Cognitive and Computational Accounts of Adaptive Inference in Dynamic Social Networks

By

Alice (Ting) Xia

B.A., University of California, Irvine, 2011
M.A., Tsinghua University, 2016
M.Cog.Sc., Carleton University, 2018

Dissertation

Submitted in partial fulfilment of the requirements for the
Degree of Doctor of Philosophy in the Department of Cognitive and Psychological
Sciences at Brown University

PROVIDENCE, RHODE ISLAND

MAY 2026

This dissertation by Alice (Ting) Xia is accepted in its present form

by the Department of Cognitive and Psychological Sciences as satisfying the

dissertation requirements for the degree of Doctor of Philosophy.

Date______________ ____

Oriel FeldmanHall, Advisor

Recommended to the Graduate Council

Date______________ ____

Michael J. Frank, Reader

Date______________ ____

David Levari, Reader

Approved by the Graduate Council

Date______________ ____

David P. Lindstrom, Dean of the Graduate School

Alice (Ting) Xia

Department of Cognitive and Psychological Sciences
Brown University, Providence, R.I.
alice_xia@brown.edu

EDUCATION

2021 – 2026	**Ph.D., Cognitive Science**, Brown University
2016 – 2018	**M.A., Cognitive Science**, Carleton University (Ottawa, Canada)
2012 – 2015	**M.A., Public Administration**, Tsinghua University (Beijing, China)
2007 – 2011	**B.A.**, **Political Science**, University of California, Irvine

EMPLOYMENT

Aug 2020 – Aug 2021	**Research Specialist**, Jenkins Lab, University of Pennsylvania
Jul 2018 – Jul 2020	**Research Specialist**, Thompson-Schill Lab, University of Pennsylvania

PUBLICATIONS (*denotes equal contribution)

1. Teoh, Y. Y., Son, J-Y., **Xia, A.**, Bhandari, A., FeldmanHall, O. (2026). Neural encoding of an abstract
2. cognitive map supports social navigation in a real-world network. *Proceedings of the National Academy of Sciences, 123(*14*).*
3. Hu, Y-F., **Xia, A.**, FeldmanHall, O. (2025). Mental abstraction aids group cohesion in large social networks [Commentary]. *Behavioral and Brain Sciences*, 48, e170.
4. **Xia, A.**, Teoh, Y. Y., Nassar, M. R., Bhandari, A., FeldmanHall, O. (2025). Knowledge of information cascades through social networks facilitates strategic gossip. *Nature Human Behaviour*, 1-14.
5. Aslarus, I., Son, J-Y., **Xia, A.**, FeldmanHall, O. (2025). Climbing the social ladder requires early insight into social network structure. *Science Advances*, *11*(25), eads2133.
6. **Xia, A.***, Solomon, S. H.*, Thompson-Schill, S. L., & Jenkins, A. C. (2023). Constructing complex social categories under uncertainty. *Cognition*, 234, 105363.
7. Tompary, A., **Xia, A.,** Coslett, B. H., & Thompson-Schill, S. L. (2023). Disruption of anterior temporal lobe reduces distortions in memory from category knowledge. *Journal of Cognitive Neuroscience*, *35*(12), 1899-1918.
8. **Xia, A.**, Teoh, Y. Y., Bhandari, A., FeldmanHall, O. (*under review*). *Errors in social network knowledge predict how ties evolve in the future.*
9. **Xia, A.***, Fan, H.*, Nassar, M.R., FeldmanHall, O. (*under review*). *Asymmetric social ties are dynamic, detectable, and consequential for inference.*

GRANTS & HONORS

2023-2025 NIH T32	Training Grant in Interactionist Cognitive Neuroscience (ICoN) ($54,288)
Fall 2017, Summer 2018	Travel Bursary, Institute of Cognitive Science, Carleton ($500 CAD)
Fall 2014	Beijing City International Student Tuition Scholarship (~$4,000 USD)
2008 – 2010	Dean's Honor List, University of California, Irvine
2007 – 2008	Academic Competitive Grant, University of California, Irvine ($750)

