

UNIVERSITY OF CALIFORNIA

Los Angeles

Tracking relationships:

Uncovering how people acquire, represent, use, and predict
social network information

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

by

Miriam Elynn Schwyck

2023

© Copyright by

Miriam Elynn Schwyck

2023

ABSTRACT OF THE DISSERTATION

Tracking relationships:
Uncovering how people acquire, represent, use, and predict
social network information

by

Miriam Elynn Schwyck

Doctor of Philosophy in Psychology

University of California, Los Angeles, 2023

Professor Carolyn Parkinson, Chair

Humans are incredibly social creatures, and every relationship is part of an individual's broader social network. Human social networks are large and complex structures (the number of relationships to track increases exponentially with every additional person), posing a huge computational challenge for the mind. Moreover, there are often limitations and biases in information availability (e.g., when entering a new community); it is also typically impossible to monitor the interactions and relationships between all of the people in one's surroundings. Yet, people rather accurately and seemingly effortlessly remember, track, and make inferences about others' relationships and patterns thereof (e.g., who is very well-connected). Indeed, the ability to successfully navigate one's social network is essential for avoiding social missteps and accessing necessary information or resources. Because little research has integrated individual cognition with

the social networks one inhabits, very little is known about how people are able to track and use the immense amount of information contained within these structures. Through my dissertation research, I seek to understand (i) how different people mentally represent their social networks through the use of priors regarding how relationships are formed, (ii) how the human brain supports this process, (iii) how one's mental map of relationships shapes one's interactions with others, and (iv) how one's assumptions of behavioral homophily (i.e., that similarly behaving people are likely to become friends with one another) shape predictions of future friendships. To this end, I created novel, experimentally controlled paradigms in which participants learn new networks with minimal extraneous information to rigorously and systematically examine these processes. That is, I decouple social network knowledge from other factors that may covary along with it in real-world contexts to examine how people acquire, neurally encode, and use social network knowledge in their everyday life, and how various contexts shape these processes. My research helps explain a vital and universal skillset that humans possess: the ability to understand and navigate their complex social worlds.

The dissertation of Miriam Elynn Schwyck is approved.

Naomi I. Eisenberger

Kerri Johnson

Matthew D. Lieberman

Carolyn M. Parkinson, Committee Chair

University of California, Los Angeles

2023

This dissertation is dedicated to my own social network
for all the support that made this work possible.

TABLE OF CONTENTS

Abstract of the Dissertation	ii
Table of Contents	vi
List of Figures, Tables, and Supplementary Materials	vii
Acknowledgments	ix
Vita	x
General Introduction	1
Chapter 1 The role of one’s own social network position in learning new networks:	
Brokerage is associated with better network learning	4
Chapter 2 Neural encoding of novel social networks: Evidence that perceivers prioritize	
others’ centrality	39
Chapter 3 Similarity among friends serves as a social prior: The assumption that “Birds of	
a feather flock together” shapes social decisions and relationship beliefs	68
Chapter 4 Predicting that birds of a feather will flock together: Expectations of homophily	
for others but not the self	119
General Discussion	142
References	147

LIST OF FIGURES, TABLES,
AND SUPPLEMENTARY MATERIALS

Chapter 1 The role of one’s own social network position in learning new networks: Brokerage is associated with better network learning	
Figure 1.1	10
Table 1.1	11
Figure 1.2	13
Figure 1.3	16
Table 1.2	17
Figure 1.4	18
Table 1.3	19
Figure 1.5	22
Table 1.4	24
Figure 1.6	25
Supplementary Materials S1	30
Table S1.1	31
Figure S1.1	33
Table S1.2	33
Figure S1.2	34
Table S1.3	35
Figure S1.3	36
Table S1.4	37
Figure S1.4	38
Chapter 2 Neural encoding of novel social networks: Evidence that perceivers prioritize others’ centrality	
Figure 2.1	45
Figure 2.2	51
Figure 2.3	55
Table 2.1	55
Table 2.2	56
Figure 2.4	58
Supplementary Materials S2	64
Figure S2.1	64
Table S2.1	64
Table S2.2	65

Chapter 3 | Similarity among friends serves as a social prior: The assumption that “Birds of a feather flock together” shapes social decisions and relationship beliefs

Figure 3.1	77
Figure 3.2	81
Figure 3.3	88
Table 3.1	90
Table 3.2	94
Figure 3.4	98
Supplementary Materials S3	107
Table S3.1	107
Figure S3.1	108
Table S3.2	110
Figure S3.2	110

Chapter 4 | Predicting that birds of a feather will flock together: Expectations of homophily for others but not the self

Figure 4.1	125
Figure 4.2	131
Table 4.1	131
Figure 4.3	133
Table 4.2	134
Supplementary Materials S4	139
Table S4.1	139
Figure S4.1	141

ACKNOWLEDGMENTS

This journey would not have been possible without the support of countless people along the way. First, to my spouse, Anne Schwyck, thank you for your unshakable belief in me, your unending encouragement and support, your loving openness to sharing this experience, and your partnership in all things. To my family, especially my parents, Shelly and Curtis Weaverdyck, thank you for your unconditional love, support, and understanding. In particular, thank you for modeling conscientious choices throughout your careers and challenging me to be thoughtful and intentional in all that I do. To my friends, thank you for helping me through the day-to-day. I love our conversations, work sessions, movie and game nights, happy hours, and shared meals. I have learned so much from each and every one of you.

To my advisor, Carolyn Parkinson, thank you for your patient guidance and support over the past six years. Thank you especially for always encouraging and trusting me to follow my own ideas. Your mentorship has helped me grow into a more confident, critical, and compassionate scientist. To the rest of my dissertation committee, Matthew Lieberman, Naomi Eisenberger, and Kerri Johnson, thank you for your thoughtful consideration and feedback about my work. Your advice has been, and will continue to be, invaluable as I navigate my academic career. Finally, thank you to all my colleagues, including my amazing research assistants and coauthors, who have helped me conduct this research.

VITA

Education

Ph.D., Social Psychology, University of California, Los Angeles, CA Projected 2023

Advisor: Dr. Carolyn Parkinson

M.A., Social Psychology, University of California, Los Angeles, CA 2018

Advisors: Drs. Carolyn Parkinson, Matthew Lieberman

B.A., Psychology, Mathematical Sciences, Bethel College, North Newton, KS 2013

Advisors: Drs. Dwight Krehbiel, Lisa Thimm

Publications

*Equal contributions

Schwycck, M.E.*, Du, M.*, Li, Y., Chang, L., & Parkinson, C. (2023). Similarity among friends serves as a social prior: The assumption that "Birds of a feather flock together" shapes social decisions and relationship beliefs. *Personality and Social Psychology Bulletin*, 014616722211402.

Schwycck, M.E., Du, M., Natarajan, P., Chwe, J.A., & Parkinson, C. (2023). Neural encoding of novel social networks: Evidence that perceivers prioritize others' centrality. *Social Cognitive and Affective Neuroscience*, 18(1) 1-9.

Sahi, R.S.*, **Schwycck, M.E.***, Parkinson, C., & Eisenberger, N.I. (2021). Having more virtual interaction partners during COVID-19 physical distancing measures may benefit mental health. *Scientific Reports*. 11(18273).

- Weaverdyck, M.E.**, Thornton, M.A., & Tamir, D.I. (2021). The representational structure of mental states generalizes across target people and stimulus modalities. *NeuroImage*, 238(118258) 1-9.
- Weaverdyck, M.E.**, Lieberman, M.D., & Parkinson, C. (2020). Multivoxel pattern analysis in fMRI: A practical introduction for social and affective neuroscientists. *Social Cognitive and Affective Neuroscience*, 15(4) 487-509.
- Thornton, M.A., **Weaverdyck, M.E.**, & Tamir, D.I. (2019). The brain represents people as the mental states they habitually experience. *Nature Communications*, 10(2291) 1-10.
- Thornton, M.A., **Weaverdyck, M.E.**, Mildner, J.N., & Tamir, D.I. (2019). People represent their own mental states more distinctly than others'. *Nature Communications*, 10(2117) 1-9.
- Thornton, M.A., **Weaverdyck, M.E.**, & Tamir, D.I. (2019). The social brain automatically predicts others' future mental states. *Journal of Neuroscience*, 39(1) 140-148.
- Weaverdyck, M.E.**, & Parkinson, C. (2018). The neural representation of social networks. *Current Opinion in Psychology*, 24, 58-66.
- Jones, M.A., Shelton, B.C., & **Weaverdyck, M.E.** (2014). On God's Number(s) for Rubik's Slide. *The College Mathematics Journal*, 45(4), 267-275.
- Anderson, A.L.*, **Weaverdyck, M.E.***, & Krehbiel, D. (2011, March). *Discovering GEMS in music: Armonique digs for music you like*. Paper presented at the National Conference for Undergraduate Research, Ithaca College, NY.

GENERAL INTRODUCTION

Human social connection is ubiquitous and essential for people to live happy, healthy lives. Indeed, for decades, psychologists from many disciplines have linked social connectedness to health outcomes (e.g., Cole et al., 2007; Eisenberger, 2013; Kawachi, 2001; Sahi*, Schwyck*, Parkinson, & Eisenberger, 2021). In social psychology, researchers have long understood that the mere presence of another person in a room can have drastic effects on attention, behavior, and performance (e.g., Augustinova & Ferrand, 2012; Risko & Kingstone, 2011; Satow, 1975; Wühr & Huestegge, 2010). Furthermore, these effects are mitigated or exacerbated depending on who that person is, how physically close they are, and the relationship that exists between them. The psychological proximity of another, whether it is temporal distance, similarity, or social closeness (or one of many other dimensions; Trope & Liberman, 2010), changes how we interact with them, how we perceive them, and how our brains respond when we see them (e.g., Liberman & Trope, 2014; Parkinson, Liu, & Wheatley, 2014; Rim, Uleman, & Trope, 2009; Thornton, Weaverdyck, Mildner, & Tamir, 2019; Wakslak, Trope, & Liberman, 2007). In particular, there are distinct brain regions that consistently track and process social information (Mars et al., 2012; Meyer, Davachi, Ochsner, & Liberman, 2019). These regions are responsible for reasoning about oneself and others, and, in particular, the ability to consider what is happening in others' minds (i.e., theory of mind, mentalizing; e.g., Mars et al., 2012; Meyer et al., 2019; Spreng, Mar, & Kim, 2009). Furthermore, this process happens spontaneously when people are mind-wandering alone as well as when they encounter others (Poerio et al., 2017). The brain's prioritization of social information is further evidenced by the automatic social categorization of others and groups (Alt & Phillips, 2022; Goodale, Alt, Lick, & Johnson, 2018; Johnson, Lick, & Carpinella, 2015). That is, humans are almost constantly affected by or thinking about others and how they relate to oneself. Thus,

social thought and interactions do not take place in a vacuum, but rather in a broader social context. One highly consequential and chronically relevant aspect of this social context is the networks of social relationships surrounding each person. Recently, researchers have begun to look beyond the individual or dyad to examine how aspects of these webs of relationships (e.g., who is friends with whom; who has a lot of friends; who bridges between communities) impact human thought and behavior.

Throughout history, our species has survived through building and maintaining complex social structures that shape every aspect of our lives, from national governments to small friend groups. The prominent theory known as the social brain hypothesis posits that the evolution of the human brain was in part shaped by the need to form, track, and navigate these large complex social networks (Dunbar, 2014). Indeed, understanding the structure of one's social networks is key to having successful interactions. Misunderstanding how individuals are connected could lead to embarrassing remarks or missteps, which may in turn weaken one's own reputation, diminish access to resources, or shift one's standing in the hierarchy. People with accurate knowledge of the relationships comprising their social networks, however, may be better able to avoid unwanted conflict, acquire and control information and other resources, manage reputations, and move themselves into more central positions, thereby increasing their control and status within the network (Hahl, Kacperczyk, & Davis, 2016; Simpson, Markovsky, & Steketee, 2011). "Networking" in professional settings is often touted as one of the best ways to find and secure jobs, learn about new opportunities or methods, and acquire funding. Thus, it is vital for humans to be able to cognitively track their broader social networks.

Yet, the size and complexity of these networks means that tracking all these relationships poses a serious computational challenge. Even individuals' local social networks are very large

and complex, consisting of approximately 150 people on average (Dunbar, 2014). This means there are 11,175 possible relationships that would need to be navigated and tracked. Still, people seem to be able to recall, track, and make inferences about their social networks relatively accurately and easily, despite the computational issues. While there is increasing research interest in social networks, it is still unknown how people are able to represent such large amounts of social information in a way that allows them to use it in their daily life, and how this context then shapes social thought and behavior.

Here, I conducted a series of experiments that seek to understand how people learn about and mentally represent their complex social networks, how the human brain supports this process, and how our mental maps of relationships shape our interactions with others, and vice versa. To test how the human mind and brain represents social network information, I created novel, experimentally controlled paradigms in which participants learn new networks with minimal extraneous information, thereby allowing for the decoupling of social network knowledge from factors that covary with it in real-world contexts. In Chapter 1, I tested how individuals' own position in their personal network shapes learning of new network structures through the use of priors learned from real-world networks. In Chapter 2, I tested if and where the brain encodes information about new social networks and how it prioritizes specific aspects of one's position in the network. In Chapter 3, I examined how priors about social network structures are used to inform behavior and shape recall of others' relationships. Finally, in Chapter 4, I tested how people use such priors to predict future relationships between others and oneself. In other words, in this work, I examine how people acquire (Chapter 1), represent (Chapter 2), use (Chapter 3), and predict (Chapter 4) social network information.

CHAPTER 1

The role of one's own social network position in learning new networks:

Brokerage is associated with better network learning

Miriam E. Schwyck, Meng Du, Carolyn Parkinson

University of California, Los Angeles

Abstract

Navigating our complex social lives requires understanding where others sit in our social networks. Some individuals may be more attuned to the structure of their social world, and thus, better able to learn new social networks due to having accumulated accurate priors about social network structure in their own lives. Correspondingly, such individuals may acquire more advantageous positions in their own social networks. In four studies ($N_{total} = 1,768$), brokers (people who connect otherwise disparate people in their own networks) were especially good at learning and remembering new networks that were structured like typical real-world social networks, but not unnaturally-structured ones, suggesting that brokers are attuned to the structure of real-world networks. Additionally, we found that brokers were able to learn networks better by focusing on ties that exist in those networks (as opposed to focusing on ties that were missing) and other brokers. We found no differences when the network was framed as a social network of friends or a non-social network of flights between airports. This work illuminates the mechanisms of network learning based on one's own experiences, and is one of the first studies that we know of to link one's own social network position to one's ability to learn new networks.

Introduction

When entering a new school, workplace, or other community, one needs to learn the social structure underlying every interaction. Misunderstanding how people are connected to one another could lead to making embarrassing remarks to a boss or about a neighbor's friend. That is, an accurate understanding of where everyone sits in their broader social network is socially advantageous (Brands, 2013). While humans in general appear to be remarkably adept at navigating their complex social networks, some people seem to be more adept at navigating new social structures than others. Here, we examine if and how individuals' own social experiences may shape their ability to learn new social network knowledge.

Human social networks are large and complex. Tracking and encoding patterns of relationships in these networks poses a significant computational problem for the human mind, as the number of possible relationships increases exponentially with each additional network member (Basyouni & Parkinson, 2022). The social brain hypothesis argues that our brains evolved, in part, to support the cognitive demands of navigating and managing the vast amount of information related to the relationships comprising our large, intricate social worlds (Dunbar, 2014; Dunbar & Shultz, 2007). Indeed, people are remarkably adept at learning the global structure of incomplete networks (Lynn & Bassett, 2020). Still, it would be extremely demanding and inefficient to learn and remember each relationship (and where there is no relationship) individually.

To solve this problem, it has been suggested that humans use "compression heuristics" (Brashears, 2013) to reduce the cognitive load of remembering whole network structures. Specifically, people may use familiar or naturally occurring features of social networks as mental shortcuts to reduce the amount of information they need to represent (e.g., Brands, 2013; De Soto, 1960; De Soto et al., 1968; Janicik & Larrick, 2005). For instance, social networks tend to be

“small-world” structures (i.e., densely connected groups with few inter-group connections) that consist of balanced triads (if two people share a friend, then they are also friends). Indeed, people often assume that two individuals who share a close friend are also friends themselves, thereby overperceiving balanced (i.e., closed) triads in their own networks (De Soto, 1960; Heider, 1946, 1958; Kilduff, Crossland, Tsai, & Krackhardt, 2008). In other words, assuming that all three people are friends negates the need to remember each of the three ties individually. This heuristic is only effective, however, because it matches naturally occurring phenomena in real-world social networks. The tendency to over-perceive many of these naturally occurring features in one’s own social networks (Kilduff et al., 2008) could reflect the integration of such features into schemas regarding how social networks tend to be structured, which would subsequently impact how one learns new networks. Thus, someone who internalizes the properties of real-world social networks will likely be better able to learn new relationships that adhere to these expectations.

The taxing cognitive demands of navigating social networks may have led the human brain to consider social networks in a fundamentally different way than other types of networks. Indeed, there is evidence that social and non-social information is differentially processed in the brain (Meyer & Lieberman, 2016; Rusch & Charpentier, 2021; Stanley, 2016) and that individuals’ abilities to learn social and non-social relations are independent (Tompson, Kahn, Falk, Vettel, & Bassett, 2019). On the other hand, even if the brain evolved mechanisms to track specifically social networks, such mechanisms may have been co-opted to allow us to represent other, non-social networks in a similar manner, or vice versa. Consistent with this latter possibility, there is considerable overlap in the brain regions that encode social and spatial relations (Du et al., 2021; Parkinson & Wheatley, 2013, 2015), and cognitive processing of linguistic and social networks are similarly impacted by common structural properties (Lynn & Bassett, 2020). These findings

suggest that domain-general mechanisms may be used to represent social and non-social connections.

While most research on the mental representations of social networks, often referred to as cognitive social structures (Krackhardt, 1987), has focused on how humans in general track relationships, there is considerable variability in the accuracy of these representations across people (Brands, 2013; Marineau, Labianca, Brass, Borgatti, & Vecchi, 2018). It may be that people who hold certain positions within their own social networks are better able to learn new network structures. For example, brokers (i.e., individuals who act as bridges that connect otherwise unconnected others) may be particularly adept at understanding the overall structure of their social networks due to their unique role as the link between disparate groups (or sets of individuals) within the network. Indeed, Janicik & Larrick (2005) suggest that exposure to missing relationships in one's own social network facilitates easier learning of incomplete networks (i.e., not every person is connected to every other person). This may be because brokers are more likely to be aware of the structural hole that they occupy compared to their connections (Hahl et al., 2016).

Brokers hold privileged positions in their networks (Burt, 2021), are often high in self-monitoring (Oh & Kilduff, 2008), and have opportunities to wield exceptional power as they have access to information and resources from distinct groups (Burt, Kilduff, & Tasselli, 2013). The ability to exploit their structural advantage to control information flow and aggregation, however, relies on at least a rough understanding of where others sit in the network. Otherwise, brokers would not know who is already privy to, for example, specific pieces of information, and which resources they can control or leverage between groups. In other words, brokers would benefit greatly from maintaining accurate social network knowledge. It could also be that people who are

particularly attuned to their social environment may use this ability to navigate themselves into an advantageous position with high brokerage capacity. In either case, they are likely to be particularly attuned to how social networks tend to be structured and use that information to scaffold their learning of new networks.

Here, across four studies, we systematically tested the relationship between brokerage and network learning, as well as potential mechanisms underlying this relationship. In Study 1, we tested if brokerage was associated with enhanced learning of novel social networks, and if brokers achieve better learning by relying on schemas about how social networks tend to be structured. In Study 2, we tested if this association between brokerage and social network learning generalized to non-social networks (specifically, flight networks between airports). In Study 3, we tested the robustness of our results from Study 2 by controlling for task order and self-monitoring (a potential correlate of brokerage). In Study 4, we conducted a pre-registered test of the generalizability of the previous findings beyond college students in a larger sample. Across all four studies, we tested if brokers were better able to learn new networks, what mechanisms they used to learn new networks, and whether they approached social and non-social networks similarly.


Methods

In each of four studies, participants learned novel networks that were framed as consisting of either friendships between people (social condition) or flights between airports (non-social condition). Flight networks share many of the same structural properties that human social networks exhibit: they both tend to be small-world structures (Amaral, Scala, Barthelemy, & Stanley, 2000), which include closed triads, and all nodes (e.g., people, airports) can be reached from one another via a small number of hubs even though most nodes have relatively few connections to other nearby nodes (e.g., connections among regional airports; Barabási &

Bonabeau, 2003; Chi et al., 2003). All networks in Study 1 were framed as social networks of friends, while each participant in Studies 2-3 was assigned to either the social or non-social framing condition. Each participant learned two networks consisting of nine nodes (i.e., people or airports) and 8-15 edges (i.e., friendships between people or flights between airports) via the network learning task, then recalled the network by drawing it using custom software. Participants also completed self-report surveys, including an ego-network survey that was used to calculate their personal brokerage capacity, referred to here as participant brokerage. In all studies, the order in which participants learned the two networks were counter balanced¹. The order in which participants learned the networks and completed the ego-network survey were counterbalanced in Studies 3-4. For a summary of the methodological differences across studies, see Table 1.1.

Figure 1.1





Learned Network Structures



Note. Each participant learned one naturally-structured network (Network A or Network B) and one unnaturally-structured network (Network C, Network D). Naturally-structured networks (referred to here as “Natural Networks”) exhibited properties found in real-world social networks, such as closed triads, while unnaturally-structured networks (referred to here as “Unnatural Networks”) did not.

¹ Due to a technical error, the order in which Study 1 participants learned the networks was not fully counterbalanced across participants: 148 participants learned the unnatural network first and 92 learned the natural network first.

Table 1.1*Methods Across Studies*

	Study 1	Study 2	Study 3	Study 4
<i>N</i>	240	177	471	841
Participant Sample	Undergraduates	Undergraduates	Undergraduates	General Population
Format	In-lab	In-lab	Online	Online
Stimuli				
Framing Conditions	Social	Social, Non-social	Social, Non-social	Social, Non-social
Natural Network(s)	Network A	Network A	Network A	Networks A, B
Unnatural Network(s)	Network D	Network C	Network D	Networks C, D
Order of tasks	Ego-network after network learning and drawing tasks	Ego-network after network learning and drawing tasks	Ego-network and learning tasks counterbalanced	Ego-network and learning tasks counterbalanced
Individual differences	Ego-network, ASQ	Ego-network, ASQ	Ego-network, Self-monitoring	Ego-network, Self-monitoring

Participants

Participants were recruited from the University of California, Los Angeles (UCLA) undergraduate population using a human participant management platform (SONA systems, Studies 1-3), or from the general population of the United States using an online research recruitment platform, Prolific (<https://www.prolific.co>; Study 4). To maximize statistical power, we collected as many participants as possible through our departmental subject pool within one year for Studies 1-3. In Study 4, we recruited 1,000 participants (sample size, methods, and hypotheses were preregistered at <https://osf.io/34cru>). Participants were required to be fluent in

English (to ensure that they understood all study materials) and to be between the ages of 18 and 40. All recruited participants were consented in accordance with UCLA’s Institutional Review Board requirements and compensated with university credit (Studies 1-3) or paid \$10.50/hr for the 20-min study (Study 4). For a full description of sample demographics and exclusions, see Supplementary Materials S1.

Stimuli

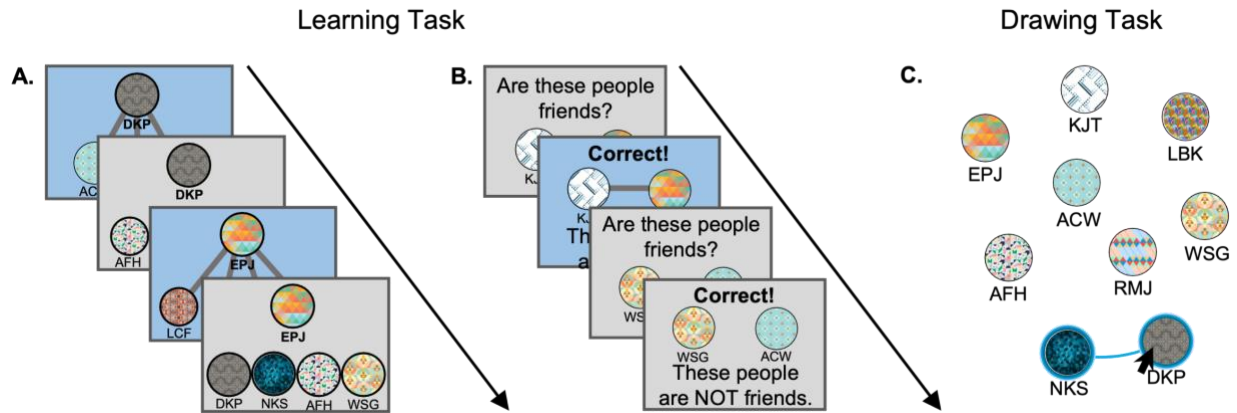
Network Structures

The artificial networks differed in how closely they resembled features of real-world social networks (Fig. 1.1). Specifically, some structures were more “natural” in that they were small-world networks (i.e., a network containing dense clusters that were sparsely interconnected, and each node is connected to each other node via a small number of edges), and consisted of closed triads (i.e., three people who are all friends). The “unnatural” network structures, on the other hand, consisted of only one community with no closed triads and minimal degree variation. Each participant learned one natural network structure and one unnatural network structure in a counterbalanced order (Table 1.1).

Images were randomly assigned to network positions across participants to disassociate image features from network positions. Three-letter identifiers (referred to as “initials” in the social condition and “airport codes” in the non-social condition) were randomly selected and assigned to each node per network per participant.

Figure 1.2

Experimental Tasks



Note. All studies used similar experimental paradigms, in which participants learned one naturally-structured network and one unnaturally-structured network structure (Fig. 1.1) via the learning task (A-B). The learning task consisted of a viewing phase (A) and a feedback phase (B). After learning each network, participants drew the network (C). All participants in Study 1 were told they would learn about a network of friendships between people (i.e., social condition), each of whom were represented by a facial photograph and initials. Participants in Studies 2-4 were assigned to either the social condition, as in Study 1, or the non-social condition in which the networks were framed as flights between airports. In Studies 2-4, all nodes (i.e., people or airports depending on condition) were represented by abstract patterns and three letter codes that were framed as initials in the social condition and airport codes in the non-social condition.

Procedure

Network Learning Task

Participants learned each network structure via a simple task consisting of a viewing phase (Fig. 1.2A) and a feedback phase (Fig. 1.2B). Participants repeated the viewing and feedback phases twice in Studies 1-2 and once in Studies 3-4 before completing the drawing task.

Viewing Phase. Participants were shown each node at the top of the screen along with that node's connections (i.e., alters) at the bottom of the screen (Fig. 1.2A). To avoid showing people who have many friends more often than people who have few friends (and thus confounding social network centrality with visual familiarity), each of these screens was followed by a similar screen showing everyone who is not friends with that person. This task was self-paced.

Feedback Phase. During the feedback phase, participants were asked to recall which pairs were connected and which were not (Fig. 1.2B). They saw every pair once per round in a fully randomized order and responded to the question, “Are these two people friends?” (social condition) or “Are these two airports connected by a direct flight?” (non-social condition). After each trial, participants received feedback telling them whether they were correct or incorrect as well as whether the two people were friends. To motivate learning, participants earned 1 point for every correct answer and lost 3 points for every incorrect answer; the total number of points earned was displayed in the top right corner of the screen.

Network Drawing Task

Participants were then told to draw the set of relationships that they had just learned to the best of their ability. They were presented with all nine nodes randomly spread out on the screen. Participants were able to rearrange the nodes by dragging the images around. Participants could draw or delete lines between two nodes to indicate if they were connected or not (Fig. 1.2C).

Individual Differences

Ego-Network Survey. Participants provided information about the structure of their own social networks—and thus, their personal brokerage capacity—using an ego-network approach. This consisted of a name generator adapted from past research (e.g., Parkinson et al., 2017, 2018); participants responded to the prompt, “Consider the people with whom you like to spend your time. Since you arrived at UCLA, who are the people with whom you socialize and/or discuss important matters?” (In Study 4, the phrase “Since you arrived at UCLA” was omitted.) Participants entered each individual’s first name, gender, and relationship type (in Studies 1-2, participants wrote in the relationship type, whereas in Studies 3-4, they selected from 11 possible labels such as “friend”, “parent”, or “schoolmate”; see Supplementary Materials S1). They

indicated which of these individuals were friends with each other using a matrix format in which each row and column corresponded to one of the people listed. Participants indicated whether two people were connected by checking or unchecking the cell corresponding to those two people's row and column. Two raters categorized the relationships reported in Studies 1-2 into the 11 categories of Studies 3-4. In Studies 1-2, participants also answered two items from the Norbeck Social Support Questionnaire (Norbeck, 1984) that asked, "How frequently do you contact this person?" (5-point Likert scale from "once a year or less" or "daily") and "How close do you feel to this person?" (3-point Likert scale from "a little close or not close" to "very close").

Other Self-Report Surveys. Participants completed other surveys, including a basic demographic survey (age, gender, race), the Autism Spectrum Quotient (ASQ) survey (Baron-Cohen, Wheelwright, Skinner, Martin, & Clubley, 2001) in Studies 1-2, and the Self-Monitoring Scale (Snyder, 1974) in Studies 3-4. The ASQ and demographic responses were included to facilitate exploratory analyses to inform future research and were not analyzed for this manuscript. The Self-Monitoring Scale was included so that we could control for the effects of self-monitoring, since it tends to be correlated with brokerage (Oh & Kilduff, 2008).

Analyses and Software

The learning task was presented using PsychoPy (Peirce, 2009) and the drawing task was presented using custom javascript code. All analyses were completed in R (Version 4.2.1; R Core Team, 2022). Linear mixed models were performed with the nlme package (Pinheiro, Bates, DebRoy, Sarkar, & R Core Team, 2019). Marginal slopes were estimated with the emmeans package (Lenth, 2019) and adjusted for multiple comparisons using false discovery rate correction (Benjamini & Hochberg, 1995). Results were plotted with ggplot2 (Wickham, 2016).

The igraph package (Csardi & Nepusz, 2006) was used to calculate each participant’s constraint score based on their responses to the personal social network survey. Constraint is an inverse measure of brokerage and is not normally distributed in the population or our sample. As such, participants’ constraint values were log transformed, then multiplied by -1. We refer to this transformed value as the participant’s brokerage throughout. To calculate participants’ accuracies in recalling the learned networks, we calculated the percent of all possible ties (i.e., 36 possible edges between all 9 nodes) that they correctly reported during the drawing task.

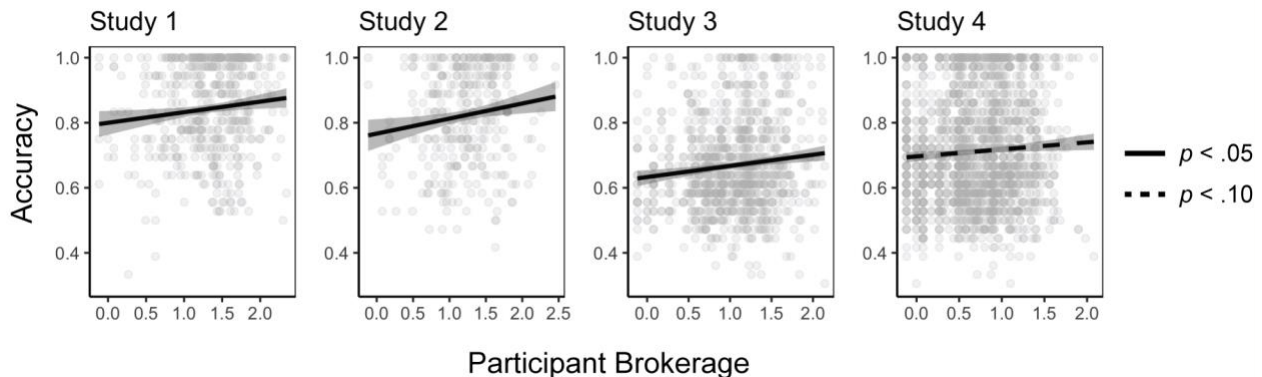
Results

Brokers are Better at Learning New Networks

To test if brokerage was associated with better learning of new network structures, we ran linear regressions with participant brokerage predicting drawing accuracy. We found that brokerage was significantly associated with higher accuracy in Study 1, $\beta = 0.03$, $F(1, 238) = 6.02$, $p = .015$, Study 2, $\beta = 0.05$, $F(1, 175) = 6.94$, $p = .009$, and Study 3, $\beta = 0.03$, $F(1, 492) = 8.58$, $p = .004$, and marginally so in Study 4, $\beta = 0.02$, $F(1, 855) = 3.08$, $p = .079$. This suggests that brokers are better at learning new network structures in general.

Figure 1.3

Brokers Are Better at Learning New Networks



Note. In Studies 1-3, brokerage was significantly predictive of recall accuracy as measured by the percent of all

possible ties correctly reported during the network drawing task. Study 4 also showed a positive association but was only marginally significant. 95% CI shown.

Brokers Focus on Ties That Exist Rather Than Missing Ties

The overall accuracy score consists of (i) ties that existed in the network (e.g., two people listed as friends) that participants either correctly drew (true positive, TP) or incorrectly left blank (false negative, FN) and (ii) missing ties (e.g., two people listed as not friends) that participants either correctly left blank (true negative, TN) or incorrectly drew (false positive, FP). To test if brokers were especially good at correctly reporting existing or missing ties, we calculated the true positive rate (TPR) and true negative rate (TNR) per participant per network. TPR was calculated as $\frac{TP}{TP+FN}$, or the proportion of all existing ties that were correctly drawn, and TNR was calculated as $\frac{TN}{TN+FP}$, or the proportion of all missing ties that were correctly left blank. We ran a linear mixed effects model with accuracy measure (TPR, TNR), participant brokerage, and their interaction predicting accuracy rate with by-participant intercepts. Indeed, we found a significant interaction between brokerage and accuracy measure, such that brokerage was associated with TPR but not TNR in every study (Fig. 1.4, Table 1.2).

Table 1.2

Linear Mixed Model Results Testing Type of Accuracy Measure Effects

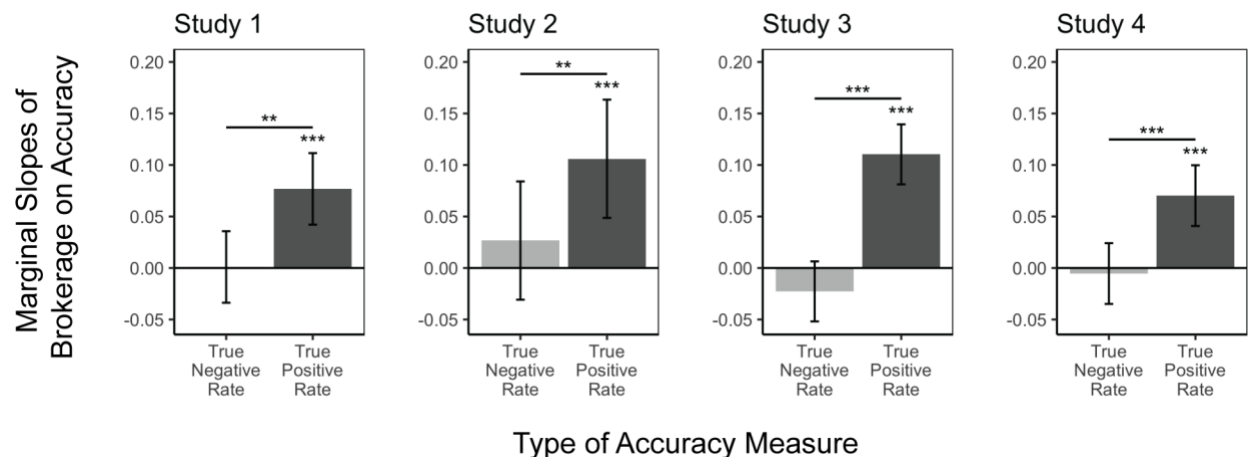
Effect	<i>df</i>	β	95% CI	<i>F</i>	<i>p</i>	
Study 1						
Brokerage	1, 238	0.04	[0.01, 0.06]	8.84	.003	**
Accuracy Measure	1, 718	0.12	[0.09, 0.16]	54.52	< .001	***
Brokerage x Accuracy Measure	1, 718	-0.04	[-0.06, -0.01]	10.28	.001	**
Study 2						
Brokerage	1, 175	0.07	[0.02, 0.12]	6.82	.010	**
Accuracy Measure	1, 529	0.17	[0.13, 0.20]	81.05	< .001	***
Brokerage x Accuracy Measure	1, 529	-0.04	[-0.07, -0.01]	7.85	.005	**

Effect	<i>df</i>	β	95% CI	<i>F</i>	<i>p</i>	
Study 3						
Brokerage	1, 492	0.04	[0.02, 0.07]	12.94	< .001	***
Accuracy Measure	1, 1480	0.27	[0.25, 0.29]	795.97	< .001	***
Brokerage x Accuracy Measure	1, 1480	-0.07	[-0.08, -0.05]	61.73	< .001	***
Study 4						
Brokerage	1, 855	0.03	[0.01, 0.06]	6.05	.014	*
Accuracy Measure	1, 2561	0.21	[0.20, 0.23]	1228.88	< .001	***
Brokerage x Accuracy Measure	1, 2561	-0.04	[-0.05, -0.02]	27.50	< .001	***

Note. Results from a linear mixed effects model with percent recall accuracy as the dependent variable, participant brokerage, accuracy measure (two levels: true positive rate, true negative rate) and their interaction as predictors, and random by-participant intercepts. Estimated marginal slopes of the interaction terms are illustrated in Fig. 1.4. *** $p < .001$, ** $p < .01$, * $p < .05$

Figure 1.4

Brokers Are Better at Learning Existing Ties Only



Note. Marginal estimated slopes of brokerage on accuracy for ties that exist in the network (true positive rate) and missing ties in the network (true negative rate). In all four studies, brokerage was significantly predictive of true positive rate but not true negative rate, suggesting brokers focus on ties that are present rather than those that are not. Error bars reflect 95% CI. *** $p < .001$, ** $p < .01$

Brokers' Network Learning Advantage Extends Beyond Social Networks

Do brokers learn social and non-social network differently? Are brokers better able to learn new networks by relying on their understanding of how networks tend to be structured in the real world? To test these questions, we ran linear mixed models with participant brokerage, network structure type (natural, unnatural), framing condition (social, non-social), and all their interactions

as predictors of accuracy, with by-participant random intercepts. Note, in Study 1, all networks were framed as social networks so framing condition was not included in the Study 1 model.

To test if brokers' advantage in learning new networks is specific to social networks, we examined the interactions between brokerage and framing condition. We found no significant two-way interactions between brokerage and framing condition, nor three-way interactions between brokerage, framing condition, and structure type (Table 1.3). That is, brokers' ability to learn new networks was not significantly impacted by the framing of the learned networks as social networks of friends or as flight networks between airports. This suggests that the network learning advantage related to brokerage is not domain-specific and extends at least to other kinds of networks that have similar structural characteristics to typical human friendship networks.

