within foci (i.e., in structural proximity to) to which this individual is connected (McPherson 2004; Kossinets and Watts 2009).

Along these lines, sociologists investigating the structural determinants of homophily in organizations show that variation in presence of opportunities for mixing pattern the extent to which homophilous dyads form, and show that constraining such opportunities by maximizing heterogeneity within explicit foci such as teams, extracurricular activities, and even seat assignments at lunch tables or classrooms can lead to more ties between people who possess different social categories through transitive closure (Moody 2001; Rohrer et al. 2021).

These studies indicate that manipulating the structure of a social network should affect – and could attenuate – levels of induced homophily. In most situations, focal arrangements and the broader network structure that lead to induced homophily are inherited, not generated, by individuals who enter an established network context such as a school or professional organization (Kossinets and Watts 2009; Fiel 2021). But

manipulating the structure of an organization's network in a way that alters these otherwise-inherited and minimal inter-group mixing opportunities through explicit foci (such as diverse cliques and friendship circles formed prior to their incorporation into the broader network) could lead to different levels of induced homophily.

In turn, persistent contact with dissimilar others through manipulations that constrain and maximize mixing opportunities within networks may in turn affect the degree of choice homophily – in-group preferences and inter-group behaviors – that individuals possess within themselves. While historically viewed as a fixed attribute, recent work emphasizes the relationally constructed nature of the in-group preferences that cause choice homophily (e.g., Fiel 2021). Indeed, adolescents with more racially diverse ego networks in secondary school contexts are far more likely to be married to racially dissimilar partner as adults (Kao et al. 2019); White college students assigned to racially dissimilar roommates develop more diverse friend groups and report lower levels of anxiety when interacting with Black strangers (Gaither and Sommers 2013); and repeated exchange in categorically diverse triads has a positive impact on future trustworthiness behavior toward unmet out-group alters (Harrell and Quinn 2022).

In short, while choice homophily and induced homophily provide a compelling account of a durable social fact, taking these explanations seriously as micro-level mechanisms requires sociologists to jointly operationalize them as such, while simultaneously linking them to the structural features of networks that are used to

explain homophily – such as inherited segregation (McPherson 2004; Kossinets and Watts 2009) or pre-existing segregation into focal domains that correlate with social category group membership (Blau 1977; McPherson et al. 2001; Wimmer and Lewis 2010). Such an approach allows researchers to investigate homophily as an emergent and property of social systems that is both malleable as an outcome and relevant to the relational construction of in-group preferences – and subsequent network-level homophily.

Doing so is very difficult, if not presently impossible, in an empirical context. While longitudinal network data has allowed researchers to identify patterns linking social structure to homophily in the form of downstream segregation, virtually all observable networks are homophilous to some degree at the outset of inquiry. As a result, plausible initial network configurations – such as a network in which no agents have yet formed ties, or where the only ties formed prior to the network’s evolution are assigned by an authority figure who presides over a given organizational space – may have substantial consequences for the emergence of homophilous ties, but lack a zone of common support for exploration in empirical longitudinal network data. Moreover, in-group preferences are typically assumed to be the excess attachment left unexplained by induced homophily terms in a model, and are not captured as the individual-level (or relational – see Fiel 2021) psychological constructs that they are. Present empirical studies are not only limited by a researcher’s capacity to operationalize the mechanisms

with available data; they also provide top-down explanation for a bottom-up process out of necessity. In short, the data that are required to develop an account of how social structures co-evolve from choice and induced homophily do not currently exist.

This paper develops an agent-based model to investigate homophily as a ground-up emergent phenomenon of social systems. The model embeds agents from two different social groups within the network of a hypothetical organization, and aims to manipulate three groups of parameters. First, the model manipulates the *initial* structure of the network space in which agents interact so that it (a) includes or excludes diverse cliques as social foci, and (b) allows the graph to be fully disconnected connected or segregated – allowing for transitive closure opportunities to either align with or conflict with opportunities to form in-group ties. Second, it varies the propensity of all agents in a given condition to engage in choice homophily (i.e., one’s affinity for in-group ties) and the closure component of induced homophily (i.e., one’s propensity to become friends with the friends of their friends). Third, and absent from the present set of results, the model will manipulate whether or not agents who form and maintain ties with out-group alters develop attenuated in-group preferences.

Across the suite of simulation conditions, I find that the structure of the initial network can be highly influential in determining the degree to which homophily emerges in an organization. When initial network configurations contain social foci in the form of diverse cliques of agents, tie formation occurs more often between dissimilar

alters through the transitivity mechanism – even when in-group preference is high. I also find that, while the choice homophily parameter (subsequently called in-group preference) and the transitivity component of induced homophily (subsequently called closure preference) interact to produce homophilous systems, closure preference is most influential only when in-group preferences are also high, or when the social system has a great deal of pre-existing segregation in its initial network configuration. Consistent with the implications of contact theory (Pettigrew and Tropp 2006), my results provide evidence to suggest that leaders of organizations can attenuate segregation in their constituent networks by assigning heterogeneous teams – or by partitioning actors into cohorts to decrease the probability that transitivity preferences lead to homophilous ties at the outset of the cohort's integration. These findings also provide a ground-up demonstration of the cognitive and structural plausibility of mechanisms used to explain the homophily principle by top-down models of social ecology.

In a future phase of the study, I hope to show that operationalizing in-group preference as a relationally learned construct further attenuates emergent homophily through socialization, or the retooling of agents' cultural meanings and affective commitments about out-group interaction partners. Upon completion, this agent-based model will include updateable in-group attachment preferences, and represent an emergent account of homophily that conceives of in-group preference as a dynamic

property shaped by ongoing ties and structurally induced opportunities for inter-group interaction.