INVITED TALKS

Apr 2026	Research Matters, Brown University
Oct 2025	Curiosity, Information Seeking & Exploration Conference Blitz Talk, Brown
Sep 2025	Multimodal Social Interaction Group (Host: Dr. Chen Zhou), Univ. of Glasgow
June 2025	Hartley Lab, NYU
May 2025	Computational Cognitive Neuroscience Lab (PI: Dr. Anne Collins), UC Berkeley
April 2023	Mende-Siedlecki Lab, University of Delaware

CONFERENCE PRESENTATIONS (*denotes equal contribution)

1. Fan, H.*, **Xia, A.***, FeldmanHall, O., Nassar, M.R. (2025, April). *Resolving uncertainty fosters tie formation in real-world social networks.* Poster accepted for the 2025 Annual Meeting of the Social Affective Neuroscience Society (SANS). Chicago, IL.
2. **Xia, A.**, Son, J.-Y., Aslarus, I., FeldmanHall, O. (2024, April). *The brain indexes social network structure to predict likelihood of information flow*. Poster presented at the 2024 Annual Meeting of the Social Affective Neuroscience Society (SANS). Toronto, Canada.
3. **Xia, A.**, Bhandari, A., FeldmanHall, O. (2023, October). *Strategic gossiping is aided by latent social network properties.* Poster presented at the 2023 Annual Meeting of the Society for Neuroeconomics (SNE). Vancouver, Canada.
4. **Xia, A.**, Solomon, S.H., & Thompson-Schill, S. L., Jenkins, A.C. (2020, July, virtual). *Constructing complex social categories from distinct membership information.* Poster presented at the 42nd Annual Meeting of the Cognitive Science Society.
5. Tompary, A., **Xia, A.**, & Thompson-Schill, S. L. (2019, October). *Semantic knowledge distorts episodic memory: Behavioral and neural investigations.* Poster presented at the 50th Society for Neuroscience. Chicago, IL.
6. Familiar, A.M., **Xia, A.**, & Thompson-Schill, S. L. (2019, March). *Learned social values modulate neural representations of faces.* Poster presented at the 26th Annual Meeting of the Cognitive Neuroscience Society, San Francisco, CA.
7. **Xia, A.**, Murasugi, K. (2017, December). *Effect of proficiency on Chinese character processing strategy in foreign language learners.* Poster presented at the International Conference on Bilingualism: Language and Heritage, Hong Kong.

TEACHING EXPERIENCE

Fall 2023	Teaching Assistant, Language and Gender (CLPS 0340), Brown University
Spring 2023	Teaching Assistant, Social Psychology (CLPS 0700), Brown University
Fall 2022	Teaching Assistant, Core I Seminar (CLPS 2001), Brown University
Fall 2019	Teaching Assistant, Perception (PSYC 111), University of Pennsylvania

SKILLS

• Social network analysis (modeling network dynamics using graph theory and statistical physics frameworks)
• Computational modeling of social networks (reinforcement learning and graph-based approaches)
• fMRI analysis (univariate GLM, RSA, and searchlight methods using SPM12 and Nilearn)
• Programming and visualization (R – fluent, MATLAB – intermediate, python – novice; Adobe Illustrator)

AD HOC REVIEW

Royal Society Open Science
PLOS Computational Biology

Acknowledgements

No amount of words will do justice to how much gratitude I have for Oriel, but I shall try my best here. If you have heard through the grapevine what it's like to be mentored by Oriel, or experienced it firsthand in all its colorful glory, then my admission that it was a tough learning curve might not come as a surprise. What I was even less prepared for, though, was the scientist and person I would become at the end. There are some who believe that people don't change, but regardless of whether that's true under this, that, or some other circumstance, it would've been a shame to be unchanged by Oriel's approach to science and to life more generally. She truly does it all, and with the kind of energy I once only dreamed of but have now somehow been able to fold into my own everyday mentality. During our first in-person meeting, I hesitantly said that I believe science is more about the people you work with. I didn't know whether she agreed. But then, her eyes opened up wide and she nodded vehemently. I still didn't know, at the time, just how much running a lab with that principle as its compass would really mean—not only in terms of the scientific output, but *how* we get there. She knows better than anyone that I did not always enjoy the process; she didn't give up on me until, through some alchemy of her wonderfully unusual mentorship, the process just simply became enjoyable. I'm sorry if that is an unsatisfactory account, but relational dynamics are not meant to be easily understood. And now, I can't see the endeavor of doing science in any other light. One of the best ways to get through anything difficult (and a PhD is up there, I think?) is to let yourself be changed by wonderful mentors. We don't get many of them throughout our lives, and I would do everything over again the exact same way—barring procrastinating on a R&R—as long as it was guaranteed that one of mine would be Oriel FeldmanHall (aka ).