Table 1.3

Linear Mixed Model Results Testing Type of Network and Framing Condition Effects

Effect	<i>df</i>	β	95% CI	<i>F</i>	<i>p</i>	
Study 1						
Brokerage	1, 238	0.03	[0.01, 0.06]	7.49	.007	**
Structure	1, 238	0.05	[0.03, 0.08]	19.02	< .001	***
Structure x Brokerage	1, 238	0.02	[0.00, 0.04]	6.22	.013	*
Study 2						
Brokerage	1, 173	0.06	[0.01, 0.10]	6.80	.010	**
Structure	1, 173	-0.06	[-0.09, -0.03]	13.19	< .001	***
Condition	1, 173	0.02	[-0.04, 0.07]	0.42	.517	
Brokerage x Structure	1, 173	0.03	[0.00, 0.05]	5.60	.019	*
Brokerage x Condition	1, 173	-0.01	[-0.05, 0.03]	0.27	.605	
Structure x Condition	1, 173	0.02	[-0.02, 0.05]	0.89	.346	
Structure x Brokerage x Condition	1, 173	-0.01	[-0.04, 0.01]	1.30	.256	
Study 3						
Brokerage	1, 490	0.03	[0.01, 0.06]	8.58	.004	**
Structure	1, 490	0.01	[-0.01, 0.03]	1.03	.310	
Condition	1, 490	0.00	[-0.02, 0.03]	0.01	.897	
Brokerage x Structure	1, 490	0.02	[0.01, 0.04]	8.02	.005	**

Effect	<i>df</i>	β	95% CI	<i>F</i>	<i>p</i>
Brokerage x Condition	1, 490	0.00	[-0.02, 0.02]	0.02	.873
Structure x Condition	1, 490	0.00	[-0.02, 0.02]	0.00	.983
Structure x Brokerage x Condition	1, 490	-0.01	[-0.02, 0.01]	0.48	.489
Study 4					
Brokerage	1, 853	0.02	[0.00, 0.04]	2.87	.091 +
Structure	1, 849	-0.01	[-0.02, 0.00]	1.69	.194
Condition	1, 853	-0.01	[-0.03, 0.01]	1.53	.216
Brokerage x Structure	1, 849	0.00	[-0.01, 0.01]	0.05	.824
Brokerage x Condition	1, 853	0.00	[-0.02, 0.03]	0.09	.763
Structure x Condition	1, 849	0.00	[-0.01, 0.01]	0.15	.700
Structure x Brokerage x Condition	1, 849	0.00	[-0.01, 0.01]	0.02	.875

Note. Results from a linear mixed effects model with percent recall accuracy as the dependent variable. Participant brokerage, network structure (natural, unnatural), and framing condition (social, non-social) are included as fixed effects (when applicable) with random by-participant intercepts. These patterns of results hold when controlling for self-monitoring and effort (see Supplementary Materials S1). *** $p < .001$, ** $p < .01$, * $p < .05$, + $p < .1$

Brokers Rely on Schemas Consistent with Real-World Networks

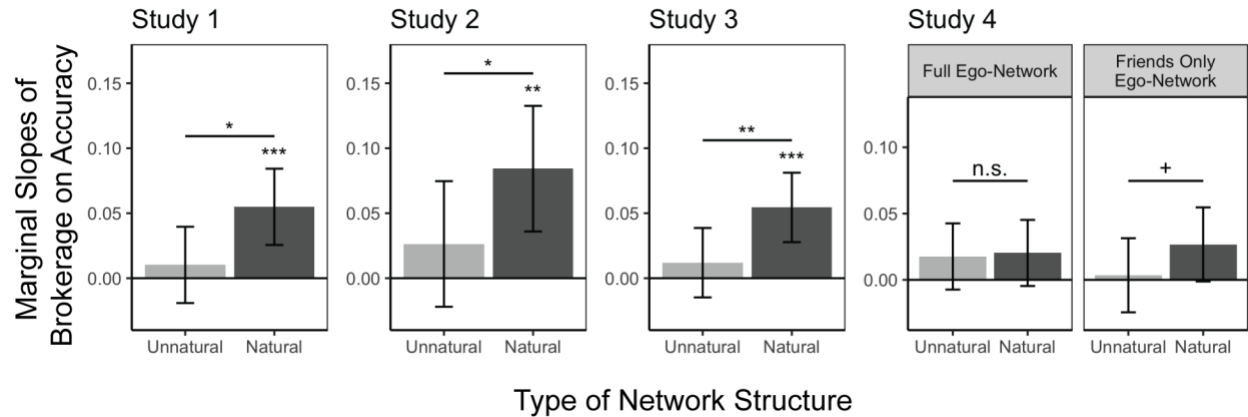
In the above models, we also tested if participants' level of brokerage predicted their recall of two types of network structures: those that were naturally-structured and others that were unnaturally-structured. Indeed, we found significant interactions between brokerage and structure type in Studies 1-3 but not in Study 4 (Table 1.3). We calculated the estimated marginal slopes of brokerage on accuracy for each type of structure and found that brokerage was only associated with better learning of the natural network type, consistent with the hypothesis that brokers incorporate structural properties of real-world networks into schemas that scaffold their learning and representation of new network structures (Fig. 1.5). In other words, people who held high-brokerage positions in their own social networks were particularly good at learning new social structures if and only if they were characterized by naturally occurring social network properties. This suggests that brokers understand how social networks tend to be structured in the real world, and that they use this knowledge to help them learn new social networks, resulting in brokers having an advantage at learning structures that adhere to these priors. It is perhaps unsurprising

that we do not see any moderating effect of framing condition (i.e., social or non-social) because flight networks of airports are structured very similarly to social networks (e.g., they have hubs that act as brokers between clusters of other airports). As such, brokers' learning of social network structures would translate well to flight networks, making this a very conservative test of domain specificity.

To ensure that these results were not due to other factors, we re-ran the models while controlling for self-monitoring and a measure of the amount of effort that participants exerted in the task. Given that brokerage is often associated with self-monitoring (Oh & Kilduff, 2008), it could be that self-monitors are particularly attuned to how networks are structured, and that brokerage is not in itself associated with network learning. When controlling for self-monitoring in the above models, however, we find no significant changes in our results (Table S1.2). It could also be that our measure of participant brokerage measured attention instead: participants who were more attentive may have been more likely to carefully report their ego-networks. This heightened attentiveness could also lead them to both learn the network better and to take more time in the learning and drawing tasks. To test this, we ran the above models again while controlling for the average amounts of time participants spent on each trial of the self-paced learning and drawing tasks. Again, we found no significant changes (Table S1.3). These additional tests further suggest that our findings specifically reflect the enhanced ability of brokers to learn new networks.

Figure 1.5

Brokers Are Especially Good at Learning Naturally Structured Networks



Note. Marginal estimated slopes of brokerage on recall accuracy of networks categorized as Natural and Unnatural. In the first three studies, brokerage was significantly predictive of better learning in the natural network but not the unnatural network. There was no difference, however, in Study 4, which sampled a different population. When calculating brokerage based on participants' friendships only (the vast majority of participants' reported relationships were friendships in Studies 1-3 but not 4), however, we find a trending difference in line with results from Studies 1-3. Error bars reflect 95% CI. *** $p < .001$, ** $p < .01$, * $p < .05$, + $p < .1$, n.s. = not significant

To better understand why this pattern of results was so consistent across Studies 1-3 but not present in Study 4, we examined the types of relationships that participants reported in their ego-networks. Some types of relationships are more elective (e.g., friendships), and thus the patterns of these relationships are more likely to be driven by one's intrinsic social tendencies, while other types of relationships (e.g., family ties) are more driven by external factors. As such, one's position in a network comprised mainly of the latter type of relationships would presumably be less related to how that individual navigates and constructs their social network. Indeed, we found a striking difference between Studies 1-3 and Study 4 in the types of relationships reported. We found that in Studies 1-3, which consisted of college student samples, most people they listed in their ego-networks fell in the "friend" category (Study 1: $M = 72\%$, $SD = 23\%$; Study 2: $M = 72\%$, $SD = 23\%$; Study 3: $M = 68\%$, $SD = 25\%$; Study 4: $M = 43\%$, $SD = 31\%$), while participants in Study 4, which was sampled from the general population via Prolific, reported more

family members (Study 1: $M = 19\%$, $SD = 17\%$; Study 2: $M = 21\%$, $SD = 21\%$; Study 3: $M = 26\%$, $SD = 24\%$; Study 4: $M = 51\%$, $SD = 31\%$) than friends (see Supplementary Materials S1). To test the possibility that brokerage within friendship networks drove the above results, we calculated Study 4 participants' brokerage capacities among their friends only and re-ran the mixed model. In this post-hoc model, we found a trending interaction between friendship-only brokerage and structure type such that brokerage was more positively associated with learning the natural networks than the unnatural networks (Fig. 1.5). This suggests that there may be distinct differences between brokerage calculated based on freely chosen friendships compared to predetermined familial ties.

Brokers Are Particularly Attuned to Other Brokers

Thus far, we have found that brokers are especially good at learning new networks, that they do so by focusing on ties that exist rather than missing ties, and by using real-world properties as domain-general schemas. To further understand the cognitive mechanisms used to learn these naturally-structured networks, we tested if participants focused on specific people in those networks. Specifically, we tested if brokers were particularly attuned to other brokers. To do so, we calculated the brokerage of each node in the two naturally structured networks and scaled them to be between 0 and 1 (Fig. 1.6A). We then ran a linear mixed model with by-participant random intercepts, and participant brokerage, target brokerage, and their interaction predicting recall accuracy of that target's connections (both existing and missing). In all four studies, we found a significant interaction such that higher participant brokerage was associated with better learning of target nodes with higher brokerage capacity (Table 1.4, Fig. 1.6B). That is, brokers were especially good at learning other brokers' ties in naturally structured networks.

Importantly, as in real-world networks, target brokerage was highly correlated with target degree (i.e., the number of connections they have). Given that participant brokerage is associated with better learning of existing ties rather than missing ties, it may be that this is driving the positive interaction found above. Thus, we ran the mixed model again while controlling for target degree and found the same pattern of results (Table S1.4, Fig. S1.4).

Table 1.4

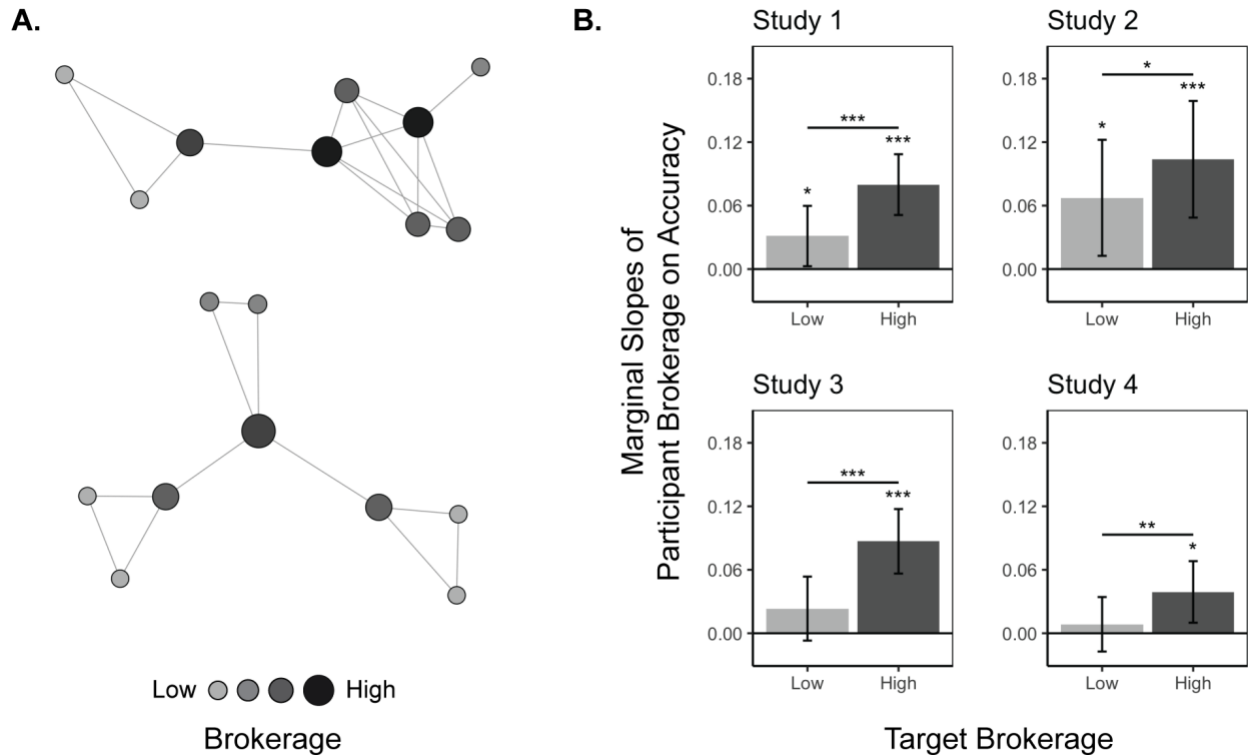
Linear Mixed Model Results Testing Target Brokerage Effects

Effect	<i>df</i>	β	95% CI	<i>F</i>	<i>p</i>	
Study 1						
Participant Brokerage	1, 238	0.03	[0.00, 0.06]	4.91	.028	*
Target Brokerage	1, 1918	-0.20	[-0.24, -0.15]	77.54	< .001	***
Participant Brokerage x Target Brokerage	1, 1918	0.07	[0.04, 0.10]	20.38	< .001	***
Study 2						
Participant Brokerage	1, 175	0.07	[0.01, 0.12]	6.01	.015	*
Target Brokerage	1, 1414	-0.25	[-0.32, -0.18]	47.79	< .001	***
Participant Brokerage x Target Brokerage	1, 1414	0.05	[0.00, 0.11]	3.95	.047	*
Study 3						
Participant Brokerage	1, 469	0.02	[-0.01, 0.05]	2.50	.114	
Target Brokerage	1, 3766	-0.36	[-0.40, -0.32]	302.84	< .001	***
Participant Brokerage x Target Brokerage	1, 3766	0.09	[0.06, 0.13]	26.34	< .001	***
Study 4						
Participant Brokerage	1, 837	0.01	[-0.02, 0.03]	0.35	.556	
Target Brokerage	1, 6710	-0.16	[-0.17, -0.14]	375.88	< .001	***
Participant Brokerage x Target Brokerage	1, 6710	0.03	[0.01, 0.05]	8.74	.003	**

Note. Results from a linear mixed effects model with percent recall accuracy of each target node’s connections in the natural networks as the dependent variable, participant brokerage, target brokerage, and their interaction as predictors, and random by-participant intercepts. In Study 4, where there were two naturally structured networks, target brokerage was scaled to [0, 1] within each network. Estimated marginal slopes of the interaction terms are illustrated in Fig. 1.6. *** $p < .001$, ** $p < .01$, * $p < .05$

Figure 1.6

Brokers Are Especially Good at Learning Other Brokers' Ties



Note. Within the natural networks (A), the association between participant brokerage and better recall was significantly stronger for targets with higher brokerage capacity compared to targets with lower brokerage capacity (B). This was true in all four studies. Error bars reflect 95% CI. *** $p < .001$, ** $p < .01$, * $p < .05$

Discussion

The challenges associated with organizing into and navigating large, complex social networks were likely driving forces in the evolution of the human brain, and the ability to successfully meet such challenges may have fostered humans' ability to thrive across the planet (Dunbar, 2014; Dunbar & Shultz, 2007). Still, the cognitive demands of learning, remembering, and tracking large numbers of relationships would likely be unmanageable if people did not rely on general frameworks regarding how social networks tend to be structured. Here, we found that brokers (i.e., those who connect disparate people or groups that would not otherwise be connected to each other) were especially good at learning new networks that resemble the structure of

naturally occurring networks (i.e., those that adhere to schemas that could be learned through observation of real-world networks). This advantage did not persist when the network was unnaturally-structured, and it was not significantly affected by framing the network as a social network of friends vs. a non-social flight network between airports. Furthermore, we found that brokers were able to learn new networks by focusing on existing connections, as opposed to non-existent ties, and the connections to especially important nodes, like other brokers.

We did not find any significant effect of framing the networks as social or non-social on brokerage and learning, suggesting this advantage that brokers have is not domain-specific. It is difficult to interpret the null effect beyond this, but recent behavioral and neural evidence suggests that generally people can learn non-social network structures that are modular in structure (such as our naturally-structured networks) better than lattice structures (such as our unnaturally-structured networks) (Kahn, Karuza, Vettel, & Bassett, 2018; Lynn & Bassett, 2020; Lynn, Kahn, Nyema, & Bassett, 2020; Lynn, Papadopoulos, Kahn, & Bassett, 2020). While this recent work does not investigate participants' own network position, it suggests that the overall schemas used by humans (especially by brokers, as we show here) may be applicable across many domains. Importantly, we used a very conservative test when comparing social and non-social networks given that flight networks tend to have the same type of structure as social networks. Thus, brokers' schema for social network structures would translate well to flight networks. Future research could further test for domain specificity or domain generality using networks that have distinct real-world properties.

Research examining how humans and other highly social animals learn linear status hierarchies suggests that they are learned through observing and participating in encounters between pairs of individuals, then using transitive inference to ascertain the relative status of those

whom one has not seen interacting directly (Cheney & Seyfarth, 1990; Gazes, Hampton, & Lourenco, 2017; Paz-y-Miño C, Bond, Kamil, & Balda, 2004; White & Gowan, 2013). Similarly, people may learn about others' positions in social networks (which are more complex in that they cannot be reduced to a simple ordinal hierarchy) by observing some social relationships and making inferences regarding other social relationships about which they are uncertain (i.e., based on prior schemas regarding how social networks tend to be structured). Brokers seem to be particularly adept at doing so, given that brokerage was associated with (i) higher rates of recalling edges that exist in social networks—suggesting that they are exceptionally adept at learning what relationships *do* exist—and (ii) better learning of social networks with natural, compared to unnatural, structures—suggesting that brokers effectively use schemas about how social networks tend to be structured to fill in gaps in their knowledge.

There are many possible explanations for why brokers may have this advantage: Brokers could have generally heightened awareness of and attention to social information, either because of their social network position, or vice versa. Additionally, people may expect new social worlds to resemble their own and be particularly good at navigating new social spaces that fit their expectations. Thus, for brokers, who bridge gaps between others in a network, the naturally-structured networks may match their own experiences particularly closely, because they include groups of people with minimal inter-group connections. Indeed, when we examined which people and relationships they were particularly adept at remembering, we found that brokers were especially attuned to other brokers and their ties, further suggesting that it is easier to learn about social networks that match one's own experiences. It could also be that people who are particularly good at learning new social networks are able to navigate into powerful positions in social networks, and thus become brokers *because* of this ability. Future research is needed to fully

identify and test the reasons why brokers demonstrate this advantage in learning new social networks and the mechanisms they use to do so.

In Study 4, we tested the generalizability of our results beyond undergraduate students and found some key differences. For one, recall accuracy was only marginally associated with participants' brokerage capacity. Second, we found no difference in the strength of this association for the naturally- and unnaturally-structured networks. After a post-hoc examination of the differences in the types of relationships participants included in their ego-networks, we found that studies 1-3 nominated friends as a vast majority of their relationships, while Study 4, which sampled a larger age range from the general population, nominated friends as less than half of their relationships, instead nominating family more often. When we calculated brokerage based on friendships only in Study 4, we saw a pattern emerge that was closer to the previous studies. This suggests that the ties on which certain positional characteristics are calculated may drastically shape the results found.

This may be particularly true of brokerage in which a structural hole among friends may be qualitatively different than one among family. Furthermore, here, and in much of past research, brokerage was calculated based on one's ego-network, whereas brokerage based on full network models may provide different insights (Kwon, Rondi, Levin, De Massis, & Brass, 2020). Finally, we believe these findings may reflect the difference between elective ties (e.g., friendships) and ties that are formed beyond one's control (e.g., familial relationships) or based on external factors (e.g., work relationships). One's intrinsic social tendencies, abilities, and preferences may have more opportunities to impact an individual's position (e.g., brokerage) in networks consisting of almost exclusively of friendships that people select themselves compared to networks based on external factors. Thus, cases in which intrinsic factors (e.g., being particularly socially

attuned/skilled) are shaping network position more strongly may provide a more sensitive test of our hypotheses. Future research examining these possibilities would be valuable in further elucidating links between the structure of one's own social world and how one learns about and represents novel social information.

Future research would also benefit from looking at how people learn larger networks. Here, we presented participants with small networks, consisting of only nine nodes. The average individual's social network consists of many more people (Gonçalves, Perra, & Vespignani, 2011; Hill & Dunbar, 2003; Roberts, Dunbar, Pollet, & Kuppens, 2009), with exponentially more possible ties. It may be that, in our study, the networks were not difficult enough for participants to learn for us to be able to pick up on subtle differences stemming from, for example, framing the connections as social or non-social in nature. If the networks were much larger, brokers may be forced to rely on heuristics more heavily, resulting in a greater number of systematic errors that could then be analyzed more thoroughly, potentially illuminating the mechanisms of this network learning advantage.

Conclusion

As humans, we lead highly social lives. Navigating and tracking the intricate social networks in which we each live can be a computationally daunting task. To our knowledge, this is the only study directly linking one's personal social network position to their ability to learn new networks. Here, we illuminated how one's own social network experience uniquely relates to the ability to learn entirely new networks, examined the mechanisms through which this occurs, and generated questions for future research to pursue.

Supplementary Materials S1

Sample Characteristics

In Study 1, 250 participants were recruited from the University of California, Los Angeles (UCLA) undergraduate population. Ten participants were excluded due to technical issues or researcher error that prevented complete data collection. As such, our final sample size was 240 (178 women, 61 men, 1 non-binary; ages 18-32 years, $M = 20.23$, $SD = 1.56$).

In Study 2, 196 participants were recruited from the UCLA undergraduate population. Fourteen participants were excluded due to technical issues that prevented complete data collection. Five participants were excluded for responding in the same way on all or all-but-one of the trials in the network learning task, suggesting they were not trying to answer correctly. Thus, our final sample size was 177 (126 women, 46 men, 1 non-binary, 4 did not report gender;).

In Study 3, 600 online participants were recruited from the UCLA undergraduate population. Twenty-three participants were excluded due to technical issues that prevented complete data collection. Eighty-three participants were excluded for reporting that they paid little or no attention during the task. Twenty-three participants were excluded for taking more than 3x the estimated length of the study. Thus, our final sample size was 471 (373 women, 91 men, 5 non-binary, 2 did not report gender; ages 18-36 years, $M = 20.25$, $SD = 2.13$).

In Study 4, 1,000 online participants were recruited from the Prolific (www.prolific.co) population. Nineteen participants were excluded for reporting that they paid little or no attention during the task. Nine participants were excluded for responding in the same way on all or all-but-one of the trials in the network learning task, suggesting they were not trying to answer correctly. Two participants were excluded pressing the attention check key twice as often as it was needed. Sixteen participants were excluded for taking more than more than three times the estimated length

of the study. Thus, our final sample size was 841 (355 women, 456 men, 22 non-binary, 8 did not report gender; ages 18-42 years, $M = 29.85$, $SD = 5.92$).

Table S1.1

Racial/Ethnic Demographics of Participants

Study	Asian	Black	Latinx	Middle Eastern	Native American	White	Multi-Racial	Not Listed	Prefer Not to Answer
Study 1	81	5	53	13	1	59		23	5
Study 2	89	5	33	6	0	28		11	1
Study 3	175	13	96	27	3	95	45	6	11
Study 4	82	82	104	2	3	510	43	1	14

Study 1 Images

In Study 1, each node in the network was represented by a neutral face from the Chicago Face Database (Ma et al., 2015). These images were cropped to a circle and edited such that the gray shirts were colored to aid in distinguishing between people. Across participants, images were semi-randomly assigned to network structures while controlling for gender and race across structures. Each network consisted of one female and one male of each of the four ethnic categories in the Chicago Face Database: Asian, Black, Latinx, White. The ninth node was represented by either a Black man or a Latina woman. This semi-random allocation was used to prevent any potential confounds that inter-image similarity may have on participants' perceived connections. Within each network, images were randomly assigned to network positions. As such, images were not consistently associated with a given position within a network.

Initials/Airport Codes

There were 28 three-letter codes that were framed as either initials (social condition) or airport codes (non-social condition). These were selected from all possible combinations such that

they could realistically represent initials, were dissimilar from any major airports (to prevent the learned information potentially competing with any participants' prior knowledge in Study 2) and were not common abbreviations for English words or phrases. The full list is: WSG, SBJ, JGC, ENM, DWM, MBF, NWB, RLA, LBK, AFH, LWP, LCF, DBG, RMJ, KJT, TDH, HGP, EBH, EKF, GFT, NFD, NKS, EAG, MLH, RNP, EPJ, ACW, DKP.

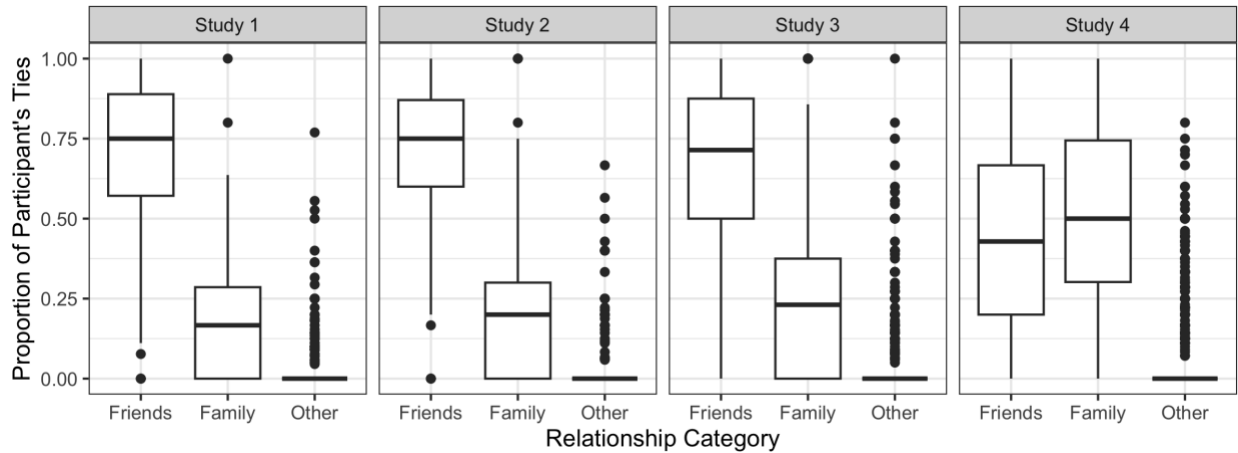
Ego-Network Relationships

In Studies 3-4, participants listed the people in their ego-network and selected one of the following 11 labels (modified from Cohen, Doyle, Skoner, Rabin, & Gwaltney, 1997) that best reflected their relationship with that person: friend, spouse, child, parent, parent-in-law, other close family member, schoolmate, workmate, close neighbors, member of religious group, member of group without religious affiliations.

Two research assistants coded the free responses that were given in Studies 1-2 according to these 11 labels. The two coders were in high agreement ($IRR = .97$). We then categorized these labels into three categories: friends, family, and other. We then calculated the proportion of each participant's ties that fell into each of those categories (Fig. S1.1)

Figure S1.1

Proportions of Relationships in Ego-Network



Note. Boxplots illustrating the proportion of each participant's ties that were friends, family, or other.

The Potential Confound of Self-Monitoring

Table S1.2

Linear Mixed Model Results Testing Type of Network and Framing Condition Effects

Controlling for Self-Monitoring

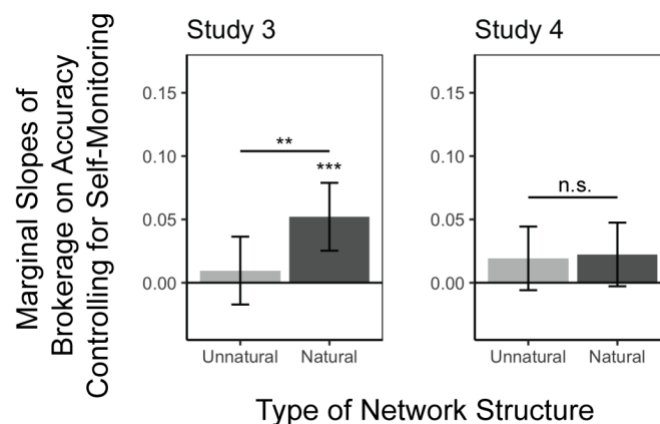
Effect	<i>df</i>	β	95% CI	<i>F</i>	<i>p</i>	
Study 3						
Brokerage	1, 489	0.03	[0.01, 0.05]	7.33	.007	**
Structure	1, 490	0.01	[-0.01, 0.03]	1.03	.310	
Condition	1, 489	0.00	[-0.02, 0.03]	0.02	.885	
Self-Monitoring	1, 489	0.00	[0.00, 0.01]	3.00	.084	+
Brokerage x Structure	1, 490	0.02	[0.01, 0.04]	8.02	.005	**
Brokerage x Condition	1, 489	0.00	[-0.02, 0.02]	0.02	.890	
Structure x Condition	1, 490	0.00	[-0.02, 0.02]	0.00	.983	
Structure x Brokerage x Condition	1, 490	-0.01	[-0.02, 0.01]	0.48	.489	
Study 4						
Brokerage	1, 847	0.02	[0.00, 0.04]	3.38	.066	+
Structure	1, 848	-0.01	[-0.02, 0.00]	1.95	.163	
Condition	1, 847	-0.01	[-0.03, 0.01]	1.78	.182	
Self-Monitoring	1, 847	0.00	[0.00, 0.00]	0.33	.564	

Effect	<i>df</i>	β	95% CI	<i>F</i>	<i>p</i>
Brokerage x Structure	1, 848	0.00	[-0.01, 0.01]	0.07	.798
Brokerage x Condition	1, 847	0.00	[-0.02, 0.03]	0.17	.680
Structure x Condition	1, 848	0.00	[-0.01, 0.01]	0.24	.622
Structure x Brokerage x Condition	1, 848	0.00	[-0.01, 0.01]	0.07	.795

Note. Results from a linear mixed effects model with percent recall accuracy as the dependent variable. Participant brokerage, network structure (natural, unnatural), and framing condition (social, non-social) are included as fixed effects with random by-participant intercepts. Self-monitoring was included as a covariate. $**p < .01$, $+p < .1$

Figure S1.2

Brokerage Association With Accuracy When Controlling for Self-Monitoring



Note. Marginal estimated slopes of brokerage on recall accuracy of networks categorized as Natural and Unnatural when controlling for self-monitoring. The pattern of results is the same when controlling for self-monitoring and when it is not included in the model. Error bars reflect 95% CI. $***p < .001$, $**p < .01$, n.s. = not significant

Accounting for the Potential Confound of Effort

It is possible that the association of brokerage with network learning is not due to actual network position, but rather the amount of time and effort participants are willing to put into the network position, but rather the amount of time and effort participants are willing to put into the study. That is, participants who are more attentive and spend more time filling in the ego network survey may have spent more time learning and recalling the network, and therefore performed better when recalling the network in the drawing task. To test this possibility, we ran the same linear mixed models with the following covariates: average time spent on the viewing trials of the learning task, average time spent on the feedback trials of the learning task, and time spent drawing the network. The results show the same pattern as seen without including these covariates. That is,

even when controlling for the amount of time the participant spent on the task at each stage, we still see a strong association between brokerage and ability to learn naturally-structured social networks, but no other network structure.

Table S1.3

Linear Mixed Model Results Testing Type of Network and Framing Condition Effects

Controlling for Effort

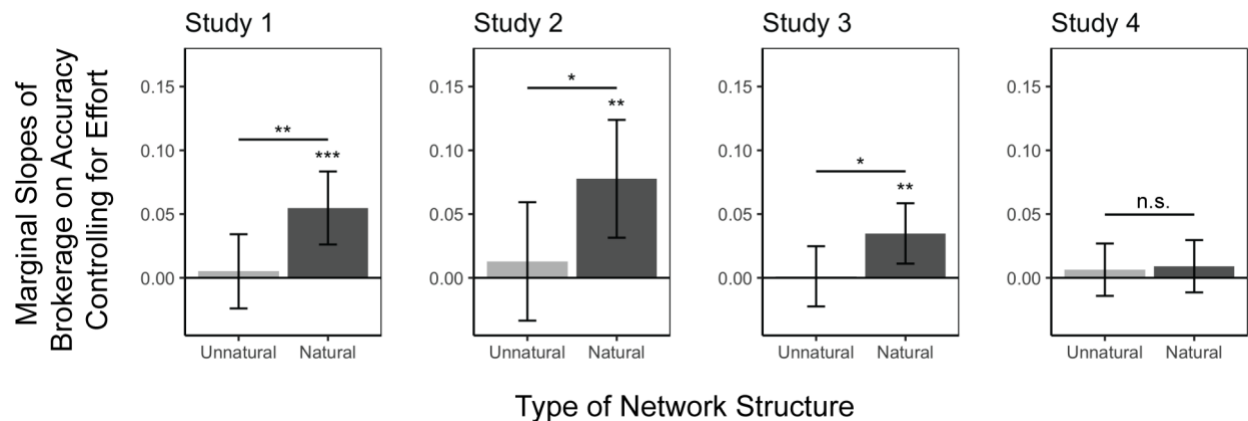
Effect	<i>df</i>	β	95% CI	<i>F</i>	<i>p</i>	
Study 1						
Brokerage	1, 238	0.03	[0.01, 0.05]	6.87	.009	**
Structure	1, 235	0.05	[0.04, 0.09]	22.43	< .001	***
Time on Viewing Phase	1, 235	0.00	[0.00, 0.01]	8.83	.003	**
Time on Feedback Phase	1, 235	0.00	[-0.01, 0.01]	0.00	.948	
Time on Drawing Task	1, 235	0.00	[0.00, 0.00]	1.99	.159	
Structure x Brokerage	1, 235	0.02	[0.01, 0.04]	7.31	.007	**
Study 2						
Brokerage	1, 173	0.05	[0.01, 0.08]	5.21	.024	*
Structure	1, 170	-0.06	[-0.10, -0.03]	15.02	< .001	***
Condition	1, 173	0.03	[-0.02, 0.08]	1.09	.298	
Time on Viewing Phase	1, 170	0.01	[0.00, 0.01]	8.67	.004	**
Time on Feedback Phase	1, 170	0.01	[0.00, 0.02]	2.40	.123	
Time on Drawing Task	1, 170	0.00	[0.00, 0.00]	5.19	.024	*
Brokerage x Structure	1, 170	0.03	[0.01, 0.06]	6.68	.011	*
Brokerage x Condition	1, 173	-0.02	[-0.06, 0.02]	0.85	.359	
Structure x Condition	1, 170	0.02	[-0.01, 0.05]	1.34	.249	
Structure x Brokerage x Condition	1, 170	-0.02	[-0.04, 0.01]	1.80	.181	
Study 3						
Brokerage	1, 490	0.02	[0.00, 0.04]	3.90	.049	*
Structure	1, 487	0.01	[0.00, 0.03]	2.08	.150	
Condition	1, 490	0.00	[-0.02, 0.02]	0.01	.919	
Time on Viewing Phase	1, 487	0.00	[0.00, 0.00]	1.13	.288	
Time on Feedback Phase	1, 487	0.01	[0.01, 0.01]	29.38	< .001	***
Time on Drawing Task	1, 487	0.00	[0.00, 0.00]	116.94	< .001	***
Brokerage x Structure	1, 487	0.02	[0.00, 0.03]	4.57	.033	*
Brokerage x Condition	1, 490	0.00	[-0.01, 0.02]	0.11	.744	
Structure x Condition	1, 487	0.00	[-0.02, 0.01]	0.27	.601	

Effect	<i>df</i>	β	95% CI	<i>F</i>	<i>p</i>
Structure x Brokerage x Condition	1, 487	0.00	[-0.02, 0.01]	0.03	.855
Study 4					
Brokerage	1, 853	0.01	[-0.01, 0.02]	0.82	.365
Structure	1, 842	-0.01	[-0.02, 0.00]	1.75	.186
Condition	1, 853	0.00	[-0.02, 0.01]	0.40	.529
Time on Viewing Phase	1, 842	0.00	[0.00, 0.00]	2.40	.121
Time on Feedback Phase	1, 842	0.02	[0.02, 0.02]	317.64	< .001 ***
Time on Drawing Task	1, 842	0.00	[0.00, 0.00]	54.74	< .001 ***
Brokerage x Structure	1, 842	0.00	[-0.01, 0.01]	0.05	.823
Brokerage x Condition	1, 853	0.00	[-0.02, 0.01]	0.11	.744
Structure x Condition	1, 842	0.00	[-0.01, 0.01]	0.32	.574
Structure x Brokerage x Condition	1, 842	0.00	[-0.01, 0.02]	0.33	.569

Note. Results from a linear mixed effects model with percent recall accuracy as the dependent variable. Participant brokerage, network structure (natural, unnatural), and framing condition (social, non-social) are included as fixed effects (when applicable) with random by-participant intercepts. Time spent on the viewing and feedback phases of the learning task and the drawing task were included as covariates. *** $p < .001$, ** $p < .01$, * $p < .05$

Figure S1.3

Brokerage Association With Accuracy When Controlling for Time Spent on Tasks



Note. Marginal estimated slopes of brokerage on recall accuracy of networks categorized as Natural and Unnatural when controlling for effort. Effort was measured by the time spent on the self-paced portions of the network learning task (viewing phase and feedback phase) and the drawing task. The pattern of results is the same when controlling for effort and when it is not included in the model. Error bars reflect 95% CI. *** $p < .001$, ** $p < .01$, * $p < .05$, n.s. = not significant

Controlling for Target Degree

Table S1.4

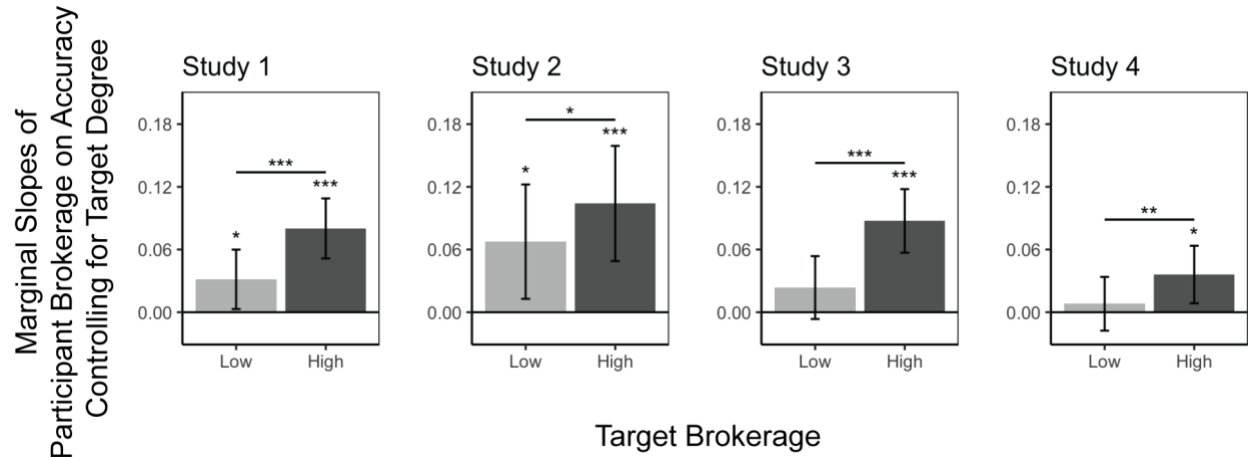
Linear Mixed Model Results Testing Target Brokerage Effects Controlling for Target Degree

Effect	<i>df</i>	β	95% CI	<i>F</i>	<i>p</i>	
Study 1						
Participant Brokerage	1, 238	0.03	[0.00, 0.06]	4.74	.030	*
Target Brokerage	1, 1917	-0.07	[-0.10, -0.03]	13.24	< .001	***
Target Degree	1, 1917	-0.08	[-0.11, -0.06]	35.29	< .001	***
Participant Brokerage x Target Brokerage	1, 1917	0.05	[0.03, 0.07]	20.75	< .001	***
Study 2						
Participant Brokerage	1, 175	0.07	[0.01, 0.12]	5.93	.016	*
Target Brokerage	1, 1413	-0.06	[-0.12, -0.01]	4.68	.031	*
Target Degree	1, 1413	-0.13	[-0.17, -0.08]	32.58	< .001	***
Participant Brokerage x Target Brokerage	1, 1413	0.04	[0.00, 0.07]	4.04	.045	*
Study 3						
Participant Brokerage	1, 469	0.02	[-0.01, 0.05]	2.37	.124	
Target Brokerage	1, 3765	-0.07	[-0.11, -0.04]	15.87	< .001	***
Target Degree	1, 3765	-0.22	[-0.25, -0.18]	188.45	< .001	***
Participant Brokerage x Target Brokerage	1, 3765	0.06	[0.04, 0.09]	27.66	< .001	***
Study 4						
Participant Brokerage	1, 837	0.01	[-0.02, 0.03]	0.38	.538	
Target Brokerage	1, 6709	-0.01	[-0.04, 0.02]	0.53	.466	
Target Degree	1, 6709	-0.16	[-0.19, -0.13]	129.10	< .001	***
Participant Brokerage x Target Brokerage	1, 6709	0.03	[0.01, 0.05]	8.72	.003	**

Note. Results from a linear mixed effects model with percent recall accuracy of each target node's connections in the natural networks as the dependent variable, participant brokerage, target brokerage, and their interaction as predictors, and random by-participant intercepts. In Study 4, where there were two naturally structured networks, target brokerage and target degree were scaled to [0, 1] within each network. Estimated marginal slopes of the interaction terms are illustrated in Fig. S1.2. *** $p < .001$, ** $p < .01$, * $p < .05$

Figure S1.4

Brokers Are Especially Good at Learning Other Brokers' Ties Even When Controlling for Degree



Note. Within the natural networks, the association between participant brokerage and better recall was significantly stronger for targets with higher brokerage capacity compared to targets with lower brokerage capacity when controlling for target's degree. This was true in all four studies. Error bars reflect 95% CI. *** $p < .001$, ** $p < .01$, * $p < .05$

CHAPTER 2

Neural encoding of novel social networks:
Evidence that perceivers prioritize others' centrality

Miriam E. Schwyck¹, Meng Du¹, Pratishta Natarajan², John Andrew Chwe², Carolyn Parkinson¹

¹University of California, Los Angeles

²New York University

This is a reprint of the published manuscript:

Schwyc, M.E., Du, M., Natarajan, P., Chwe, J.A., & Parkinson, C. (2023). Neural encoding of novel social networks: Evidence that perceivers prioritize others' centrality. *Social Cognitive and Affective Neuroscience*, 18(1) 1-9.