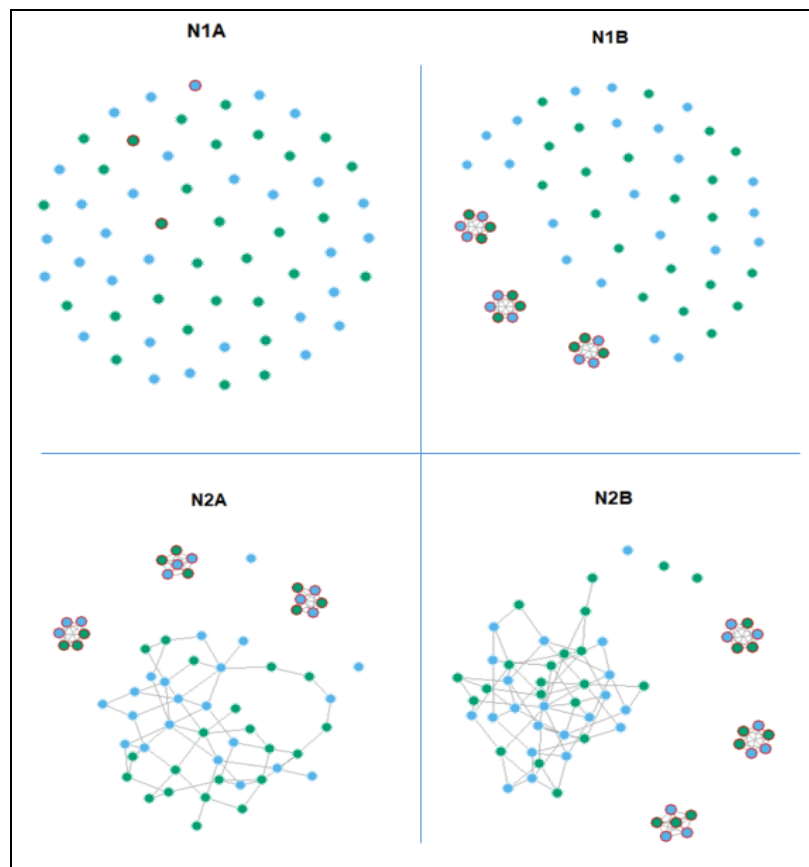
### **4.3. Methods**

#### **4.3.1 The Organizational Setting**

Prior agent-based simulations modeling the emergence of segregation conceive of the social space as a lattice where positions represent physical proximity (Schelling 1971; Bruch and Mare 2006) – primarily because these models attempt to imitate the process of residential decision-making, or choosing neighbors. By contrast, the agents in the computational experiment in this study choose social ties. The environment in the model affects the selection of ties based on the ongoing structure of relations in the environment at any given state. The environment in the model does *not* affect the selection of these ties through proximity in physical space: the setting is constrained to 60 agents, allowing me to plausibly assume that each agent has the ability to reach any other agent as a potential interaction partner.

I vary the structure of the organization in two ways (shown in Figure 5): the broader patterns of connectivity in the initial (i.e., when time=0) configuration of the social space, and the presence or absence of explicit foci in the form of diverse cliques embedded within that initial configuration. In the N1A factorial condition shown in Figure 5, the initial network space (i.e., the structure of the network prior to any decision-making or evolution) is a fully-disconnected graph. The N1B initial network

condition is also a fully disconnected graph, with the exception of diverse cliques: three fully-connected clusters of six agents, three from one identity group and three from another, have already formed ties prior to the start of the simulated organizational dynamics. This condition is analogous to a strategic decision on the part of an organization such as an orientation program for a subset of the broader population.



**Figure 5. Initial Network Configurations for the Simulation Experiment.**

By contrast, the N2 factorial conditions include nominal degree of initial connectivity in the social space. This connectivity is randomly generated through either an Erdos-Rényi model with a new tie probability of 0.01, or a conditional Erdos-Rényi

that induces in-group favoritism in the initial set of ties with a mixing matrix of 0.05, 0.025, 0.025, 0.05. The N2B condition contains cliques that exist in junction with initially segregated connectivity; the N2A condition includes cliques that co-occur with a randomly connected graph.

### 4.3.2 The Agents

Each of the 60 agents make a set of decisions in each state, or interval of time, over which the network evolves. All agents are assigned to one social category or another. Within any given run of the simulation, all agents within the run are given one constant value for in-group tie formation preference, and another constant value for the closure preference term that varies independently from the in-group preference term. These values are described in more detail in the following section; the unique values used to produce the results reported in this study are shown in Table 6.

**Table 6. Parameterization of the Agents and Social Space**

Group	Parameter	Level of Variation	Values
Agent Parameters	In-group preference	Network	0.0   0.5   0.9
	Closure preference	Network	0.0   0.5   0.9
	Social category	Agent	0 [blue]   1 [green]
	Maximum number of ties	Agent	$X \sim N(6, 1.75)$
	Pr seeking out an unmet interaction partner	Agent	$1 - \text{logit}(\text{centered "max n ties"})$
Space Parameters	Number of agents	Network	60
	Number of agents from social category 1	Network	30
	Diverse focal cliques	Network	0 [no]   1 [yes]
	Focal clique size	Network	6
	Any initial connectivity	Network	0 [no]   1 [yes]
	Segregated initial connectivity	Network	0 [no]   1 [yes]
	Pr initial within-group ties (if segregated = 1)	Network	0.05
	Pr initial between-group ties (if segregated = 1)	Network	0.025
Pr initial ties (if segregated = 0)	Network	gnp(0.01)	

Each agent was assigned a maximum number of ties before they cease to seek out new ties in their network. This value was determined for each simulation run by 60 draws from a random normal distribution with a mean of 6 and a standard deviation of 1.75. The result was rounded so that each agent's maximum number of ties was a whole number.

### **4.3.2 Decisions Linking Agents to the Organization**

The first decision agents make in each iteration of a simulation run is whether to interact with one of their existing ties or a new interaction partner during this state. This decision is determined by a logistic function modeling the centered difference between the maximum number of ties an agent may want in the system (i.e., their extroversion or sociability) and the number of ties they currently have in a given state. When agents have very few ties relative to their personal maximum number of ties, they are much more likely to opt to interact with an alter with whom they have not yet formed a tie. When agents have a number of ties that is close to their maximum number of desired ties, they are much more likely to seek a potential interaction partner within their own set of existing ties.

Once an agent has determined whether they are going to meet someone new or interact with their existing partners in a state, they make their second decision: who within the set of potential interaction partners (i.e., the batch determined by decision 1) will they interact with. This is determined by two parameters: in-group preference

("ingrpboost") and preference for triadic closure ("closureboost"). I vary the values for both across three different levels: 0.0, 0.5, and 0.9. A value of 0.5 means that, within the set of potential interaction partners, all of the in-group agents are 50% more likely to be selected than the out-group agents (relative to the group identity of the agent making the decision). This principle is also true for the closure term, biasing agents in favor of seeking transitive ties. I then combine these probabilities, so that each potential interaction partner agent is represented by one joint probability of being selected, and then I choose based on these probabilities within the decision-making agent's choice set.

Finally, agents make a third decision – whether to form a tie with their potential interaction partner. Conditional on choosing an interaction partner that one has not yet met, this probability is fixed at 0.5. For agents who opted to interact with an existing tie, the probability of maintaining the tie is set at 0.95. This value allows the model to converge on a stable outcome while maintaining sociometric plausibility, as ties tend to be more readily made than dissolved<sup>7</sup>. To prevent the model from converging on a maximally connected hairball after one state of the simulation, I randomly select half of the agents to implement the decision, while the other half opt not to interact with any agents during the present state.

---


<sup>7</sup> Note that tie dissolution may be marginally asymmetric when modeled as a function of pairwise similarity in empirical networks (e.g., Kossinets and Watts 2009). The present model abstracts this difference away for simplicity, as the difference in dissolution frequency found in empirical friendship networks tends to be extremely small when comparing highly dissimilar pairs to highly similar pairs.