Speaking of mentors, I was lucky to have had more than one during my time here. Thank you, Apoorva Bhandari, whose questions and critiques have raised the bar for what I consider to be "good thinking" in general; who suggested testing a model that eventually explained almost all of my data, and more down the line; who sat through each iteration of model simulations with me until one day, after a months-long pursuit, our expected pattern finally emerged, culminating in one of the most joyous moments of this whole research journey. Matt Nassar, who, from my faulty yet vivid memory, was the first to casually point out that information flows in a way that is very much *unlike* walking through a maze! Thank you for all the times you let me crash your lab's group work sessions and for your seemingly endless patience with all my computational modeling questions. Thank you so much to Michael J. Frank and David Levari, who served on my dissertation committee, and pushed me to consider the implications of my work in ways I won't soon forget—and in doing so, have made me feel like a scientist worth taking seriously. I am extremely indebted to Haoxue Fan, who made the third chapter of this dissertation possible. I loved that our meetings would be equal parts socializing and science-ing and I will miss working with such a spirited, smart, and engaged person.

To the students in our DormNets study—part of the Class of 2026, same as me! My intellectual journey here has, thankfully, collided with all of yours, and it is a beautiful thing to again take in just how vast our social networks really are. Thank you for your precious data and for letting us examine your social world. You've contributed to our understanding

of how people perceive the real, messy, and complicated relationships around them—and I hope your networks only grow richer and messier with time, because what's life without a bit of chaos?

To the members of the FeldmanHall Lab, past and present:

Alex: Thank you for bringing a positive energy to everything you do. I loved getting to know you, Heather, and Emily, and even being neighbors for a short stint. It was awesome to see you and Daantje take on a huge project, and I can't wait to see all the fruits of that labor.

Amrita: To the queen of computational modeling, thank you for making it immensely easier to step into that foreign and intimidating world by sharing your knowledge and wisdom. I especially appreciate you comforting me with anecdotes about your own struggles with learning those same concepts, even though you made it look so effortless. I love your passion, not just for science but for all the other things in the world that truly matter for well-being. None of us will ever have the high rate that you did for presenting ideas that were so irrefutably good, but I hope I'll get a little closer someday and if I ever do, I'll be thinking of you.

Daantje: To the most presidential lady of all, thank you for showing me the value of being optimistic. You have no idea how much your views have impacted me over the years. You are just the best combination of traits—strong, caring, sharp, funny, kind, ambitious, and always able to look on the bright side, which is not an easy thing to do and not something I cared about as much before getting to know you. Not only are you an accomplished researcher, but perhaps more importantly, you know what you want. That takes a lot of courage and practice in listening to your own voice instead of those of other people. I deeply respect the values you bring to this odd, messed up world. I wish everyone could have the privilege of spending time with you because they would almost certainly become better versions of themselves.

Izzy: You really made DormNets possible, and you have all the ingredients of a great colleague: clear communicator, quick learner, fun collaborator. But that's only the research side. I'll never forget the ski trips, especially the ones you so generously organized. On and off the slopes, you gave people one of the most precious things anyone can offer—your time, and plenty of it.

Jae: Thank you for being *the* giant on whose shoulders I've stood since our first conversation about your work. I truly believe this dissertation would not have happened without the foundational knowledge you built. You blew the door wide open so that we could saunter through. We first talked about the SR sitting outside over coffee and I thought, "This makes sense"—but to help me actually understand, you patiently walked through your code and every level of analysis, from technical details to larger conceptual questions. Your steadfast dedication to research and to discovering knowledge more generally has been infectious. Also, nearly everything I have learned about how to put together effective and visually pleasing slides, I have learned from you.