Abstract

Knowledge of someone's friendships can powerfully impact how one interacts with them. Past research suggests that information about others' real-world social network positions—e.g., how well-connected they are (centrality), 'degrees of separation' (relative social distance)—is spontaneously encoded when encountering familiar individuals. However, many types of information covary with where someone sits in a social network. For instance, strangers' face-based trait impressions are associated with their social network centrality, and social distance and centrality are inherently intertwined with familiarity, interpersonal similarity, and memories. To disentangle the encoding of social network position from other social information, participants learned a novel social network in which social network position was decoupled from other factors, then saw each person's image during functional magnetic resonance imaging scanning. Using representational similarity analysis, we found that social network centrality was robustly encoded in regions associated with visual attention and mentalizing. Thus, even when considering a social network in which one is not included and where centrality was unlinked from perceptual and experience-based features to which it is inextricably tied in naturalistic contexts, the brain encodes information about others' importance in that network, likely shaping future perceptions of and interactions with those individuals.

Introduction

When encountering a stranger, the human brain spontaneously encodes specific pieces of information about that person. Information related to inferences of trustworthiness, dominance, and other socially relevant characteristics based on facial features is encoded in a set of regions associated with social cognitive processes, often referred to as the default mode network (e.g., Cao, Li, Todorov, & Wang, 2020; Engell, Haxby, & Todorov, 2007; Gobbini & Haxby, 2007; Parkinson et al., 2017; Su, Luo, Tan, & Qu, 2021; Wagner, Haxby, & Heatherton, 2012; Winston, Strange, O’Doherty, & Dolan, 2002). Recent evidence suggests that people also encode where familiar others sit in their broader social networks, even when there is no task directing their attention to this information (Parkinson et al., 2017; Peer, Hayman, Tamir, & Arzy, 2021; Zerubavel, Bearman, Weber, & Ochsner, 2015). Such evidence stems from functional magnetic resonance imaging (fMRI) studies on real-world social networks in which participants viewed images of their fellow network-members (e.g., members of the same community). Brain regions associated with mentalizing and attentional allocation encoded how well-connected, or central, the individual was in the participant’s own social network (Parkinson et al., 2017; Zerubavel et al., 2015). Additionally, brain regions implicated in encoding spatial and abstract distances encoded how proximal perceived individuals were in the network (friends, friends-of-friends, friends-of-friends-of-friends, etc.), either to the participant or to each other (Parkinson et al., 2017; Peer et al., 2021). Thus, the human brain appears to prioritize information about familiar others’ positions in one’s real-world social networks and spontaneously retrieves this information when encountering them.

There are many confounding pieces of information, however, that are inextricably tied to where people sit in their social networks. Indeed, when encountering familiar friends, there is a plethora of information immediately available, including personal history, personality, and shared

experiences, all of which are inherently linked with that person's social network position. For instance, people who are exceptionally well-connected (e.g., people who have many friends, or in other words, are high in "degree centrality") will likely be seen at more social gatherings and be discussed more frequently, and thus become more visually and socially familiar. Furthermore, people are not randomly assigned to positions in their real-world social networks, and thus, there are a variety of factors that may lead individuals to hold their respective places (e.g., more extraverted individuals are likely to have more connections). Recent evidence suggests that, even without first-hand experience with others or knowledge of their personalities, face-based trait impressions (e.g., apparent trustworthiness, warmth, and attractiveness) are associated with actual and perceived social network centrality (Alt, Parkinson, Kleinbaum, & Johnson, 2022). That is, naïve observers were able to accurately identify characteristics of others' social network positions based solely on their facial features. Furthermore, observers' impressions of targets' personality traits (again based only on facial features) were linked to where those targets sat in their social network. Thus, familiarity, person knowledge, shared experiences, and physical features may systematically covary with real-world network position characteristics, such as relative social distance and social network centrality. Given these potentially confounding factors in real-world social networks, it is difficult to determine if perceivers truly spontaneously encode knowledge of others' social network positions when encountering them, rather than features that covary with where someone sits in their social network.

Additionally, it is unknown how context shapes the encoding of information related to where people sit in their social network. Different aspects of this information may be more relevant in one context and less so in another. For instance, if the goal is to spread information about an event as quickly as possible, then one would likely seek out the most well-connected individuals

who are able to efficiently disseminate the message to as many people as possible. On the other hand, when planning a wedding seating chart, one needs to consider how closely people are connected so that individuals who are nearer to each other in the couple's social network will be seated together (e.g., a table for a bride's college friends, another for her partner's cousins). In the first scenario, an individual's number of connections, or degree centrality (one measure of how well-connected an individual is in a network), is particularly relevant while in the second, the geodesic distance between two people (i.e., the number of steps between them in the network) is more relevant. Does this contextual relevancy affect how the mind encodes their social network position when encountering them? It is possible that certain brain regions incorporate the relevancy of information to the current situation and encode information like degree centrality to a greater extent when it is relevant than when it is not. We can thus examine neural patterns elicited by others to shed light on if and where social network information is encoded, as well as how the mind integrates situational factors with person knowledge.

In the current experiment, we taught participants a novel network structure and used fMRI to measure the neural encoding of others' social network positions in this network. This allowed us to examine the encoding of social network knowledge decoupled from other potential confounding factors that are inherently linked to social network characteristics in real-world contexts. We also systematically varied the contextual relevance of two different facets of others' network positions: how many friends someone has (degree centrality) and how close people are to one another in the network (relative geodesic distance). Network members (represented by their names and faces) were randomly assigned to positions in a learned social network across participants. In doing so, we were able to dissociate social network position characteristics from confounding variables that exist in real-world social networks, and to dissociate degree centrality

from relative geodesic distance. Finally, by using a novel network that participants were not a part of, we were able to test if the brain encodes allocentric social distance (distance between two others) rather than egocentric (distance from oneself), further isolating social network knowledge from potential feelings of affiliation or preferences for individuals closer to oneself. Through this controlled paradigm, we were able to empirically test if the human brain encodes various aspects of social network position over and above other aspects of person knowledge, social experiences, familiarity, and facial features, and to explore how contextual relevancy shapes this encoding.

Methods

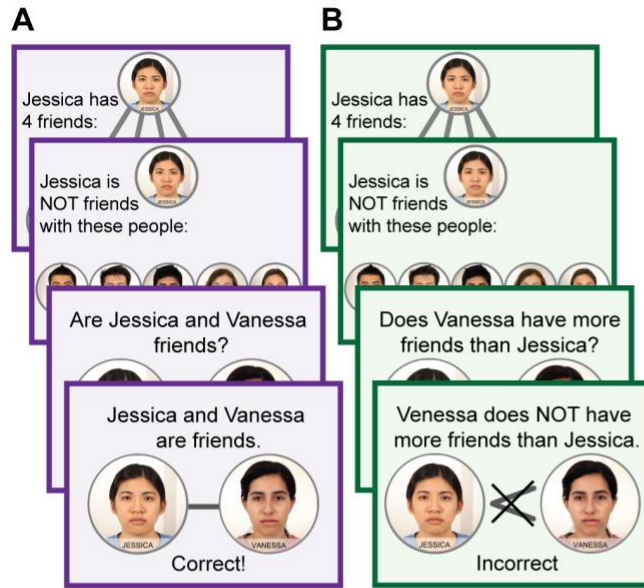
Participants

Participants were recruited from the University of California, Los Angeles (UCLA) campus via flyers and were required to be fluent in spoken and written English, 18-35 years old, right-handed, and have no metal in their body. Additionally, participants had to sufficiently learn a social network during session 1 to be eligible for the fMRI session. To reach our target sample size of 30 (determined *a priori*), we recruited 78 participants for session 1, 31 of whom passed (see Procedure section) and participated in session 2. One subject was excluded due to technical issues with the projection system in the scanner. As such, our final sample size was 30 (13 female, 17 male; ages 18-35 years, $M = 21.27$, $SD = 3.33$). Participants were paid \$15/hour for session 1 and \$20/hour for session 2. All recruited participants were consented in accordance with UCLA Institutional Review Board requirements.

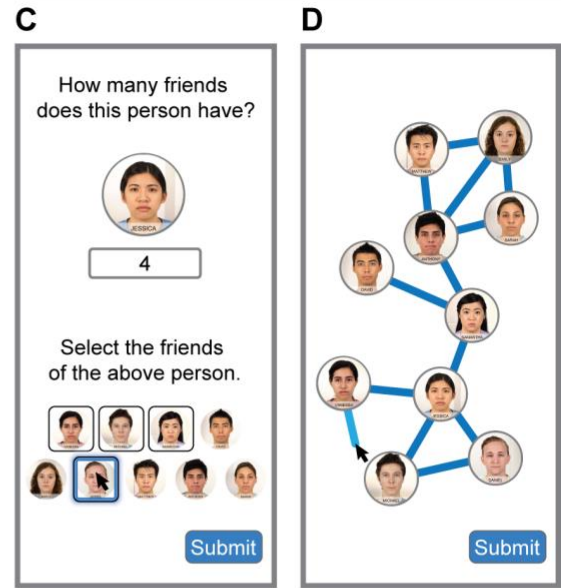
Figure 2.1

Experimental Paradigm


Learning Network Structure



Evaluating Network Learning



FMRI Task



Note. (A and B) Participants learned the social network by viewing each person at the top the screen with their number of friends (degree) and the pictures of those friends at the bottom of the screen. Next, they saw the same individual at the top of the screen with everyone who is not their friend at the bottom in order to ensure that every person is seen the same number of times. (A) On half of the rounds, participants then saw every pair and answered whether or not they were friends. (B) On the other rounds, participants saw each pair and answered whether the person on the right had more friends than the person on the left. They were told if they were correct or not immediately following each trial. After eight rounds, they were evaluated on their knowledge of the network through two tasks. (C) In the first, they reported how many friends and who those friends were for each person, one at a time. (D) In the second task, they drew the network by drawing lines between friends. Participants who knew the network sufficiently completed a practice version of the task they would complete in the scanner. (E) In the fMRI task, they saw each person for 1.5 s followed by jittered fixation time. Randomly throughout each run were a small number of catch trials that were immediately followed by a second image on a blue screen. (F) On half of the runs, participants had to answer if the person on the blue screen had more friends than the last person (centrality condition). On the other half of the rounds, participants answered if the two were friends or not (distance condition). Participants who could successfully complete the scanner task were scheduled for the fMRI session 1–6 days later.

Procedure

The study was completed in two sessions, one to six days apart. In session 1, participants learned two aspects of a friendship network consisting of 13 friendships among 10 individuals. Specifically, they learned how many friends each individual had (i.e., their degree centrality; Fig. 2.1B) and who was friends with whom (Fig. 2.1A). We then evaluated their knowledge. If participants failed to recall each individual's degree and friendships with 100% accuracy (Fig. 2.1C) or were less than 70% accurate in their drawing of the network (Fig. 2.1D), then they did not pass the evaluation task, their participation was terminated, and they were paid for their participation. If they did pass the evaluation task, then they completed a practice version of the task they would complete in the scanner (Fig. 2.1E), which required at least 80% accuracy to be eligible to participate in the fMRI session (session 2). We used strict passing thresholds to ensure that participants knew the network well and were able to complete the scanner task. In session 2, participants completed a shortened version of the network learning task before entering the scanner and then completed the fMRI task while in the scanner.

Each node in the network was represented by an emotionally neutral face from the Chicago Face Database (CFD; Ma et al., 2015). To aid in distinguishing between people, the shirts were colored and names were added. Images were randomly assigned to network position across participants. The network was symmetric (Fig. 2.1D) to maximally dissociate degree centrality from relative distance between people.

Session 1

Learning the Network Structure. Participants learned the two social network features (centrality: number of friends; friendship/distance: friends' identities) in discrete blocks of a learning paradigm, which was presented using PsychoPy (Peirce, 2009). Participants saw each

network member at the top of the screen along with how many friends that person had (i.e., their degree centrality) and who those friends were (Fig. 2.1A-B). To avoid showing well-connected people more often than people who have fewer friends (and thus confounding social network centrality with visual familiarity to participants), this was followed by a trial showing everyone who is not friends with that person.

Next, participants saw each pair in a randomized order and were either asked if the two people were friends with each other (friendship blocks; Fig. 2.1A), or if the person on the right had more friends than the person on the left (centrality blocks; Fig. 2.1B). They were given immediate feedback on whether or not they were correct for a minimum of 0.25 seconds. This was repeated over eight rounds that were blocked such that the first half of the rounds were of one type, and the second half was of the other type, counterbalanced across subjects.

Evaluating Network Learning. To evaluate how well participants learned each feature of the network, participants were shown each person in the network and asked how many friends they had, followed by who those friends were (Fig. 2.1C). To pass, 100% accuracy was required. Next, participants were asked to draw the full network (Fig. 2.1D). All 10 people were presented and participants drew lines between those they remembered as friends. To pass this task, participants needed at least 70% accuracy. They were told which ones were wrong (if any) and to fix them before continuing.

Practice fMRI Task. To ensure all participants that participated in session 2 would be able to do the task in the scanner, those who passed the evaluation tasks practiced the fMRI task (presented using PsychoPy) at the end of session 1. During the fMRI task, every trial consisted of one network member being displayed for 1.5 seconds followed by 4.5-10.5 seconds of jittered fixation time. Randomly spaced throughout each run were catch trials in which a second image on

a blue background was shown immediately after the first for 1 second (Fig. 2.1E). At the beginning of each run, participants were told to answer one of two questions whenever they saw a blue background: (i) “Does this person have more friends than the last person?” or (ii) “Is this person friends with the last person?” At the end of each round, participants were told how many catch trials they answered correctly, how many they answered incorrectly, and how many they missed. Participants needed to reach 80% accuracy in session 1 to be eligible for session 2.

Session 2

If participants passed all tasks in session 1, they participated in session 2 one to six days later. During this session, they completed a shortened version of the learning and evaluation tasks described above. During the evaluation task, they were given immediate feedback on each trial and told to correct their mistake to ensure participants knew the network as well as possible before entering the scanner. Participants then completed eight runs of the fMRI task in the scanner (approximately one hour). Per run, each target person was shown four times as non-catch trials, one time as a catch trial (on a blue screen), and one time as the person shown immediately before the catch trial. Importantly, catch trials were only used to focus participants’ attention on individuals’ relative centrality or distance from others (if they were friends – i.e., separated by a geodesic distance of 1– or not – i.e., separated by a geodesic distance greater than 1). We only analyzed the four occurrences of each image that were not part of a catch trial to test if people encoded social network information even when they were not asked about it directly.

FMRI Data Acquisition

MRI data were collected on a Siemens 3-Tesla Prisma Fit MRI Scanner with a 32-channel head coil. Functional scans were obtained using a gradient echo sequence with 64 interleaved slices (2.0 mm isotropic voxels, TR = 750 ms, TE = 35 ms, flip angle = 52°, FOV = 184 mm).

Participants used a 2-button response box to make choices during the task. For each subject, two echo-planar field maps were obtained after functional scans began in order to correct for the effects of field inhomogeneity. Finally, a T1-weighted (T1w) MPRAGE sequence (1 mm isotropic voxels, 208 slices, TR = 1900 ms, TE = 2.48 ms, flip angle = 9°, FOV = 256 mm) was acquired after the field maps.

FMRI Analyses

Image Preprocessing

Preprocessing was performed using fMRIPrep 1.4.0 (Esteban et al., 2019), which is based on Nipype 1.2.0 (Gorgolewski et al., 2019). The preprocessing descriptions provided here are taken from the recommended citation boilerplate text generated by fMRIPrep (released under a CC0 license with the intention that researchers reuse the text to facilitate clear, consistent descriptions of preprocessing steps, thereby enhancing reproducibility).

Anatomical Data Preprocessing. The T1w image was corrected for intensity non-uniformity with N4BiasFieldCorrection, distributed with ANTs 2.1.0 (Avants, Epstein, Grossman, & Gee, 2008) and used as T1w-reference throughout the workflow. The T1w-reference was skull-stripped with a Nipype implementation of the antsBrainExtraction.sh workflow, using OASIS30ANTs as target template. Brain tissue segmentation of cerebrospinal fluid (CSF), white matter (WM) and gray matter (GM) was performed on the brain-extracted T1w using fast (S. M. Smith et al., 2004). Brain surfaces were reconstructed using recon-all (FreeSurfer 6.0.0; Dale, Fischl, & Sereno, 1999).

Functional Data Preprocessing. For each of the eight BOLD runs per subject, the following preprocessing was performed. First, a reference volume and its skull-stripped version were generated. A deformation field to correct for susceptibility distortions was estimated based

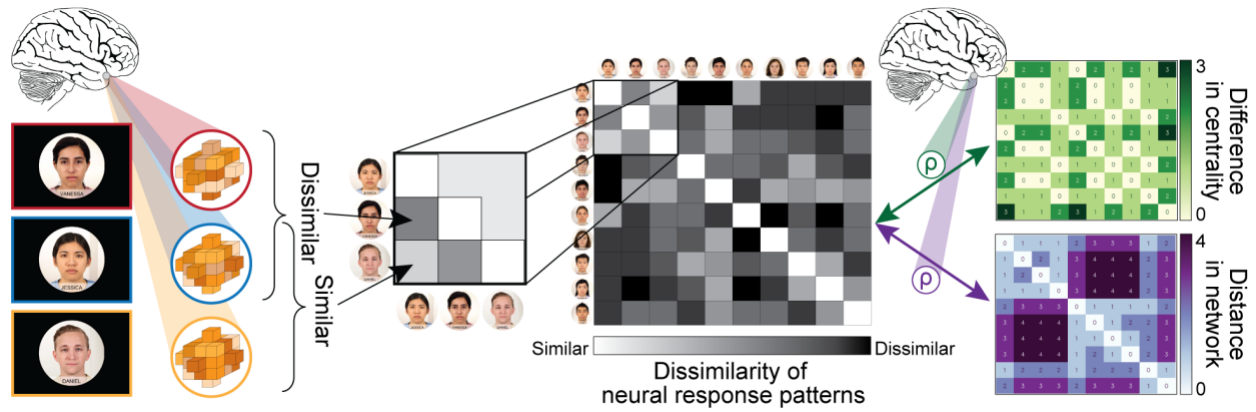
on two echo-planar imaging references with opposing phase-encoding directions, using 3dQwarp in AFNI (Cox, 1996; Cox & Hyde, 1997). Based on the estimated susceptibility distortion, an unwarped BOLD reference was calculated for a more accurate co-registration with the anatomical reference. The BOLD reference was co-registered to the T1w reference using `bbregister` in FreeSurfer with nine degrees of freedom to account for distortions remaining in the BOLD reference. Head-motion parameters with respect to the BOLD reference (transformation matrices; six corresponding rotation and translation parameters) were estimated before any spatiotemporal filtering using FSL's `mcfliirt` (Jenkinson, Bannister, Brady, & Smith, 2002). The BOLD time series were resampled onto their original, native space by applying a single, composite transform to correct for head motion and susceptibility distortions. The first six volumes of each scan were removed from data prior to subsequent analyses.

First-Level Analysis

We fit a general linear model within each catch trial condition (focused on either degree centrality or friendship) of the fMRI data using Nistats (Abraham et al., 2014) to estimate the BOLD response evoked for each of the 10 nodes in the network (represented by different images across participants). The following confounding variables were included in the model as nuisance regressors: three translational motion parameters, three rotational motion parameters, three global signals extracted within the CSF, WM, and whole-brain mask. All regressors of interest were convolved with a Glover hemodynamic response function. The t -statistic maps (i.e., maps of beta coefficients divided by their standard error estimates) resulting from these analyses were used for subsequent pattern similarity analyses.

Figure 2.2

Representational Similarity Analysis Procedure



Note. Using only the non-catch trials in the fMRI task, patterns of neural activation were extracted for each network member in a searchlight procedure. At each searchlight center (i.e. each voxel), the response pattern within a sphere centered on that voxel was extracted for each person seen by the participant while in the scanner. The Euclidean distances between each pattern were then calculated and arranged into a RDM in which each row and column are associated with a network member and the corresponding cell is the dissimilarity in patterns of activity elicited by those two people. This RDM was then Spearman rank-correlated with RDMs representing the difference in network members' centralities (top right matrix) and the distance between people in the network (bottom right matrix). At the group level, we tested where each correlation coefficient (ρ) at each point in the brain (i.e. at each searchlight center) was significantly greater than zero.

Overall Encoding of Social Network Position Characteristics. We tested if and where each facet of network position was encoded throughout the fMRI task, regardless of condition. To do so, we first averaged the t -maps from the two conditions, resulting in one overall distributed neural response pattern for each target person participants encountered in the study. Using a searchlight procedure, we iteratively extracted the multi-voxel pattern of t -values evoked by each target person within “spheres” (radius = 4 voxels) centered at each voxel. We conducted representational similarity analysis (RSA) using these neural response patterns (Kriegeskorte, Mur, & Bandettini, 2008), which allowed us to compare neural representations to models based on each facet of social network position (degree centrality and social distance) to test if a brain region encoded that particular feature (Fig. 2.2; Weaverdyck, Lieberman, & Parkinson, 2020).

This was achieved through the creation and comparison of representational dissimilarity matrices (RDMs). First, we created the model centrality RDM in which each row and column was associated with an individual in the network and the corresponding cell in the matrix was the absolute value of the difference in degree centrality between those two individuals (Fig. 2.2, green matrix). Similarly, we created a model distance RDM in which each cell represented the geodesic distance between two individuals in the network (Fig. 2.2, purple matrix).

We then compared these model RDMs with neural RDMs. To create the neural RDMs, we calculated the Euclidean distance between the neural patterns elicited by different nodes. Euclidean distance was used as it reflects differences in both overall response magnitudes and in topological response patterns between conditions. For results from parallel analyses using Pearson correlations, which reflect differences in topological response patterns only, see the Supplementary Materials S2 (Fig. S2.1, Table S2.1, Table S2.2). That is, each cell of the neural RDM reflected how similar a brain region represented the people corresponding with that cell's row and column. Next, to determine the extent to which each facet of social network knowledge was encoded (independent of the other), we calculated the Spearman rank correlation coefficient, ρ , between the lower off-diagonal triangles of the neural RDMs and each model RDM (Fig. 2.2). (Spearman correlation was used instead of Pearson correlation because it does not assume a linear relationship between the neural and model RDMs.) In other words, we tested if similarity in neural representations were correlated with similarities in degree centrality or proximity in the network. The correlation coefficients were then mapped back onto the central voxel of the searchlight sphere. The two resulting whole-brain maps demonstrated the extent to which distributed neural response patterns in each region (area surrounding each voxel) reflected the degree centrality and the relative social distance of the people being viewed.

Encoding Based on Contextual Relevancy. How does context shape the encoding of this information? To begin to answer this question, we conducted the same analysis described above within each of the conditions: the centrality condition when participants were focused on individuals' degree centralities, and the distance condition when participants were focused on individuals' relationships (i.e., their "degrees of separation" from one another). To test if degree centrality was encoded more when it was relevant than when it was irrelevant, we subtracted the correlation coefficients between the neural RDMs and the centrality model RDM in the irrelevant condition from the relevant condition (i.e., centrality-relevant condition > distance-relevant condition). We ran the same analysis for distance, testing where the correlation coefficient between the neural RDM and the distance model RDM was higher in the relevant condition than in the irrelevant condition (i.e., distance-relevant condition > centrality-relevant condition).

Second-Level Analysis

All first-level analyses were conducted in participants' T1w space. For group-level analyses, we transformed individuals' first-level maps to The ICBM 152 Nonlinear Asymmetrical template version 2009c (Collins, Zijdenbos, Baaré, & Evans, 1999; Fonov et al., 2011) space using ANTs and the mapping generated by fMRIPrep. All first-level maps underwent smoothing (6mm FWHM gaussian kernel). To determine where the brain encoded centrality or distance at the group level, we ran nonparametric permutation testing (10,000 iterations) using FSL's randomise function (Winkler, Ridgway, Webster, Smith, & Nichols, 2014) within a 5mm-dilated brain mask, with 10mm variance smoothing. Results underwent threshold-free cluster enhancement (S. M. Smith & Nichols, 2009) to correct for multiple comparisons.

Parcellation Analyses

In addition to the whole-brain searchlight analyses described above, we ran the same analyses in each parcel of the 200-region Schaefer parcellation (Schaefer et al., 2018) to test for convergence. We conducted both searchlight and parcellation analyses because the searchlight approach provides continuous statistical maps of social network encoding, but defines regions as artificial spheres that are unlikely to resemble the size or shape of functionally or anatomically-defined brain regions, which could lead, for example, to collapsing response patterns across functionally distinct areas. The parcellation approach, however, results in a coarser map of encoding, but defines regions based on their functional response profiles (or anatomy, depending on the parcellation chosen). We transformed the parcellation to each participant's T1w space using the mapping generated by fMRIPrep and ANTs. We extracted the patterns of t -values within each region to create the neural RDMs. The correlations between the neural and model RDMs were then mapped back onto each region. To determine if a region encoded centrality or distance at the group level, we conducted one-sample one-sided t -tests ($p > 0$) within each region in R (Version 3.6.1; R Core Team, 2018). P -values were corrected for multiple comparisons across the 200 parcels using FDR-correction.

Results


Neural Encoding of Degree Centrality

Using RSA, we tested if and where two facets of social network position (degree centrality and social distance from others) were encoded overall (i.e., across conditions). First, we tested which regions encoded network members' degree centrality (i.e., how many friends they had). Results from the searchlight analysis show significant encoding of others' degree centrality in large swaths around bilateral temporoparietal junction (TPJ), superior and inferior parietal lobules

(SPL and IPL), and superior and middle temporal gyri (STG, SMG; Fig. 2.3A; Table 2.1). We found convergent results using the Schaefer parcellation (Fig. 2.3B; Table 2.2).

Figure 2.3

Overall Neural Encoding of Degree Centrality



Note. Regions that showed significant encoding of degree centrality across conditions as measured by correlations between neural RDMs and the degree centrality RDM. (A) Results using a searchlight procedure with a four-voxel radius. All p-values from searchlight-based analyses were corrected for multiple comparisons using TFCE. (B) Results using the 200-region Schaefer parcellation. All p-values from parcellation-based analyses were FDR-corrected for multiple comparisons. Only regions that surpass a corrected threshold of $p < .05$ are shown.

Table 2.1

Searchlight Clusters That Encoded Degree Centrality

General Region	N Voxels	Peak t - Value	Peak Coordinate (x, y, z)	Center of Gravity (x, y, z)
Posterior Lateral Temporal Cortex, Posterior Parietal Cortex, and Occipital Cortex	36,194	5.36	(63.5, -28.5, -8.5)	(-4.3, -56, 34.9)
Left Dorsolateral Prefrontal Cortex	865	4.12	(-38.5, 15.5, 59.5)	(-44.9, 12.1, 50.6)
Left Dorsomedial Prefrontal Cortex	128	3.64	(-4.5, 11.5, 61.5)	(-3.67, 9.64, 61.4)

Note. Significant clusters from the searchlight analysis (TFCE-corrected, $p < .05$). General regions are named based on approximate location of the cluster. Localization of clusters is depicted more precisely in Fig. 2.3A.

Table 2.2

Parcels That Encoded Degree Centrality

General Region	Network	Index	β	t	df	p
Left Posterior Parietal Cortex	Control	61	0.15	4.66	29	.006**
Left Premotor Cortex	Control	70	0.16	4.20	29	.012*
Left Posterior Parietal Cortex	Dorsal Attention	37	0.11	3.77	29	.019*
Left Posterior Parietal Cortex	Control	62	0.15	3.85	29	.019*
Left Superior Temporal Cortex	Default	78	0.12	3.57	29	.022*
Right Superior Frontal Cortex	Default	196	0.10	3.56	29	.022*
Left Lateral Occipital Cortex	Visual	8	0.08	3.50	29	.022*
Left Posterior Parietal Cortex	Dorsal Attention	34	0.11	3.32	29	.030*
Left Posterior Parietal Cortex	Saliency/Ventral Attention	46	0.12	3.28	29	.030*
Left Posterior Parietal Cortex	Default	82	0.11	3.22	29	.031*
Left Posterior Parietal Cortex	Dorsal Attention	36	0.10	3.00	29	.037*
Left Dorsolateral Prefrontal Cortex	Default	93	0.11	3.00	29	.037*
Left Medial Premotor Cortex	Default	95	0.10	3.04	29	.037*
Right Somatomotor Cortex	Somatomotor	127	0.11	3.02	29	.037*
Right Posterior Parietal Cortex	Default	184	0.10	3.09	29	.037*
Left Posterior Parietal Cortex	Saliency/Ventral Attention	45	0.10	2.97	29	.037*
Left Lateral Occipital Cortex	Dorsal Attention	33	0.13	2.93	29	.038*
Left Posterior Parietal Cortex	Control	63	0.12	2.87	29	.040*
Right Somatomotor Cortex	Somatomotor	126	0.09	2.88	29	.040*
Right Premotor Cortex	Control	175	0.09	2.80	29	.045*
Left Ventrolateral Prefrontal Cortex	Default	85	0.08	2.73	29	.046*

Left Dorsomedial Prefrontal Cortex	Default	91	0.10	2.74	29	.046*
Right Somatomotor Cortex	Somatomotor	124	0.09	2.75	29	.046*
Left Middle Temporal Gyrus	Dorsal Attention	32	0.08	2.57	29	.048*
Left Posterior Parietal Cortex	Dorsal Attention	35	0.06	2.52	29	.048*
Left Ventrolateral Prefrontal cortex	Saliency/Ventral Attention	50	0.09	2.67	29	.048*
Left Dorsolateral Prefrontal Cortex	Control	69	0.09	2.55	29	.048*
Left Dorsomedial Prefrontal Cortex	Default	92	0.09	2.53	29	.048*
Left Dorsolateral Prefrontal Cortex	Default	94	0.10	2.68	29	.048*
Right Somatomotor Cortex	Somatomotor	125	0.07	2.60	29	.048*
Right Somatomotor Cortex	Somatomotor	128	0.10	2.52	29	.048*
Right Somatomotor Cortex	Somatomotor	130	0.07	2.51	29	.048*
Right Posterior Parietal Cortex	Dorsal Attention	137	0.10	2.55	29	.048*
Right Precuneus	Dorsal Attention	140	0.09	2.53	29	.048*
Right Middle Temporal Gyrus and Superior Temporal Sulcus	Saliency/Ventral Attention	148	0.09	2.62	29	.048*
Right Posterior Parietal Cortex	Control	166	0.09	2.52	29	.048*
Right Middle Temporal Gyrus and Superior Temporal Sulcus	Default	188	0.09	2.52	29	.048*
Right Posterior Parietal Cortex	Control	167	0.08	2.50	29	.049*
Right Posterior Parietal Cortex	Dorsal Attention	142	0.08	2.48	29	.049*
Right Posterior Parietal Cortex	Dorsal Attention	139	0.09	2.47	29	.050*

Note. Indices and network names are provided by the database corresponding to the Schaefer et al. (2018) parcellation. General regions are named based on approximate location of parcel. All p -values are corrected for multiple comparisons using FDR-based correction. ** $p < .01$, * $p < .05$

Neural Encoding of Social Distance


Next, we tested if and where the social distance between network members was encoded.

We did not find any significant ($p < .05$) encoding of distance. However, we did find that the

overall encoding of targets' social distances to one another was trending ($p < .10$) in a 520-voxel cluster in the left middle temporal gyrus (peak value = 4.59, peak coordinate = (-62.5, -24.5, -10.5), center of gravity = (-61.2, -24.4, -11.7)) in the searchlight analysis (Fig. 2.4A) and, in the parcellation-based analysis, in the right medial prefrontal cortex (mPFC), $\rho = 0.07$, $t(29) = 3.53$, $p = .070$, and anterior STG, $\rho = 0.09$, $t(29) = 3.58$, $p = .070$ (Fig. 2.4B). Since none of these results reached significance, these findings must be interpreted with caution and future research is required to confirm and clarify them.

Figure 2.4

Overall Neural Encoding of Distance Between Network Members



Note. Parcellation results showing regions that trended towards encoding distance between network members. Note, there were no significant regions after correcting for multiple comparisons using FDR-correction. Only regions that surpass a corrected threshold of $p < .1$ are shown.

Effects of Contextual Relevancy on the Encoding of Social Network Information

Lastly, we tested if and where contextual goals modulated the encoding of centrality. That is, we examined if centrality was encoded *more* when it was relevant than when it was irrelevant. In both the searchlight and parcellation analyses, we did not find any regions that survived correction for multiple comparisons. Similarly, we tested if and where contextual goals modulated

the encoding of distance. We did not find any regions that significantly encoded distance *more* when it was relevant than when it was irrelevant.

Discussion

In this study, we tested if and where the human brain encodes information about others' social network positions when that information is disassociated from other features that covary with it in real-world contexts. Specifically, we tested if and where relative degree centrality and social distance were tracked by the brain when viewing others' faces. To decouple these social network features from other typically confounding types of information (e.g., trait impressions, person knowledge, visual characteristics, familiarity, memories), we taught participants a new social network where network members' identities were randomly assigned to positions in the network.

We found that degree centrality was robustly encoded in broad regions surrounding the superior and inferior parietal lobules, bilateral temporoparietal junction, and superior and middle temporal gyri. That is, the brain prioritized information regarding others' centrality in the social network even though the participant was not directly involved in that network, and even when other facets of social network position were more relevant. This suggests that social network centrality may be chronically important to monitor. Indeed, measures of social network centrality capture the importance of a person in a social network and can be considered to comprise a facet of social status (Basyouni & Parkinson, 2022; Weaverdyck & Parkinson, 2018). Here, we found that centrality was encoded in regions that support attentional modulation (Corbetta & Shulman, 2002; Kastner & Ungerleider, 2000; Uncapher & Wagner, 2009) and social cognitive processes, such as understanding others' mental states (Morelli, Leong, Carlson, Kullar, & Zaki, 2018; Tamir, Thornton, Contreras, & Mitchell, 2016; Van Overwalle & Baetens, 2009). Additionally, our results

overlap significantly with previous findings regarding the spontaneous encoding of other's centrality in their real-world social networks (Parkinson et al., 2017; Zerubavel et al., 2015). Thus, it may be that social network centrality modulates one's overall attention towards others, because it signals those individuals' importance in the community and/or as people who are particularly valuable to attend to for ascertaining group norms (Basyouni & Parkinson, 2022; Paluck & Shepherd, 2012). This may also increase attention to and consideration of high-status individuals' mental states, which could then shape downstream thoughts and behaviors. Given the potential implications for social influence and reputation management (Weaverdyck & Parkinson, 2018), future research should test these possibilities by examining how social network-based status shapes how much perceivers attend to others and to what they appear to be thinking.

While we did not find any significant encoding of relative distance in the network, we found a trend suggesting that allocentric distances may be encoded in aspects of the lateral temporal cortex and medial prefrontal cortex (mPFC). Importantly, previous research has primarily focused on if and where egocentric distance (distance from oneself) is encoded in the brain (Parkinson et al., 2017; Zerubavel et al., 2015). Here, however, we taught participants a new network and tested the extent to which allocentric distance (distance between others) was encoded. Because participants were not members of this network, we cannot directly test the extent to which egocentric distance was encoded in this controlled setting, and are thus unable to directly compare our results to previous literature examining egocentric distance; given that allocentric distance is distinct from egocentric distance, its neural representation may differ. It could simultaneously be that allocentric distance – particularly in a novel network in which one has no part – is encoded less robustly than more self-relevant information (e.g., egocentric distance), and we were not

sufficiently powered to detect it in our paradigm. Thus, it is unclear if our trending result is due to a lack of effect or a lack of power in our sample to detect the encoding of allocentric distance.

One reason to suspect that the trending results in the lateral temporal cortex and mPFC are due to a lack of power is that these areas overlap with and neighbor regions that are known to support person models (Hassabis et al., 2014; Wagner et al., 2012; Wang et al., 2017; Welborn & Lieberman, 2015). It could be, then, that people who are close to each other in the network are assumed to be more similar to each other due to phenomena such as homophily (i.e., the tendency for similar others to become friends) and social influence (Schwyck*, Du*, Li, Chang, & Parkinson, 2023; Son, Bhandari, & FeldmanHall, 2021). This would result in the brain representing more proximate individuals in a network as more similar, which is consistent with the trending results. Secondly, there is recent evidence that allocentric distance in one's real-world social media network is encoded in the default mode network, including similar regions to those found in the present study, and that allocentric distance is encoded distinctly from egocentric distance (Peer et al., 2021). Thus, there are several possible reasons why the current results differ from previous findings. First, previous findings studying real-world social networks may partially reflect *similarity* of person knowledge or associated memories. Second, egocentric distance might be encoded more robustly than allocentric distance because it connotes self-relevance. Finally, egocentric and allocentric social distance may be qualitatively different types of information that are processed differently in the brain.