The simulation then repeats for 20 states. I run the simulation 10 times within each parameter specification of the initial network configuration, and for differences in the in-group bias and closure-bias terms. I then compute mean assortativity for the initial and final states of the graph, take the difference to estimate the change in assortativity, and average this difference across all runs of a simulation for a given parameter specification set to compute the mean assortativity – or homophily – that emerges from each state space. Minor differences in the initial specification of the network yield a resulting data set that includes the outcomes for 691,200 runs of 20-state simulations repeated 10 times each for the varied initial parameter specifications of in-group bias, transitivity bias, and initial network structure. The analysis file contains one observation for each parameter specification at which the simulation was run. Runs contain 3 in-group preference levels, 3 closure-propensity levels, 2 conditions for the presence or absence of diverse focal cliques, 2 conditions for connectivity in the initial network, 2 conditions for whether that initial connectivity is random or segregated (if it exists) – 720 summary observations in total.

#### **4.4. Results**

First, I compare the effects of the in-group bias term to the transitivity bias term across all runs of the simulation. Figure 6A shows that, even in the presence of all types of variation in the transitivity bias and initial network configuration in the simulation, in-group preference has a strong naïve signal. By contrast, the parameter specifications

representing slight variations of the networks from figure 1 nearly indistinguishable differences at different levels of transitivity bias (Figure 6B).




**Figure 6. In-Group Bias (left) and Transitivity Bias (right) Plotted Against Change in Mean Assortativity in Simulated Networks.**

Notes: Blue = 0.9; green = 0.5; red = 0.0

Transitivity preferences appear to have a weaker effect on emergent change in assortativity, in part because the N1 conditions do not provide an initial network structure that allows for transitivity to operate at the outset. In graphs where initial ties between actors outside of the focal cliques do not yet exist, transitivity preferences cannot shape the formation of ties, as no agents (outside of the diverse focal cliques) are

in the risk set of structural opportunities for closure. In other words, there are not yet any friends of friends with whom an agent can become friends. Looking at networks where in-group preferences are maximized allows us to begin to see differences in the effects of different levels of the transitivity preferences. This is likely due to the fact that high levels of in-group preference rapidly lead to highly segregated ties that transitivity can then exploit to compound segregation.




**Figure 7. Transitivity Bias Levels Plotted Against Change in Mean Assortativity in Simulated Networks, Where In-Group Bias Levels = 0.9**

Notes: Blue = 0.9; green = 0.5; red = 0.0

I specify a generalized linear model to estimate the effects of the different computational conditions on changes in mean assortativity, focusing on the role of initial

network configurations on the change in assortativity – the emergence of homophily in a network. The results of the model are represented visually in Figure 8 to demonstrate the hypothesis that the presence of diverse cliques allows transitivity bias to suppress some of the effects of closure bias.



**Figure 8. Difference in Mean Assortativity Change Across In-Group and Transitivity Preference Levels, for Networks With and Without Diverse Cliques**

The x-axis of Figure 8 shows three discrete categories from left to right: simulation runs where in-group preferences did not exist; runs where in-group preference levels were set to 0.5; and runs where in-group preference levels were set to 0.9. The y-axis values represent the average change in mean assortativity. Unsurprisingly, network assortativity does not occur in the absence of in-group preferences, and increases alongside increases to in-group preference.

The three lines shaded red – one for each level of transitivity preference – represent the average mean assortativity change differences across in-group preference levels for simulation runs that do *not* contain diverse focal cliques (i.e., N1A networks at low, medium, and high values of transitivity preference). The blue lines represent the same metric at three different levels of transitivity preference in networks that are initialized with diverse focal cliques (N1B, N2A, and N2B variants). The plot shows that incorporating diverse cliques into the social space prior to the evolution of the network (i.e., the networks represented by the blue lines) reduces the emergence of homophily compared to graphs that do not, doing so even more as in-group bias increases. In other words, I find support for the hypothesis that initial network configurations can affect the emergence of homophily by exploiting the induced homophily mechanism of transitivity propensity. This finding persists even when the degree of in-group preference in attachment – the choice homophily mechanism – is set to arbitrarily high values.

#### **4.5. Discussion**

Many studies of homophily emphasize the ubiquity and inevitability of homophilous clustering in social systems. Prior research that investigates how structural mechanisms can moderate the degree to which in-group preferences lead to segregated outcomes still use empirical network data from networks that are already segregated to varying degrees at their initial state of analysis.

My simulation looks at the same mechanisms from a different perspective: structural mechanisms of induced homophily are neutral, and merely amplify existing patterns of connectivity and preference. As organizational leadership often have the capacity to manipulate these patterns, networks of actors within them could consequently have more friends who have more out-group friends (or fewer friends who have more in-group friends) and make transitive tie formation choices that result in lower levels of network segregation.

The results of my simulation support this proposition. Transitive closure and the presence of foci cause lower levels of network assortativity when I arrange networks to be connected differently before I begin to observe their evolution. In other words, structural mechanisms are agnostic to identity-preferences. They amplify in-group preferences, but they also amplify the patterns that organizations create when they assign teams, create groups for activities, or decide to limit opportunities for actors to connect with clusters that might already be segregated within them. Moreover, transitivity preferences only become relevant to homophilous clustering in more connected graphs, and do not systematically pattern homophily across a variety of initial network configurations in the same way that in-group preferences do.

These findings echo prior empirical work on network ecology and segregation, which show that organizations which restrict the scale and identity composition of explicit foci like clubs, teams, activities, and academic tracks converge in ways that

include more ties between people from different social categories (e.g., Moody 2001; Kossinets and Watts 2009; Wimmer and Lewis 2010). My simulation also expands on these insights with treatment conditions that explicitly manipulate network structures in homophily-attenuating ways. These conditions vary levels of initial network segregation or connectivity that are, while plausible for many organizations to implement, not observed in available longitudinal empirical network data. Doing this in a ground-up simulation allows me to show that reductions in homophily found in the network ecology literature could be even more substantial, while also eliding concerns about causal order from non-manipulated empirical contexts.

This study of emergent homophily emphasizes that viable manipulations to initial network structure – assignments that place actors in diverse cliques, and minimizing the extent to which the initial network is already segregated (through, for example, partitioning new-coming actors into discrete cohorts) – can reduce the extent to which homophilous sorting emerges in an organization. The results hold even when the simulated agents in the organization possess very strong preferences for in-group tie formation. Organizations looking to increase connections between people from different groups – because of a mission and vision of creating integrated contexts, or because of an interest in the benefits of diverse ties for new idea formation, search, and discovery, may be able to do so by creating heterogeneous focal cliques, and by limiting the extent to which new agents can reach segregated clusters to begin with.

In the next phase of this ongoing project, I plan to include a social learning component to the model that reflects the lessons from Chapter 3, which shows that behavioral tendencies toward unmet out-group alters can change based on the social category composition of a person's ties in a network where they repeatedly interact. This addition will allow me to investigate the implications of a relationally constructed in-group preference within network contexts, and how organizations may have the capacity to act as sources of cultural belief change about social categories.