Jerry: Oh Jerry, Jerry, Jerry. You are a special one for sure. I could tell from the first little ponytail I saw you wearing in class in our first year. I didn't know if we could become friends, even after you joined the lab, but then I very quickly realized that you are also crazy. We're crazy in somewhat different ways, but we're fundamentally the same too. And that makes me upset sometimes, but I also accept it. You know what I mean. You always know what I mean! Thank you for considering me a friend, and for showing me that you are my friend by telling me secrets. On a more serious note, you are truly a deep thinker and someone who I will always be able to learn from. Your passion for research is inspiring but honestly, I know that whatever you do (even though I know you'll be staying in academia), you will win, so just relax and let us be the ones who worry about you taking over the world one day. Again, I am glad we are friends—don't have paranoia, thanks! Wait, I have more to say. I really appreciate your mind. You can bring a funny, ironic, and oftentimes weird perspective to almost anything and I think that's a very special skill to have. You are definitely one of the most interesting people I have ever met and I cannot wait to "exchange" labs one day.

Joey: Thank you for taking me under your wing from my first wobbly-stepped days in the lab and for continuing to be an inspiring presence, even from afar. You were always interested in chatting science and so naturally good at mentoring, which I could appreciate even as someone who *really* didn't know what they were doing then. Looking back, I remember you giving me last minute tips on how to make my first-year talk clearer and then, afterward, offering not just generic compliments but specific observations about the work itself that made me feel genuinely more confident in it. That's going to be something your students will relish. Lucky, lucky them.

Marc: Even though you never took me to Claudine in the end, I know that you have rooted for me every step of the way during this journey and I'm so grateful that the world is small enough for us to end up closer to each other again. As soon as I met you on my first day, I could tell you were warm and caring and obviously FUN. You taught me what Flanderization is. You also selflessly shared things about your experiences as a postdoc that helped me through my floundering early days in the lab. I'll never forget that you took me out for coffee after I dramatically announced on Zoom, in so many words, that I was feeling unmotivated. You always heard me out. Thank you for being you.

Molina: My desk buddy, I am amazed at how quickly time flies. In two short years, I've watched you really grow into your own and handle challenges like a pro. In my mind, there is only one way this will go, which is that you will become even stronger and more capable over the years. Thank you for supporting me through the ups and downs towards the end of my time here, and also for giving me opportunities to tell *you* that you are good at this thing called research and that you not only can run, but soar.

MZ: MZ, thank you for showing me what commitment and grit look like, in high concentrations of both. You have made a lot of sacrifices to be here, and you've definitely never shied away from the more difficult path—that's the kind of energy that I think everyone should adopt. I hope that you find what you love and that you wake up everyday remembering what your strengths are. I certainly will not forget.

Yang: Whether you accept my request or not, we are friends, I hate to break it to you. Putting that aside, there is absolutely no denying that your contribution to this lab and to my work in particular has been incredibly valuable—and that's the mildest possible way of putting it. You've taught me how to think, how to code, how to approach certain kinds of statistical problem. You are the whole package, and I really wish I could shout your worth from the rooftops. You always bring more than expected to the table, which is strange, because you'd think I'd become desensitized over time—but my theory of why that hasn't happened is that you become more remarkable week after week, at a rate that far outpaces my ability to correct my prediction errors. You bring the best balance of fun and intellectual rigor, and when I grow up (i.e., become a postdoc), I want to be just like you—as I've told you many times before. And I'll tell you again when I eventually fail because you've set the bar too high, which is your own fault.

To my family: My parents are a bit wacky, as all parents are, more or less—but it is an unobjectionable truth that they sacrificed more than I ever have, and perhaps ever will, to give me the opportunities and freedom that allowed me to end up here. I am grateful that they always saw so much potential in me, even when I didn't see it myself. I hope I have made them proud. There are not enough words in English or Chinese to properly thank my younger sister and brother, Mona and Evan, who I laugh harder with than anyone else in the world. I am so lucky to have you both by my side as we shine the same light on all the ridiculous circumstances we find ourselves in and then chuckle our way out of them together. Our shared sense of humor is one of the most treasured things in my life. Mona has taught me how to be brave and embrace my own values, regardless of whether other people approve. And Evan has made me feel immensely proud—simply to know someone who is such a brilliant goofball. This is the ultimate triad, and it's not just closed, it is *sealed shut forever*.