The current study implemented a controlled task in which participants learned a new pattern of relationships. This was an intentional departure from previous research which used participants' real-world social networks, but where social network positions were inextricably linked to other types of social knowledge (e.g., memories) and perception (e.g., visual familiarity;

face-based trait impressions; Parkinson et al., 2017; Peer et al., 2021; Zerubavel et al., 2015). These previous studies along with the current research present strong evidence that the human brain prioritizes the recall of social network knowledge when encountering others.

Future research is needed to replicate these findings and extend this work in several ways. In the current study, participants were not directly part of the learned social network, and the task directed people to specific network features. This limited us to examining the encoding of allocentric (and not egocentric) social distances and may have rendered participants less likely to call to mind information about others' network positions when viewing them than they would be in a more personally relevant social network. Future research should work to further disentangle how egocentric and allocentric distances are encoded, how contextual goals modulate the encoding of social networks in which the participant is included, and if this information is spontaneously encoded (i.e., even when the task does not direct participants' attention towards this information), as previous research suggests (Parkinson et al., 2017). Additionally, future research should examine how these facets of social network positions are neurally encoded for different types of relationships other than friendships (e.g., kinship, work hierarchies), and how other measures of node importance (e.g., eigenvector centrality, betweenness centrality) are encoded and used to facilitate inferences about other traits (e.g., competence), likely shaping downstream processes and behaviors. Finally, it is important to examine if and how these phenomena differ across cultures and age groups, and as a function of other individual differences (e.g., in patients with disorders characterized by atypical social functioning). Such studies (necessitating much larger sample sizes; Marek et al., 2022) could focus on understanding how the encoding and application of network knowledge differs in perceivers with different social cognitive abilities and from different backgrounds. These expansions of the current work would provide valuable insight into how the


brain represents, uses, and integrates information about the social networks in which everyone is embedded.

As it has been said time and time again: humans are social animals. The people with whom we regularly interact do not exist in a vacuum, but rather, in the broader context of our social networks. As such, our perceptions of others are not only defined by our impressions and knowledge of them as individuals, but also by the patterns of social relationships that surround them. Understanding this social structure is impactful in everyday life yet it is not well-understood. It is important to characterize how healthy brains support the capacity to learn and represent social networks to understand how such social processes may be compromised in disorders characterized by deficits in social cognition and behavior. One's ability to learn and process new social network information, and apply it in different contexts, likely has serious consequences for downstream behavioral interactions in all aspects of one's social life. Here, we found evidence that the human brain prioritizes specific aspects of social network knowledge that signal the relative importance of others in a community.

Supplementary Materials S2

Figure S2.1

Neural Encoding of Degree Centrality Using Pearson Correlation



Note. Regions that showed significant encoding of degree centrality across conditions as measured by correlations between neural RDMs (comprised of Pearson correlation coefficients) and the degree centrality RDM. (A) Results using a searchlight procedure with a four-voxel radius. All searchlight p -values were corrected for multiple comparisons using threshold-free cluster enhancement. (B) Results using the 200-region Schaefer parcellation. All parcellation p -values were FDR-corrected for multiple comparisons. Only regions that surpass a corrected threshold of $p < .05$ are shown.

Table S2.1

Searchlight Clusters That Encoded Degree Centrality Using Pearson Correlation

General Region	N Voxels	Peak t - Value	Peak Coordinate (x, y, z)	Center of Gravity (x, y, z)
Posterior Lateral Temporal Cortex, Posterior Parietal Cortex, and Occipital Cortex	36,194	5.36	(63.5, -28.5, -8.5)	(-4.3, -56, 34.9)
Left Dorsolateral Prefrontal Cortex	865	4.12	(-38.5, 15.5, 59.5)	(-44.9, 12.1, 50.6)
Left Dorsomedial Prefrontal Cortex	128	3.64	(-4.5, 11.5, 61.5)	(-3.67, 9.64, 61.4)

Note. Significant clusters from the searchlight analysis (TFCE-corrected, $p < .05$). General regions are named based on approximate location of the cluster. Localization of clusters is depicted more precisely in Fig. 2.3A.

Table S2.2*Parcels That Encoded Degree Centrality Using Pearson Correlation*

General Region	Network	Index	β	t	df	p
Left Posterior Parietal Cortex	Control	71	0.14	5.39	29	.001***
Left Posterior Parietal Cortex	Control	61	0.16	4.89	29	.002**
Left Posterior Parietal Cortex	Default	82	0.13	4.75	29	.002**
Left Posterior Parietal Cortex	Dorsal Attention	34	0.16	4.49	29	.002**
Left Posterior Parietal Cortex	Control	62	0.15	4.44	29	.002**
Left Posterior Parietal Cortex	Dorsal Attention	35	0.15	4.15	29	.004**
Right Posterior Parietal Cortex	Default	184	0.14	4.10	29	.004**
Left Premotor Cortex	Control	70	0.14	3.87	29	.007**
Left Posterior Parietal Cortex	Control	63	0.15	3.68	29	.010**
Right Precuneus	Dorsal Attention	140	0.13	3.67	29	.010**
Right Middle Temporal Gyrus and Superior Temporal Sulcus	Default	188	0.12	3.56	29	.012*
Left Posterior Parietal Cortex	Saliency/Ventral Attention	45	0.13	3.49	29	.012*
Right Posterior Parietal Cortex	Control	167	0.13	3.48	29	.012*
Left Somatomotor Cortex	Somatomotor	25	0.11	3.44	29	.013*
Left Somatomotor Cortex	Somatomotor	23	0.13	3.38	29	.014*
Left Middle Temporal Gyrus	Dorsal Attention	32	0.11	3.28	29	.016*
Left Dorsolateral Prefrontal Cortex	Control	68	0.10	3.23	29	.016*
Right Lateral Occipital Cortex	Visual	115	0.09	3.22	29	.016*
Right Premotor Cortex	Control	175	0.11	3.24	29	.016*
Right Middle Temporal Sulcus	Default	186	0.08	3.26	29	.016*
Right Medial Occipital Cortex	Visual	114	0.12	3.19	29	.016*
Left Superior Temporal Cortex	Default	78	0.11	3.17	29	.016*
Left Ventrolateral Prefrontal cortex	Saliency/Ventral Attention	50	0.12	3.09	29	.017*

Left Cingulate Cortex	Saliience/Ventral Attention	52	0.1	3.06	29	.017*
Left Middle Temporal Gyrus and Superior Temporal Sulcus	Default	77	0.08	3.11	29	.017*
Right Somatomotor Cortex	Somatomotor	124	0.10	3.08	29	.017*
Right Somatomotor Cortex	Somatomotor	128	0.13	3.06	29	.017*
Right Posterior Parietal Cortex	Control	177	0.10	3.12	29	.017*
Left Occipital Pole	Visual	7	0.13	3.02	29	.018*
Right Posterior Parietal Cortex	Dorsal Attention	142	0.11	3.02	29	.018*
Right Precuneus Cortex	Default	200	0.10	2.98	29	.019*
Left Dorsomedial Prefrontal Cortex	Default	92	0.10	2.95	29	.019*
Right Posterior Parietal Cortex	Control	166	0.11	2.89	29	.022*
Right Posterior Parietal Cortex	Default	182	0.09	2.87	29	.022*
Left Posterior Parietal Cortex	Dorsal Attention	38	0.12	2.86	29	.022*
Right Calcarine Cortex	Visual	109	0.10	2.84	29	.022*
Left Fusiform Gyrus	Visual	2	0.12	2.82	29	.023*
Left Dorsomedial Prefrontal Cortex	Default	91	0.09	2.80	29	.024*
Right Inferior Temporal Gyrus and Middle Temporal Sulcus	Control	168	0.10	2.78	29	.024*
Left Lateral Occipital Cortex	Visual	3	0.09	2.77	29	.024*
Left Somatomotor Cortex	Somatomotor	19	0.09	2.71	29	.027*
Right DorsoAttn_Post	Dorsal Attention	141	0.08	2.68	29	.028*
Left Posterior Parietal Cortex	Dorsal Attention	37	0.10	2.69	29	.028*
Left Dorsolateral Prefrontal Cortex	Default	93	0.11	2.68	29	.028*
Right Posterior Parietal Cortex	Default	183	0.09	2.64	29	.028*
Left Dorsolateral Prefrontal Cortex	Default	94	0.09	2.64	29	.028*
Left Lateral Occipital Cortex	Visual	14	0.11	2.64	29	.028*
Right Dorsolateral Prefrontal Cortex	Control	174	0.07	2.57	29	.032*

Right Occipital Pole	Visual	112	0.10	2.55	29	.034*
Left Somatomotor Cortex	Somatomotor	27	0.08	2.52	29	.035*
Left Posterior Parietal Cortex	Dorsal Attention	36	0.09	2.45	29	.040*
Left Posterior Parietal Cortex	Saliency/Ventral Attention	46	0.10	2.4	29	.041*
Left Inferior Temporal Gyrus and Middle Temporal Sulcus	Control	64	0.09	2.39	29	.041*
Left Posterior Parietal Cortex	Default	81	0.10	2.40	29	.041*
Right Occipital Pole	Visual	106	0.07	2.41	29	.041*
Right Superior Frontal Cortex	Default	196	0.09	2.39	29	.041*
Right Superior Frontal Cortex	Default	197	0.10	2.40	29	.041*
Right Posterior Parietal Cortex	Dorsal Attention	144	0.07	2.44	29	.041*
Left Fusiform Cortex	Dorsal Attention	31	0.08	2.32	29	.047*

Note. Indices and network names are provided by the database corresponding to the Schaefer et al. (2018) parcellation. General regions are named based on approximate location of parcel. All *p*-values are corrected for multiple comparisons using FDR-based correction. ****p* < .001, ***p* < .01, **p* < .05

CHAPTER 3

Similarity among friends serves as a social prior:
The assumption that “Birds of a feather flock together” shapes
social decisions and relationship beliefs

Miriam E. Schwyck*¹, Meng Du*¹, Yuchen Li¹, Luke J. Chang², Carolyn Parkinson¹

¹University of California, Los Angeles

²Dartmouth College

This is a reprint of the published manuscript:

Schwyck, M.E.*, Du, M.*, Li, Y., Chang, L., & Parkinson, C. (2023). Similarity among friends serves as a social prior: The assumption that "Birds of a feather flock together" shapes social decisions and relationship beliefs. *Personality and Social Psychology Bulletin*, 014616722211402.

Abstract

Social interactions unfold within networks of relationships. How do beliefs about others' social ties shape—and how are they shaped by—expectations about how others will behave? Here, participants joined a fictive online game-playing community and interacted with its purported members, who varied in terms of their trustworthiness and apparent relationships with one another. Participants were less trusting of partners with untrustworthy friends, even after they consistently showed themselves to be trustworthy, and were less willing to engage with them in the future. To test if people not only expect friends to behave similarly, but also expect those who behave similarly to be friends, an incidental memory test was given. Participants were exceptionally likely to falsely remember similarly-behaving partners as friends. Thus, people expect friendship to predict similar behavior and *vice versa*. These results suggest that knowledge of social networks and others' behavioral tendencies reciprocally interact to shape social thought and behavior.

Introduction

Sayings like "birds of a feather flock together" reflect the widespread intuition that we tend to be surrounded by those who are similar to us. Indeed, similarity between people who are socially close to one another is a common feature of human social networks and can arise from various mechanisms. For example, social influence can result in friends becoming more similar (Altermatt & Pomerantz, 2003; Brechwald & Prinstein, 2011; de Klepper, Sleebos, van de Bunt, & Agneessens, 2010) and people who are similar (e.g., in terms of demographic characteristics) are more likely to become friends due to both structural factors that constrain whom individuals are likely to encounter (i.e., "induced homophily") and individuals' preferences (i.e., "choice homophily"; McPherson et al., 2001). Indeed, there is a large body of sociological research examining this phenomenon, particularly focusing on the existence and effects of demographic homophily in real-world social networks (e.g., Lawrence & Shah, 2020; McPherson et al., 2001; Shrum et al., 2016; J. A. Smith et al., 2014). Additional work bridging sociology and psychology has begun to examine how cognitive, affective, and behavioral tendencies are related to friendship, and more generally, proximity in social ties (e.g., Apicella et al., 2012; Dehghani et al., 2016; Ehlert et al., 2020; Fowler & Christakis, 2010; Parkinson et al., 2018; Pradel et al., 2009; K. M. Smith et al., 2018). This body of research complements related work in psychology regarding similarity-based attraction and social influence (e.g., Henderson & Furnham, 1982; Kuwabara et al., 2022; Wetzel & Insko, 1982). Thus, convergent evidence spanning multiple disciplines has linked social closeness to interpersonal similarity with respect to a variety of factors, including how people tend to think, feel, and behave, as well as demographic characteristics.

Given that similarity among friends is commonly observed in human social networks, it may be that people *assume* such similarities exist, and that this assumption serves as a heuristic to

inform predictions of how others will behave and to scaffold mental representations of friendships between others. Although people's perceptions and reality are not always aligned, there have been decades of research on how people cognitively represent their social networks, finding that people often overperceive common characteristics of social networks, including for example, balanced triads (two people who have a friend in common are often assumed to be friends), and the likelihood that people of the same race are friends compared to individuals of different racial backgrounds (Brands, 2013). Assuming such characteristics would be beneficial, as it would reduce the amount of information that humans need to remember. Moreover, because people cannot exhaustively observe others' social behaviors and relationships, they may use strategies to fill in the gaps. By characterizing the systematic errors people make in their inferences about others, we can gain insight into such strategies. It is unknown, however, if an assumption of behavioral similarity among friends exists and how such an assumption might shape social thought and behavior. Does knowing how a stranger's friend behaved affect how we treat that stranger? Furthermore, does this assumption shape our perception of others and our memory of others' relationships?

Prosocial tendencies provide a useful lens through which to examine these questions for several reasons. Whereas much work on homophily has focused on relatively coarse and often readily observable characteristics (e.g., race, age), some research has argued for homophily in prosocial behavioral tendencies (Apicella et al., 2012), which may be possible to the extent that cooperative tendencies are observable (e.g., learned through one's own social interactions, observing others' interactions, or gossip). Indeed, people can accurately learn about others' tendencies to cooperate or freeload, this information shapes individuals' reputations, and people use such reputations to maximize their own gain (Krasnow, Cosmides, Pedersen, & Tooby, 2012;

Price, 2006). Therefore, it appears that people are keenly attuned to and readily learn about how prosocial others are, and cooperation thus provides a context where people are motivated to do so. In particular, when playing cooperative games in which they themselves have a stake in the outcome, people are motivated to do their best to learn about others' behaviors and accurately predict their decisions. Thus, cooperative games provide a sensitive lens for examining these mental processes in a controlled context.

Here, we use prosocial tendencies to test if similarity among friends acts as a social prior (i.e., as a belief people have before taking new evidence into account), leading people to believe that friends are likely to behave similarly (i.e., those who “flock together” are probably “birds of a feather”), and people who behave similarly are more likely to be friends (i.e., “birds of a feather” probably “flock together”). That is, we test two hypotheses about people's assumption of a similarity-friendship association: 1) assumptions of similarity among friends lead people to believe that friends will behave similarly to one another, and 2) assumptions of similarity among friends lead people to believe that those who behave similarly are particularly likely to be friends.

Do People Expect Friends to Behave Similarly?

Effective, beneficial interactions with others are critical to individuals' wellbeing, as well as to their personal and professional success. However, people also often need to interact with individuals whom they have not previously encountered. Decisions about how to interact with such strangers can lead to impactful consequences. For example, if you ask a stranger to watch your laptop for a few minutes, you may end up mourning the loss of your computer if the stranger proves to be untrustworthy. Many researchers have examined how people infer others' trustworthiness through direct interactions (Chang, Doll, van 't Wout, Frank, & Sanfey, 2010; Fareri, Chang, & Delgado, 2012; Fouragnan et al., 2013), through character knowledge and

reputation (e.g., explicit character descriptions, gossip; Delgado et al., 2005; Feinberg et al., 2014; Fouragnan et al., 2013), and by interpreting what physical appearance might signal (e.g., ingroup status, face-based trait attributions; Stanley et al., 2011; Todorov et al., 2009; van 't Wout & Sanfey, 2008; Wilson & Eckel, 2006). Yet, many of the strangers with whom we interact also have relationships with individuals we know, and knowledge of such associations may provide valuable clues regarding how unfamiliar others will behave in the absence of previous interactions or character information. That is, people may make assumptions about how a stranger will behave based on their knowledge of how that person's friends tend to act. Indeed, there is growing research interest in how we think about and are affected by the social networks that surround us (E. R. Smith & Collins, 2009; Weaverdyck & Parkinson, 2018). Knowing that a stranger has a trustworthy friend may increase the likelihood that you would trust this stranger to watch your laptop. In other words, we may believe that a stranger will behave similarly to how their friend behaves due to an assumption of similarity among friends.

It is also unclear how third-party relationship knowledge (i.e., knowledge of others' relationships) affects our perceptions of someone we *have* previously encountered (i.e., people about whom we have direct knowledge). When encountering people we know (e.g., acquaintances, friends), perceivers spontaneously retrieve information about those people's relationships, presumably to prepare for appropriate, beneficial interactions (Parkinson et al., 2017). Thus, it is possible that people use information about others' relationships when deciding how to behave in an upcoming interaction. For instance, if a recent acquaintance has consistently behaved trustworthily, but you know they have a deceitful friend, you may be less trusting of that individual, compared with another recent acquaintance who behaves identically but has friends whom you

trust. Thus, when interacting with a familiar other, people may continue to refer to their knowledge of that individual's relationships.

Do People Expect Those Who Behave Similarly to Be Friends?

While beliefs about a partner's relationships may shape predictions about how that partner will behave in an upcoming interaction, others' behavioral tendencies (i.e., how they typically behave in particular situations) may also shape our beliefs about those individuals' relationships with one another. In human social networks, friendship is linked to similarity in a wide variety of factors, including some forms of prosocial behavioral tendencies (Apicella et al., 2012). Thus, people may expect their interaction partners to be friends with others who think and behave like they do, and accordingly, this expectation may shape mental representations of social networks. For example, you may be unsurprised to learn that two people who both value honesty or rule-following are friends with one another, and you may even *expect* or *assume* that these two people would be friends with each other if you learned that they both belonged to the same small community. In the same vein, you may also expect two people within that community who are both rebellious to be friends with one another. On the other hand, you may be surprised to learn that an exceptionally rebellious person is friends with someone who is a stickler for the rules, or that an honest, trustworthy person is friends with someone who is conniving and deceitful. More generally, if two individuals tend to behave similarly, people may be more likely to believe that those individuals are friends with one another, even when they are not, compared to a pair of individuals who behave dissimilarly to one another.

The Current Study

In the current study, we sought to test how knowledge of others' social relationships shapes people's subjective expectations about how those individuals will behave, both when interacting

with strangers and when repeatedly interacting with the same people. Additionally, we examined how knowledge of others' behavior shapes beliefs about friendship ties. That is, we also tested if people are more likely to assume that others are friends² with one another if they behave similarly to each other. Thus, rather than testing what is true (e.g., whether or not friends behave exceptionally similarly to each other), in the current study, we test what people *assume* to be true (e.g., whether or not people believe that friends will behave exceptionally similarly to each other) and the consequences of such assumptions for social interactions and memory.

To this end, we constructed a fictive online game-playing community in which participants played games with various partners (see Figs. 1 and 2). By creating this fictive online game-playing community, we were able to measure and manipulate participants' beliefs about both their gaming partners' behavior and the friendships between those partners in a fully balanced design. We thus examined how beliefs about others' behavioral tendencies shape, and are shaped by, beliefs about their relationship ties. More specifically, we tested two complementary hypotheses: assumptions of similarity among friends will lead people to expect partners who are friends to behave similarly (Hypothesis 1) and assumptions of similarity among friends will lead people to expect that two partners are more likely to be friends if they behave similarly than if they behave dissimilarly (Hypothesis 2).

² In this experiment, we refer to "friend" as one of the partner's "Top 3 Friends" on the gaming site where members could interact with one another in a variety of ways. Participants were told that members of the community regularly answered questions about how much they favored one another and how much they would like to engage with one another in the future; an algorithm allegedly computed each member's "Top 3 Friends" based on who that member consistently favored and was favored by.

Methods

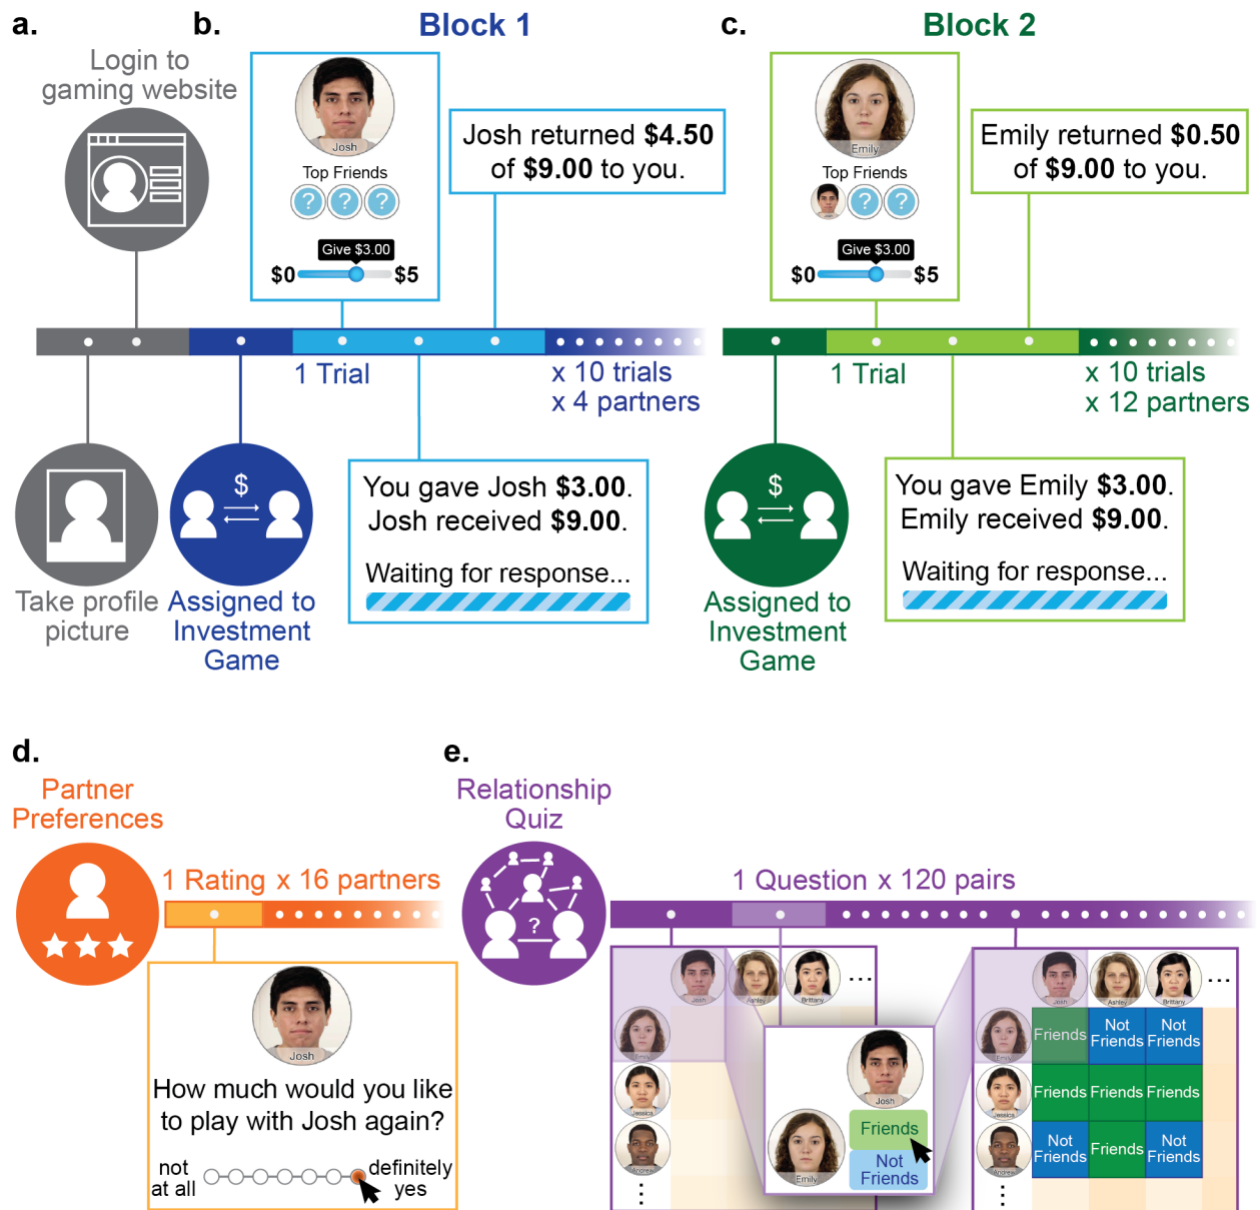
Materials and code used to conduct the experiment are publicly available (<https://github.com/meng-du/TGame>), as are all data and analysis scripts (<https://osf.io/nx87u/>).

Participants

Participants were 80 undergraduate students (50 female, 29 male, 1 did not specify) at University of California, Los Angeles (UCLA), between 18 and 31 years old ($M = 20.51$ years, $SD = 2.37$). Two additional participants were excluded from the analysis because, after finishing the experiment, they explicitly told the experimenter that they either did not understand the instructions or fell asleep during the experiment (all measures, manipulations and exclusions are reported in this manuscript). Based on the medium to large effect sizes found in studies utilizing similar paradigms (Chang et al., 2010; van 't Wout & Sanfey, 2008), we determined that we would need 63 participants to detect the smallest of these effects ($\eta^2 = 0.05$; Chang et al., 2010) with 80% power. These calculations were made in the R package *pwr* (Champely, 2018) and were based on detecting a main effect friend knowledge in a balanced one-way ANOVA (although a more appropriate linear mixed effects model was implemented). Sample size was determined before any data analysis. All participants provided written informed consent in accordance with the policies of the UCLA ethical review board.

Figure 3.1

Overview of Paradigm



Note. Participants were told that they would be testing out a new online gaming community in which members had profiles and regularly played a variety of games with one another. As shown in (a), the experimenter first created a profile for the participant. Next (b), participants were “randomly assigned” to play a series of 10 Investment Games (i.e., trust games) with each of 4 partners. When playing with each partner, that individual’s name and profile picture were shown, along with their “Top 3 Friends” in the community, who were hidden if the participant had not already “met” them on the website, ostensibly due to the website’s privacy controls (see Fig. 3.2). Thus, in Block 1 (b), the identities of participants’ partners’ “Top 3 Friends” were hidden. In Block 2 (c), participants were assigned to play the Investment Game again with 12 new partners, some of whom were indicated to have a “Top Friend” with whom the participant had interacted during Block 1, and thus, whose identity was visible. Afterwards (d), participants rated each partner on how much they would like to play with them again in a new, cooperative game. Finally, participants

completed a surprise memory test (e) of who appeared as friends and who did not by filling out a matrix with “Friends” or “Not Friends” for every pair of partners.

Procedure

Introduction to The Online Game-Playing Community

Before the study began, participants were told that the experimenter needed to make profiles for them, so that they could join an online game-playing community. The experimenter took a facial photograph of the participant (Fig. 3.1a). The photograph would ostensibly serve as the participant’s profile picture on the website, which would be displayed to partners during game play (participants later saw analogous photographs of their partners when interacting with them). Participants were told that this online game-playing community was being developed and tested at multiple college campuses, and that members can win small amounts of money by playing a variety of simple interactive games. Specifically, participants were informed that they would receive a portion of the total money that they earned while playing on the site. They were also told that their partners would be players from other campuses who regularly engage with the website and its associated social networking features; these individuals had ostensibly played a variety of games on the website several times before and continue to do so regularly for fun and profit. In reality, participants were playing with a computer that simulated different behavioral tendencies (see Manipulating Interaction Partners’ Behavior section). Additionally, participants were told that the gaming website was also rolling out a new social network feature, in which each player’s “Top 3 Friends” (described in the Manipulating Social Network Knowledge section) would be displayed along with their profile photo. Deception was used in order to manipulate our key constructs of interest and to maximize the equivalence in experience across participants. Participants were fully debriefed regarding all aspects of the procedure immediately after finishing the study.

Next, participants were told that they would be randomly assigned to play one game out of a set of many possible games during each of three consecutive blocks. In reality, all participants were assigned to play trust games (described in the Manipulating Interaction Partners' Behavior section; Berg et al., 1995) in the first two blocks. Participants were then told that the third block would be skipped due to time constraints. This was done to further instill the idea that members of the website (including the participant) could have played a variety of games, and to make participants believe, during the second block, that they might not be done with interacting with their partners just yet. The trust game was always referred to as "The Investment Game" to minimize demand characteristics. As motivation, participants were informed that they would be paid at the end of the experiment based on their earnings during the game. Participants were also informed that as a member of this website, they would later have the chance to rate each partner, and that these ratings contribute to the determination of each member's "Top 3 Friends," as described in more detail in the next section.

Manipulating Social Network Knowledge

To test if friendship knowledge shapes how participants expected people to behave (Hypothesis 1), we systematically manipulated friendships between partners. For each interaction, participants viewed their partners' first names and profile photos, along with photos of their partners' "Top 3 Friends" (Fig. 3.1b-c). Prior to the start of the experiment, it was explained that members regularly answered questions about how much they favored one another and how much they would like to engage with one another in the future; an algorithm allegedly computed each member's "Top 3 Friends" based on who that member consistently favored and was favored by. Participants were also told that, due to the default privacy settings on the website, the faces of a given partner's "Top 3 Friends" would not be visible if the participant had not previously played

































































a game with that friend on the website. Instead, question marks would be displayed in place of those individuals' faces (Figs. 1b, 2). Thus, participants were unable to view any of their partners' friends in Block 1 (Fig. 3.2), but were then able to see the face and name of up to one friend for Block 2 partners (i.e., friends of Block 2 partners who were people the participant had encountered in Block 1; Fig. 3.2). Half of the Block 2 partners' visible friends were known by the participant to be trustworthy (based on the Block 1 games), and half were known to be untrustworthy. For a full description of all partners' relationships and behavioral tendencies, see Fig. 3.2.

Manipulating Interaction Partners' Behavior

In each round of a trust game, the participant is endowed with a sum of money (\$5 in the current study). They must choose a portion of that money for the experimenter to triple and send to their partner, who can then choose to return any amount of the tripled sum to the participant (Fig. 3.1b-c; Berg et al., 1995). It is maximally advantageous for the participant to invest all of their endowment if their current partner can be trusted to return more than one third of what they receive; investing in an untrustworthy partner (i.e., someone who returns less than one third of the tripled sum), however, leads to a net loss. Participants always played the role of the initial allocator (Player 1) in the trust games, and never the role of the player who received the initial offers (Player 2).

Figure 3.2

Partner Characteristics

	Partner	Partner Trustworthiness	Partner Variance	Friend Knowledge	Top 3 Friends
Block 1		Trustworthy	Consistent	No Known	  
		Trustworthy	Consistent	No Known	  
		Untrustworthy	Consistent	No Known	  
		Untrustworthy	Consistent	No Known	  
Block 2		Trustworthy	Consistent	No Known	  
		Trustworthy	Consistent	Trustworthy	  
		Trustworthy	Consistent	Untrustworthy	  
		Trustworthy	Inconsistent	No Known	  
		Trustworthy	Inconsistent	Trustworthy	  
		Trustworthy	Inconsistent	Untrustworthy	  
		Untrustworthy	Consistent	No Known	  
		Untrustworthy	Consistent	Trustworthy	  
		Untrustworthy	Consistent	Untrustworthy	  
		Untrustworthy	Inconsistent	No Known	  
		Untrustworthy	Inconsistent	Trustworthy	  
		Untrustworthy	Inconsistent	Untrustworthy	  

Note. Participants played repeated trust games with 4 partners in Block 1 (top) and 12 partners in Block 2 (bottom). *Partner:* Each partner was identified with a profile picture and name. *Partner Trustworthiness:* Each partner's reciprocation rate was drawn from a gaussian distribution that had a mean of either 50% (for "trustworthy" partners) or 5% (for "untrustworthy" partners). *Partner Variance:* The distribution from which a partner's reciprocation rate was drawn on each trial had a variance of either 0.01 (for "consistent" partners) or 0.12 (for "inconsistent" partners). *Top 3 Friends:* If the participant had already played with a partner's friend, then they would be able to see who that friend was. If they had not played with a partner's friend yet, then that friend was kept anonymous and only a question mark was shown. Participants did not know any of their partners' friends in Block 1. In Block 2, participants knew either zero or one of each partner's friends. *Friend Knowledge:* Participants either did not know any of the partner's friends ("no known"), or they knew that they had a trustworthy friend ("trustworthy") or untrustworthy friend ("untrustworthy") whom they had played in Block 1 (images with colored borders). In Block 2, there was one partner for each cell of the 2x2x3 factorial design. Images and names were randomly assigned to partners/conditions for each participant. Note: colored borders are only shown here to clearly depict correspondences between Block 1 partners and Block 2 partners' "Top 3 Friends", and did not appear in the experiment.

In the current study, unbeknownst to participants, a computer algorithm determined the average proportion of the tripled sum each of their partners would return to them on each trial (i.e., their trustworthiness). Trustworthy partners' return rates on each trial were drawn from a gaussian distribution with a mean of 50%, whereas untrustworthy partners' return rates were drawn from a gaussian distribution with a mean of 5%. Furthermore, some partners were consistent (their return rates were drawn from a gaussian distribution with a relatively low variance of 0.01), while others were inconsistent (their return rates were drawn from a gaussian distribution with a relatively high variance of 0.12). This offer variance was used for exploratory analyses examining how the relative consistency of a social partner's behavior impacts the extent to which one relies on knowledge of that partner's friends when forming an impression of them. In addition to manipulating partner trustworthiness (2 levels: untrustworthy, trustworthy) and the variance of their offers (2 levels: inconsistent, consistent), we also controlled participant's friend knowledge (i.e., apparent friendships between partners) as described above (3 levels: untrustworthy friend, trustworthy friend, no known friends).

This paradigm provides a quantifiable behavioral measure of how much an individual trusts their partner, while allowing the experimenter to maintain experimental control over the partners' behavior and apparent relationships. By manipulating partners' return rates and known associates,

it is possible to measure how knowledge gained through direct experiences with a partner and knowledge of that partner's relationships shape interpersonal trust, as well as how partners' behaviors shape beliefs about their relationships. That is, the current experimental approach allowed us to (1) decouple social behavioral tendencies from other variables that often covary in real-world situations, (2) ensure that we have data in each cell of the experimental design (e.g., trustworthy partners with no known, trustworthy, and untrustworthy friends) to allow us to test our hypotheses (e.g., to test how first interactions between partners with untrustworthy friends differ from those with trustworthy or no known friends), (3) provide a strong test of our hypotheses by introducing social partners with precisely equivalent behavioral tendencies paired with differing friendship information, and testing for the impact of that friendship information on future partner preferences and interactions within the trust game.

Block 1 of Trust Games. In the first block, participants played with two consistently trustworthy partners and two consistently untrustworthy partners (all Block 1 partners behaved in a consistent manner to facilitate learning; Fig. 3.2). Participants played 10 rounds of the trust game with each of their Block 1 partners in a random, interleaved order, allowing them to learn the relative trustworthiness of each of these four individuals (Chang et al., 2010; Fareri et al., 2012; Fareri, Chang, & Delgado, 2015).

Block 2 of Trust Games. In Block 2, participants played with 12 new partners. As in Block 1, participants played 10 rounds of the trust game with each partner in a random, interleaved order. Participants' Block 2 partners varied on all three possible dimensions: partner trustworthiness, variance, and apparent friendships (as described above). Of the 12 partners in Block 2, there were six trustworthy and six untrustworthy partners, six consistent and six inconsistent partners, and

four partners each with no known, trustworthy, and untrustworthy friends. Thus, each partner filled a different cell of the 2x2x3 factorial design (Fig. 3.2).

Stimuli. Great care was taken to make the online game-playing experience as naturalistic as possible. This was done to ensure that participants believed they were playing with human partners, while still maintaining sufficient experimental control to be able to manipulate participants' beliefs about their partners' behavior, and the friendships between those partners, in a fully balanced design (which allowed us to systematically test the relationships between these variables). On each trial, participants saw their partner's picture with at most one other picture of that partner's top friends. The 16 photos (see Fig. 3.2) were selected from the Chicago Face Database (Ma et al., 2015) and consisted of neutral photos of male and female Asian, Black, Latinx, and White individuals. The images were altered to look more naturalistic (e.g., partners appeared to be wearing various colored shirts in front of natural white walls, rather than the same gray t-shirt in front of artificially removed backgrounds), and their ostensible first name was overlaid on the photo (Fig. 3.2). Below the images was a slider from \$0 to \$5 which the participant used to make their offers (Fig. 3.1b-c). After making each offer, a progress bar appeared as the partner was ostensibly making their decision. For each participant, the partner photos and names were randomly assigned to partner condition (i.e., consistent/inconsistent behavior, trustworthy/untrustworthy behavior, trustworthy/untrustworthy/no known friends). This was done to ensure that there would be no systematic relationships between the behavioral tendencies of participants' partners and aspects of those partners' physical appearance (e.g., race, gender, apparent trustworthiness, attractiveness).

Measuring Generalized Partner Preferences

As described above, the amount that a participant offered on each round of the trust game provided a measure of how much that participant trusted that partner, which, in turn, provided a way to test if participants expected friends to behave similarly (Hypothesis 1). To further test if participants expected friends to behave similarly beyond the immediate context, we measured their general preferences for playing with each partner. After the two blocks of trust games, participants rated each partner from both blocks in terms of how much they would like to play with them again in a different context in the future by responding to the following prompt (adapted from Hackel et al., 2015): “You may have the chance to be invited back to complete a cooperative puzzle-solving game with a partner. If this happens, we’ll do our best to follow your preferences in assigning you a partner. Please rate how much you would like to be paired with each partner you played with today” (Fig. 3.1d). Responses to this question provided a measure of participants’ more general preferences for engaging with each partner in the future, beyond the immediate context of the trust game.

Measuring Perceived Friendships Between Interaction Partners

To test if participants expected similarly behaving partners to be friends (Hypothesis 2), participants completed an incidental memory test assessing their knowledge of their partners’ social network at the end of the study. Participants saw a grid with all 16 partners’ profile pictures along the rows and columns (Fig. 3.1d). They responded to the prompt, “For each pair of players, please indicate whether they appeared as friends in your previous games.” This was done by selecting either the label “Friends” or “Not Friends” in each cell of the matrix, indicating whether the participant thought the partners in the corresponding row and column were in each other’s “Top 3 Friends.”

At the end of the experiment, participants provided basic demographic information (age, gender, ethnicity) and feedback regarding their enjoyment of the game to maintain the cover story regarding the purpose of the experiment. All tasks were completed alone in a private room.