## 5. Conclusion

The studies in this dissertation collectively suggest that the forces that shape out people perceive, interact with, and form ties with dissimilar others are malleable, and manipulating these mechanisms in settings over extended periods of time could attenuate in-group biases and patterns of segregation that are often viewed as social problems in modern civil society.

Broadly speaking, my dissertation explores how environments and network contexts can shape individual-level beliefs about and interactions between actors who inhabit different social identity categories, and how these beliefs and behaviors can in turn shape broader patterns of social structure. Each study contributes to sociological areas of inquiry in social psychology, culture and cognition, and network theory; when taken in combination, they demonstrate the role of beliefs in shaping micro-level behaviors and social interactions, the viability and potential causes of belief change about people from different groups, and the extent to which beliefs about and behaviors toward others influence, and are influenced by, the structure of social systems.

The results of the study presented in Chapter 2 show that belief change about social identities can accompany an event which places individuals into persistent structural conditions that make certain identities – and our potentially-poor-fitting prior beliefs about those identities – particularly salient. To sociologists of culture and cognition, this finding suggests that, while the cohort replacement control model of

belief stability likely dominates in stable social contexts, explanations of cultural change that involve individuals updating their beliefs cannot be dismissed without more thought about the mechanisms by which we might expect such changes to occur. To answer the question of whether cultural belief change can occur in adulthood, social scientists must test psychologically informed hypotheses about the *causes* of such changes rather than characterizing them as edge cases. More broadly, this finding reflects the power of persistent social conditions in shaping our beliefs about social identity categories in ways that persist beyond the presence of the conditions themselves. Given that the salient social identities in the data were primarily working class occupations, the study also suggests that environmental changes which make us persistently depend upon, valorize, or experience concern for people who inhabit social identity categories that are not in high status positions in society can impact the ways that we perceive and interact with these actors.

The results of the experiment in Chapter 3 provide a mechanism for changes in beliefs and prosocial behavior toward different social category groups that does not require a manipulation at the magnitude of a global pandemic. Persistent repeated exchange in joint tasks leads to higher levels of cooperation than tasks that are more separable and fail to inculcate a sense of shared responsibility amongst interaction partners, the pattern holds in both homogeneous and diverse networks. Most importantly, being exposed to multiple unmet out-group others in repeated exchange

led people to find members of their out-group generally more trustworthy – even people they had not yet met.

The study fills an important gap in the literature on social exchange by incorporating well-established findings on in-group preferences from across the social sciences. Results show that actors make behaviorally consequential meaning from the minimal group identity information in ways that affect the outcome of repeated network exchange. Within each type of exchange structure, I find significant differences in the emergence of cooperation frequency between partners embedded in triads with homogeneous social identity and diverse social identity compositions. At the same time, I also show that structural features of the triadic context of each exchange structure still has strong causal effects, even for diverse triads: ATSE's predicted differences in the degree of cooperation between the different network exchange structures are robust to all social identity information conditions, and the inclusion of identity information does not moderate or suppress this powerful structural effect.

In the trust game following each phase of the experiment, people give significantly more monetary units when making trustworthiness decisions to unmet alters who do *not* share their social category after phases of repeated exchange in diverse-identity triads compared to homogeneous-identity phases. Both findings imply that the structure of repeated exchange may impact the extent to which people in diverse groups may cooperate, and suggest more broadly that people can learn to

cooperate with, cultivate social cohesion with, and develop trust toward people they are not categorically similar to. In other words, the results provide the foundation for a set of questions about whether people can learn to cooperate with trust people who inhabit different social categories just as much as people within their in-group through emotional commitment and social learning.

Beyond the contribution to the literature, these results have important substantive implications. People cooperate the most with each other when their efforts are centered around a task in which their individual contributions are minimally apparent, and they all feel a sense of shared responsibility – even when these people are members of different identity groups. Moreover, being persistently involved in any form of diverse repeated network arrangement with dissimilar others leads people to become more likely to see members of their outgroup as more trustworthy. Coaches and teachers, for example, can leverage strategies informed by these results not only to cultivate cohesion and community within their classrooms and on their fields, but to cultivate broader notions of trust and cooperation that their students and athletes may transpose when interacting with people from different social identity groups beyond the school and sport domains.

The computational experiment in Chapter 4 investigates how the structural conditions and mechanisms that sociologists use to explain homophily can be leveraged to reduce homophily in organizational networks. By manipulating initial network

connectivity between agents of different social categories in an agent-based model, I demonstrate that extreme levels of homophily tend to emerge more often in systems where social ties are already highly segregated at the outset, and find that placing pre-established heterogeneous cliques into an existing network increases the probability that triadic closure will produce new ties between alters with different identities, thereby decreasing the degree of emergent homophily. This finding has important implications for school leaders, directors of programs in workplaces, and policy-makers whose decisions impact the ways in which people from different groups mix in schools, offices, and other organizational contexts. Seeding these spaces with diverse social ties in the form of structures such as orientation groups and long-term project teams, in conjunction with limiting the exposure that not-yet-segregated parts of an organization have for connecting with components that are already connected in segregated ways, could attenuate the degree to which homophilous sorting emerges in these spaces.

The contents of these chapters connect to my broader interest in developing sociometrically plausible models of the processes through which individual-level beliefs and behaviors can impact the structure of social systems – structures that themselves may stabilize and reproduce new stable beliefs.

## References

- Abascal, Maria. 2015. "Us and Them: Black-White Relations in the Wake of Hispanic Population Growth." *American Sociological Review* 80(4):789–813. doi: 10.1177/0003122415587313.
- Aksoy, Ozan. 2015. "Effects of Heterogeneity and Homophily on Cooperation." *Social Psychology Quarterly* 78(4):324–44. doi: 10.1177/0190272515612403.
- Aksoy, Ozan. 2019. "Crosscutting Circles in a Social Dilemma: Effects of Social Identity and Inequality on Cooperation." *Social Science Research* 82:148–63. doi: 10.1016/j.ssresearch.2019.04.009.
- Allport, Gordon W. 1954. "The Nature of Prejudice." *The Nature of Prejudice* 537 p.-537 p.
- Ambrasat, J., C. von Scheve, M. Conrad, G. Schauenburg, and T. Schroder. 2014. "Consensus and Stratification in the Affective Meaning of Human Sociality." *Proceedings of the National Academy of Sciences* 111(22):8001–6. doi: 10.1073/pnas.1313321111.
- Anon. 2020. "Coronavirus Disease 2019." *Centers for Disease Control and Prevention*. Retrieved October 1, 2020 (<https://www.cdc.gov/media/releases/2020/s0528-coronavirus-death-toll.html>).
- Anon. n.d. "A Timeline of COVID-19 Developments in 2020." *AJMC*. Retrieved October 1, 2020a (<https://www.ajmc.com/view/a-timeline-of-covid19-developments-in-2020>).
- Anon. n.d. "Employment Situation News Release." Retrieved October 1, 2020b ([https://www.bls.gov/news.release/archives/empsit\\_05082020.htm](https://www.bls.gov/news.release/archives/empsit_05082020.htm)).
- Baldassarri, Delia, and Maria Abascal. 2020. "Diversity and Prosocial Behavior." *Science* 369(6508):1183–87. doi: 10.1126/science.abb2432.
- Balliet, Daniel, Junhui Wu, and Carsten K. W. De Dreu. 2014. "Ingroup Favoritism in Cooperation: A Meta-Analysis." *Psychological Bulletin* 140(6):1556–81. doi: 10.1037/a0037737.
- Barrera, Davide. 2007. "The Impact of Negotiated Exchange on Trust and Trustworthiness." *Social Networks* 29(4):508–26. doi: 10.1016/j.socnet.2007.03.004.