Lastly, I'd like to thank the universe for all the good fortune that has come my way. Duly noted.

Table of Contents

List of Tables

List of Figures

Abstract of "Cognitive and Computational Accounts of Adaptive Inference in Dynamic Social Networks" by Alice (Ting) Xia, Ph.D., Brown University, May 2026

Social success depends not only on who we interact with in the moment, but also on the many connections those individuals have within the community. When deciding whether to share sensitive information, for example, we often anticipate how it might travel through others—gauging the risks of diffusion across the broader network. How do people represent social networks, and how do those representations guide social inference and behavior?

Across three chapters, in both artificial and real-world environments, I show that a representational format capturing patterns of indirect connectivity underlies adaptive inference. Specifically, a representation that encodes global topology via accumulated indirect paths (Katz communicability) best explains strategic communication behavior and accurate predictions about which relationships form and dissolve up to six months into the future, providing converging evidence that knowledge of global network topology meets the cognitive demands of different inferential problems with important social consequences.

At the same time, indirect connectivity does not capture everything people know about relational structure. Relational asymmetry—cases in which two individuals hold divergent views of their connection—contains information that global connectivity alone cannot explain. I show that asymmetric ties are prevalent yet volatile, resolving toward symmetry over time, and that relational asymmetry shapes how the broader community reasons about the network. Together, these findings reframe social network cognition as active and adaptive—and reveal that the relationship between mental representations and network structure is not one-way: how people model their social world shapes how that world changes.

Chapter 1

General Introduction

1. Familiar chaos

Sociologist George Homans described social behavior as "familiar chaos"—familiar because of our shared experience navigating group dynamics since childhood, but chaotic because a systematic understanding of how individuals contend with ever-changing social landscapes has remained incomplete[1]. More than fifty years later, many open questions remain about what drives our ability to weave through situations that appear simple on the surface, from deciding whether to pass along a piece of gossip, to playing matchmaker for two friends. And yet, when considering how the consequences of these actions play out in the larger web of relationships that surround each person, the mundaneness of these scenarios obscures the underappreciated question of how individuals anticipate social outcomes that extend beyond one or two others. Our social worlds are vast, relative to what we are able to observe at a given moment in time. The larger and more dynamic the network, the less any snapshot of it can be trusted, as friendships cool or new alliances emerge. Given our evident success in cultivating large societies[2], it may be tempting to treat this as a solved problem—but *how* do humans make adaptive inferences in a highly interconnected social world that is constantly changing?

This problem is easy to overlook because social navigation so often feels fluent. We quickly develop a working sense of who is close to whom when starting a new job, who to approach for advice given our read on expertise and hierarchy, and how information is likely to travel through the people around us. Much of what we gather about the social ties surrounding us is informative not just for mapping out a community, but for making predictions with real consequences, about reputation and influence, or who will hear what and when. But how accurate are these assessments? Empirical data suggest that our confidence is not always warranted. When people are asked to report on the relationships around them, they tend to be accurate in recalling gist-like information, such as cliques, but frequently inaccurate in the specifics of who actually interacts with whom[3]. Mental portraits of social environments are not just fuzzy and incomplete, but systematically distorted: people tend to believe that relationships are more reciprocated than they really are, and consistently place themselves closer to the center than they actually are[4]. These distortions are understandable, as the mind likely prioritizes forming a coherent and actionable model of a complex social world[5]. However, in a network that is large, interconnected, and constantly shifting, the gap between the model a person carries and the structure that actually surrounds them can be substantial, with consequences that only become visible when an inference goes wrong.