Analyses

Analyses of all data were implemented in R (version 3.6.1). Linear mixed models were implemented using the package nlme (Pinheiro et al., 2019). All means and standard errors reported for linear mixed models are estimated marginal means (i.e., least-squares means) and standard errors using the package emmeans (Lenth, 2019). The reported p -values for all pairwise t -tests have been corrected for multiple comparisons using the Holm correction (Holm, 1979).

Results

Do People Expect Friends to Behave Similarly? (Testing Hypothesis 1)

Block 1 was used for the sole purpose of manipulating participants' beliefs about the trustworthiness of future partners' friends. As such, only offers made in Block 2 were used to test the main hypotheses of the current study. Please refer to the Supplementary Materials S3 for a comprehensive analysis of participants' behavior during Block 1 (Fig. S3.1, Table S3.1).

Interacting With Partners for the First Time

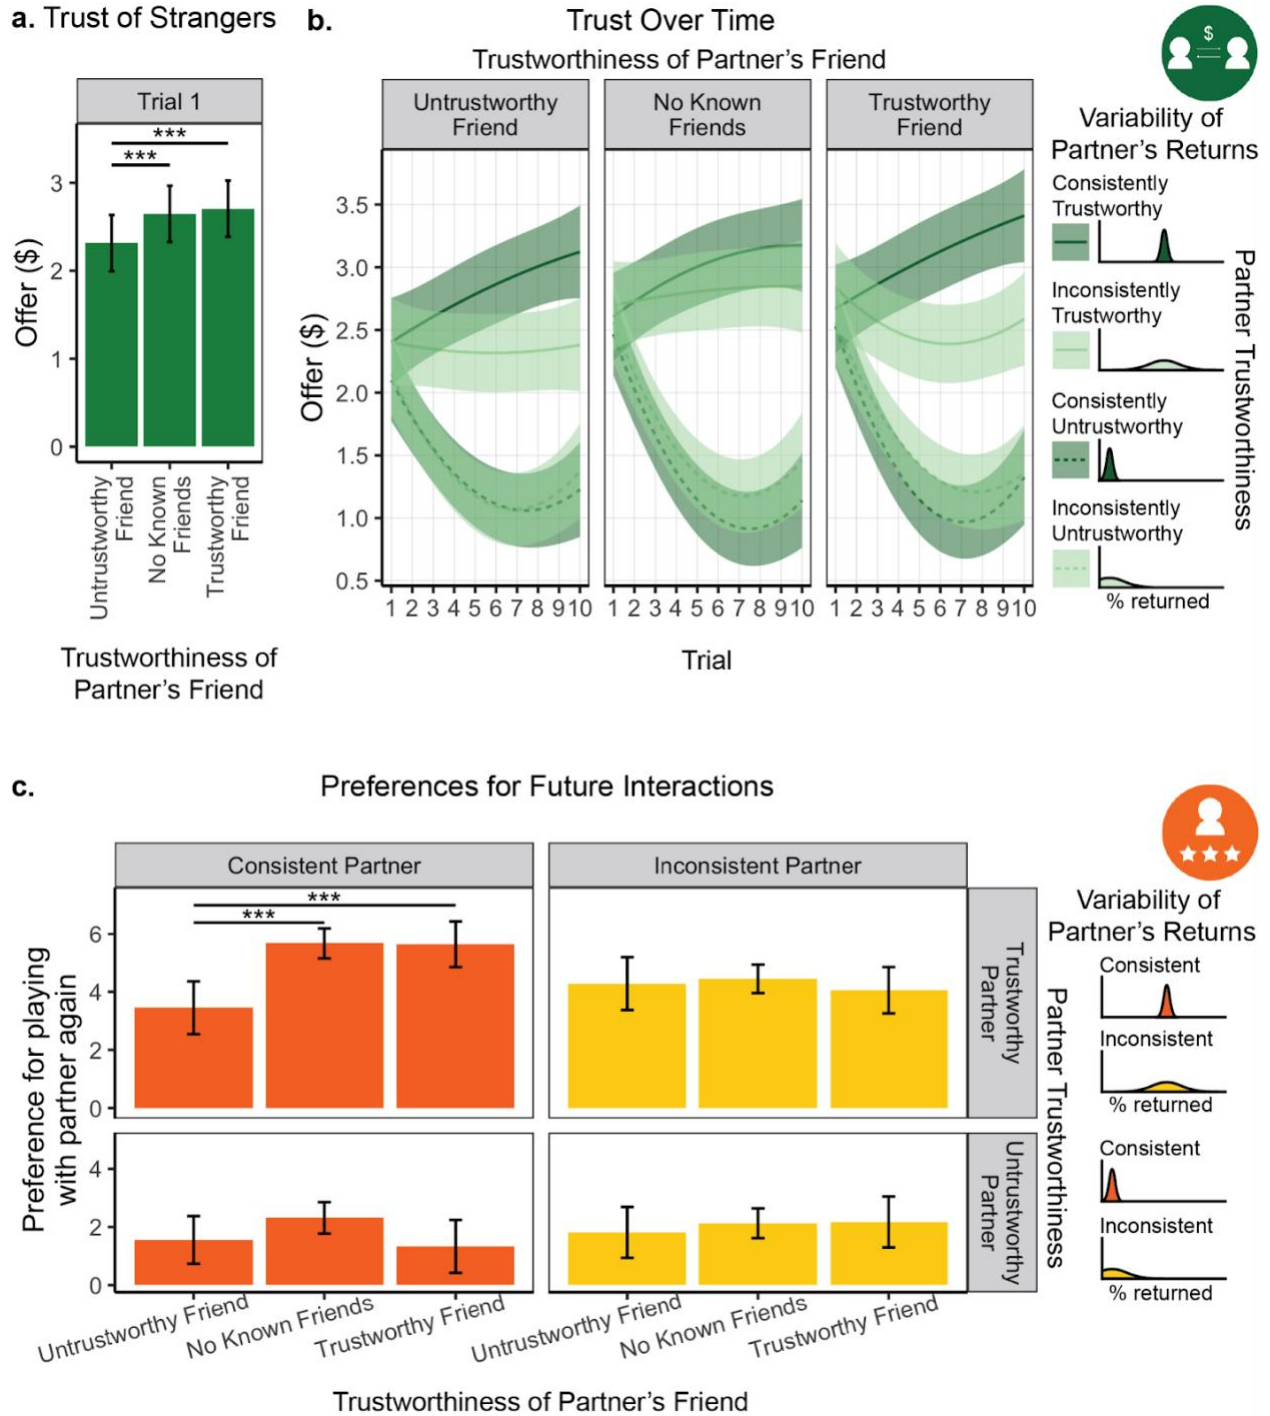
Participants' initial offers to their Block 2 partners provided a window into how third-party relationship knowledge may shape expectations about how strangers will behave. We ran a linear mixed effects model on offers made to each *new* partner in Block 2. Friend knowledge was included as a fixed effect predictor, along with random intercepts for each participant (which account for non-independence among observations of the same participant). This analysis revealed a significant main effect of friend knowledge on participants' initial trust in strangers, $\beta_{\text{trustworthy-unknown}} = 0.15$, 95% CI [0.08, 0.22], $\beta_{\text{untrustworthy-unknown}} = -0.24$, 95% CI [-0.31, -0.17], $F(2,$

878) = 23.62, $p < .001$. Subsequent pairwise t -tests on the least-squares estimation of the means revealed that participants offered partners with untrustworthy friends ($M = 2.32$, $SE = 0.16$) significantly less than those with either trustworthy friends ($M = 2.71$, $SE = 0.16$), $\Delta = 0.39$, 95% CI [0.24, 0.54], $t(878) = 6.37$, $p < .001$, $d = 0.43$, or no known friends ($M = 2.65$, $SE = 0.16$), $\Delta = 0.33$, 95% CI [0.19, 0.48], $t(878) = 5.42$, $p < .001$, $d = 0.37$ (Fig. 3.3a). However, there was no significant difference between offers made to partners with trustworthy friends and those with no known friends, $\Delta = 0.06$, 95% CI [-0.09, 0.21], $t(878) = 0.95$, $p = .343$, $d = 0.06$. The pattern of significance of these effects was not impacted by taking into account other extraneous factors that could shape participants' initial interactions with new partners, such as whether those partners belonged to the same or different demographic categories than the participant (see Supplementary Materials S3).

These findings support our first hypothesis, as they suggest that participants were wary of new partners who had untrustworthy friends, deciding to risk less in those encounters than when playing with partners who had trustworthy friends or unknown friends. Interestingly, it did not matter whether the participant knew nothing about a partner's friends (i.e., no known friends) or if they knew they were trustworthy. If we consider *not* knowing who someone's friends are as the baseline, then it appears that having trustworthy friends provided no benefit in new encounters. However, having an untrustworthy friend resulted in less trust being bestowed.

Figure 3.3

Knowing Who is Friends with Whom Shapes Social Behavior



Note. In Block 2, **(a)** participants initially (i.e., before they had direct knowledge of each partner's trustworthiness) offered smaller amounts to partners with untrustworthy friends than to those with either trustworthy or no known friends. Even after gaining direct knowledge of their partners' behavior, **(b)** the same main effect of friend knowledge held, such that partners with untrustworthy friends were offered less than partners with no known or trustworthy

friends. This main effect of friend knowledge appeared to be driven by cases where the partners themselves were trustworthy (top panel) and consistent (dark green), who were offered more and more over time. Participants were less willing to interact in a (c) new, cooperative context in the future with partners who had untrustworthy friends than those with either no known or trustworthy friends, but only if those partners were consistently trustworthy themselves. Trial variable was centered around 0 in the model and relabeled with 1-10 here for clarity. Asterisks indicate significant differences between friend knowledge conditions (***) $p < .001$. Error bars show 95% CI.

Repeatedly Interacting With Partners

Next, we tested if and how participants' beliefs about their partners' trustworthiness were shaped by (i) knowledge of their partners' social relationships and (ii) directly acquired knowledge of their partners' behavioral tendencies. We examined how friend knowledge affected participants' offers to their partners over time and their preferences for playing with those partners again in the future in different contexts.

Trust Game Offers. We again used a linear mixed effects model to test how offers were affected by friend knowledge (Table 3.1). This method was chosen because it accounts for several sources of non-independence among observations. Specifically, this method accounts for the fact that participants may differ in the overall amounts that they tend to invest in partners who behave in a trustworthy or untrustworthy manner, and in the rate at which those offers change over time, through the inclusion of random by-participant intercepts and slopes, as described below. Further, this approach also allows us to account for temporal autocorrelation by specifying a first-order autoregressive (AR1) covariance structure reflecting that correlations between offers made by a participant decrease as they become farther removed from one another in time. We included the linear effect of time (centered at 0) as well as its quadratic transformation to account for the possibility of a non-linear trend of offers made over time. In total, we included the following fixed effects in this model: friend knowledge, partner trustworthiness, partner variance, and linear and quadratic transformations of time, along with all of their interactions. We also included random intercepts and linear slopes per participant for each level of partner trustworthiness. There was a significant main effect of friend knowledge (Table 3.1, Fig. 3.3b), suggesting that participants

continued to use information about their partners' friendships even after having directly acquired information on which to draw. Additionally, there was a significant interaction effect with the linear transformation of time. This suggests that knowledge of how someone's friend tends to behave (i.e., their friend's trustworthiness) affected how participants learned about that partner's own behavioral tendencies (i.e., their own trustworthiness).

Table 3.1

Linear Mixed Model on Block 2 Offers Over Time

Effect	β	95% CI	<i>F</i>	<i>df1</i>	<i>df2</i>	<i>p</i>	
Partner Trustworthiness	0.79	[0.67, 0.90]	192.55	1	9485	<.001	***
Friend Knowledge			9.10	2	9485	<.001	***
Trustworthy – No Known Friends	0.02	[-0.03, 0.07]					
Untrustworthy – No Known Friends	-0.11	[-0.16, -0.05]					
Variance	0.08	[0.04, 0.12]	17.72	1	9485	<.001	***
Trial	-0.04	[-0.07, -0.02]	13.21	1	9485	<.001	***
Trial ²	0.02	[0.01, 0.02]	91.61	1	9485	<.001	***
Partner Trustworthiness : Friend Knowledge			6.63	2	9485	.001	**
Trustworthy – No Known Friends	-0.02	[-0.07, 0.03]					
Untrustworthy – No Known Friends	-0.07	[-0.12, -0.02]					
Partner Trustworthiness : Variance	0.16	[0.12, 0.19]	71.21	1	9485	<.001	***
Friend Knowledge : Variance			4.52	2	9485	.011	*
Trustworthy – No Known Friends	0.02	[-0.03, 0.07]					
Untrustworthy – No Known Friends	0.06	[0.00, 0.11]					
Partner Trustworthiness : Trial	0.08	[0.06, 0.09]	93.39	1	9485	<.001	***
Partner Trustworthiness : Trial ²	-0.02	[-0.02, -0.01]	93.71	1	9485	<.001	***
Friend Knowledge : Trial			3.58	2	9485	.028	*
Trustworthy – No Known Friends	-0.01	[-0.02, 0.00]					

Untrustworthy – No Known Friends	0.02	[0.00, 0.03]					
Friend Knowledge : Trial ²			1.56	2	9485	.210	
Trustworthy – No Known Friends	0.00	[0.00, 0.01]					
Untrustworthy – No Known Friends	0.00	[-0.01, 0.00]					
Variance : Trial	0.02	[0.01, 0.03]	15.55	1	9485	<.001	***
Variance : Trial ²	0.00	[-0.01, 0.00]	1.29	1	9485	.256	
Partner Trustworthiness : Friend Knowledge : Variance			5.11	2	9485	.006	**
Trustworthy – No Known Friends	0.08	[0.03, 0.14]					
Untrustworthy – No Known Friends	-0.04	[-0.09, 0.01]					
Partner Trustworthiness : Friend Knowledge : Trial			2.38	2	9485	.093	+
Trustworthy – No Known Friends	0.00	[-0.01, 0.01]					
Untrustworthy – No Known Friends	-0.01	[-0.02, 0.00]					
Partner Trustworthiness : Friend Knowledge : Trial ²			1.51	2	9485	.220	
Trustworthy – No Known Friends	0.00	[0.00, 0.01]					
Untrustworthy – No Known Friends	0.00	[0.00, 0.01]					
Partner Trustworthiness : Variance : Trial	0.02	[0.01, 0.03]	28.61	1	9485	<.001	***
Partner Trustworthiness : Variance : Trial ²	0.00	[-0.01, 0.00]	3.18	1	9485	.075	+
Friend Knowledge : Variance : Trial			2.18	2	9485	.113	
Trustworthy – No Known Friends	0.01	[0.00, 0.02]					
Untrustworthy – No Known Friends	0.00	[-0.01, 0.01]					
Friend Knowledge : Variance : Trial ²			0.15	2	9485	.859	
Trustworthy – No Known Friends	0.00	[0.00, 0.01]					
Untrustworthy – No Known Friends	0.00	[-0.01, 0.00]					
Partner Trustworthiness : Friend Knowledge : Variance : Trial			0.48	2	9485	.622	
Trustworthy – No Known Friends	0.00	[-0.01, 0.02]					
Untrustworthy – No Known Friends	0.00	[-0.01, 0.01]					
Partner Trustworthiness : Friend Knowledge : Variance : Trial ²			1.80	2	9485	.165	
Trustworthy – No Known Friends	0.00	[-0.01, 0.00]					
Untrustworthy – No Known Friends	0.00	[0.00, 0.01]					

Note. Results from a linear mixed effects model with offer amount as the dependent variable, and friend knowledge, partner trustworthiness, partner variance, and both linear and quadratic transformations of time (centered Trial) as fixed effect predictors. Random intercepts and by-trial linear slopes were included for each participant for each level of trustworthiness. For effects including Friend Knowledge, which consists of three levels, two estimates and confidence intervals are provided. The first value reflects the difference between offers to partners with trustworthy friends and those with no known friends, and the second reflects the difference between offers made to partners with untrustworthy friends and those with no known friends. *** $p < .001$, ** $p < .01$, * $p < .05$, + $p < .1$

To better understand the three-way interaction effect of friend knowledge, partner trustworthiness, and variance (see Table 3.1), we ran pairwise t -tests on the estimated marginal means. That is, we compared each of the three levels of friend knowledge at each of the combined levels of partner trustworthiness and variance. This approach showed that the above effects were mostly driven by trustworthy partners, as described in more detail below. The only significant pairwise comparison for untrustworthy partners involved *inconsistently* untrustworthy partners; participants offered inconsistently untrustworthy partners who had trustworthy friends ($M = 1.26$, $SE = 0.14$) more than those who had untrustworthy friends ($M = 1.13$; $SE = 0.14$), $\Delta = 0.23$, 95% CI [0.01, 0.45], $t(9485) = 2.50$, $p = .037$, $d = 0.05$. When a participant's direct experience with a partner showed that they were *consistently* not to be trusted, however, that participant did not consider how trustworthy their friends were when deciding whether or not to trust them.

For trustworthy partners, on the other hand, there was a significant difference between partners who acted in a consistent way across trials (i.e., low variance) and those who behaved inconsistently across trials (i.e., high variance). When a partner was consistently trustworthy, the effects of friend knowledge were exactly the same as what was found with the trial 1 data, supporting Hypothesis 1: participants offered significantly less to those with untrustworthy friends ($M = 2.83$, $SE = 0.16$), than those with either trustworthy friends ($M = 3.09$, $SE = 0.16$), $\Delta = 0.26$, 95% CI [0.04, 0.48], $t(9485) = 2.85$, $p = .013$, $d = 0.06$, or no known friends ($M = 3.04$, $SE = 0.16$), $\Delta = 0.21$, 95% CI [0.00, 0.44], $t(9485) = 2.32$, $p = .041$, $d = 0.05$. This can be seen clearly in Fig. 3.3b, where consistent, trustworthy partners (top panel, dark green lines) have

similar slopes, irrespective of participants' knowledge about their friends, but show the lowest intercept when they have an untrustworthy friend, rather than a trustworthy friend, or no known friends.

If a trustworthy partner is inconsistent in their return rates, however, participants offered those with no known friends ($M = 2.81$, $SE = 0.16$) considerably more than those with either trustworthy ($M = 2.40$, $SE = 0.16$), $\Delta = 0.41$, 95% CI [0.19, 0.63], $t(9485) = 4.53$, $p < .001$, $d = 0.09$, or untrustworthy friends, ($M = 2.32$, $SE = 0.16$), $\Delta = 0.50$, 95% CI [0.28, 0.72], $t(9485) = 5.45$, $p < .001$, $d = 0.11$ (Fig. 3.3b, top panel, light green; see Supplementary Materials S3 for further analysis and discussion on how variability in partners' returns modulates the use of friendship knowledge). Together, these results suggest that people were wary of potentially negative interactions: if their partner behaved in an untrustworthy way, or if they had untrustworthy friends, they were offered less money. If, however, they tended to be trustworthy and they had a trustworthy friend, but they were inconsistent (for example, sometimes they returned much less than what was expected for a trustworthy partner with a trustworthy friend), then people trusted them significantly less than they would have if they did not have these priors (i.e., no known friends). For further analyses and discussion of how behavioral variability may influence how and when people use third party knowledge, see the Supplementary Materials S3 (Fig. S3.2).

Generalized Partner Preferences. After the game-playing portion of the experiment had concluded, participants rated how much they would like to play with each partner in a cooperative, and thus new, context in the future. Since it is possible that participants did not accurately remember who was friends with whom, we tested how participant's *true* knowledge about their Block 2 partners' relationships affected their expectations of how those partners would behave in

the future. We used the results of the incidental memory test participants completed at the very end of the experiment to only select rating data for partners whose relationships were correctly recalled. We then used a linear mixed model to test the effects of friend knowledge, partner trustworthiness, and partner variance on partner preference ratings (Table 3.2). The model included friend knowledge, partner trustworthiness, and partner variance as fixed effects, as well as random by-participant intercepts. Here, we saw similar results to what we found in the trust game data. There was a significant three-way interaction, which appeared to be driven by consistent, trustworthy partners who were rated lower if they had an untrustworthy friend ($M = 3.45$, $SE = 0.46$) than if they had a trustworthy friend ($M = 5.64$, $SE = 0.40$), $\Delta = 2.20$, 95% CI [0.74, 3.65], $t(206) = 3.63$, $p = .001$, $d = 0.51$, or no known friends ($M = 5.68$, $SE = 0.27$), $\Delta = 2.23$, 95% CI [0.96, 3.50], $t(206) = 4.24$, $p < .001$, $d = 0.59$. No other pairwise comparisons of friend knowledge were significant. In other words, participants' ratings of their partners suggest that people used relationship knowledge when considering future interactions with partners who had behaved in a consistently trustworthy way, but not when considering future interactions with other partners (Fig. 3.3c). Again, this suggests that participants did not trust people who acted in an untrustworthy manner, no matter who their friends were, and that they were wary of people who seemed trustworthy, but had untrustworthy friends.

Table 3.2

Linear Mixed Model on Preferences for Future Interactions with Partners

Effect	β	95% CI	F	$df1$	$df2$	p	
Friend Knowledge			5.82	2	206	.003	**
Trustworthy – No Known Friends	0.06	[-0.26, 0.39]					
Untrustworthy – No Known Friends	-0.46	[-0.80, -0.13]					

Partner Trustworthiness	1.35 [1.14, 1.57]	150.59	1	206	<.001	***
Partner Variance	0.09 [-0.13, 0.31]	0.66	1	206	.417	
Friend Knowledge : Partner Trustworthiness		1.23	2	206	.296	
Trustworthy – No Known Friends	0.20 [-0.13, 0.52]					
Untrustworthy – No Known Friends	-0.26 [-0.59, 0.07]					
Friend Knowledge : Variance		3.04	2	206	.050	*
Trustworthy – No Known Friends	0.10 [-0.23, 0.42]					
Untrustworthy – No Known Friends	-0.36 [-0.69, -0.03]					
Partner Trustworthiness : Variance	0.24 [0.02, 0.46]	4.80	1	206	.030	*
Friend Knowledge : Partner Trustworthiness : Variance		3.06	2	206	.049	*
Trustworthy – No Known Friends	0.37 [0.04, 0.69]					
Untrustworthy – No Known Friends	-0.39 [-0.72, -0.06]					

Note. Results from a linear mixed model on participants' preferences for interacting with partners again in a new context. Friend knowledge, partner trustworthiness, and partner variance were included as fixed effects, and random by-participant intercepts were also included in the model. Follow-up tests show that friend knowledge is particularly predictive of preferences when partners behave in a consistent and trustworthy manner. For effects including Friend Knowledge, which consists of three levels, two estimates and confidence intervals are provided. The first value reflects the difference between ratings of partners with trustworthy friends and those with no known friends, and the second reflects the difference between preferences for partners with untrustworthy friends and those with no known friends. *** $p < .001$, ** $p < .01$, * $p < .05$

Summary

Analyses of participants' trust game offers and post-game preference ratings suggest that the assumption of similarity among friends shapes expectations about how others will behave (supporting Hypothesis 1). When initially interacting with strangers, people shape their behavior based on knowledge of how those individuals' friends behave. Further, in some cases (e.g., when a trustworthy-seeming individual is known to have untrustworthy friends), social network knowledge continues to shape expectations about a partner's behavior, such that people remain wary of those who associate with untrustworthy people, even in the presence of countervailing evidence about that individual's own behavior.

Do People Expect Those Who Behave Similarly to Be Friends (Testing Hypothesis 2)?

The results in the previous section suggest that assumptions of similarity among friends lead knowledge of others' social relationships to shape expectations about how they will behave. Next, we tested if the converse would also be true. In other words, we tested if knowledge of others' behavior would shape how participants recall their social relationships. Specifically, we tested if participants would expect similarly behaving individuals in a community to be friends with one another, even when they are not.

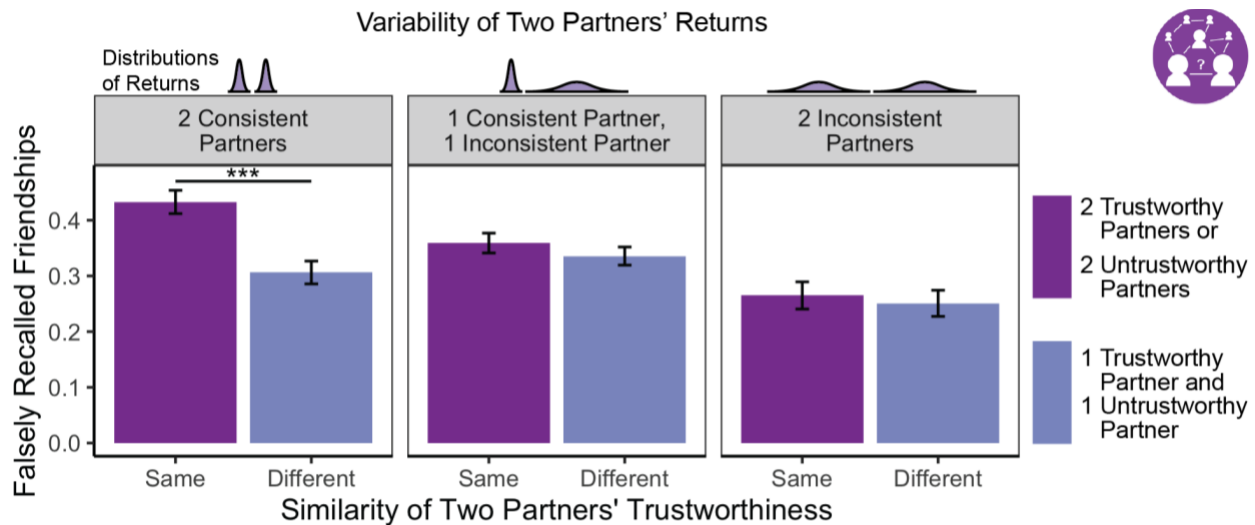
To test the effect of behavioral similarity on beliefs about friendship ties, we created two dyad-level variables to describe each pair of partners that participants were quizzed on (i.e., all possible unique pairs). The first such variable, dyad variance, had three levels: dyads composed of (i) one consistent and one inconsistent partner, (ii) two consistent partners, or (iii) two inconsistent partners. The second dyad-level variable, dyad trustworthiness, had two levels: each pair either consisted of two partners who (i) behaved similarly (i.e., both trustworthy partners or both untrustworthy partners), or (ii) behaved differently (i.e., one trustworthy partner and one untrustworthy partner). Consistent with other work investigating how schemas influence memory in non-social domains (Brewer & Treyens, 1981; De Brigard, Brady, Ruzic, & Schacter, 2017; Lampinen, Copeland, & Neuschatz, 2001), we examined false alarms in the relationship quiz. Specifically, we calculated each participant's false-positive rate for each type of dyad. False positive rate refers to the number of partner pairs that a participant reported as "Top 3 Friends" who were never presented as such, compared to the total number of pairs who were never presented as friends:

$$\text{False Positive Rate} = \frac{\text{Number of pairs falsely remembered as friends}}{\text{Total number of pairs who are not friends}}$$

We used a linear mixed model to test the effects of dyad variance similarity (three levels: two consistent partners; two inconsistent partners; one consistent and one inconsistent partner) and dyad trustworthiness similarity (two levels: similar overall trustworthiness; dissimilar overall trustworthiness) on false-positive rate in the friendship memory quiz. The model included the dyad partners' variance similarity and trustworthiness similarity as fixed effects, as well as random intercepts for each participant. Because the data were bounded by 0 and 1, false positive rates were arcsine transformed before being entered into the model (Freeman & Tukey, 1950). There were significant main effects of both dyad trustworthiness similarity, $\beta = 0.03$, 95% CI [0.01, 0.04], $F(1, 395) = 10.39$, $p = .001$, and dyad variance similarity, $\beta_{\text{consistent-mixed}} = 0.04$, 95% CI [0.01, 0.04], $\beta_{\text{inconsistent-mixed}} = -0.07$, 95% CI [-0.09, -0.04], $F(2, 395) = 16.18$, $p < .001$, as well as their interaction, $\beta_{\text{consistent-mixed}} = -0.04$, 95% CI [-0.06, -0.01], $\beta_{\text{inconsistent-mixed}} = 0.02$, 95% CI [0.00, 0.04], $F(2, 395) = 4.52$, $p = .011$. This was also true when we included a covariate indicating whether or not two partners were of the same race and another accounting for shared gender between partners, $\beta_{\text{trustworthiness}} = 0.02$, 95% CI [0.01, 0.04], $F(1, 1645) = 9.05$, $p = .003$; $\beta_{\text{consistent-mixed}} = 0.03$, 95% CI [0.01, 0.05], $\beta_{\text{inconsistent-mixed}} = -0.06$, 95% CI [-0.08, -0.04], $F_{\text{variance}}(2, 1645) = 12.81$, $p < .001$; $\beta_{\text{consistent-mixed}} = -0.03$, 95% CI [-0.06, -0.01], $\beta_{\text{inconsistent-mixed}} = 0.04$, 95% CI [0.02, 0.06], $F_{\text{interaction}}(2, 1645) = 6.87$, $p = .001$.

Figure 3.4

Behavioral Similarity Shapes Beliefs About Who Is Friends With Whom



Note. Participants were more likely to falsely believe that two partners who behaved similarly (i.e., both trustworthy or both untrustworthy) were friends than two partners who behaved dissimilarly (i.e., one trustworthy, one untrustworthy). However, this was only true if both partners behaved in a consistent manner (i.e., both low variance). Note, for clarity, only the significance of effects of central interest to the current study (i.e., the effects of Friend Knowledge on participants' behavior) are indicated visually with asterisks (***) ($p < .001$). The y-axis reflects the arcsine transformation of false positive rates. Error bars show 95% CI.

Follow-up paired *t*-tests on the estimated marginal means suggest that the main effect of dyad trustworthiness similarity was driven by consistently behaving pairs. This effect was significant when both partners behaved consistently ($M_{similar} = 0.43$, $SE = 0.04$; $M_{dissimilar} = 0.31$, $SE = 0.04$), $\Delta = 0.13$, 95% CI [0.07, 0.18], $t(395) = 4.31$, $p < .001$, $d = 0.43$, but not when both partners behaved inconsistently, or when one partner behaved in a consistent manner and the other did not (Fig. 3.4). That is, participants were more likely to falsely remember a pair of partners as friends if they consistently behaved similarly to one another—i.e., both trustworthy or both untrustworthy—supporting our second hypothesis that people who behave similarly are expected to be friends. This pattern of results also held when looking at trustworthy and untrustworthy pairs separately, such that there was no difference in false positive rates between consistently trustworthy ($M = 0.40$, $SE = 0.04$) and untrustworthy pairs ($M = 0.41$, $SE = 0.04$), but both had

significantly greater false positive rates than pairs that behaved dissimilarly ($M = 0.31$, $SE = 0.04$), $\Delta_{trustworthy} = 0.09$, 95% CI [0.00, 0.17], $t(632) = 2.53$, $p = .023$, $d = 0.20$; $\Delta_{untrustworthy} = 0.11$, 95% CI [0.02, 0.19], $t(632) = 2.98$, $p = .009$, $d = 0.24$.

Summary

Beliefs about others' behavioral tendencies impact how people recall social relationships between others, such that individuals who behave similarly are more likely to be falsely remembered as friends (Hypothesis 2). This suggests that the assumption of similarity among friends shapes mental representations of social networks.

Discussion

In real-world social networks, people tend to be socially connected with those who are similar to themselves. Here, we tested if and how assumptions of similarity among friends shape our beliefs about other people's relationships and behavioral tendencies by having participants play a series of trust games in a fictive online community. More specifically, we tested two hypotheses: that assumptions of similarity among friends will lead people to expect partners who are friends to behave similarly (Hypothesis 1), and that assumptions of similarity among friends will lead people to expect that two partners are more likely to be friends if they behave similarly than dissimilarly (Hypothesis 2).

We frequently make judgments about how to interact with other people, including how much to trust them. Sometimes, we must decide whether or not to trust a complete stranger without any knowledge about that individual. Even in the absence of direct experience with someone, there is a wealth of information on which one can draw when making such decisions. One source of information is the web of relationships in which we, and all of our interactions, are embedded. Here, we found that knowledge of others' social ties shapes evaluations of their behavioral

tendencies. That is, people assumed similarity among friends, leading them to expect that partners who were friends with each other were more likely to behave similarly to each other (supporting Hypothesis 1) and that partners who behaved similarly were more likely to be friends (supporting Hypothesis 2).

Social Network Knowledge Shapes Social Behavior

On the first encounter (i.e., first trust game trial) with each new partner, participants made trust decisions with no prior information about their partner's own behavior. Using third-party relationship knowledge, participants offered less money to those with untrustworthy friends than to other partners, consistent with recent work demonstrating that knowledge of social affiliations biases decisions when interacting with strangers for the first time (Martinez, Mack, Gelman, & Preston, 2016). In the current study, even after gaining direct experience with each partner, participants continued to partially rely on third-party relationship knowledge when making trust-based decisions, offering less money to partners with untrustworthy associates. Interestingly, this was only true when the third-party information signaled a potentially untrustworthy partner: in both the initial offers and subsequent trials, participants did not treat partners with trustworthy friends differently than those with no known friends, but consistently treated those with untrustworthy friends as less trustworthy. This pattern of results was also reflected in participants' post-game preference ratings of each partner. People were significantly less inclined to cooperate in the future with partners who had untrustworthy friends than those with either trustworthy or no known friends. This suggests that people are especially avoidant of cues to potentially risky future social interactions. Taken together, these results provide convergent evidence that knowledge of a social partner's relationships with others shapes one's immediate and long-term expectations of that partner's behaviors.

Knowledge of Others' Behavioral Tendencies Shapes Social Network Beliefs

We also found that a partner's behavioral tendencies (i.e., their trustworthiness in the game) shaped participants' perceptions of third-party ties. Specifically, we tested participants' recall of each partner's friendship ties at the end of the experiment and found that participants were more likely to misremember non-existent friendships between partners who behaved similarly (i.e., both were predictably trustworthy or untrustworthy) compared with those who behaved dissimilarly. That is, knowledge of others' behavioral tendencies shaped memories of their social ties.

Past research on the occurrence (McPherson et al., 2001), causes (Carrarini, Jackson, & Pin, 2009) and assumption (Flynn, Reagans, & Guillory, 2010) of similarity among friends has focused on relatively coarse variables (e.g., demographics). More recent research provides evidence for similarity among friends in terms of behavioral tendencies and how individuals process the world around them (Apicella et al., 2012; Centola, 2011; Parkinson et al., 2018; Pradel et al., 2009; K. M. Smith et al., 2018). The current results suggest that people may internalize such associations between behavioral similarity and social network proximity, and that this in turn may shape, and sometimes distort, their mental representations of social networks.

That said, it is important to note that the evidence for links between social affiliation and interpersonal similarity in cooperative behavior, has been somewhat mixed (i.e., it is not always observed; Ehlert et al., 2020; Simpson et al., 2014; Vernarelli, 2016). Given that participants were led to believe that other players' friendships with one another were based on numerous and diverse interactions on the social gaming website, it is possible that participants inferred that similarly-behaving players had broader interpersonal similarities (e.g., in terms of their general underlying personality traits and values) that led them to both befriend one another and to behave similarly to one another when playing the investment game. It is possible that different results would be

obtained (e.g., the assumption of similarity among friends may have been weaker) had participants been led to believe that their partners only interacted with one another in investment games, given that the link between similarity and social closeness in cooperative tendencies is not always observed. We suggest that future research examine if the assumption of similarity among friends is applied uniformly across traits and contexts, or if the existence and/or magnitude of this assumption, and its impact on people's thoughts and behaviors, varies depending on the trait and/or context at hand (e.g., depending on the extent to which the assumption of similarity among friends reflects the ground truth for a particular attribute).

Potential Mechanisms

Future work could build on the current findings by elucidating the underlying mechanisms that may give rise to people assuming a link between friendship and interpersonal similarity. It may be that people observe this association in their own social networks, and thus assume it exists when interacting with new individuals. Indeed, homophily (i.e., the increased tendency for similar people to become socially linked compared to dissimilar individuals; Lazarsfeld & Merton, 1954) based on readily observable socio-demographic characteristics is commonly observed in real-world social networks. It may be that people who have similar cooperative tendencies are more likely to become friends, thus leading to this social prior. The well-documented effect of social contagion in cooperation (Christakis & Fowler, 2013; Fowler & Christakis, 2010) may also be a driving force in expectations of behavioral similarity. For instance, it could be that people believe partners with untrustworthy friends are more likely to become untrustworthy themselves over time—which could be addressed with future research on how people anticipate networks to form. Finally, there could also be biases in preferences and opportunities (Carrarini et al., 2009), including similarity among socially close individuals that is driven by broader structural forces

rather than solely by individual-level social choices (Schelling, 1978). Future research could integrate these lines of inquiry by examining if people who are more likely to interact with similar others in their real-world social networks also have a stronger tendency to expect people who are connected to one another to behave similarly, and vice versa. Similarly, we suggest that future research examine the match or mismatch between expectations and reality. More specifically, future research should test if people's expectations about how similarity among friends arises, along with the extent to which similarity is expected among friends, differ depending on (1) how likely that similarity is to exist in reality and (2) how much they themselves prefer to befriend similar others.

Prior work suggests that similarities among friends in their cooperative tendencies could arise through direct and indirect reciprocity (Chiang & Takahashi, 2011; Takahashi & Mashima, 2006), as well as mechanisms consistent with biological market theory (i.e., everyone prefers the best or most cooperative partners, resulting in each person being paired with others at their own level; Baumard et al., 2013; Chiang, 2010; Noë & Hammerstein, 1995). To decisively rule out the possibility that the current results reflect strategies related to indirect reciprocity or assumptions about biological market-related mechanisms, future work could extend the current experimental approach to similarity in attributes that are not inherently beneficial or detrimental to others (e.g., similarity in tastes or preferences, rather than trustworthiness). Additionally, as previously noted, in the current study, we were careful to present the gaming website and its members as an established community with a wide variety of ways of interacting with each other. The friendships viewed on this website, therefore, were ostensibly developed over time through a wide variety of games and interactions. This setup makes it somewhat unlikely that participants were assuming mechanisms such as reciprocity or biological markets: two players who behaved in a selfish

manner toward the participant in the investment game would not necessarily have regularly betrayed one another's trust before (since their interactions could have been primarily in other kinds of game-playing contexts); further, the qualities that make a desirable (or undesirable) partner in the investment game would not necessarily generalize to other games. Given that participants were led to believe that other players' friendships were based on numerous and diverse interactions, similarly behaving people could have been expected to be friends with one another because broader similarities between them (e.g., in terms of their tastes, traits, values, etc.) led both to friendship and to similar behavior within the context of the investment game or because broad interpersonal similarities among friends resulted from social influence. That said, we suggest that future work examine people's beliefs about such phenomena (e.g., biological markets, reciprocity), the extent to which those beliefs reflect reality, and how such beliefs shape people's expectations regarding how others will behave.

Rather than assuming similarity among friends, some of our results may reflect that participants expect relationships among three "nodes" in a triad to be balanced (Heider, 1958). In a triad where two nodes are people (e.g., the participant's current partner and that partner's friend) with a positive link between them (they like each other) and the third node is "being untrustworthy," if the partner likes being untrustworthy, the triad is balanced if their friend does too. Thus, if a participant assumes that triads are balanced and knows about their current partner's friend's preference, they would assume that the partner would also prefer to be untrustworthy. Importantly, however, people do not always behave in ways consistent with their attitudes (one can regret a past behavior or dislike things about oneself), and we did not directly manipulate partners' ostensible attitudes. Thus, further research is needed to delineate the mechanisms through

which friendship knowledge impacts presumed behavioral similarities (and attitudes about such behaviors).