- Bettencourt, B. Ann, Marilynn B. Brewer, Marian Rogers Croak, and Norman Miller. 1992. "Cooperation and the Reduction of Intergroup Bias: The Role of Reward Structure and Social Orientation." *Journal of Experimental Social Psychology* 28(4):301–19. doi: 10.1016/0022-1031(92)90048-O.
- Billig, Michael, and Henri Tajfel. 1973. "Social Categorization and Similarity in Intergroup Behaviour." *European Journal of Social Psychology* 3(1):27–52. doi: 10.1002/ejsp.2420030103.
- Blau, Peter M. 1977. "A Macrosociological Theory of Social Structure." *American Journal of Sociology* 83(1):26–54.
- Bruch, Elizabeth E., and Robert D. Mare. 2006. "Neighborhood Choice and Neighborhood Change." *American Journal of Sociology* 112(3):667–709. doi: 10.1086/507856.
- Buchan, Nancy R., Rachel T. A. Croson, and Robyn M. Dawes. 2002. "Swift Neighbors and Persistent Strangers: A Cross-Cultural Investigation of Trust and Reciprocity in Social Exchange." *American Journal of Sociology* 108(1):168–206. doi: 10.1086/344546.
- Centola, Damon, Joshua Becker, Devon Brackbill, and Andrea Baronchelli. 2018. "Experimental Evidence for Tipping Points in Social Convention." *Science* 360(6393):1116–19. doi: 10.1126/science.aas8827.
- Cheshire, Coye, Alexandra Gerbasi, and Karen S. Cook. 2010. "Trust and Transitions in Modes of Exchange." *Social Psychology Quarterly* 73(2):176–95. doi: 10.1177/0190272509359615.
- Coleman, James S. 1988. "Social Capital in the Creation of Human Capital." *American Journal of Sociology* 94:S95–120.
- Collins, Randall. 1981. "On the Microfoundations of Macrosociology." *American Journal of Sociology* 86(5):984–1014.
- Cook, Karen S., Coye Cheshire, and Alexandra Gerbasi. n.d. "Power, Dependence and Social Exchange." *Contemporary Social Psychological Theories* 166–93.
- Cook, Karen S., and Richard M. Emerson. 1984. "Exchange Networks and the Analysis of Complex Organizations." *Research in the Sociology of Organizations* 3(4):1–30.

- Cunningham, George. 2005. "The Importance of a Common In-Group Identity in Ethnically Diverse Groups." *Group Dynamics: Theory, Research, and Practice* 9:251–60. doi: 10.1037/1089-2699.9.4.251.
- Davis, James. 1992. "Changeable Weather in a Cooling Climate atop the Liberal Plateau: Conversion and Replacement in Forty-Two General Social Survey Items, 1972–1989." *Public Opinion Quarterly* 56(3):261–306.
- Dovidio, John F., Samuel L. Gaertner, and Gladys Kafati. 2000a. "Group Identity and Intergroup Relations The Common In-Group Identity Model." in *Advances in group processes*. Emerald Group Publishing Limited.
- Dovidio, John F., Samuel L. Gaertner, and Gladys Kafati. 2000b. "Group Identity and Intergroup Relations The Common In-Group Identity Model." Pp. 1–35 in *Advances in Group Processes*. Vol. 17, *Advances in Group Processes*. Emerald Group Publishing Limited.
- Ekeh, Peter. 1974. "Social Exchange Theory." *The Two Traditions*.
- Emerson, Richard M. 1972. "Exchange Theory, Part I: A Psychological Basis for Social Exchange." *Sociological Theories in Progress* 2:38–57.
- Emerson, Richard M. 1976. "Social Exchange Theory." *Annual Review of Sociology* 335–62.
- Fehr, Ernst, and Herbert Gintis. 2007. "Human Motivation and Social Cooperation: Experimental and Analytical Foundations." *Annual Review of Sociology* 33(1):43–64. doi: 10.1146/annurev.soc.33.040406.131812.
- Feld, Scott L. 1981. "The Focused Organization of Social Ties." *American Journal of Sociology* 86(5):1015–35.
- Feld, Scott L. 1982. "Social Structural Determinants of Similarity among Associates." *American Sociological Review* 47(6):797–801. doi: 10.2307/2095216.
- Festinger, Leon. 1957. *A Theory of Cognitive Dissonance*. Vol. 2. Stanford university press.
- Fiel, Jeremy E. 2021. "Relational Segregation: A Structural View of Categorical Relations." *Sociological Theory* 39(3):153–79. doi: 10.1177/073527512111029979.
- Fine, Gary Alan. 2012. "Group Culture and the Interaction Order: Local Sociology on the Meso-Level." *Annual Review of Sociology* 38(1):159–79. doi: 10.1146/annurev-soc-071811-145518.

- Fischbacher, Urs. 2007. "Zurich Toolbox for Ready-Made Economic Experiments. Forthcoming in *Experimental Economics*."
- Foddy, Margaret, Michael J. Platow, and Toshio Yamagishi. 2009. "Group-Based Trust in Strangers: The Role of Stereotypes and Expectations." *Psychological Science* 20(4):419–22. doi: 10.1111/j.1467-9280.2009.02312.x.
- Fowler, James H., and Cindy D. Kam. 2007. "Beyond the Self: Social Identity, Altruism, and Political Participation." *The Journal of Politics* 69(3):813–27. doi: 10.1111/j.1468-2508.2007.00577.x.
- Freeland, Robert E., and Jesse Hoey. 2018. "The Structure of Deference: Modeling Occupational Status Using Affect Control Theory." *American Sociological Review* 83(2):243–77. doi: 10.1177/0003122418761857.
- Gaertner, Samuel L., John F. Dovidio, Phyllis A. Anastasio, Betty A. Bachman, and Mary C. Rust. 1993. "The Common Ingroup Identity Model: Recategorization and the Reduction of Intergroup Bias." *European Review of Social Psychology* 4(1):1–26. doi: 10.1080/14792779343000004.
- Gaertner, Samuel L., John F. Dovidio, and Betty A. Bachman. 1996. "Revisiting the Contact Hypothesis: The Induction of a Common Ingroup Identity." *International Journal of Intercultural Relations* 20(3–4):271–90.
- Gaertner, Samuel L., John F. Dovidio, Jason A. Nier, and Brenda S. Banker. 2000. "The Common Ingroup Identity Model for Reducing Intergroup Bias: Progress and Challenges." Pp. 133–48 in *Social Identity Processes: Trends in Theory and Research*. London: SAGE Publications Ltd.
- Gaertner, Samuel L., John F. Dovidio, Mary C. Rust, Jason A. Nier, Brenda S. Banker, Christine M. Ward, Gary R. Mottola, and Missy Houlette. 1999. "Reducing Intergroup Bias: Elements of Intergroup Cooperation." *Journal of Personality and Social Psychology* 76(3):388.
- Gaither, Sarah E., and Samuel R. Sommers. 2013. "Living with an Other-Race Roommate Shapes Whites' Behavior in Subsequent Diverse Settings." *Journal of Experimental Social Psychology* 49(2):272–76.
- Giddens, Anthony. 1986. *The Constitution of Society: Outline of the Theory of Structuration*. Vol. 349. Univ of California Press.
- Goodenough, Ward H. 1961. "Comment on Cultural Evolution." *Daedalus* 90(3):521–28.

- Habyarimana, James, Macartan Humphreys, Daniel N. Posner, and Jeremy M. Weinstein. 2009. *Coethnicity: Diversity and the Dilemmas of Collective Action*. Russell Sage Foundation.
- Hamm, Jill V. 2000. "Do Birds of a Feather Flock Together? The Variable Bases for African American, Asian American, and European American Adolescents' Selection of Similar Friends." *Developmental Psychology* 36(2):209–19. doi: 10.1037/0012-1649.36.2.209.
- Hamm, Jill V., and Beverly S. Faircloth. 2005. "The Role of Friendship in Adolescents' Sense of School Belonging." *New Directions for Child and Adolescent Development* 2005(107):61–78.
- Harrell, and Joseph M. Quinn. 2022. "Shared Identities and the Structure of Exchange Distinctly Shape Cooperation."
- Harwood, Jake. 2020. "Social Identity Theory." *The International Encyclopedia of Media Psychology* 1–7.
- Heise, D., and N. MacKinnon. 2010. *Self, Identity, and Social Institutions*. Springer.
- Heise, David R. 1979. *Understanding Events: Affect and the Construction of Social Action*. CUP Archive.
- Heise, David R. 2007. *Expressive Order: Confirming Sentiments in Social Actions*. Springer Science & Business Media.
- Heise, David R. 2010. *Surveying Cultures: Discovering Shared Conceptions and Sentiments*. John Wiley & Sons.
- Hobbes, Thomas. 1651. *Leviathan*. Cambridge: Cambridge University Press.
- Holland, Paul W., and Samuel Leinhardt. 1971. "Transitivity in Structural Models of Small Groups." *Comparative Group Studies* 2(2):107–24. doi: 10.1177/104649647100200201.
- Homans, George C. 1974. "Social Behavior: Its Elementary Forms."
- Iyengar, Shanto, and Sean J. Westwood. 2015. "Fear and Loathing across Party Lines: New Evidence on Group Polarization." *American Journal of Political Science* 59(3):690–707. doi: 10.1111/ajps.12152.

- Kao, Grace, Kara Joyner, and Kelly Stamper Balistreri. 2019. *The Company We Keep: Interracial Friendships and Romantic Relationships from Adolescence to Adulthood*. Russell Sage Foundation.
- Kiley, Kevin, and Stephen Vaisey. 2020. "Measuring Stability and Change in Personal Culture Using Panel Data." *American Sociological Review* 85(3):477–506. doi: 10.1177/0003122420921538.
- Kiyonari, Toko. 2002. "Expectations of a Generalized Exchange System and Ingroup Favoritism: An Experimental Study of Bounded Reciprocity." *Japanese Journal of Psychology*.
- Kiyonari, Toko, and Toshio Yamagishi. 2004. "Ingroup Cooperation and the Social Exchange Heuristic." *Contemporary Psychological Research on Social Dilemmas* 269–86.
- Kollock, Peter. 1994. "The Emergence of Exchange Structures: An Experimental Study of Uncertainty, Commitment, and Trust." *American Journal of Sociology* 100(2):313–45.
- Kollock, Peter. 1998. "Social Dilemmas: The Anatomy of Cooperation." *Annual Review of Sociology* 24(1):183–214. doi: 10.1146/annurev.soc.24.1.183.
- Kossinets, Gueorgi, and Duncan J. Watts. 2009. "Origins of Homophily in an Evolving Social Network." *American Journal of Sociology* 115(2):405–50.
- Kuwabara, Ko. 2011. "Cohesion, Cooperation, and the Value of Doing Things Together: How Economic Exchange Creates Relational Bonds." *American Sociological Review* 76(4):560–80. doi: 10.1177/0003122411414825.
- Lawler, Edward J. 2001. "An Affect Theory of Social Exchange." *American Journal of Sociology* 107(2):321–52. doi: 10.1086/324071.
- Lawler, Edward J., Shane R. Thye, and Jeongkoo Yoon. 2000. "Emotion and Group Cohesion in Productive Exchange." *American Journal of Sociology* 106(3):616–57. doi: 10.1086/318965.
- Lawler, Edward J., Shane R. Thye, and Jeongkoo Yoon. 2008. "Social Exchange and Micro Social Order." *American Sociological Review* 73(4):519–42. doi: 10.1177/000312240807300401.
- Lawler, Edward J., Shane R. Thye, and Jeongkoo Yoon. 2009. *Social Commitments in a Depersonalized World*. Russell Sage Foundation.

- Lawler, Edward J., Shane R. Thye, and Jeongkoo Yoon. 2016. "The Problem of Social Order in Nested Group Structures." Pp. 149–66 in *Handbook of contemporary sociological theory*. Springer.
- Lawler, Edward J., Shane R. Thye, and Jeongkoo Yoon. 2021. "Theorizing Nested Group Ties." *Theoretical Sociology: The Future of a Disciplinary Foundation* 153–72.
- Lawler, Edward J., and Jeongkoo Yoon. 1996. "Commitment in Exchange Relations: Test of a Theory of Relational Cohesion." *American Sociological Review* 61(1):89–108. doi: 10.2307/2096408.
- Lazarsfeld, Paul F., and Robert K. Merton. 1954. "Friendship as a Social Process: A Substantive and Methodological Analysis." *Freedom and Control in Modern Society* 18(1):18–66.
- Lee, Sangseok, and Dong Kyu Lee. 2018. "What Is the Proper Way to Apply the Multiple Comparison Test?" *Korean Journal of Anesthesiology* 71(5):353–60. doi: 10.4097/kja.d.18.00242.
- Lemay, Edward P., and Joshua E. Ryan. 2021. "Common Ingroup Identity, Perceived Similarity, and Communal Interracial Relationships." *Personality and Social Psychology Bulletin* 47(6):985–1003. doi: 10.1177/0146167220953984.
- Lerner, Melvin J. 1971. "Observers Evaluation of a Victim: Justice, Guilt, and Veridical Perception." *Journal of Personality and Social Psychology* 20(2):127–35. doi: 10.1037/h0031702.
- Lévi-Strauss, Claude. 1969. "The Elementary Structures of Kinship, Trans." *James Harle Bell and John Richard von Sturmer (Boston: Beacon, 1969)* 30.
- Lincoln, Alan, and George Levinger. 1972. "Observers' Evaluations of the Victim and the Attacker in an Aggressive Incident." *Journal of Personality and Social Psychology* 22(2):202–10. doi: 10.1037/h0032708.
- Lizardo, Omar, and Michael Strand. 2010. "Skills, Toolkits, Contexts and Institutions: Clarifying the Relationship Between Different Approaches to Cognition in Cultural Sociology." *Poetics* 38(2):204–27.
- MacKinnon, Neil J., and Alison Luke. 2002. "Changes in Identity Attitudes as Reflections of Social and Cultural Change." *The Canadian Journal of Sociology / Cahiers Canadiens de Sociologie* 27(3):299–338. doi: 10.2307/3341547.