Limitations and Future Directions

More generally, it remains to be seen how these results would generalize to other forms of social behavior (e.g., gossip), which may be affected differently by third-party ties due to social consequences (e.g., reputation management; Jolly & Chang, 2018; Smith & Collins, 2009). It is also unclear how these results would generalize across cultures. Given that homophily is thought to pervade social networks across diverse cultures and to have shaped the emergence and disappearance of ties in social networks from an early point in human history (Apicella et al., 2012), it may be the case that assumptions of similarity among friends serve as a social prior across cultures. That said, although the current study's sample was ethnically diverse (see Supplementary Materials S3), this research was limited to undergraduate students at a single American institution. The heuristics that guide social thought and behavior do not necessarily generalize across cultures; many psychological phenomena that were previously thought to generalize across cultures have since been shown to be specific to the population most often studied: people from WEIRD (western, educated, industrialized, rich, and democratic) cultures (Henrich, Heine, & Norenzayan, 2010; Rad, Martingano, & Ginges, 2018). As such, further research is needed to examine if and how the assumption of similarity among friends, as well as its impact on how people think and behave, differs across cultures and populations.

A growing body of research suggests that tracking the structure of the social world is vital to many aspects of everyday social thought and behavior (Brands, 2013; Brent, 2015; E. R. Smith & Collins, 2009; Weaverdyck & Parkinson, 2018). This requires accurately tracking others' relationships (i.e., ties between third parties). The current research advances our understanding of

these phenomena, and could be integrated with research on other kinds of social knowledge, such as the inference of in-group/out-group boundaries (Lau, Gershman, & Cikara, 2020; Lau, Pouncy, Gershman, & Cikara, 2018), to gain a fuller picture of how we infer the structure of the social world in which we are embedded (Parkinson & Du, 2020). Additionally, future research should explore the extent to which people assume that others will be connected with individuals who are similar to themselves with respect to additional behavioral tendencies and traits; such assumptions may be particularly strong with respect to cooperative tendencies, as in the current study, given that the tendency for cooperators to sever ties with exploitative individuals may underlie broader cooperation (Izquierdo, Izquierdo, & Vega-Redondo, 2014).

Conclusion

All social interactions unfold within networks of social relationships. Here, we found evidence that people use third-party relationship knowledge when making trust decisions, even after having direct experience on which to draw. Additionally, we found that expectations of behavioral similarity among friends shape mental representations of social networks. Taken together, these results suggest that similarity among friends serves as a social prior, causing knowledge of social networks and of others' behavioral tendencies to reciprocally interact to shape social thought and behavior. This reciprocal relationship between beliefs about others' behavioral tendencies and their relationships with one another suggests that the social decisions that people make are fundamentally intertwined with the networks of social relationships that they inhabit.

Supplementary Materials S3

Learning Partner Trustworthiness in Block 1 Trust Games

Block 1 trust games were used for the sole purpose of manipulating participants' beliefs about the trustworthiness of future partners' friends. To verify that this manipulation was effective, we conducted a linear mixed effects model with offer amount as the dependent variable, and partner trustworthiness, both linear and quadratic transformations of time (Trial), and interactions between time and partner trustworthiness, as fixed effect predictors. Random intercepts and random by-trial linear slopes were included for each participant for each level of trustworthiness (Table S3.1). The main effect of partner trustworthiness showed that participants invested less in untrustworthy partners ($M = 0.95$, $SE = 0.13$) than in trustworthy partners ($M = 3.56$, $SE = 0.15$), $\Delta = 2.61$, 95% CI [2.33, 2.89], $t(3115) = 18.17$, $p < .001$, $d = 0.65$. Additionally, the interaction effects of partner trustworthiness with both the linear and quadratic transformations of time were significant, suggesting that participants successfully learned, over the course of the block, which Block 1 partners were trustworthy and which were not (Figure S3.1). Thus, we were able to assess the impact of this knowledge on their behavior towards new partners during Block 2.

Table S3.1


Linear Mixed Model on Block 1 Offers Over Time

Effect	β	95% CI	F	$df1$	$df2$	p	
Partner Trustworthiness	1.30	[1.16, 1.44]	330.07	1	3115	<.001	***
Trial	0.00	[-0.03, 0.02]	0.01	1	3115	.926	
Trial ²	0.01	[0.00, 0.01]	9.42	1	3115	.002	**
Partner Trustworthiness : Trial	0.14	[0.12, 0.16]	153.43	1	3115	<.001	***
Partner Trustworthiness : Trial ²	-0.03	[-0.03, -0.02]	150.80	1	3115	<.001	***

Note. Results from a linear mixed model on offers made in Block 1. The model included partner trustworthiness and both linear and quadratic transformations of time (centered Trial) as fixed effects, and random by-participant intercepts and by-participant linear slopes over time at each level of partner trustworthiness. *** $p < .001$, ** $p < .01$

Figure S3.1

Participants Successfully Learned Partner Trustworthiness in Block 1



Note. Participants offered more money to partners who behaved in a trustworthy manner than to those who behaved in an untrustworthy manner, suggesting that they successfully learned which partners were trustworthy and which were not in Block 1. Error bars show 95% CI.

Are People More Influenced by Knowledge of Their Partners' Friends When Their Partners Themselves Are Less Predictable?

In our main analyses, we tested how the consistency of partners' behavior affected the use of friendship knowledge in trust game offers and partner ratings by including it in the full linear mixed effects models. The results of such analyses indicated, for example, that people were impacted by knowledge of their current partners' friends when interacting with *inconsistently* untrustworthy partners, but not when interacting with *consistently* untrustworthy partners, as described in the main text. That said, these analyses are limited to testing directed effects of partner variance and therefore would not detect an effect if participants varied in *how* they use friendship knowledge (see below for further discussion). Thus, we also conducted an exploratory analysis

examining if people are generally more impacted by friendship knowledge when their partner is relatively unpredictable (i.e., when they have behaved inconsistently) than when their partner is more predictable (i.e., when their return rates have been very consistent). To examine the variability of offers made to Block 2 partners as a function of friend knowledge, we calculated the coefficient of variation (i.e., the ratio of the standard deviation to the mean) across mean offers to the three partners (one at each level of friend knowledge) for each participant within each unique combination of partner trustworthiness and variance. This yielded an estimate of how much each participant's offers varied as a function of friend knowledge within each unique combination of the remaining two conditions, which pertain to their partners' own behavior (partner trustworthiness and variance). We ran a linear mixed effect model on these data, with partner trustworthiness and variance as fixed effects, and random by-participant intercepts (Table S3.2). This analysis revealed a significant main effect of partner variance, such that participants varied their offers based on friend knowledge more when their partner was inconsistent ($M = 0.39$, $SE = 0.03$) than when their partner behaved consistently ($M = 0.33$, $SE = 0.03$; Table S3.2). Subsequent analyses of the estimated marginal means for each combination of condition levels revealed that this was true only in the trustworthy condition, $\Delta = 0.08$, 95% CI [0.01, 0.15], $t(237) = 2.14$, $p = .033$, $d = 0.28$. A non-significant difference in the same direction was found when analyzing variability in the partner preference ratings in the same way (Table S3.2, Fig. S3.2). Thus, knowledge of how a partner's friend had behaved influenced participants' behavior more when that partner was inconsistently, rather than consistently, trustworthy.

Table S3.2

Linear Mixed Model on Coefficients of Variation for Trust Game Offers and Partner Preference

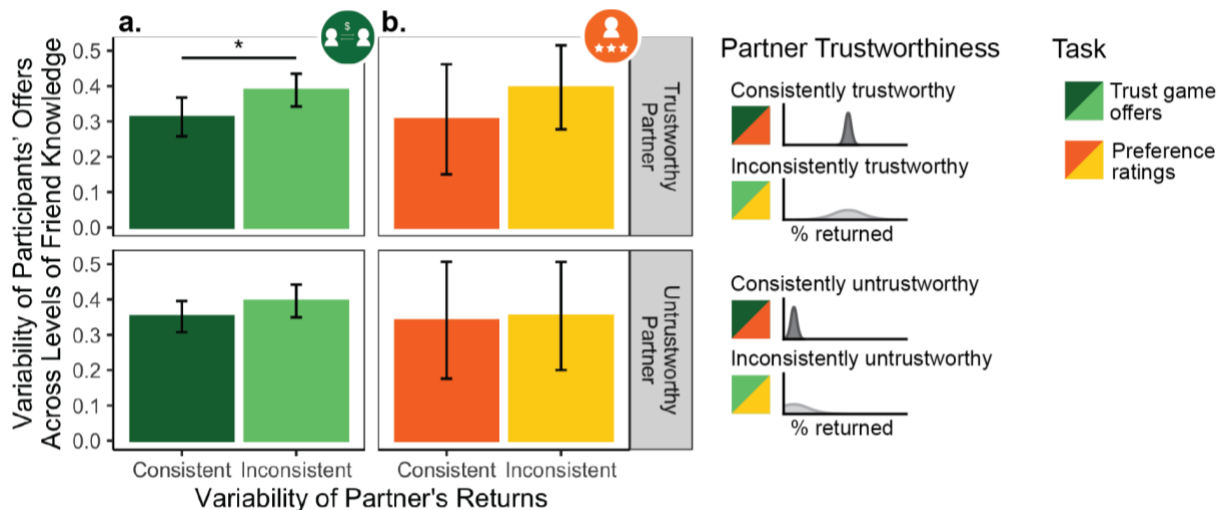
Ratings Across Levels of Friendship Knowledge in Block 2

Effect	β	95% CI	<i>F</i>	<i>df1</i>	<i>df2</i>	<i>p</i>
Trust Game Offers						
Partner Trustworthiness	-0.01	[-0.04, 0.01]	0.86	1	237	.355
Partner Variance	-0.03	[-0.05, -0.01]	5.75	1	237	.017 *
Partner Trustworthiness : Partner Variance	-0.01	[-0.03, 0.02]	0.40	1	237	.529
Partner Preference Ratings						
Partner Trustworthiness	0.00	[-0.08, 0.08]	0.00	1	34	.998
Partner Variance	-0.02	[-0.10, 0.06]	0.23	1	34	.632
Partner Trustworthiness : Partner Variance	-0.01	[-0.09, 0.06]	0.15	1	34	.700

Note. The coefficient of variation (i.e., the ratio of the standard deviation to the mean) across mean offers to (i.e., Trust Game Offers) and ratings of (i.e., Partner Preference Ratings) the three partners within each unique combination of partner trustworthiness and variance was calculated per participant, which provides an estimate of how much each participant's offers varied as a function of friend knowledge. One linear mixed model was run per task, and included partner trustworthiness and partner variance as fixed effects, along with random by-participant intercepts. Follow-up *t*-tests show that the friend knowledge influenced participants' offers more (i.e., participants' offers had a higher coefficient of variation) when playing with inconsistently compared to consistently trustworthy partners. **p* < .05

Figure S3.2

People Are Influenced More by Friend Knowledge When Their Partners Are Unpredictable



Note. The coefficient of variation (i.e., the ratio of the standard deviation to the mean) across (a) mean offers to and (b) preference ratings of the three Block 2 partners within each unique combination of partner trustworthiness and variance was calculated per participant. This provides an estimate of how much each participant's offers varied as a function of friend knowledge within each of the four cells and is plotted on the y-axis. As shown in (a), friend knowledge influenced participants' offers more (i.e., higher coefficient of variation) when playing with inconsistent

(compared to consistent) partners, $F(1, 237) = 5.75, p = .017$. This effect was driven by trustworthy partners, such that when participants played with trustworthy partners, their offers varied more as a function of their knowledge of a partner's friends if the partner behaved inconsistently (i.e., high-variability), rather than consistently (i.e., low-variability), $\Delta = 0.08$, 95% CI [0.01, 0.15], $t(237) = 2.14, p = .033, d = 0.28$. A similar, but non-significant, pattern of results is seen in partner preference ratings, as shown in **(b)**. Asterisks indicate significant differences between friend knowledge conditions ($*p < .050$), and error bars show 95% CI.

These results (which reflect how much, in general, participants' offers were influenced by friend knowledge for each kind of partner they played with, irrespective of how this influence unfolded within each participant) differ somewhat from the effects of partner variance noted in our main analyses. It should be noted that the analyses of participants' coefficients of variation described above are sensitive to the extent to which friend knowledge shaped behavior, but not the direction of such effects, whereas the effects of partner variance in the analyses described in the main text could only detect the influence of friend knowledge if it unfolded in a consistent direction across participants. As such, one possibility is that participants' behavior varied more as a function of friend knowledge when their partners were inconsistently trustworthy (i.e., mean return rate = 50%; variance = 0.12) than when their partners were consistently trustworthy (i.e., mean return rate = 50%, variance = 0.01; Figure S3.2), but the precise nature of this relationship varied somewhat across participants. For example, when interacting with inconsistent but generally trustworthy partners, some participants may have been strongly influenced by expectations of similarity among friends that persisted across interactions, expecting people with trustworthy friends to be trustworthy themselves, even when their own behavior was variable. Contrastingly, other participants may have had particularly aversive responses to sometimes receiving relatively low offers from these same partners (as return rates for high-variance trustworthy partners dipped below 33% on some trials), especially when they had a strong prior expectation of fair behavior (i.e., especially when a trustworthy partner who was known to have trustworthy friends offered less than expected on a given trial). We also note that it may have been relatively difficult for some or all participants to learn that the inconsistently trustworthy partners were in fact trustworthy,

given that although their return rates were drawn from a distribution centered on 50%, this distribution was relatively wide and included offers that were considerably lower than 50%. Thus, 10 trials with each interaction partner in Block 2 may not have been sufficient for all participants to learn that high-variance, trustworthy partners were in fact trustworthy (consistent with Figure 3b-c). It is also possible that for many people, the concepts of trustworthiness and behavioral consistency are fundamentally intertwined, and thus, two partners with equivalent return rates (e.g., 50%) but different levels of variability across trials may not always be perceived as equivalently trustworthy, even if the central tendency of their behavior is equivalent across trials. Thus, further research is needed to better understand the relationship between behavioral variability, social network knowledge, and expectations about how other people will behave.

Examining the effects of shared demographic characteristics between participants and their partners

All stimulus images were randomly assigned to partner profiles to negate any potential effects of demographics on trust behavior. For completeness's sake, and to explore any possible influence these incidental social categories may have had on participants' behavior, we reran all of the linear mixed models examining trust game offers with covariates accounting for shared race and gender between the participant and each partner. The stimuli were categorized as Asian, Black, Latinx, or White according to the Chicago Face Database. Participants were categorized into the same four categories ($N_{asian} = 39$, $N_{black} = 3$, $N_{latinx} = 14$, $N_{white} = 16$), along with a fifth category ($N = 8$) for participants who did not self-identify as one of these. For each participant, a binary variable was created to code for whether or not the participant and the partner were of the same race. Similarly, the stimuli were categorized as male or female in the Chicago Face Database, and

a second binary variable was created to code for whether or not the participant and the partner identified as the same gender.

Having the same race as one's partner did not significantly impact offers made in block 1 ($M_{same} = 2.23$, $SE = 0.13$; $M_{different} = 2.25$, $SE = 0.12$), $\Delta = 0.02$, 95% CI [-0.13, 0.08], $t(3113) = 0.47$, $p = .638$, $d = 0.02$, or block 2, for either initial offers ($M_{same} = 2.64$, $SE = 0.16$; $M_{different} = 2.53$, $SE = 0.16$), $\Delta = 0.11$, 95% CI [-0.01, 0.23], $t(876) = 1.75$, $p = .080$, $d = 0.12$, or offers over time ($M_{same} = 1.96$, $SE = 0.13$; $M_{different} = 1.96$, $SE = 0.13$), $\Delta = 0.01$, 95% CI [-0.05, 0.07], $t(9483) = 0.28$, $p = .785$, $d = 0.01$. Similarly, having the same gender as a partner did not significantly predict offers made in block 1 ($M_{same} = 2.22$, $SE = 0.12$; $M_{different} = 2.26$, $SE = 0.12$), $\Delta = 0.03$, 95% CI [-0.05, 0.12], $t(3113) = 0.74$, $p = .460$, $d = 0.03$, or initial offers in block 2 ($M_{same} = 2.58$, $SE = 0.16$; $M_{different} = 2.59$, $SE = 0.16$), $\Delta = 0.01$, 95% CI [-0.09, 0.11], $t(876) = 0.19$, $p = .847$, $d = 0.01$; however, there was a significant effect on offers over time, such that participants offered more to players of different genders ($M_{same} = 1.92$, $SE = 0.13$; $M_{different} = 2.00$, $SE = 0.13$), $\Delta = 0.08$, 95% CI [0.03, 0.13], $t(9483) = 2.97$, $p = .003$, $d = 0.06$. Furthermore, the patterns of results discussed in the main text of the manuscript remained the same with and without the added covariates.

It is of particular interest that neither gender nor race impacted offers made to new partners in block 2, but friendship knowledge did. Importantly, the current study was not designed to test effects of demographic knowledge on trust behavior, and thus we interpret these results with extreme caution. These findings do suggest, however, that when encountering someone for the first time, knowledge of that person's relationships with others may play a particularly important role in impression formation. As such, future research examining effects of perceived social

categories on cognition and behavior may benefit from considering the effects of social network information.

Examining if People Tended to Forget Ties Among Dissimilar Partners

In our main analyses, we examined if the assumption of a link between behavioral similarity and friendship colored mental representations of social relationships by examining false alarms, consistent with previous work that has examined the impact of schemas on memory of non-social contents (Brewer & Treyens, 1981; De Brigard et al., 2017; Lampinen et al., 2001). For completeness, we also tested if assumptions of a link between behavioral similarity and friendship shaped the extent to which participants forgot friendships that did exist (i.e., false negative rate). We ran a linear mixed model with false negative rate (arcsine transformed) as the dependent variable, dyad trustworthy similarity, dyad variance similarity, and their interaction as fixed effects, and random by-participant intercepts. We did not find any significant main or interaction effects in this analysis, $\beta_{\text{trustworthiness similarity}} = 0.04$, 95% CI [-0.01, 0.10], $F(1, 237) = 2.19$, $p = .140$; $\beta_{\text{variance similarity}} = 0.02$, 95% CI [-0.04, 0.07], $F(1, 237) = 0.37$, $p = .542$; $\beta_{\text{interaction}} = -0.02$, 95% CI [-0.07, 0.04], $F(1, 237) = 0.37$, $p = .542$. That said, consistent with the pattern of results discussed in the main text, we found that the marginal mean false negative rates tended to be slightly higher for dyads composed of people who were dissimilar in their trustworthiness compared to dyads composed of people who were similar to each other. This was true for both pairs of people who behaved in a consistent manner ($M_{\text{dissimilar}} = 0.85$, $M_{\text{similar}} = 0.81$), and those composed of one consistently-behaving and one inconsistently-behaving partner ($M_{\text{dissimilar}} = 0.85$, $M_{\text{similar}} = 0.74$). (Note, since all Block 1 partners were consistent, there were no known friendships that consisted of two inconsistent partners.) That is, people tended to be slightly more likely to forget that two people were friends if they behaved dissimilarly. Importantly, these effects did not reach

significance and should thus be interpreted with caution; thus, further research is needed to confirm and clarify this effect.

Examining Whether Similarity-Attraction Was Present

The primary question of this manuscript asks if people *assume* that friends will be exceptionally similar to one another when interacting with others. A separate question is: Does similarity-attraction (i.e., preference for similarly-behaving others) exist in the context of our experiment? We are unable to directly test this question with the current data because participants always played the role of Player 1 in the trust game, but they only observed others play the role of Player 2. Therefore, we have data on how much participants tended to place *their* trust in others (i.e., how *wary/trusting* they were of others), but participants did not learn how *wary/trusting* their partners were of others. Similarly, while participants learned how *trustworthy* their partners were, we do not have data on how trustworthy participants were themselves. One's own trustworthiness and how wary one is of trusting others are distinct qualities: For example, someone could be both cautious when it comes to placing their trust in new people and still be unlikely to betray others' trust themselves (i.e., wary but trustworthy); someone could also be both cautious when it comes to placing their own trust in new people and likely to betray others trust (i.e., wary and untrustworthy). We suggest that future work augment the current paradigm by (1) also having participants observe their partners be Player 1 in the trust game, and/or (2) by measuring participants' behavior when they assume the role of Player 2 in the trust game. This would provide data on (1) both participants' own wariness and others' apparent wariness (i.e., how readily they trust people) and/or (2) both participants' own trustworthiness and the apparent trustworthiness of others. Such data would allow future work to build on the current study, which tests if people assume that others will be similar to their friends, by testing the separate but complementary

questions of (1) if this assumption reflects reality and (2) if the degree to which people expect others to behave similarly to their friends is related to the extent to which they themselves prefer to interact with and befriend similar others.

All of that being said, we conducted exploratory analyses testing if the extent to which people were wary of trusting others was related to the extent to which they preferred partners who are trustworthy (i.e., examine associations between participants' behavioral tendencies regarding how *wary/trusting* they were of others, and their preference ratings of others who were differentially *trustworthy*). We ran a linear mixed model with participant wariness (calculated as their average offer to all partners, such that higher average offers reflect less wariness), partner trustworthiness, and their interaction predicting partner preference ratings (*z*-scored within each participant) with random by-participant slopes. If participants who are less wary (i.e., more trusting) prefer trustworthy partners *and* more wary participants (i.e., less trusting) prefer untrustworthy partners, then we should see a significant interaction effect. However, only partner trustworthiness significantly predicts partner preferences, $\beta = 0.63$, 95% CI [0.55, 0.72], $F(1, 1198) = 208.90$, $p < .001$, such that trustworthy partners are preferred above untrustworthy partners by everyone. As noted above, people who are trusting/wary of others are not necessarily trustworthy/untrustworthy themselves, and vice versa. As such, we cannot speak confidently as to whether or not participants preferred those who were more similar to themselves.

Although it is not possible to draw conclusions about the existence of similarity-attraction based on these results, given the limitations of the data noted above, this data analytic approach could be adopted in future work to test for similarity-attraction (e.g., to test if the trustworthiness of participants is related to their preference for trustworthy partners and/or if participants' wariness

of placing their trust in others is related to their preference for others who are similarly wary of placing their trust in others).

Experimental Instructions

Excerpts from the onscreen instructions provided to participants are included below. The full paradigm is available online at <https://anonymous.4open.science/r/c4145239-6a86-47b9-a681-a063184ec182>.

“Welcome to our social gaming website! You are participating in a study that examines how various factors, like the amount of reward that’s at stake, impact players’ enjoyment of online games. You’ll be playing a series of simple games with other people on this website. We’re currently testing out this website with collaborators at other colleges in the U.S. The players you’ll be partnered with today are students at other colleges who regularly play a variety of games with one another for fun and the chance to earn prizes (e.g., money). Users of the site regularly have the chance to rate each other in terms of how much they prefer playing with one another. Players who consistently choose one another as their favorite partners to play with on the site are called one another’s "Top Friends". When members of this website play with one another, their 3 Top Friends are displayed below their own profile photo.... Since this is your first time using this website, you won’t yet have any Top Friends displayed alongside your profile picture. Due to most users’ privacy settings, new users like you can only see those friends’ faces if you’ve interacted with them before on the site.

....

Please wait while we assign games for you...

In Block X, you are assigned to play the Investment Game with other players. In each round of the Investment Game, you'll be partnered with one person.

Please wait while we assign your role...

You've been assigned the role of first player.

....

You have now completed Block 2! We really want to hear about your enjoyment of the games you played today. To ensure that your session today doesn't run overtime, we're now going to move on to the post-game feedback session rather than having you complete Block 3."

CHAPTER 4

Predicting that birds of a feather will flock together:

Expectations of homophily for others but not the self

Miriam E. Schwyck, Carolyn Parkinson

University of California, Los Angeles

Abstract

Similarity among friends – or other socially connected individuals – is a common and ubiquitous characteristic of social networks. There are several mechanisms through which such similarity arises, including homophily, or the tendency for similar people to attract one another. While past research has found that people use similarity heuristics to structure their mental representations of social networks (predicting that similar people are likely to be friends), it is unknown if people assume such similarity arises through homophily or other processes (e.g., social influence). Here, we tested if people assume that homophily will govern their own and others' future friendships. Through a series of repeated trust games, participants learned how trustworthy various partners were towards the participant (i.e., amount of money the partner returned) and how trusting other partners were of the participant (i.e., amount of money the partner invested). Participants predicted which partners would become friends with one another and which would become friends with themselves were they to meet in person. Across two studies and both trait measures, we found that participants were significantly more likely to predict that partners who behaved similarly would become friends compared to those who behaved dissimilarly. Interestingly, we found that participants were significantly more likely to predict that they would become friends with highly trustworthy and highly trusting partners compared to highly untrustworthy or highly untrusting partners, regardless of their own behavioral tendencies or even their own self-perceptions. These findings suggest that participants assume homophily for others' relationships but not their own. Such expectations likely shape how people approach or foster new friendships for themselves and between others.

Introduction

Every social interaction takes place within the context of social relationships. One common feature of real-world social networks is that people who are directly connected (e.g., friends) tend to be more similar than those who are more distantly connected, as reflected in the common phrase, “birds of a feather flock together”. Research in psychology, sociology, and network science has studied this phenomenon of similarity among friends, how it arises over time, and the social implications of it (Altermatt & Pomerantz, 2003; Apicella et al., 2012; Bahns, Pickett, & Crandall, 2012; Brechwald & Prinstein, 2011; Centola, 2011; Christakis & Fowler, 2013; de Klepper et al., 2010; Dehghani et al., 2016; Ehlert et al., 2020; Fowler & Christakis, 2010; Lawrence & Shah, 2020; McPherson et al., 2001; Parkinson et al., 2018; J. A. Smith et al., 2014; K. M. Smith et al., 2018). Interestingly, this body of research covers similarity among socially close people in terms of a wide range of features, including relatively coarse demographic characteristics (e.g., age, race, gender), abstract ideologies (e.g., religion, political ideation), and behavioral choices.

The ubiquity of similarity among friends makes it a potentially powerful tool for people to use when inferring new information about others. Indeed, humans have been shown to capitalize on this phenomenon by assuming such similarity exists when viewing or interacting with others: when predicting others’ likely friendships based on knowledge of their individual characteristics, people often assume that people who are similar in terms of social values (Goel, Mason, & Watts, 2010), demographic categories (Flynn et al., 2010), and even preferences and behaviors (Z. Liberman, Kinzler, & Woodward, 2021; Schwyck*, Du* et al., 2023) are more likely to be friends. Conversely, when predicting others’ characteristics based on knowledge of their friends, there is evidence that people also assume friends are more likely to be similar (Schwyck*, Du* et al., 2023). This is especially helpful given that social networks are large, complex structures such that

it would be infeasible to directly observe and cognitively track each relationship individually (Basyouni & Parkinson, 2022; Brands, 2013). Thus, people appear to use this similarity heuristic to infer and track information about others, thereby reducing the cognitive load associated with representing social network information.

There are many possible mechanisms through which similarity among close others (e.g., friends) arises in real-world social networks. People may become more similar to their friends over time due to both social influence and exposure to shared environments, or similar others may be attracted to each other (Altermatt & Pomerantz, 2003; Brechwald & Prinstein, 2011; Christakis & Fowler, 2009, 2013; de Klepper et al., 2010; McPherson et al., 2001). This last possibility is known as homophily, or the tendency for people to seek out or associate with similar others. While there is research showing that people assume similarity among people who are currently friends, as discussed above, it is not clear if people assume that similar people will *become* friends through homophily. Here we test this possibility using online games in which participants learn about others' social traits through repeated interactions, then predict who would become friends with whom. In addition to asking which gaming partners would become friends, we also examined participants' predictions for their own relationships.

When testing the accuracy of individuals' mental representations of real-world social networks, researchers usually determine the ground truth based on whether the two people themselves report having a relationship (Brands, 2013). Thus, much of the research in this area has focused on perceptions of others' relationships rather than one's own, and it is unknown if people apply the same heuristics that they use to represent others' relationships to their own. It may be that people predict all relationships similarly, irrespective of whether they themselves are involved, or they may use a different method for predicting their own versus others' relationships.

When considering such possibilities in the context of socially desirable traits (e.g., trustworthiness), it is important to consider not only objective measures of people's behavior, but also their own self-perceptions. This is because people tend to have a positive bias when considering their own levels of desirable traits, such that most people believe themselves to be better than average (Zell, Strickhouser, Sedikides, & Alicke, 2020), a statistical impossibility. This phenomenon could distort people's assessments of similarities between themselves and others with respect to desirable traits. For example, an untrustworthy person could assume that they will become friends with trustworthy others, not because they use different methods for predicting their own future relationships, but rather, because they perceive themselves to be more trustworthy than they actually are. Alternatively, people may have qualitatively different expectations about the factors that will guide their own and others' friendships. For example, people may be overly optimistic about the traits their future friends would possess, illustrating a positivity bias about their own future. They may thus base their predictions of their own friendships on different criteria (e.g., the assumption that their own friends will have socially desirable traits, irrespective of their own self-perceptions) than their predictions of others' friendships (e.g., the assumption that others' friendships will be governed primarily by homophily). Do people assume they will become friends with those who are (in reality) similar to themselves, those whom they believe to be similar to themselves even when they are not (e.g., due to inflated self-perceptions), or those who are most desirable regardless of their own traits?

In the current set of studies, participants played repeated online trust games through which they could earn small rewards with various partners. Trust games (Berg et al., 1995) are a useful tool that allows researchers to manipulate how *trusting* a gaming partner is (how much money did they invest in the other player) and how *trustworthy* that partner is (how much of the investment

did they return to the other player). Additionally, this paradigm allows us to measure participants' own levels of trustworthiness and trustingness, as well as their perceptions of their own and others' levels of these traits. Across two studies, we used these two measures of behavioral similarity to systematically test if people assume homophily for others and for themselves.

Methods

Participants

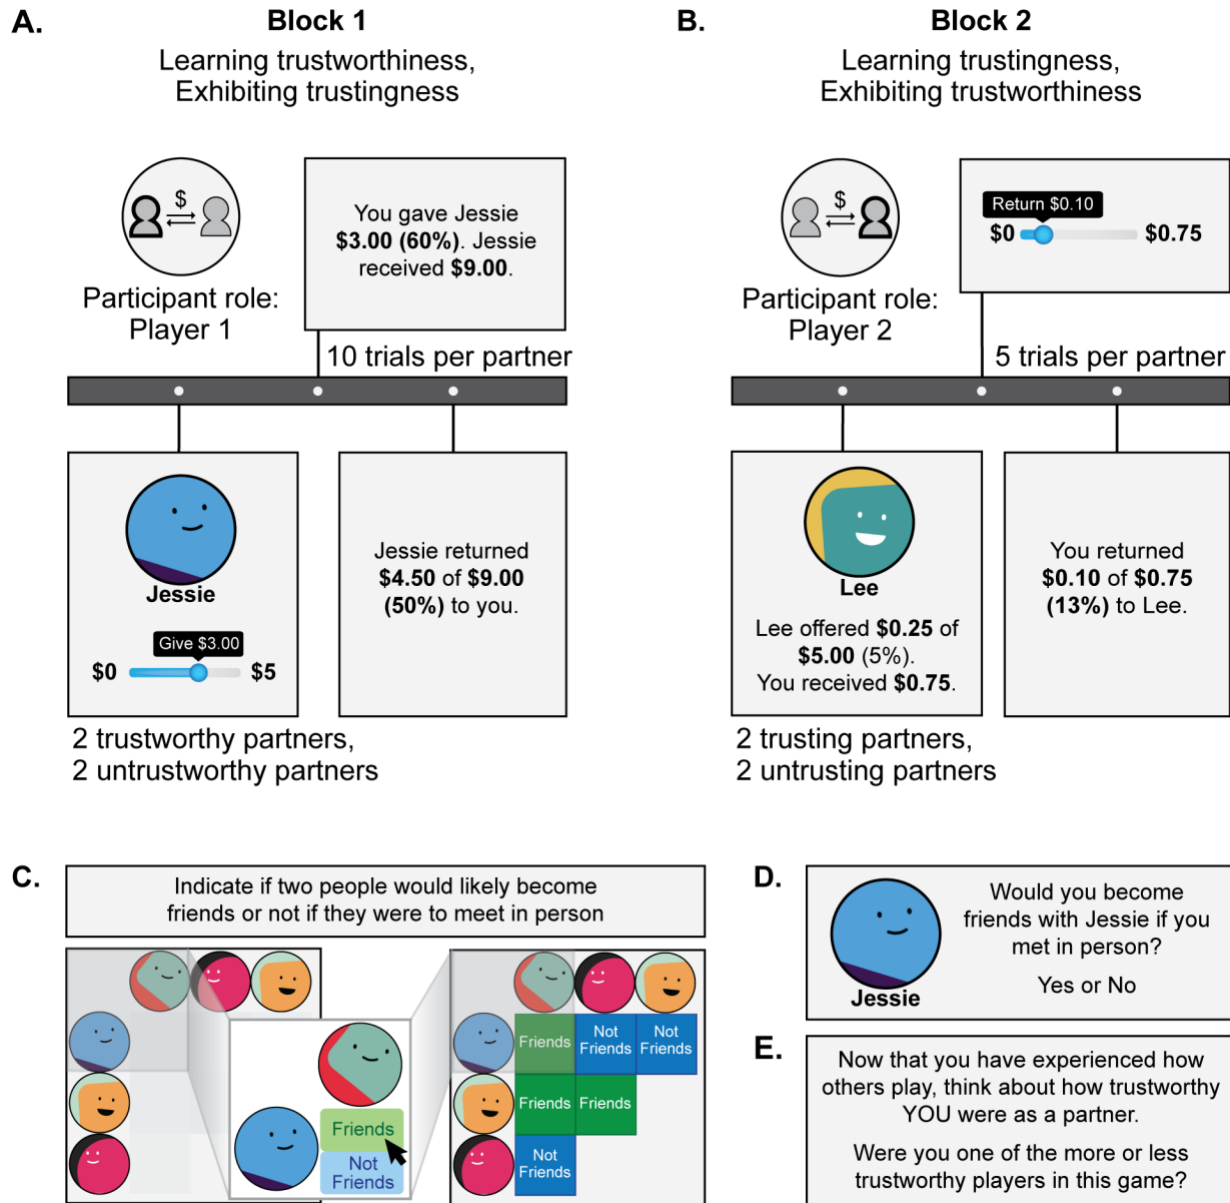
Participants were undergraduate students recruited through the University of California, Los Angeles (UCLA) Sona Systems and provided written informed consent in accordance with the policies of the UCLA ethical review board. For both studies, we recruited as many participants as possible within one term. We excluded people for reporting that they paid little to no attention during the task ($N_{\text{Study1}} = 8$; $N_{\text{Study2}} = 4$) and for taking more than three times the expected length of the study ($N_{\text{Study1}} = 5$; $N_{\text{Study2}} = 4$). In Study 1, this resulted in a final sample size of 345 (280 female, 62 male, 1 non-binary, 3 did not specify, ages = 18-52 years, $M_{\text{age}} = 20.22$, $SD_{\text{age}} = 3.14$). In Study 2, we had a final sample size of 215 (182 female, 26 male, 5 non-binary, 2 did not specify, ages = 18-44 years, $M_{\text{age}} = 20.37$, $SD_{\text{age}} = 2.67$).

Procedure

Both studies followed very similar paradigms that were administered online. Participants played repeated trust games (Berg et al., 1995) with various partners who were ostensibly past participants (Fig. 4.1). That is, participants were told that their partners had already made choices based on various scenarios and that the behavior the participant would observe was based on these choices. Note, because the partners were not playing in real time, they could not react to the participants' choices. In reality, a computer algorithm determined the response that each partner gave.

Figure 4.1

Trust Game and Friendship Prediction Paradigms



Note. Participants played repeated trust games with various partners in two blocks. (A) In Block 1, participants played the role of Player 1 with four partners, two of whom returned high percentages (trustworthy partners) and two of whom returned low percentages (untrustworthy partners). The offers made by participants in Block 1 served as a measure of their level of trust in their partners. (B) In Block 2, participants played the role of Player 2 with four new partners, two of whom offered high percentages (trusting partners) and two of whom offered low percentages (untrusting partners). The return rates by participants in Block 2 served as a measure of their level of trustworthiness. (C) After each block, participants predicted which partners would likely become friends and (D) which partners they would become friends with if they were to meet in person. (E) In Study 2, participants reported if they were one of the more trustworthy players and if they were one of the more trusting players.

In each round of a trust game, Player 1 is endowed with a sum of money (\$5 in the current studies). They must choose a portion of that money to be tripled and sent to their partner, Player 2, who can then choose to return any amount of the tripled sum to Player 1 (Berg et al., 1995). It is maximally advantageous for Player 1 to invest all of their endowment if their current partner can be trusted to return more than one third of what they receive; investing in an untrustworthy partner (i.e., someone who returns less than one third of the tripled sum), however, leads to a net loss. Thus, the amount Player 1 offers to Player 2 indicates the extent to which Player 1 is trusting of Player 2, and the amount Player 2 returns to Player 1 indicates the extent to which Player 2 is trustworthy. In Block 1 of the current studies, participants were in the Player 1 role, and in Block 2, participants were in the Player 2 role.

To avoid potential inferences based on facial features, each partner was represented by a simple smiling avatar and a gender-neutral name. Furthermore, all images and names were randomized across participants. Before playing, participants chose an avatar for themselves and entered their own display name. In Block 1, participants (as Player 1) played with four partners, two exhibiting trustworthy behaviors (their returns were drawn from a gaussian distribution with a mean 50%) and two exhibiting untrustworthy behaviors (their returns were drawn from a gaussian distribution with a mean of 5%). Participants played with each partner 10 times in interleaved trials. Next, they were asked to predict who would become friends with whom. Specifically, they were asked to “indicate if two people would likely become friends or not if they were to meet in person” and to answer the question, “Would you become friends with <partner name> if you met in person?” for each partner they played in Block 1. In Study 2, an additional question was added before making these friendship predictions that quizzed participants knowledge of their partners’ behavioral patterns. Specifically, participants responded to the

prompt, “People who are more trustworthy returned more of the money they received back to you compared to those who returned very little. Was <partner name> one of the more trustworthy partners that you played with in block 1?”

In Block 2, participants (as Player 2) played four new partners, two exhibited trusting behaviors (they offered on average 50% of the endowment) and two exhibiting untrusting behaviors (they offered on average 5% of the endowment). Participants played with each partner 5 times in interleaved trials. Note that the Player 2 role does not intrinsically require participants to learn the overall patterns of their partners in order to perform optimally. This is because participants can simply respond to the amount offered to them on a given trial, whereas in the Player 1 role, participants must predict what that partner will do next with their investment in order to perform optimally. Thus, we quizzed them on their partners’ relative levels of trustiness in both studies. Specifically, participants responded to the prompt, “People who are more trusting invest more money with the hope that you will return more than they invested. Was <partner name> one of the more trusting partners that you played with in block 2?” Next, participants were asked to predict who would become friends with whom using the same prompts as in Block 1. In Study 2, participants were also asked to reflect on their own behavior and report if they were one of the more trusting players in the game, and if they were one of the more trustworthy players in the game.

At the end of the experiments, participants completed a personal social network survey which was to be used for exploratory analysis and will not be included here. Finally, they provided basic demographic information (age, gender, ethnicity) and were then debriefed regarding the deception that was involved in the design.