- MacKinnon, Neil Joseph. 1994. *Symbolic Interactionism as Affect Control*. SUNY Press.
- Marsden, Peter V. 1987. "Core Discussion Networks of Americans." *American Sociological Review* 52(1):122–31. doi: 10.2307/2095397.
- Marsden, Peter V. 2012. *Social Trends in American Life: Findings from the General Social Survey since 1972*. Princeton University Press.
- Marsden, Peter V., Tom W. Smith, and Michael Hout. 2020. "Tracking US Social Change Over a Half-Century: The General Social Survey at Fifty." *Annual Review of Sociology* 46(1):109–34. doi: 10.1146/annurev-soc-121919-054838.
- McPherson, M. 2004. "A Blau Space Primer: Prolegomenon to an Ecology of Affiliation." *Industrial and Corporate Change* 13(1):263–80. doi: 10.1093/icc/13.1.263.
- McPherson, Miller, Lynn Smith-Lovin, and James M. Cook. 2001. "Birds of a Feather: Homophily in Social Networks." *Annual Review of Sociology* 27(1):415–44.
- McPherson, Miller, Lynn Smith-Lovin, and Craig Rawlings. 2021. "The Enormous Flock of Homophily Researchers: Assessing and Promoting a Research Agenda." *Personal Networks: Classic Readings and New Directions in Egocentric Analysis* 459.
- Melamed, David, Brent Simpson, Ashley Harrell, Christopher W. Munn, Jered Z. Abernathy, and Matthew Sweitzer. 2020. "Homophily and Segregation in Cooperative Networks." *American Journal of Sociology* 125(4):1084–1127. doi: 10.1086/708142.
- Molm, Linda D. 1991. "Affect and Social Exchange: Satisfaction in Power-Dependence Relations." *American Sociological Review* 56(4):475–93. doi: 10.2307/2096269.
- Molm, Linda D. 1994. "Dependence and Risk: Transforming the Structure of Social Exchange." *Social Psychology Quarterly* 57(3):163–76. doi: 10.2307/2786874.
- Molm, Linda D. 2010. "The Structure of Reciprocity." *Social Psychology Quarterly* 73(2):119–31. doi: 10.1177/0190272510369079.
- Molm, Linda D., Jessica L. Collett, and David R. Schaefer. 2006. "Conflict and Fairness in Social Exchange." *Social Forces* 84(4):2331–52.
- Molm, Linda D., Jessica L. Collett, and David R. Schaefer. 2007. "Building Solidarity through Generalized Exchange: A Theory of Reciprocity." *American Journal of Sociology* 113(1):205–42. doi: 10.1086/517900.

- Molm, Linda D., David R. Schaefer, and Jessica L. Collett. 2009. "Fragile and Resilient Trust: Risk and Uncertainty in Negotiated and Reciprocal Exchange." *Sociological Theory* 27(1):1–32. doi: 10.1111/j.1467-9558.2009.00336.x.
- Molm, Linda D., Nobuyuki Takahashi, and Gretchen Peterson. 2000. "Risk and Trust in Social Exchange: An Experimental Test of a Classical Proposition." *American Journal of Sociology* 105(5):1396–1427. doi: 10.1086/210434.
- Moody, James. 2001. "Race, School Integration, and Friendship Segregation in America." *American Journal of Sociology* 107(3):679–716. doi: 10.1086/338954.
- Osgood, Charles E. 1962. "Studies on the Generality of Affective Meaning Systems." *American Psychologist* 17(1):10–28. doi: 10.1037/h0045146.
- Osgood, Charles E. 1964. "Semantic Differential Technique in the Comparative Study of Cultures." *American Anthropologist* 66(3):171–200.
- Osgood, Charles Egerton, George J. Suci, and Percy H. Tannenbaum. 1957. *The Measurement of Meaning*. University of Illinois Press.
- Pettigrew, Thomas F., and Linda R. Tropp. 2006. "A Meta-Analytic Test of Intergroup Contact Theory." *Journal of Personality and Social Psychology* 90(5):751–83. doi: 10.1037/0022-3514.90.5.751.
- Portes, Alejandro, and Julia Sensenbrenner. 1993. "Embeddedness and Immigration: Notes on the Social Determinants of Economic Action." *American Journal of Sociology* 98(6):1320–50.
- Rapoport, Anatol. 1953. "Spread of Information through a Population with Socio-Structural Bias: I. Assumption of Transitivity." *The Bulletin of Mathematical Biophysics* 15(4):523–33.
- Rho, Hye Jin, Shawn Fremstad, and Hayley Brown. 2020. *A Basic Demographic Profile of Workers in Frontline Industries*. Washington, DC: Center for Economic and Policy Research.
- Ridgeway, Cecilia L., Elizabeth Heger Boyle, Kathy J. Kuipers, and Dawn T. Robinson. 1998. "How Do Status Beliefs Develop? The Role of Resources and Interactional Experience." *American Sociological Review* 63(3):331–50. doi: 10.2307/2657553.
- Ridgeway, Cecilia L., and Kristan Glasgow Erickson. 2000. "Creating and Spreading Status Beliefs." *American Journal of Sociology* 106(3):579–615. doi: 10.1086/318966.