Analyses

We used two measures of behavior in this game to test assumptions of homophily: trustworthiness (the proportion of money returned by Player 2 to Player 1) and trustingness (the proportion of money offered by Player 1 to Player 2). We calculated each participants' level of trustingness and trustworthiness in the game by calculating the mean of their offers in Block 1 and the mean percent they returned in Block 2, respectively. We then categorized each participant based on their means. Participants were categorized as untrusting if their mean offer was below 15% and trusting if their mean offer was above 40%. Similarly, participants were categorized as untrustworthy if their mean percent return was below 15% and trustworthy if their mean percent return was above 40%. These values were chosen to be within 10 percentage points of the 5% and 50% means that categorized partners as untrusting/untrustworthy and trusting/trustworthy, respectively.

To test if people expect similarly behaving people to become friends, we calculated several dyad-level variables with 3 levels each. Partner-partner trustworthiness similarity describes whether two partners were both untrustworthy, both trustworthy, or dissimilar (i.e., one untrustworthy and one trustworthy) based on their behavior. Self-partner trustworthiness similarity describes whether a partner and the participant were categorized as both untrustworthy, both trustworthy, or dissimilar based on their behavior. Analogous variables were calculated for trustingness: partner-partner trustingness similarity, and self-partner trustingness similarity. In Study 2, participants also reported whether they believed they had been trustworthy/untrustworthy and trusting/untrusting. Thus, we calculated two additional dyad-level variables for Study 2, self-report-based trustworthiness similarity and self-report-based trustingness similarity; these

variables reflect if participants' self-reported trustworthiness and trustingness matched each partner's trustworthiness and trustingness, respectively.

Analyses of all data were implemented in R (version 4.2.1; R Core Team, 2022). Mixed effects logistic regression models were implemented using the package lme4 (Bates, Maechler, Bolker, & Walker, 2015) included by-participant random intercepts and predicted the binary choice ('yes', 'no') of whether two people would likely become friends were they to meet in person (i.e., friendship prediction). All means and standard errors reported for linear mixed models are estimated marginal means (i.e., least-squares means) and standard errors using the package emmeans (Lenth, 2019). The reported *p*-values for all pairwise *t*-tests on the marginal means were corrected for multiple comparisons using the Holm correction (Holm, 1979).

For each mixed effects regression, we calculated the relevant marginal means and the relevant pairwise comparisons. Specifically, for models in which trustworthiness was considered, we calculated the mean for pairs who were both trustworthy and the mean for pairs who were both untrustworthy. For the partner-partner similarity models, we calculated the mean for pairs who behaved dissimilarly (i.e., one trustworthy partner and one untrustworthy partner), and we calculated all pairwise comparisons. For self-partner similarity models, we calculated one mean for dissimilarly behaving pairs in which the participant was trustworthy (i.e., with an untrustworthy partner), and another for which the participant was untrustworthy (i.e., with a trustworthy partner). In these models, we calculated the differences between similarly behaving pairs and dissimilarly behaving pairs with participant category. That is, we tested the difference between trustworthy participants with trustworthy partners (i.e., similarly behaving) and those with untrustworthy partners (i.e., dissimilarly behaving). Similarly, we tested the differences between untrustworthy participants with untrustworthy partners (i.e., similarly behaving) and those with trustworthy

partners (i.e., dissimilarly behaving). That is, we tested similarly behaving pairs to dissimilarly behaving pairs within trustworthy participants and within untrustworthy participants.

For models in which trustingness was considered, we ran the equivalent comparisons as described above. For partner-partner similarity models, we calculated and compared the means for pairs who were both trusting, both untrusting, and dissimilarly behaving. For self-partner similarity models, we calculated the means for pairs who were both trusting, both untrusting, dissimilarly behaving in which the participant was trusting, and dissimilarly behaving in which the participant was untrusting. We then compared pairs who were both trusting to pairs with a trusting participant and an untrusting partner, and we compared pairs who were both untrusting to pairs with an untrusting participant and a trusting partner.

Results

Predicting Others' Friendships


To test if people have expectations of homophily for others (that is, to test if people expect those who are already similar to each other to become friends in the future), we ran separate mixed effects logistical regressions with partner trustworthiness similarity and partner trustingness similarity as the predictors of predicted partner friendship. If homophily is expected, then participants will predict similarly behaving partners (e.g., both trustworthy, both untrustworthy) to be friends more frequently than dissimilarly behaving pairs (e.g., one trustworthy, one untrustworthy).

This is exactly what we found in both studies (Fig. 4.2, Table 4.1, Table S4.1). When making predictions about future friendships between others after having learned about their trustworthiness, participants were significantly less likely to predict that dissimilarly behaving partners would become friends in the future than two trustworthy partners or two untrustworthy

partners. Additionally, participants were significantly more likely to predict that two trustworthy partners would become friends than two untrustworthy partners.

Figure 4.2

Partner-Partner Friendship Predictions



Note. In both studies, people were more likely to predict that similarly behaving partners would become friends in the future compared with dissimilarly behaving partners. This was true both when participants learned about their partners' trustworthiness (left panels) and their trustingness (right panels). Additionally, participants were more likely to predict that two trustworthy partners would become friends than two untrustworthy partners, and that two trusting partners would become friends than two untrusting partners. The y-axis reflects the mean proportion of pairs that were predicted to become friends, the error bars reflect standard error, and the comparison lines reflect the pairwise difference tests on the marginal means. *** $p < .001$

A similar pattern of results was observed in both studies with respect to trustingness. Specifically, when making predictions about future friendships between other people after having learned through experience how trusting those people were, participants were significantly less likely to predict that dissimilarly behaving partners would become friends in the future than two

trusting partners or two untrusting partners. Additionally, they were significantly more likely to predict that two trusting partners would become friends in the future compared with two untrusting partners.

Table 4.1

Marginal Mean Comparisons From Logistic Regression Predicting Partner-Partner Friendships

Logistic Regression Marginal Means Comparisons	Study 1			Study 2		
	OR	SE	z	OR	SE	z
Both Trustworthy Partners (similar) / Differently Behaving Partners (dissimilar)	13.98	2.21	16.69***	46.10	12.03	14.68***
Both Untrustworthy Partners (similar) / Differently Behaving Partners (dissimilar)	7.34	1.00	14.70***	9.94	1.71	13.31***
Both Trustworthy Partners (similar) / Both Untrustworthy Partners (dissimilar)	1.90	0.36	3.39***	4.64	1.34	5.33***
Both Trusting Partners (similar) / Differently Behaving Partners (dissimilar)	23.30	4.02	18.26***	18.18	3.99	13.20***
Both Untrusting Partners (similar) / Differently Behaving Partners (dissimilar)	7.47	0.99	15.18***	5.03	0.85	9.63***
Both Trusting Partners (similar) / Both Untrusting Partners (dissimilar)	3.12	0.62	5.77***	3.61	0.88	5.28***

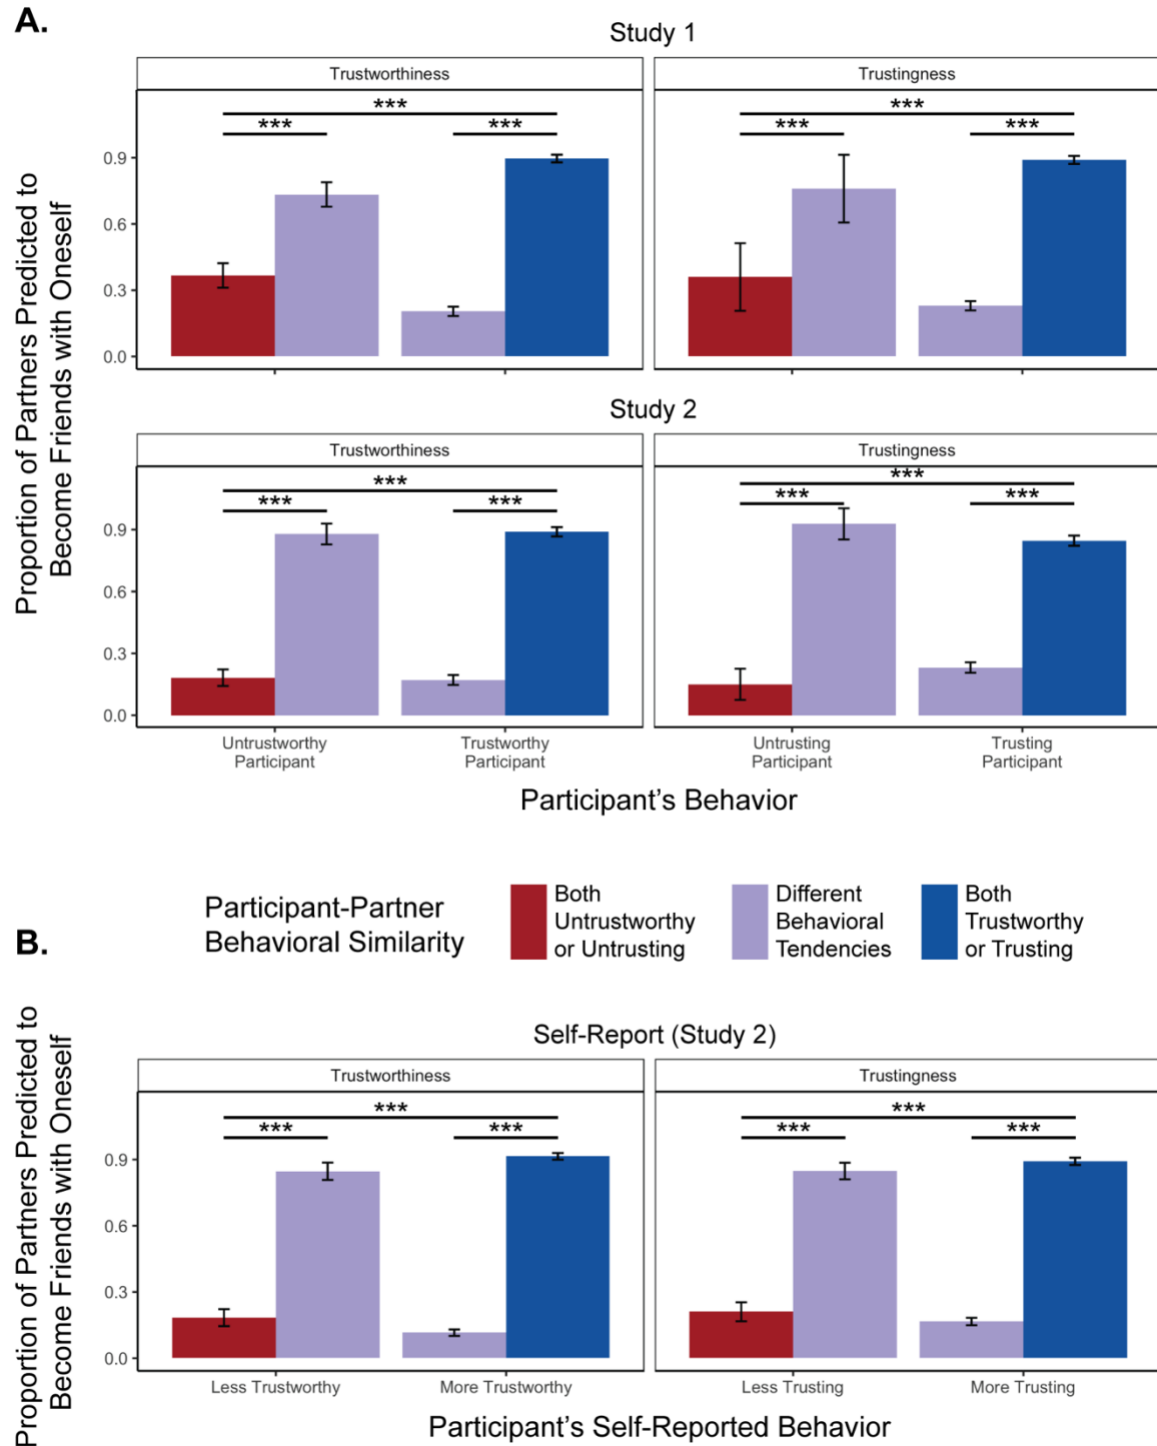
Note. We ran mixed effects logistic regressions with partner behavioral similarity predicting participants’ choices of who would become friends in the future and by-participant random intercepts. We ran one model on trustworthiness behavior and another on trustingness behavior per study. The above planned contrasts were run on the marginal means and corrected for multiple comparisons using the Holm method. *** $p < .001$

Predicting One’s Own Friendships

To test if people have expectations of homophily for themselves, we ran separate mixed effects logistical regressions involving trustworthiness and trustingness. In each model, we included self-partner behavioral similarity, partner’s behavior, and their interaction as predictors of predicted self-partner friendship. If homophily is expected for relationships involving oneself, then participants will predict that they will become friends with partners who behave similarly to themselves more often than partners who behave dissimilarly.

Figure 4.3

Self-Partner Friendship Predictions



Note. Contrary to assumptions of homophily, participants did not predict they would become friends with partners who behaved similarly to themselves (dark red and blue bars) more than dissimilarly behaving partners (light purple bars). Instead, they consistently were more likely to predict that they would become friends with those who behaved

in a trustworthy (left panels) or trusting (right panels) manner. (a) This was true for participants who were categorized based on their game play as trustworthy, untrustworthy, trusting, or untrusting. (b) It was also true for participants who categorized their own behavior as trustworthy, untrustworthy, trusting, or untrusting. The y-axis reflects the mean proportion of pairs that were predicted to become friends, the error bars reflect standard error, and the comparison lines reflect the pairwise difference tests on the marginal means. *** $p < .001$ corrected for multiple comparisons.

Table 4.2

Marginal Mean Comparisons From Logistic Regressions Predicting Self-Partner Friendships

Logistic Regression Marginal Means Comparisons	Study 1			Study 2					
	Behavioral Similarity			Behavioral Similarity			Self-Reported Behavioral Similarity		
	OR	SE	z	OR	SE	z	OR	SE	z
Both Trustworthy (similar) / Trustworthy Participant, Untrustworthy Partner (dissimilar)	43.68	9.30	17.74***	65.50	19.90	13.76***	81.69	21.10	17.05***
Both Untrustworthy (similar) / Untrustworthy Participant, Trustworthy Partner (dissimilar)	0.07	0.02	-8.43***	0.02	0.01	-9.03***	0.04	0.02	-8.10***
Both Trustworthy (similar) / Both Untrustworthy (similar)	14.92	4.85	8.32***	69.13	33.30	8.79***	47.01	15.58	11.62***
Both Trusting (similar) / Trusting Participant, Untrusting Partner (dissimilar)	36.96	7.51	17.76***	32.69	10.06	11.34***	51.76	16.78	12.18***
Both Untrusting (similar) / Untrusting Participant, Trusting Partner (dissimilar)	0.07	0.05	-4.01***	0.01	0.01	-5.52***	0.04	0.02	-6.92***
Both Trusting (similar) / Both Untrusting (similar)	12.00	8.03	3.71***	56.11	46.91	4.82***	33.94	13.51	8.86***

Note. We ran mixed effects logistic regressions with participant-partner similarity, partner behavior, and their interaction predicting friendship prediction, with by-participant random intercepts. We ran one model on trustworthiness behavior and another trustingness behavior per study. In Study 2, we ran additional models using self-reported trustworthiness/trustingness to calculate similarity. The above planned contrasts were run on the marginal means and corrected for multiple comparisons using the Holm method. *** $p < .001$

In Study 1, we found that friendship prediction was significantly predicted by partner trustworthiness, OR = 5.00, $z = 16.91$, $p < .001$, self-partner trustworthiness similarity, OR = 1.32, $z = 2.92$, $p = .003$, and their interaction, OR = 0.77, $z = -2.72$, $p = .006$. Examining the marginal means of this interaction, we find both untrustworthy and trustworthy participants were more likely to predict that they would become friends with trustworthy partners more than untrustworthy partners (Table 4.2, Fig. 4.3). In Study 2, neither self-partner trustworthiness similarity nor the

interaction term were significantly predictive of people's anticipated future friendships, whereas partner trustworthiness was highly associated with predictions of their own future friendships, $OR = 7.90$, $z = 15.17$, $p < .001$, such that participants consistently predicted that they would become friends with trustworthy partners more than untrustworthy partners.

For trustiness, we found in both studies that partner trustiness was significantly predictive of friendship prediction, $OR_{study1} = 4.82$, $z_{study1} = 8.99$, $p_{study1} < .001$; $OR_{study2} = 7.80$, $z_{study2} = 8.61$, $p_{study2} < .001$, while there was no significant effect of behavioral similarity or the interaction. Specifically, once again, we found that participants consistently predicted that they would become friends with trusting partners more than untrusting partners.

One possible reason why such a pattern of results might emerge is that participants did assume that they would become friends with others who are similar to themselves, but that their self-perceptions did not match their behavior. For example people who behave in an untrustworthy way could still perceive themselves to be trustworthy due to the better than average effect (Zell et al., 2020). Accordingly, such individuals would predict that they would become friends with trustworthy others if using assumptions of homophily to scaffold their predictions about their future friendships. In other words, these results could be congruent with expectations of homophily for the self. As such, in Study 2, we asked participants to reflect on their own behavior in the game and state if they were one of the more or less trustworthy players, and if they were one of the more or less trusting partners. We did find that self-perceptions and reality do not completely align, although we found both participants with a positive bias and those with a negative bias (see Supplementary Materials S4).

Using this self-categorization, we ran logistic regression models (one for trustworthiness, one for trustiness) with trust similarity (now calculated based on participants' self-reported

behavioral tendencies) and partner behavior. Interestingly, we found the same pattern of results in which all categories of participants predicted that they would become friends with trustworthy and trusting partners over untrustworthy and untrusting partners, respectively. Specifically, we found a significant main effect of partner trustworthiness, $OR = 6.69$, $z = 16.10$, $p < .001$, and trustworthiness similarity, $OR = 1.35$, $z = 2.54$, $p = .011$, but no interaction effect. For trustingness, we only found a main effect of partner trustingness, $OR = 6.08$, $z = 11.17$, $p < .001$, on friendship prediction. This suggests that people were not predicting that homophily would determine which relationships they themselves would be involved in, despite predicting that homophily would determine relationships that only involve others.

Discussion

People need to be able to track their own and others' relationships to successfully navigate their social networks. Given the vast size of these networks, humans rely on heuristics, such as similarity among friends, to reduce the computational load that tracking these networks requires (Brands, 2013; Brashears, 2013). Here, we test if people assume that similarity among friends arises through homophily for themselves and for others. Across two studies, participants played repeated trust games with various partners and predicted who would become friends were they to meet in person. Participants played both roles in the trust game, allowing us to test for assumptions of homophily with respect to interpersonal similarity in two distinct traits: *trustingness* was measured by the amount a player offered to their partner (in hopes their partner would return a substantial portion of their investment) and *trustworthiness* was measured by the amount that was returned by the player who received the offer. We tested if participants assumed homophily shaped the formation of friendships between their partners after participants learned about their partners' social behavioral tendencies through repeated game play. By having participants play both roles,

we could also test if people used the same assumptions of homophily for themselves as they do for others.

In line with the hypothesis that people assume homophily for others, we consistently found that participants predicted that similarly behaving partners were more likely to become friends with each other in the future than partners who behaved dissimilarly. Importantly, this was true for both partners who were consistently trustworthy and those who were consistently untrustworthy. We found the same pattern of results for participants' predictions of future friendships between people high and low in trustingness. Taken together, these results suggest that assumptions of homophily for others did not vary across levels of the traits examined here. This is consistent with past work examining the heuristic of similarity among current friends (which could have been driven by homophily, social influence, shared experiences among existing friends, or other factors), which found that behavioral similarity predicts friendship recall for both positive and negative traits (Schwyck*, Du* et al., 2023).

Interestingly, we found that participants did not assume homophily when predicting their own friendships. Participants consistently predicted that they would become friends with trustworthy partners much more frequently than untrustworthy partners, and trusting partners more frequently than untrusting partners, regardless of participants' own behavioral patterns. That is, to the extent that these qualities are desirable, people predict they will become friends with the most desirable others, not those who are most similar to themselves. We tested if this was due to the tendency for people to believe they are above average on these potentially desirable traits (Zell et al., 2020), but found the same results when using participants' self-reported categorization based on their game-playing behavior. That is, even when using participants' own self-perceptions, participants did not predict their own friendships based on homophily. One possible explanation

for this asymmetry in how people predict friendships involving themselves and how they predict friendships that only involve others stems from the reward associated with certain partners' behaviors. Since interactions with trustworthy and trusting partners both yielded more money to the participant than those with untrustworthy and untrusting partners, participants likely found these interactions more rewarding and thus associated these partners with positive experiences. Given that people likely prefer to associate with those with whom they have shared positive experiences, it may be that this mechanism of association over-rode the assumption of homophily that people apply to others' relationships.

Given that trustworthiness is a socially desirable trait to have in friends, it could also be that people did not assume homophily for their own relationships because trustworthiness is particularly desirable. However, since we also found the same pattern of results for trustingness (a trait that is not necessarily desirable – e.g., indiscriminately trusting people could be viewed as foolish and/or gullible), this suggests that our findings may not be specific to socially desirable traits. Future work could further test how homophily shapes predictions of one's own and others' friendships for traits that are unambiguously socially undesirable or neutral. By investigating positive, negative, and neutral traits, future work would expand the current findings to provide a more nuanced understanding of how people predict their own future relationships.

Here, we tested if and when people expect homophily to govern the formation of future social relationships. We found that people consistently use expectations of homophily when predicting others' future relationships but predicted only the most desirable future social partners for themselves. These findings build on and provide new generative findings for the rich literature that sits at the intersection of psychology, sociology, and network science examining the cognitive mechanisms that support the ability for humans to navigate their complex social networks.

Supplementary Materials S4

Table S4.1

Friend Prediction Means and Standard Deviations

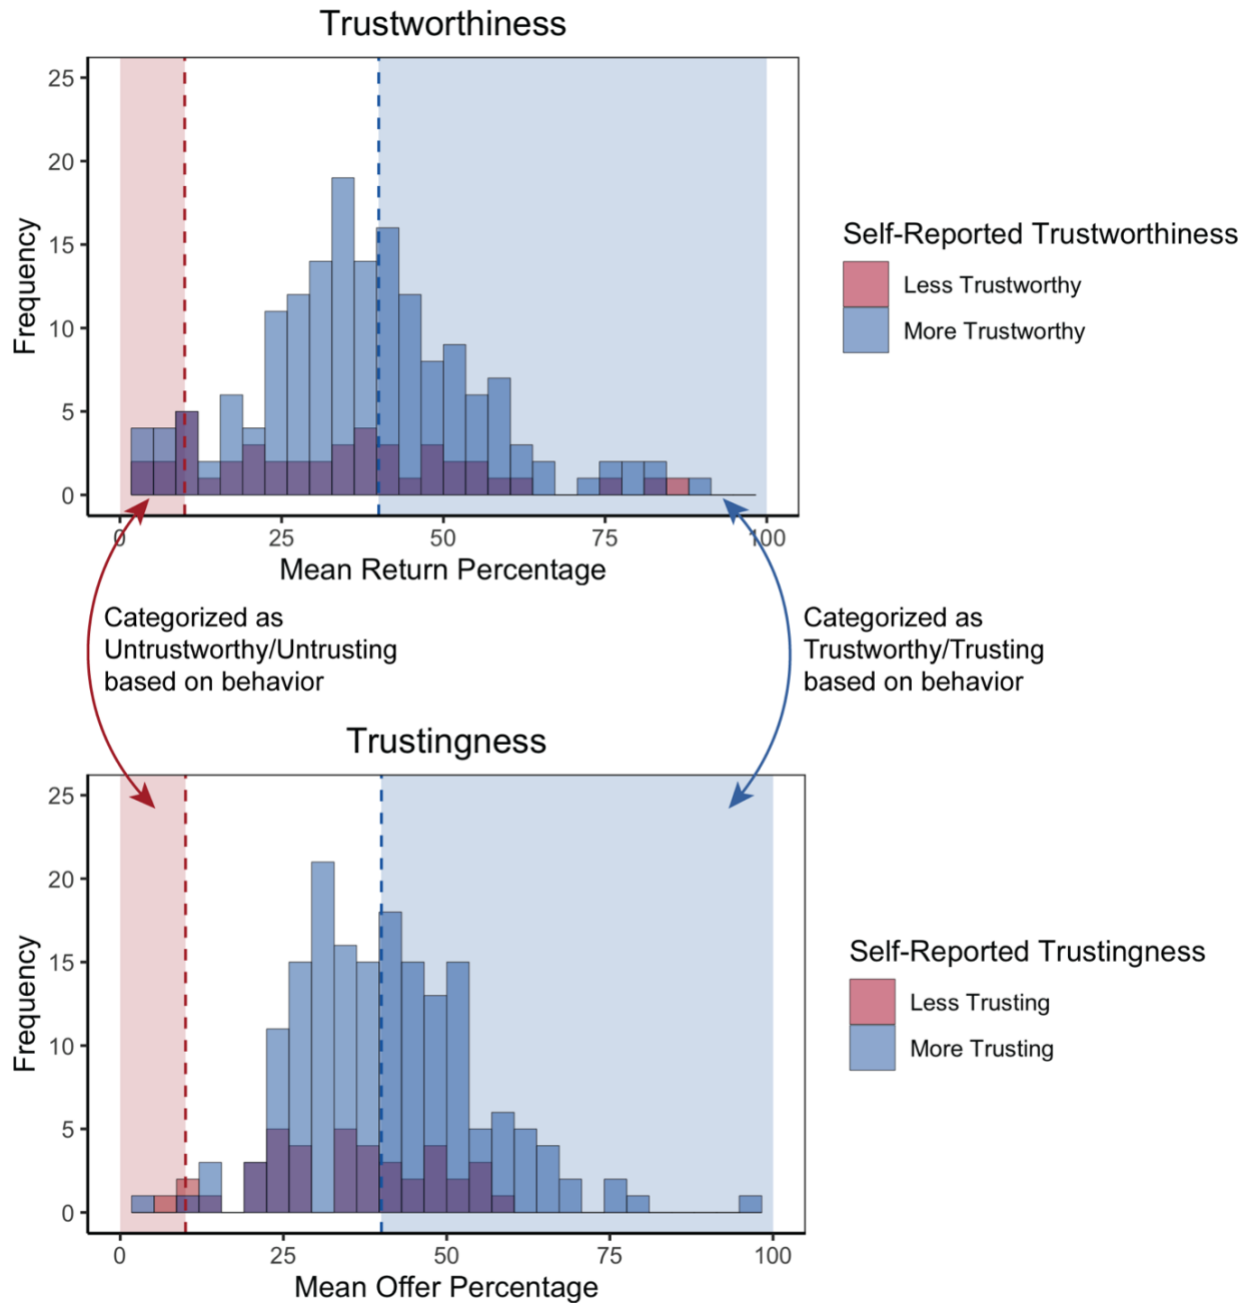
	Study 1		Study 2	
	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
Partner Behavioral Similarity				
Two Trustworthy Partners (similar)	.84	.42	.92	.32
Two Untrustworthy Partners (similar)	.73	.49	.70	.50
One Trustworthy Partner, One Untrustworthy Partner (dissimilar)	.27	.41	.19	.37
Two Trusting Partners (similar)	.87	.38	.86	.39
Two Untrusting Partners (similar)	.51	.68	.63	.52
One Trusting Partner, One Untrusting Partner (dissimilar)	.22	.39	.26	.40
Participant-Partner Behavioral Similarity				
Both Trustworthy (similar)	.90	.31	.89	.30
Both Untrustworthy (similar)	.37	.43	.18	.30
Trustworthy Participant, Untrustworthy Partner (dissimilar)	.20	.38	.17	.32
Untrustworthy Participant, Trustworthy Partner (dissimilar)	.73	.43	.88	.38
Both Trusting (similar)	.89	.32	.85	.34
Both Untrusting (similar)	.36	.48	.15	.32
Trusting Participant, Untrusting Partner (dissimilar)	.23	.38	.23	.36
Untrusting Participant, Trusting Partner (dissimilar)	.76	.48	.93	.32
Participant-Partner Self-Report Behavioral Similarity				
Both Trustworthy (similar)			.91	.28
Both Untrustworthy (similar)			.18	.37
Trustworthy Participant, Untrustworthy Partner (dissimilar)			.12	.27
Untrustworthy Participant, Trustworthy Partner (dissimilar)			.85	.38
Both Trusting (similar)			.89	.31
Both Untrusting (similar)			.21	.39
Trusting Participant, Untrusting Partner (dissimilar)			.17	.31
Untrusting Participant, Trusting Partner (dissimilar)			.85	.33

Self-Perception Accuracy

In Study 2, people reported if they thought they were one of the more trustworthy people in the game by selecting one of two buttons (“Less Trustworthy”, “More Trustworthy”). They also reported if they thought they were one of the more trusting people in the game by selecting one of two buttons (“Less Trusting”, “More Trusting”). We examined how accurate these self-perceptions were by comparing participants’ mean offers as Player 1 (their trustingness) and mean returns as Player 2 (their trustworthiness) as illustrated in Fig. S4.1. We found that people showed both a positive bias and a negative bias. That is, some people showed a positive bias such that some who thought they were trustworthy behaved in an untrustworthy manner (mean return percentage: range = 0.75-100%, $M = 37.39$) and some who thought they were trusting behaved in an untrusting manner (mean offer percentage: range = 4.90-100%, $M = 41.47$). Simultaneously, some people showed a negative bias such that some who thought they were untrustworthy behaved in a trustworthy manner (mean return percentage: range = 0.96-100%, $M = 35.17$) and some who thought they were untrusting behaved in a trusting manner (mean offer percentage: range = 7.89-59.48%, $M = 34.81$).

Figure S4.1

Self-Perception vs. Behavioral Categorization



Note. Participant's self-perception of behavioral tendencies (bars) did not perfectly align with behavioral categorization (colored background). That is, neither self-reported trustworthiness nor self-reported trustingness were accurate.

GENERAL DISCUSSION

Every interaction is embedded in the complex social systems surrounding each person. Using social network analysis to quantify this social context provides a unique opportunity to understand how the social network surrounding a perceiver, as well as the one surrounding the target of perception, shape social thought and behavior. Given the immense size and complexity of social networks, humans have developed necessary shortcuts to efficiently track and navigate their relationships. Here, I sought to uncover the neural and cognitive mechanisms through which people are able to track their social networks and how this social context then shapes their perceptions, thoughts, and decisions.

Specifically, in Chapter 1, I systematically tested the relationship between one's capacity for brokerage in one's own social network and the ability to learn new networks. I found that brokers (i.e., those who connect disparate people or groups that would not otherwise be connected to each other) were better at learning new networks than people with lower brokerage capacity, especially when the networks consist of naturally occurring network features. Framing the networks as social networks of friends versus non-social flight networks between airports did not moderate this advantage, suggesting that brokers are better at learning new networks across domains. I examined potential mechanisms through which this occurs, finding that brokers focused on existing over non-existent ties and the connections to especially important nodes, like other brokers.

I built on this work in Chapter 2 by examining the neural representations of others' positions in a social network once participants fully learned a new social network. When viewing network member's faces while undergoing fMRI, participants' brains consistently encoded others' social network centrality (i.e., the number of friends one has) in regions associated with visual

attention and mentalizing. That is, the human brain encoded information about others' importance in their social network, even when considering a network in which one is not included, and where centrality was unlinked from perceptual and experience-based features to which it is inextricably tied in naturalistic contexts.

Next, I tested how such social network information is used in interpersonal interactions in Chapter 3. Participants in this study joined a fictive online game-playing community of individuals who varied in terms of their trustworthiness and relationships with one another. I found that participants were less trusting of partners with untrustworthy friends, even after they consistently showed themselves to be trustworthy, and were less willing to engage with them in the future. Additionally, participants were exceptionally likely to falsely remember similarly behaving partners as friends. Thus, people expected friends to behave similarly and for similarly behaving people to be friends.

There are several mechanisms through which similarity among friends might arise. To test if participants assumed that friendships form through homophily (i.e., the tendency for similar people to be attracted to one another), I tested how similarity shaped participants' predictions of who would become friends in Chapter 4. I found that participants were significantly more likely to predict that partners who behaved similarly would become friends in the future compared to those who behaved dissimilarly, congruent with assumptions of homophily. Interestingly, however, participants did not predict homophily for themselves, instead consistently predicting that they would become friends with highly trustworthy and highly trusting partners compared to highly untrustworthy or highly untrusting partners. These findings suggest that participants assume homophily for others' relationships but not their own.

The information that individuals spontaneously retrieve and use when encountering someone they know illustrates what people prioritize and find most important when preparing to engage with that person. Here, I found that people consistently draw on social network information to inform their expectations and behavior. Furthermore, due to the complexity involved in representing social network information, I found that people draw on common characteristics that are observed in real-world social networks to scaffold their learning and use of this information. Together, these lines of research provide new insights into how our broader social networks shape our thoughts, behaviors, and interactions.

Future Directions

In addition to the future research suggested within each chapter, there are several directions in which I would like to expand this research. First, future research should focus on how types of relationships impact these findings. In Chapter 1, we found that participants who were drawn from the general population reported different types of relationships in their ego-network than undergraduate samples. Specifically, most of their ego-network relationships involved family ties instead of friendships. At different life stages, people form relationships in different ways. In particular, one's social ties may be primarily driven by one's own preferences and social behavioral tendencies (e.g., if one tends to introduce one's friends to each other) in contexts such as college where one has access and freedom to socialize with many different peers, whereas at other life stages, one's social ties may be shaped by opportunity and circumstances (e.g., work obligations, existing relationships among family members). Future research should investigate how different types of relationships within one's own network shape our perceptions of others' relationships. Similarly, future research should examine how people perceive these different relationships in distinct or similar ways.

In Chapter 2, we found that the human brain incorporated others' social relationships into its representation of that person. Yet in Chapter 4, we saw distinct differences in how people perceived their own and others' relationships. Thus, future work should examine how one's own relationships shape self-representations. Additionally, future research could test if the social networks of perceivers and targets interact to shape perceptions and, consequently, interactions between them. One way to approach this question is to draw on research examining the mental representations and effects of psychological distance between two people by defining psychological distance in terms of social network positions or characteristics.

Finally, future work should address the motivations behind network learning. We hypothesized in Chapter 1 that brokers' enhanced network learning ability may be linked to the opportunities their position affords. If so, people who are motivated to take advantage of these opportunities would be more likely to attend to their own networks and apply known characteristics to learning new social networks. It could also be that people are inherently motivated to learn networks and connections, as evidenced by common games such as "Six degrees of Kevin Bacon" in which people challenge each other to connect a named actor, through other actors, to Kevin Bacon through films in which they co-starred. The similar game in which people seek to identify the shortest path connecting two Wikipedia pages through hyperlinks suggests that this motivation may not be specific to social networks. Thus, future work could examine what drives people to learn networks and if this motivation changes for social and non-social connections.

Conclusion

Social networks are a ubiquitous and constant context that surround every person and interaction between them. In this dissertation, I sought to uncover how people acquire new social

network knowledge, how the brain prioritizes this information, when people use this information when interacting with others, and how people predict this information differently for themselves and others. Using novel and experimentally controlled paradigms, I decoupled social network knowledge from other factors that inevitably covary with it in real-world contexts. This work generates new questions for future research at the intersection of psychology, social network analysis, and neuroscience, and it provides novel insight into a vital and universal skillset that humans possess: the ability to understand and navigate their complex social worlds.