- Riek, Blake M., Eric W. Mania, Samuel L. Gaertner, Stacy A. McDonald, and Marika J. Lamoreaux. 2010. "Does a Common Ingroup Identity Reduce Intergroup Threat?" *Group Processes & Intergroup Relations* 13(4):403–23. doi: 10.1177/1368430209346701.
- Rohrer, Julia M., Tamás Keller, and Felix Elwert. 2021. "Proximity Can Induce Diverse Friendships: A Large Randomized Classroom Experiment." *PLOS ONE* 16(8):e0255097. doi: 10.1371/journal.pone.0255097.
- Romano, Angelo, Daniel Balliet, and Junhui Wu. 2017. "Unbounded Indirect Reciprocity: Is Reputation-Based Cooperation Bounded by Group Membership?" *Journal of Experimental Social Psychology* 71:59–67.
- Romano, Angelo, Daniel Balliet, Toshio Yamagishi, and James H. Liu. 2017. "Parochial Trust and Cooperation across 17 Societies." *Proceedings of the National Academy of Sciences* 114(48):12702–7.
- Romano, Angelo, Matthias Sutter, James H. Liu, Toshio Yamagishi, and Daniel Balliet. 2021. "National Parochialism Is Ubiquitous across 42 Nations around the World." *Nature Communications* 12(1):1–8.
- Romney, A. Kimball. 1991. "Cultural Consensus and Social Intelligence." *Mathematical Social Sciences* 22(1):98–102. doi: 10.1016/0165-4896(91)90072-Y.
- ROMNEY, A. KIMBALL, WILLIAM H. BATCHELDER, and SUSAN C. WELLER. 1987. "Recent Applications of Cultural Consensus Theory." *American Behavioral Scientist* 31(2):163–77. doi: 10.1177/000276487031002003.
- Ryder, Norman B. 1965. "The Cohort as a Concept in the Study of Social Change." *American Sociological Review* 30(6):843–61. doi: 10.2307/2090964.
- Savage, S. V., and M. M. Whitham. 2018. "The Social Exchange Perspective." *Contemporary Social Psychological Theories* 29–53.
- Schelling, Thomas C. 1969. "Models of Segregation." *The American Economic Review* 59(2):488–93.
- Schelling, Thomas C. 1971. "Dynamic Models of Segregation." *Journal of Mathematical Sociology* 1(2):143–86.
- Schneider, Andreas, and Tobias Schröder. 2012. "Ideal Types of Leadership as Patterns of Affective Meaning: A Cross-Cultural and Over-Time Perspective." *Social Psychology Quarterly* 75(3):268–87. doi: 10.1177/0190272512446755.

- Scholl, Wolfgang. 2013. "The Socio-Emotional Basis of Human Interaction and Communication: How We Construct Our Social World." *Social Science Information* 52(1):3–33. doi: 10.1177/0539018412466607.
- Schwadel, Philip. 2011. "Age, Period, and Cohort Effects on Religious Activities and Beliefs." *Social Science Research* 40(1):181–92. doi: 10.1016/j.ssresearch.2010.09.006.
- Sherif, Muzafer, O. J. Harvey, B. Jack White, and William R. Hood. 1954. "Intergroup Conflict and Cooperation: The Robbers Cave Experiment." 18.
- Simmel, Georg. 1950. *The Sociology of Georg Simmel*. Vol. 92892. Simon and Schuster.
- Simpson, Brent, and Robb Willer. 2015. "Beyond Altruism: Sociological Foundations of Cooperation and Prosocial Behavior." 23.
- Stryker, Sheldon. 1980. *Symbolic Interactionism: A Social Structural Version*. Benjamin-Cummings Publishing Company.
- Swidler, Ann. 1986. "Culture in Action: Symbols and Strategies." *American Sociological Review* 51(2):273–86. doi: 10.2307/2095521.
- Tajfel, Henri. 1981. *Human Groups and Social Categories*. Cambridge university press Cambridge.
- Tajfel, Henri, Michael G. Billig, Robert P. Bundy, and Claude Flament. 1971. "Social Categorization and Intergroup Behaviour." *European Journal of Social Psychology* 1(2):149–78.
- Tanis, Martin, and Tom Postmes. 2005. "A Social Identity Approach to Trust: Interpersonal Perception, Group Membership and Trusting Behaviour." *European Journal of Social Psychology* 35(3):413–24.
- Turner, Jonathan H. 1988. *A Theory of Social Interaction*. Stanford University Press.
- Turner, William. 1545. "The Rescuing of Romish Fox."
- Vaisey, Stephen. 2009. "Motivation and Justification: A Dual-Process Model of Culture in Action." *American Journal of Sociology* 114(6):1675–1715. doi: 10.1086/597179.
- Vezzali, Loris, Sofia Stathi, Richard J. Crisp, Dino Giovannini, Dora Capozza, and Samuel L. Gaertner. 2015. "Imagined Intergroup Contact and Common Ingroup Identity." *Social Psychology*.

- Werner, Carol, and Pat Parmelee. 1979. "Similarity of Activity Preferences among Friends: Those Who Play Together Stay Together." *Social Psychology Quarterly* 62–66.
- Wimmer, Andreas, and Kevin Lewis. 2010. "Beyond and Below Racial Homophily: ERG Models of a Friendship Network Documented on Facebook." *American Journal of Sociology* 116(2):583–642. doi: 10.1086/653658.
- Whitham, Monica M. 2018. "Paying It Forward and Getting It Back: The Benefits of Shared Social Identity in Generalized Exchange." *Sociological Perspectives* 61(1):81–98.
- Whitham, Monica M. 2021. "Generalized Generosity: How the Norm of Generalized Reciprocity Bridges Collective Forms of Social Exchange." *American Sociological Review* 86(3):503–31.
- Yamagishi, Toshio, Nobuhito Jin, and Toko Kiyonari. 1999. "Bounded Generalized Reciprocity: Ingroup Boasting and Ingroup Favoritism." *Advances in Group Processes* 16(1):161–97.
- Yamagishi, Toshio, and Toko Kiyonari. 2000. "The Group as the Container of Generalized Reciprocity." *Social Psychology Quarterly* 116–32.
- Yamagishi, Toshio, and Nobuhiro Mifune. 2009. "Social Exchange and Solidarity: In-Group Love or out-Group Hate?" *Evolution and Human Behavior* 30(4):229–37.
- Yamagishi, Toshio, and Midori Yamagishi. 1994. "Trust and Commitment in the United States and Japan." *Motivation and Emotion* 18(2):129–66.
- Zerubavel, Noam, Mark Anthony Hoffman, Adam Reich, Kevin N. Ochsner, and Peter Bearman. 2018. "Neural Precursors of Future Liking and Affective Reciprocity." *Proceedings of the National Academy of Sciences* 115(17):4375–80. doi: 10.1073/pnas.1802176115.

## Biography

Joseph Michael Quinn holds a B.S. degree in Geology and English from The College of William & Mary in Virginia and an M.A. in Sociology from Duke University. While pursuing his Ph.D. in the Sociology Department at Duke, Joe received research fellowships from the Kenan Institute for Ethics as well as departmental research awards from Duke University's Sociology Department. Joe has published pieces with his colleagues in *American Behavioral Scientist*, *Social Science and Medicine*, *Urban Education*, and *SAGE Research Methods Foundations*. He will begin an appointment as an Assistant Professor of Sociology at the University of South Carolina in the fall of 2022.