REFERENCES

- Abraham, A., Pedregosa, F., Eickenberg, M., Gervais, P., Mueller, A., Kossaifi, J., ... Varoquaux, G. (2014). Machine learning for neuroimaging with scikit-learn. *Frontiers in Neuroinformatics*, 8(14), 1–10. <https://doi.org/10.3389/fninf.2014.00014>
- Alt, N. P., Parkinson, C., Kleinbaum, A. M., & Johnson, K. L. (2022). The face of social networks: Naive observers' accurate assessment of others' social network positions from faces. *Social Psychological and Personality Science*, 13(1), 118–126. <https://doi.org/10.1177/19485506211003723>
- Alt, N. P., & Phillips, L. T. (2022). Person perception, meet people perception: Exploring the social vision of groups. *Perspectives on Psychological Science*, 17(3), 768–787. <https://doi.org/10.1177/17456916211017858>
- Altermatt, E. R., & Pomerantz, E. M. (2003). The development of competence-related and motivational beliefs: An investigation of similarity and influence among friends. *Journal of Educational Psychology*, 95(1), 111–123. <https://doi.org/10.1037/0022-0663.95.1.111>
- Amaral, L. A. N., Scala, A., Barthelemy, M., & Stanley, H. E. (2000). Classes of small-world networks. *Proceedings of the National Academy of Sciences*, 97(21), 11149–11152. <https://doi.org/10.1073/pnas.200327197>
- Apicella, C. L., Marlowe, F. W., Fowler, J. H., & Christakis, N. A. (2012). Social networks and cooperation in hunter-gatherers. *Nature*, 481(7382), 497–501. <https://doi.org/10.1038/nature10736>
- Augustinova, M., & Ferrand, L. (2012). The influence of mere social presence on Stroop interference: New evidence from the semantically-based Stroop task. *Journal of*

- Experimental Social Psychology*, 48(5), 1213–1216.
<https://doi.org/10.1016/j.jesp.2012.04.014>
- Avants, B. B., Epstein, C. L., Grossman, M., & Gee, J. C. (2008). Symmetric diffeomorphic image registration with cross-correlation: Evaluating automated labeling of elderly and neurodegenerative brain. *Medical Image Analysis*, 12(1), 26–41.
<https://doi.org/10.1016/j.media.2007.06.004>
- Bahns, A. J., Pickett, K. M., & Crandall, C. S. (2012). Social ecology of similarity: Big schools, small schools and social relationships. *Group Processes and Intergroup Relations*, 15(1), 119–131. <https://doi.org/10.1177/1368430211410751>
- Barabási, A.-L., & Bonabeau, E. (2003). Scale-free networks. *Scientific American*, 288(5), 60–69.
Retrieved from <http://www.jstor.org/stable/26060284>
- Baron-Cohen, S., Wheelwright, S., Skinner, R., Martin, J., & Clubley, E. (2001). The Autism-Spectrum Quotient (AQ): Evidence from Asperger Syndrome/high-functioning autism, males and females, scientists and mathematicians. *Journal of Autism and Developmental Disorders*, 31(1), 5–17. <https://doi.org/10.1023/A:1005653411471>
- Basyouni, R., & Parkinson, C. (2022). Mapping the social landscape: Tracking patterns of interpersonal relationships. *Trends in Cognitive Sciences*, 26(3), 204–221.
<https://doi.org/10.1016/j.tics.2021.12.006>
- Bates, D., Maechler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1), 1–48.
<https://doi.org/doi:10.18637/jss.v067.i01>

- Baumard, N., André, J.-B., & Sperber, D. (2013). A mutualistic approach to morality: The evolution of fairness by partner choice. *Behavioral and Brain Sciences*, *36*(1), 59–78. <https://doi.org/10.1017/S0140525X11002202>
- Benjamini, Y., & Hochberg, Y. (1995). Controlling the False Discovery Rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Methodological)*, *57*(1), 289–300. <https://doi.org/10.1111/j.2517-6161.1995.tb02031.x>
- Berg, J., Dickhaut, J., & McCabe, K. (1995). Trust, reciprocity, and social history. *Games and Economic Behavior*, *10*(1), 122–142. <https://doi.org/10.1006/game.1995.1027>
- Brands, R. A. (2013). Cognitive social structures in social network research: A review. *Journal of Organizational Behavior*, *34*(S1), S82–S103. <https://doi.org/10.1002/job.1890>
- Brashears, M. E. (2013). Humans use compression heuristics to improve the recall of social networks. *Scientific Reports*, *3*(1). <https://doi.org/10.1038/srep01513>
- Brechwald, W. A., & Prinstein, M. J. (2011). Beyond homophily: A decade of advances in understanding peer influence processes. *Journal of Research on Adolescence*, *21*(1), 166–179. <https://doi.org/10.1111/j.1532-7795.2010.00721.x>
- Brent, L. J. N. (2015). Friends of friends: Are indirect connections in social networks important to animal behaviour? *Animal Behaviour*, *103*, 211–222. <https://doi.org/10.1016/j.anbehav.2015.01.020>
- Brewer, W. F., & Treyens, J. C. (1981). Role of schemata in memory for places. *Cognitive Psychology*, *13*(2), 207–230. [https://doi.org/10.1016/0010-0285\(81\)90008-6](https://doi.org/10.1016/0010-0285(81)90008-6)
- Burt, R. S. (2021). Structural Holes: Capstone, Cautions, and Enthusiasms. In M. L. Small & B. L. Perry (Eds.), *Personal Networks: Classic Readings and New Directions in Egocentric*

- Analysis* (pp. 371–416). Cambridge University Press. Retrieved from <https://books.google.com/books?id=K3tGEAAAQBAJ>
- Burt, R. S., Kilduff, M., & Tasselli, S. (2013). Social Network Analysis: Foundations and frontiers on advantage. *Annual Review of Psychology*, *64*(1), 527–547. <https://doi.org/10.1146/annurev-psych-113011-143828>
- Cao, R., Li, X., Todorov, A., & Wang, S. (2020). A flexible neural representation of faces in the human brain. *Cerebral Cortex Communications*, *1*(1), 1–12. <https://doi.org/10.1093/texcom/tgaa055>
- Centola, D. (2011). An experimental study of homophily in the adoption of health behavior. *Science*, *334*(6060), 1269–1272. <https://doi.org/10.1126/science.1207055>
- Champely, S. (2018). *pwr: Basic functions for power analysis*. Retrieved from <https://cran.r-project.org/web/packages/pwr/pwr.pdf>
- Chang, L. J., Doll, B. B., van 't Wout, M., Frank, M. J., & Sanfey, A. G. (2010). Seeing is believing: Trustworthiness as a dynamic belief. *Cognitive Psychology*, *61*(2), 87–105. <https://doi.org/10.1016/j.cogpsych.2010.03.001>
- Cheney, D. L., & Seyfarth, R. M. (1990). The representation of social relations by monkeys. *Cognition*, *37*(1–2), 167–196. [https://doi.org/10.1016/0010-0277\(90\)90022-C](https://doi.org/10.1016/0010-0277(90)90022-C)
- Chi, L. P., Wang, R., Su, H., Xu, X. P., Zhao, J. S., Li, W., & Cai, X. (2003). Structural properties of US flight network. *Chinese Physics Letters*, *20*(8), 1393–1396. <https://doi.org/10.1088/0256-307X/20/8/362>
- Chiang, Y.-S. (2010). Self-interested partner selection can lead to the emergence of fairness. *Evolution and Human Behavior*, *31*(4), 265–270. <https://doi.org/10.1016/j.evolhumbehav.2010.03.003>

- Chiang, Y.-S., & Takahashi, N. (2011). Network homophily and the evolution of the pay-it-forward reciprocity. *PLoS ONE*, 6(12), e29188. <https://doi.org/10.1371/journal.pone.0029188>
- Christakis, N. A., & Fowler, J. H. (2009). *Connected: The surprising power of our social networks and how they shape our lives*. New York: Little, Brown and Company. <https://doi.org/10.1111/j.1756-2589.2011.00097.x>
- Christakis, N. A., & Fowler, J. H. (2013). Social contagion theory: Examining dynamic social networks and human behavior. *Statistics in Medicine*, 32(4), 556–577. <https://doi.org/10.1002/sim.5408>
- Cohen, S., Doyle, W. J., Skoner, D. P., Rabin, B. S., & Gwaltney, J. M. (1997). Social ties and susceptibility to the common cold. *Journal of the American Medical Association*, 277(24), 1940–1944. <https://doi.org/10.1001/jama.277.24.1940>
- Cole, S. W., Hawkley, L. C., Arevalo, J. M., Sung, C. Y., Rose, R. M., & Cacioppo, J. T. (2007). Social regulation of gene expression in human leukocytes. *Genome Biology*, 8(9), R189. <https://doi.org/10.1186/gb-2007-8-9-r189>
- Collins, D. L., Zijdenbos, A. P., Baaré, W. F. C., & Evans, A. C. (1999). ANIMAL+INSECT: Improved cortical structure segmentation. In *Lecture Notes in Computer Science* (Vol. 1716, pp. 210–223). https://doi.org/10.1007/3-540-48714-X_16
- Corbetta, M., & Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, 3(3), 201–215. <https://doi.org/10.1038/nrn755>
- Cox, R. W. (1996). AFNI: Software for analysis and visualization of functional magnetic resonance neuroimages. *Computers and Biomedical Research*, 29(3), 162–173. <https://doi.org/10.1006/cbmr.1996.0014>

- Cox, R. W., & Hyde, J. S. (1997). Software tools for analysis and visualization of fMRI data. *NMR in Biomedicine*, 10(4–5), 171–178. [https://doi.org/10.1002/\(SICI\)1099-1492\(199706/08\)10:4/5<171::AID-NBM453>3.0.CO;2-L](https://doi.org/10.1002/(SICI)1099-1492(199706/08)10:4/5<171::AID-NBM453>3.0.CO;2-L)
- Currarini, S., Jackson, M. O., & Pin, P. (2009). An economic model of friendship: Homophily, minorities, and segregation. *Econometrica*, 77(4), 1003–1045. <https://doi.org/10.3982/ECTA7528>
- Dale, A. M., Fischl, B., & Sereno, M. I. (1999). Cortical surface-based analysis. *NeuroImage*, 9(2), 179–194. <https://doi.org/10.1006/nimg.1998.0395>
- De Brigard, F., Brady, T. F., Ruzic, L., & Schacter, D. L. (2017). Tracking the emergence of memories: A category-learning paradigm to explore schema-driven recognition. *Memory and Cognition*, 45(1), 105–120. <https://doi.org/10.3758/s13421-016-0643-6>
- de Klepper, M., Sleebos, E., van de Bunt, G., & Agneessens, F. (2010). Similarity in friendship networks: Selection or influence? The effect of constraining contexts and non-visible individual attributes. *Social Networks*, 32(1), 82–90. <https://doi.org/10.1016/j.socnet.2009.06.003>
- De Soto, C. B. (1960). Learning a social structure. *The Journal of Abnormal and Social Psychology*, 60(3), 417–421. <https://doi.org/10.1037/h0047511>
- De Soto, C. B., Henley, N. M., & London, M. (1968). Balance and the grouping schema. *Journal of Personality and Social Psychology*, 8(1, Pt.1), 1–7. <https://doi.org/10.1037/h0025321>
- Dehghani, M., Johnson, K., Hoover, J., Sagi, E., Garten, J., Parmar, N. J., ... Graham, J. (2016). Purity homophily in social networks. *Journal of Experimental Psychology: General*, 145(3), 366–375. <https://doi.org/10.1037/xge0000139>

- Delgado, M. R., Frank, R. H., & Phelps, E. A. (2005). Perceptions of moral character modulate the neural systems of reward during the trust game. *Nature Neuroscience*, 8(11), 1611–1618. <https://doi.org/10.1038/nn1575>
- Du, M., Basyouni, R., & Parkinson, C. (2021). How does the brain navigate knowledge of social relations? Testing for shared neural mechanisms for shifting attention in space and social knowledge. *NeuroImage*, 118019. <https://doi.org/10.1016/j.neuroimage.2021.118019>
- Dunbar, R. I. M. (2014). The social brain: Psychological underpinnings and implications for the structure of organizations. *Current Directions in Psychological Science*, 23(2), 109–114. <https://doi.org/10.1177/0963721413517118>
- Dunbar, R. I. M., & Shultz, S. (2007). Evolution in the social brain. *Science*, 317(5843), 1344–1347. <https://doi.org/10.1126/science.1145463>
- Ehlert, A., Kindschi, M., Algesheimer, R., & Rauhut, H. (2020). Human social preferences cluster and spread in the field. *Proceedings of the National Academy of Sciences*, 117(37), 22787–22792. <https://doi.org/10.1073/pnas.2000824117>
- Eisenberger, N. I. (2013). Social ties and health: A social neuroscience perspective. *Current Opinion in Neurobiology*, 23(3), 407–413. <https://doi.org/10.1016/j.conb.2013.01.006>
- Engell, A. D., Haxby, J. V., & Todorov, A. (2007). Implicit trustworthiness decisions: Automatic coding of face properties in the human amygdala. *Journal of Cognitive Neuroscience*, 19(9), 1508–1519. <https://doi.org/10.1162/jocn.2007.19.9.1508>
- Esteban, O., Markiewicz, C. J., Blair, R. W., Moodie, C. A., Isik, A. I., Erramuzpe, A., ... Gorgolewski, K. J. (2019). fMRIPrep: A robust preprocessing pipeline for functional MRI. *Nature Methods*, 16(1), 111–116. <https://doi.org/10.1038/s41592-018-0235-4>

- Fareri, D. S., Chang, L. J., & Delgado, M. R. (2012). Effects of direct social experience on trust decisions and neural reward circuitry. *Frontiers in Neuroscience*, 6, 148. <https://doi.org/10.3389/fnins.2012.00148>
- Fareri, D. S., Chang, L. J., & Delgado, M. R. (2015). Computational substrates of social value in interpersonal collaboration. *Journal of Neuroscience*, 35(21), 8170–8180. <https://doi.org/10.1523/jneurosci.4775-14.2015>
- Feinberg, M., Willer, R., & Schultz, M. (2014). Gossip and ostracism promote cooperation in groups. *Psychological Science*, 25(3), 656–664. <https://doi.org/10.1177/0956797613510184>
- Flynn, F. J., Reagans, R. E., & Guillory, L. (2010). Do you two know each other? Transitivity, homophily, and the need for (network) closure. *Journal of Personality and Social Psychology*, 99(5), 855–869. <https://doi.org/10.1037/a0020961>
- Fonov, V., Evans, A. C., Botteron, K., Almli, C. R., McKinstry, R. C., & Collins, D. L. (2011). Unbiased average age-appropriate atlases for pediatric studies. *NeuroImage*, 54(1), 313–327. <https://doi.org/10.1016/j.neuroimage.2010.07.033>
- Fouragnan, E., Chierchia, G., Greiner, S., Neveu, R., Avesani, P., & Coricelli, G. (2013). Reputational priors magnify striatal responses to violations of trust. *Journal of Neuroscience*, 33(8), 3602–3611. <https://doi.org/10.1523/jneurosci.3086-12.2013>
- Fowler, J. H., & Christakis, N. A. (2010). Cooperative behavior cascades in human social networks. *Proceedings of the National Academy of Sciences*, 107(12), 5334–5338. <https://doi.org/10.1073/pnas.0913149107>
- Freeman, M. F., & Tukey, J. W. (1950). Transformations related to the angular and the square root. *The Annals of Mathematical Statistics*, 21(4), 607–611. <https://doi.org/10.1214/aoms/1177729756>

- Gazes, R. P., Hampton, R. R., & Lourenco, S. F. (2017). Transitive inference of social dominance by human infants. *Developmental Science*, *20*(2), 1–10. <https://doi.org/10.1111/desc.12367>
- Gobbini, M. I., & Haxby, J. V. (2007). Neural systems for recognition of familiar faces. *Neuropsychologia*, *45*(1), 32–41. <https://doi.org/10.1016/j.neuropsychologia.2006.04.015>
- Goel, S., Mason, W., & Watts, D. J. (2010). Real and perceived attitude agreement in social networks. *Journal of Personality and Social Psychology*, *99*(4), 611–621. <https://doi.org/10.1037/a0020697>
- Gonçalves, B., Perra, N., & Vespignani, A. (2011). Modeling users' activity on twitter networks: Validation of Dunbar's number. *PLoS ONE*, *6*(8). <https://doi.org/10.1371/journal.pone.0022656>
- Goodale, B. M., Alt, N. P., Lick, D. J., & Johnson, K. L. (2018). Groups at a glance: Perceivers infer social belonging in a group based on perceptual summaries of sex ratio. *Journal of Experimental Psychology: General*, *147*(11), 1660–1676. <https://doi.org/10.1037/xge0000450>
- Gorgolewski, K. J., Esteban, O., Markiewicz, C. J., Ziegler, E., Ellis, D. G., Jarecka, D., ... Ghosh, S. (2019, May 9). *nipy/nipype: 1.2.0*. <https://doi.org/10.5281/ZENODO.2685428>
- Hackel, L. M., Doll, B. B., & Amodio, D. M. (2015). Instrumental learning of traits versus rewards: Dissociable neural correlates and effects on choice. *Nature Neuroscience*, *18*(9), 1233–1235. <https://doi.org/10.1038/nn.4080>
- Hahl, O., Kacperczyk, A. “Olenka,” & Davis, J. P. (2016). Knowledge asymmetry and brokerage: Linking network perception to position in structural holes. *Strategic Organization*, *14*(2), 118–143. <https://doi.org/10.1177/1476127015624274>

- Hassabis, D., Spreng, R. N., Rusu, A. A., Robbins, C. A., Mar, R. A., & Schacter, D. L. (2014). Imagine all the people: How the brain creates and uses personality models to predict behavior. *Cerebral Cortex*, 24(8), 1979–1987. <https://doi.org/10.1093/cercor/bht042>
- Heider, F. (1946). Attitudes and cognitive organization. *Journal of Psychology*, 21(1), 107–112. <https://doi.org/10.1080/00223980.1946.9917275>
- Heider, F. (1958). *The Psychology of Interpersonal Relations*. Hillsdale, New Jersey: Lawrence Erlbaum Associates, Inc.
- Henderson, M., & Furnham, A. (1982). Similarity and attraction: the relationship between personality, beliefs, skills, needs and friendship choice. *Journal of Adolescence*, 5(2), 111–123. [https://doi.org/10.1016/S0140-1971\(82\)80040-7](https://doi.org/10.1016/S0140-1971(82)80040-7)
- Henrich, J., Heine, S. J., & Norenzayan, A. (2010). The weirdest people in the world? *Behavioral and Brain Sciences*, 33(2/3), 61–83. <https://doi.org/10.1017/S0140525X0999152X>
- Hill, R. A., & Dunbar, R. I. M. (2003). Social network size in humans. *Human Nature*, 14(1), 53–72. <https://doi.org/10.1007/s12110-003-1016-y>
- Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics*, 6(2), 65–70.
- Izquierdo, L. R., Izquierdo, S. S., & Vega-Redondo, F. (2014). Leave and let leave: A sufficient condition to explain the evolutionary emergence of cooperation. *Journal of Economic Dynamics and Control*, 46, 91–113. <https://doi.org/10.1016/j.jedc.2014.06.007>
- Janicik, G. A., & Larrick, R. P. (2005). Social network schemas and the learning of incomplete networks. *Journal of Personality and Social Psychology*, 88(2), 348–364. <https://doi.org/10.1037/0022-3514.88.2.348>

- Jenkinson, M., Bannister, P., Brady, M., & Smith, S. (2002). Improved optimization for the robust and accurate linear registration and motion correction of brain images. *NeuroImage*, 17(2), 825–841. <https://doi.org/10.1006/nimg.2002.1132>
- Johnson, K. L., Lick, D. J., & Carpinella, C. M. (2015). Emergent research in social vision: An integrated approach to the determinants and consequences of social categorization. *Social and Personality Psychology Compass*, 9(1), 15–30. <https://doi.org/10.1111/spc3.12147>
- Jolly, E., & Chang, L. J. (2018). Gossip drives vicarious learning and facilitates robust social connections. *PsyArXiv*. <https://doi.org/10.31234/osf.io/qau5s>
- Kahn, A. E., Karuza, E. A., Vettel, J. M., & Bassett, D. S. (2018). Network constraints on learnability of probabilistic motor sequences. *Nature Human Behaviour*, 2(12), 936–947. <https://doi.org/10.1038/s41562-018-0463-8>
- Kastner, S., & Ungerleider, L. G. (2000). Mechanisms of visual attention in the human cortex. *Annual Review of Neuroscience*, 23(1), 315–341. <https://doi.org/10.1146/annurev.neuro.23.1.315>
- Kawachi, I. (2001). Social ties and mental health. *Journal of Urban Health: Bulletin of the New York Academy of Medicine*, 78(3), 458–467. <https://doi.org/10.1093/jurban/78.3.458>
- Kilduff, M., Crossland, C., Tsai, W., & Krackhardt, D. (2008). Organizational network perceptions versus reality: A small world after all? *Organizational Behavior and Human Decision Processes*, 107(1), 15–28. <https://doi.org/10.1016/j.obhdp.2007.12.003>
- Krackhardt, D. (1987). Cognitive Social Structures. *Social Networks*, Vol. 9, pp. 109–134.
- Krasnow, M. M., Cosmides, L., Pedersen, E. J., & Tooby, J. (2012). What are punishment and reputation for? *PLoS ONE*, 7(9), 1–9. <https://doi.org/10.1371/journal.pone.0045662>

- Kriegeskorte, N., Mur, M., & Bandettini, P. (2008). Representational similarity analysis – connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience*, 2(4), 1–28. <https://doi.org/10.3389/neuro.06.004.2008>
- Kuwabara, K., Cao, J., Cho, S. S., & Ingram, P. (2022). Lay theories of instrumental relations: Explaining individual differences in dispositional similarity-attraction. *Academy of Management Journal*. <https://doi.org/10.5465/amj.2020.0770>
- Kwon, S. W., Rondi, E., Levin, D. Z., De Massis, A., & Brass, D. J. (2020). Network brokerage: An integrative review and future research agenda. *Journal of Management*, 46(6), 1092–1120. <https://doi.org/10.1177/0149206320914694>
- Lampinen, J. M., Copeland, S. M., & Neuschatz, J. S. (2001). Recollections of things schematic: Room schemas revisited. *Journal of Experimental Psychology: Learning Memory and Cognition*, 27(5), 1211–1222. <https://doi.org/10.1037/0278-7393.27.5.1211>
- Lau, T., Gershman, S. J., & Cikara, M. (2020). Social structure learning in human anterior insula. *ELife*, 9, 1–17. <https://doi.org/10.7554/eLife.53162>
- Lau, T., Pouncy, H. T., Gershman, S. J., & Cikara, M. (2018). Discovering social groups via latent structure learning. *Journal of Experimental Psychology: General*, 147(12), 1881–1891. <https://doi.org/10.1037/xge0000470>
- Lawrence, B. S., & Shah, N. P. (2020). Homophily: Measures and meaning. *Academy of Management Annals*, 14(2), 513–597. <https://doi.org/10.5465/annals.2018.0147>
- Lazarsfeld, P. F., & Merton, R. K. (1954). Friendship as a social process: A substantive and methodological analysis. *Freedom and Control in Modern Society*, 18(1), 18–66.
- Lenth, R. (2019). *emmeans: Estimated marginal means, aka least-squares means*. Retrieved from <https://cran.r-project.org/web/packages/emmeans/emmeans.pdf>

- Liberman, N., & Trope, Y. (2014). Traversing psychological distance. *Trends in Cognitive Sciences*, *18*(7), 364–369. <https://doi.org/10.1016/j.tics.2014.03.001>
- Liberman, Z., Kinzler, K. D., & Woodward, A. L. (2021). Origins of homophily: Infants expect people with shared preferences to affiliate. *Cognition*, *212*(March 2020), 104695. <https://doi.org/10.1016/j.cognition.2021.104695>
- Lynn, C. W., & Bassett, D. S. (2020). How humans learn and represent networks. *Proceedings of the National Academy of Sciences*, *117*(47), 29407–29415. <https://doi.org/10.1073/pnas.1912328117>
- Lynn, C. W., Kahn, A. E., Nyema, N., & Bassett, D. S. (2020). Abstract representations of events arise from mental errors in learning and memory. *Nature Communications*, *11*(1). <https://doi.org/10.1038/s41467-020-15146-7>
- Lynn, C. W., Papadopoulos, L., Kahn, A. E., & Bassett, D. S. (2020). Human information processing in complex networks. *Nature Physics*, *16*(9), 965–973. <https://doi.org/10.1038/s41567-020-0924-7>
- Ma, D., Correll, J., & Wittenbrink, B. (2015). The Chicago Face Database: A free stimulus set of faces and norming data. *Behavior Research Methods*, *47*, 1122–1135. <https://doi.org/https://doi.org/10.3758/s13428-014-0532-5>
- Marek, S., Tervo-Clemmens, B., Calabro, F. J., Montez, D. F., Kay, B. P., Hatoum, A. S., ... Dosenbach, N. U. F. (2022). Reproducible brain-wide association studies require thousands of individuals. *Nature*, *603*(7902), 654–660. <https://doi.org/10.1038/s41586-022-04492-9>
- Marineau, J. E., Labianca, G. (Joe), Brass, D. J., Borgatti, S. P., & Vecchi, P. (2018). Individuals' power and their social network accuracy: A situated cognition perspective. *Social Networks*, *54*, 145–161. <https://doi.org/10.1016/j.socnet.2018.01.006>

- Mars, R. B., Neubert, F. X., Noonan, M. A. P., Sallet, J., Toni, I., & Rushworth, M. F. S. (2012). On the relationship between the “default mode network” and the “social brain.” *Frontiers in Human Neuroscience*, *6*(189), 1–9. <https://doi.org/10.3389/fnhum.2012.00189>
- Martinez, J. E., Mack, M. L., Gelman, B. D., & Preston, A. R. (2016). Knowledge of social affiliations biases economic decisions. *PLoS ONE*, *11*(7), 1–15. <https://doi.org/10.1371/journal.pone.0159918>
- McPherson, M., Smith-Lovin, L., & Cook, J. M. (2001). Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, *27*(1), 415–444. <https://doi.org/10.1146/annurev.soc.27.1.415>
- Meyer, M. L., Davachi, L., Ochsner, K. N., & Lieberman, M. D. (2019). Evidence that default network connectivity during rest consolidates social information. *Cerebral Cortex*, *29*(5), 1910–1920. <https://doi.org/10.1093/cercor/bhy071>
- Meyer, M. L., & Lieberman, M. D. (2016). Social sorking memory training improves perspective-taking accuracy. *Social Psychological and Personality Science*, *7*(4), 381–389. <https://doi.org/10.1177/1948550615624143>
- Morelli, S. A., Leong, Y. C., Carlson, R. W., Kullar, M., & Zaki, J. (2018). Neural detection of socially valued community members. *Proceedings of the National Academy of Sciences*, *115*(32), 8149–8154. <https://doi.org/10.1073/pnas.1712811115>
- Noë, R., & Hammerstein, P. (1995). Biological markets. *Trends in Ecology & Evolution*, *10*(8), 336–339. [https://doi.org/10.1016/S0169-5347\(00\)89123-5](https://doi.org/10.1016/S0169-5347(00)89123-5)
- Norbeck, J. (1984). The Norbeck Social Support Questionnaire. *Birth Defects Original Article Series*, *20*(5), 45–57. Retrieved from <https://www.ncbi.nlm.nih.gov/pubmed/6536338>

- Oh, H., & Kilduff, M. (2008). The ripple effect of personality on social structure: Self-monitoring origins of network brokerage. *Journal of Applied Psychology, 93*(5), 1155–1164. <https://doi.org/10.1037/0021-9010.93.5.1155>
- Paluck, E. L., & Shepherd, H. (2012). The salience of social referents: A field experiment on collective norms and harassment behavior in a school social network. *Journal of Personality and Social Psychology, 103*(6), 899–915. <https://doi.org/10.1037/a0030015>
- Parkinson, C., & Du, M. (2020). How does the brain infer hidden social structures? *Trends in Cognitive Sciences, 24*(7), 497–498. <https://doi.org/10.1016/j.tics.2020.05.002>
- Parkinson, C., Kleinbaum, A. M., & Wheatley, T. (2017). Spontaneous neural encoding of social network position. *Nature Human Behaviour, 1*(5), 0072. <https://doi.org/10.1038/s41562-017-0072>
- Parkinson, C., Kleinbaum, A. M., & Wheatley, T. (2018). Similar neural responses predict friendship. *Nature Communications, 9*(1), 332. <https://doi.org/10.1038/s41467-017-02722-7>
- Parkinson, C., Liu, S., & Wheatley, T. (2014). A common cortical metric for spatial, temporal, and social distance. *Journal of Neuroscience, 34*(5), 1979–1987. <https://doi.org/10.1523/JNEUROSCI.2159-13.2014>
- Parkinson, C., & Wheatley, T. (2013). Old cortex, new contexts: Re-purposing spatial perception for social cognition. *Frontiers in Human Neuroscience, 7*. <https://doi.org/10.3389/fnhum.2013.00645>
- Parkinson, C., & Wheatley, T. (2015). The repurposed social brain. *Trends in Cognitive Sciences, 19*(3), 133–141. <https://doi.org/10.1016/j.tics.2015.01.003>

- Paz-y-Miño C, G., Bond, A. B., Kamil, A. C., & Balda, R. P. (2004). Pinyon jays use transitive inference to predict social dominance. *Nature*, 430(7001), 778–781. <https://doi.org/10.1038/nature02723>
- Peer, M., Hayman, M., Tamir, B., & Arzy, S. (2021). Brain coding of social network structure. *The Journal of Neuroscience*, 41(22), 4897–4909. <https://doi.org/10.1523/JNEUROSCI.2641-20.2021>
- Peirce, J. W. (2009). Generating stimuli for neuroscience using PsychoPy. *Frontiers in Neuroinformatics*, 2. <https://doi.org/10.3389/neuro.11.010.2008>
- Pinheiro, J. C., Bates, D., DebRoy, S., Sarkar, D., & R Core Team. (2019). *nlme: Linear and Nonlinear Mixed Effects Models*.
- Poerio, G. L., Sormaz, M., Wang, H. T., Margulies, D., Jefferies, E., & Smallwood, J. (2017). The role of the default mode network in component processes underlying the wandering mind. *Social Cognitive and Affective Neuroscience*, 12(7), 1047–1062. <https://doi.org/10.1093/scan/nsx041>
- Pradel, J., Euler, H. A., & Fetchenhauer, D. (2009). Spotting altruistic dictator game players and mingling with them: The elective assortment of classmates. *Evolution and Human Behavior*, 30(2), 103–113. <https://doi.org/10.1016/j.evolhumbehav.2008.09.003>
- Price, M. E. (2006). Monitoring, reputation, and “greenbeard” reciprocity in a Shuar work team. *Journal of Organizational Behavior*, 27(2), 201–219. <https://doi.org/10.1002/job.347>
- R Core Team. (2022). *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from <https://www.r-project.org/>
- Rad, M. S., Martingano, A. J., & Ginges, J. (2018). Toward a psychology of Homo sapiens: Making psychological science more representative of the human population. *Proceedings of*

- the National Academy of Sciences of the United States of America*, 115(45), 11401–11405.
<https://doi.org/10.1073/pnas.1721165115>
- Rim, S. Y., Uleman, J. S., & Trope, Y. (2009). Spontaneous trait inference and construal level theory: Psychological distance increases nonconscious trait thinking. *Journal of Experimental Social Psychology*, 45(5), 1088–1097. <https://doi.org/10.1016/j.jesp.2009.06.015>
- Risko, E. F., & Kingstone, A. (2011). Eyes wide shut: Implied social presence, eye tracking and attention. *Attention, Perception, and Psychophysics*, 73(2), 291–296.
<https://doi.org/10.3758/s13414-010-0042-1>
- Roberts, S. G. B., Dunbar, R. I. M., Pollet, T. V., & Kuppens, T. (2009). Exploring variation in active network size: Constraints and ego characteristics. *Social Networks*, 31(2), 138–146.
<https://doi.org/10.1016/j.socnet.2008.12.002>
- Rusch, T., & Charpentier, C. J. (2021). Domain specificity versus process specificity: The “social brain” during strategic interaction. *Neuron*, 109(20), 3236–3238.
<https://doi.org/10.1016/j.neuron.2021.09.035>
- Sahi*, R. S., Schwyck*, M. E., Parkinson, C., & Eisenberger, N. I. (2021). Having more virtual interaction partners during COVID-19 physical distancing measures may benefit mental health. *Scientific Reports*, 11(1), 1–9. <https://doi.org/10.1038/s41598-021-97421-1>
- Satow, K. L. (1975). Social approval and helping. *Journal of Experimental Social Psychology*, 11(6), 501–509. [https://doi.org/10.1016/0022-1031\(75\)90001-3](https://doi.org/10.1016/0022-1031(75)90001-3)
- Schaefer, A., Kong, R., Gordon, E. M., Laumann, T. O., Zuo, X.-N., Holmes, A. J., ... Yeo, B. T. T. (2018). Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity MRI. *Cerebral Cortex*, 28(9), 3095–3114.
<https://doi.org/10.1093/cercor/bhx179>

- Schelling, T. (1978). Micromotives and macrobehavior. In *Norton*.
- Schwyck*, M. E., Du*, M., Li, Y., Chang, L. J., & Parkinson, C. (2023). Similarity among friends serves as a social prior: The assumption that “Birds of a feather flock together” shapes social decisions and relationship beliefs. *Personality and Social Psychology Bulletin*, (014616722211402), 1–18. <https://doi.org/10.1177/01461672221140269>
- Shrum, W., Cheek, N. H., & Hunter, S. M. (1988). Friendship in school: Gender and racial homophily. *Sociology of Education*, 61(4), 227. <https://doi.org/10.2307/2112441>
- Simpson, B., Brashears, M., Gladstone, E., & Harrell, A. (2014). Birds of different feathers cooperate together: No evidence for altruism homophily in networks. *Sociological Science*, 1(December), 542–564. <https://doi.org/10.15195/v1.a30>
- Simpson, B., Markovsky, B., & Steketee, M. (2011). Network knowledge and the use of power. *Social Networks*, 33(2), 172–176. <https://doi.org/10.1016/j.socnet.2010.10.008>
- Smith, E. R., & Collins, E. C. (2009). Contextualizing person perception: Distributed social cognition. *Psychological Review*, 116(2), 343–364. <https://doi.org/10.1037/a0015072>
- Smith, J. A., McPherson, M., & Smith-Lovin, L. (2014). Social distance in the United States: Sex, race, religion, age, and education homophily among confidants, 1985 to 2004. *American Sociological Review*, 79(3), 432–456. <https://doi.org/10.1177/0003122414531776>
- Smith, K. M., Larroucau, T., Mabulla, I. A., & Apicella, C. L. (2018). Hunter-gatherers maintain assortativity in cooperation despite high levels of residential change and mixing. *Current Biology*, 28(19), 3152-3157.e4. <https://doi.org/10.1016/j.cub.2018.07.064>
- Smith, S. M., Jenkinson, M., Woolrich, M. W., Beckmann, C. F., Behrens, T. E. J., Johansen-Berg, H., ... Matthews, P. M. (2004). Advances in functional and structural MR image analysis and

- implementation as FSL. *NeuroImage*, 23, S208–S219.
<https://doi.org/10.1016/j.neuroimage.2004.07.051>
- Smith, S. M., & Nichols, T. E. (2009). Threshold-free cluster enhancement: Addressing problems of smoothing, threshold dependence and localisation in cluster inference. *NeuroImage*, 44(1), 83–98. <https://doi.org/10.1016/j.neuroimage.2008.03.061>
- Snyder, M. (1974). Self-monitoring of expressive behavior. *Journal of Personality and Social Psychology*, 30(4), 526–537. <https://doi.org/10.1037/h0037039>
- Son, J. Y., Bhandari, A., & FeldmanHall, O. (2021). Cognitive maps of social features enable flexible inference in social networks. *Proceedings of the National Academy of Sciences of the United States of America*, 118(39), 1–11. <https://doi.org/10.1073/pnas.2021699118>
- Spreng, R. N., Mar, R. A., & Kim, A. S. N. (2009). The common neural basis of autobiographical memory, prospection, navigation, theory of mind, and the default mode: A quantitative meta-analysis. *Journal of Cognitive Neuroscience*, 21(3), 489–510.
<https://doi.org/10.1162/jocn.2008.21029>
- Stanley, D. A. (2016). Getting to know you: General and specific neural computations for learning about people. *Social Cognitive and Affective Neuroscience*, 11(4), 525–536.
<https://doi.org/10.1093/scan/nsv145>
- Stanley, D. A., Sokol-Hessner, P., Banaji, M. R., & Phelps, E. A. (2011). Implicit race attitudes predict trustworthiness judgments and economic trust decisions. *Proceedings of the National Academy of Sciences*, 108(19), 7710–7715. <https://doi.org/10.1073/pnas.1014345108>
- Su, Y., Luo, Q., Tan, S., & Qu, C. (2021). The neural signature of social dominance discrimination by means of fast periodic visual stimulation. *Neuroscience*, 459, 39–49.
<https://doi.org/10.1016/j.neuroscience.2021.01.032>

- Takahashi, N., & Mashima, R. (2006). The importance of subjectivity in perceptual errors on the emergence of indirect reciprocity. *Journal of Theoretical Biology*, 243(3), 418–436. <https://doi.org/10.1016/j.jtbi.2006.05.014>
- Tamir, D. I., Thornton, M. A., Contreras, J. M., & Mitchell, J. P. (2016). Neural evidence that three dimensions organize mental state representation: Rationality, social impact, and valence. *Proceedings of the National Academy of Sciences*, 113(1), 194–199. <https://doi.org/10.1073/pnas.1511905112>
- Thornton, M. A., Weaverdyck, M. E., Mildner, J. N., & Tamir, D. I. (2019). People represent their own mental states more distinctly than others'. *Nature Communications*, 10(2117), 1–9. <https://doi.org/10.1038/s41467-019-10083-6>
- Todorov, A., Pakrashi, M., & Oosterhof, N. N. (2009). Evaluating faces on trustworthiness after minimal time exposure. *Social Cognition*, 27(6), 813–833.
- Tompson, S. H., Kahn, A. E., Falk, E. B., Vettel, J. M., & Bassett, D. S. (2019). Individual differences in learning social and nonsocial network structures. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 45(2), 253–271. <https://doi.org/10.1037/xlm0000580>
- Trope, Y., & Liberman, N. (2010). Construal-level theory of psychological distance. *Psychological Review*, 117(2), 440–463. <https://doi.org/10.1037/a0018963>
- Uncapher, M. R., & Wagner, A. D. (2009). Posterior parietal cortex and episodic encoding: Insights from fMRI subsequent memory effects and dual-attention theory. *Neurobiology of Learning and Memory*, 91(2), 139–154. <https://doi.org/10.1016/j.nlm.2008.10.011>

- van 't Wout, M., & Sanfey, A. G. (2008). Friend or foe: The effect of implicit trustworthiness judgments in social decision-making. *Cognition*, *108*(3), 796–803. <https://doi.org/10.1016/j.cognition.2008.07.002>
- Van Overwalle, F., & Baetens, K. (2009). Understanding others' actions and goals by mirror and mentalizing systems: A meta-analysis. *NeuroImage*, *48*(3), 564–584. <https://doi.org/10.1016/j.neuroimage.2009.06.009>
- Vernarelli, M. (2016). Predicting altruistic behavior and assessing homophily: Evidence from the sisterhood. *Sociological Science*, *3*, 889–909. <https://doi.org/10.15195/v3.a38>
- Wagner, D. D., Haxby, J. V., & Heatherton, T. F. (2012). The representation of self and person knowledge in the medial prefrontal cortex. *Wiley Interdisciplinary Reviews: Cognitive Science*, *3*(4), 451–470. <https://doi.org/10.1002/wcs.1183>
- Wakslak, C., Trope, Y., & Liberman, N. (2007). Construal levels and psychological distance: Effects on representation, prediction, evaluation, and behavior. *Journal of Consumer Psychology*, *17*(2), 83–95.
- Wang, Y., Collins, J. A., Koski, J., Nugiel, T., Metoki, A., & Olson, I. R. (2017). Dynamic neural architecture for social knowledge retrieval. *Proceedings of the National Academy of Sciences of the United States of America*, *114*(16), E3305–E3314. <https://doi.org/10.1073/pnas.1621234114>
- Weaverdyck, M. E., Lieberman, M. D., & Parkinson, C. (2020). Tools of the Trade: Multivoxel pattern analysis in fMRI: A practical introduction for social and affective neuroscientists. *Social Cognitive and Affective Neuroscience*, *15*(4), 487–509. <https://doi.org/10.1093/scan/nsaa057>

- Weaverdyck, M. E., & Parkinson, C. (2018). The neural representation of social networks. *Current Opinion in Psychology*, 24(1), 58–66. <https://doi.org/10.1016/j.copsyc.2018.05.009>
- Welborn, B. L., & Lieberman, M. D. (2015). Person-specific Theory of Mind in medial pFC. *Journal of Cognitive Neuroscience*, 27(1), 1–12. https://doi.org/10.1162/jocn_a_00700
- Wetzel, C. G., & Insko, C. A. (1982). The similarity-attraction relationship: Is there an ideal one? *Journal of Experimental Social Psychology*, 18(3), 253–276. [https://doi.org/10.1016/0022-1031\(82\)90053-1](https://doi.org/10.1016/0022-1031(82)90053-1)
- White, S. L., & Gowan, C. (2013). Brook trout use individual recognition and transitive inference to determine social rank. *Behavioral Ecology*, 24(1), 63–69. <https://doi.org/10.1093/beheco/ars136>
- Wickham, H. (2016). *ggplot2: Elegant Graphics for Data Analysis*. Springer-Verlag New York. Retrieved from <https://ggplot2.tidyverse.org>
- Wilson, R. K., & Eckel, C. C. (2006). Judging a book by its cover: Beauty and expectations in the trust game. *Political Research Quarterly*, 59(2), 189–202. <https://doi.org/10.1177/106591290605900202>
- Winkler, A. M., Ridgway, G. R., Webster, M. A., Smith, S. M., & Nichols, T. E. (2014). Permutation inference for the general linear model. *NeuroImage*, 92, 381–397. <https://doi.org/10.1016/j.neuroimage.2014.01.060>
- Winston, J. S., Strange, B. A., O’Doherty, J., & Dolan, R. J. (2002). Automatic and intentional brain responses during evaluation of trustworthiness of faces. *Nature Neuroscience*, 5(3), 277–283. <https://doi.org/10.1038/nn816>

- Wühr, P., & Huestegge, L. (2010). The impact of social presence on voluntary and involuntary control of spatial attention. *Social Cognition*, 28(2), 145–160. <https://doi.org/10.1521/soco.2010.28.2.145>
- Zell, E., Strickhouser, J. E., Sedikides, C., & Alicke, M. D. (2020). The better-than-average effect in comparative self-evaluation: A comprehensive review and meta-analysis. *Psychological Bulletin*, 146(2), 118–149. <https://doi.org/10.1037/bul0000218>
- Zerubavel, N., Bearman, P. S., Weber, J., & Ochsner, K. N. (2015). Neural mechanisms tracking popularity in real-world social networks. *Proceedings of the National Academy of Sciences*, 112(49), 15072–15077. <https://doi.org/10.1073/pnas.1511477112>