The costs of social miscalculations are not abstract. Information that reaches the wrong person can damage someone's reputation and reshape how they are perceived across a community, possibly closing off relationships and opportunities that could have materialized. A person who misreads their own position may invest in relationships that don't pan out or fail to cultivate the ones that would have mattered, without ever understanding why. And acting on a friendship as if it is stronger than it actually is means potentially confiding in, relying on, or publicly aligning with someone who may not reciprocate that trust, with consequences that extend well beyond two individuals. None of these are edge cases. They are woven into the familiar texture of social life in large, interconnected communities, where most of the structure that determines the consequences of our social actions is never directly visible. And yet, despite carrying models that are incomplete and systematically distorted, people manage to coordinate, cooperate, and sustain the complex social structures that define human life. This success cannot be explained by direct observation alone. Understanding the cognitive and computational mechanisms that support social inference, and how they hold up as the social world continues to change, is the central question of this dissertation.

2. Reasoning from within a changing structure

In a real-world social network, people must make inferences about others under conditions of partial observability. Unlike a physical location in space, an individual's "position" in a network is defined entirely by the configuration of relationships surrounding them. More fundamentally, an individual is part and parcel of the same structure they are trying to observe. Each person is therefore locked into an epistemic situation in which the very medium used to perceive the network is the network itself. I argue that, in practice, this means two people separated by only a few degrees—well within the short path lengths that characterize most social networks[6]—may nonetheless inhabit fundamentally different social realities, not because they are far apart, but because the local relationships that constitute each of them filter and selectively transmit different versions of the same shared world. These inherent inequalities shape what any one person can see—and therefore what they can know. Because direct observation cannot close this gap, inference is not merely useful, it is the best available tool, and the demands on social reasoning are non-trivial. A person must integrate what they know about the relationships around them, reason through paths they cannot directly observe, and do so in service of goals that change depending on the situation. But all of this assumes the structure being reasoned about holds still—and it does not.

The problem is further compounded by the fact that the structure being inferred does not hold still. Ties are constantly changing, as new connections emerge and existing ones fade. Relationships that once anchored a person's understanding of their social world may quietly and quickly transform. This is qualitatively harder than navigating a large but stable structure. A static network, however vast, offers a fixed target, while a dynamic one means that any representation a person holds is likely becoming outdated the moment it is formed. The difficulty is not merely a matter of scale or incomplete information, but that the target of the inferential problem is a moving one. One possible consequence is that the very act of navigating a social network may itself be part of what moves it. Each person is embedded in a web of ties that may still be coalescing or may be about to shift dramatically, for example, as an important relationship crumbles and opens the door for others to leverage new social opportunities. Position compounds this problem because what an individual can observe is determined by where they sit in the network, and because that network is itself in motion, the ground beneath any social inference is never quite stable. A satisfying account of social network cognition must therefore explain not just how people represent network structure under these conditions, but how they reason flexibly through those representations, and then examine how that capacity holds up when the structure itself is changing.

3. The limits of existing approaches

A tradition of research known as "cognitive social structures" has systematically documented how people perceive the networks around them[4]. Rather than treating network perception as incidental, this work established it as a distinct cognitive capacity, and one that operates through abstraction rather than faithful reproduction. People encode the global organization of a network (e.g., who is central, which clusters exist) while systematically missing details about specific ties[7], suggesting that the social networks encoded in people's minds reflect a compressed form of the global structural regularities (Brashears, 2013). This compression produces patterned misjudgments: people over-report reciprocated ties, perceive more clustering and balance than actually exists, and place themselves closer to the center than they are[7–9]. These errors are not failures of attention but rather evidence of a mind that abstracts over direct experience to encode global, structural knowledge of the social world[10–12]. What this body of work did not address, however, is what people do with these representations. The question of how people use compressed models of network structure to make predictive inferences about how that structure will change, and whether those inferences are accurate, strategic, or consequential, remained largely outside its scope.

A parallel body of work approached the same social landscape from the opposite direction. Whereas cognitive social structures research focused on the perceiver, network science focused on the structure itself. The field has produced a sophisticated account of how ties in social networks form, e.g., through triadic closure[13] or friends of friends becoming friends, and how network position shapes individual outcomes in various domains, from health to economic mobility[14,15]. It has revealed that not all positions are equal: weak ties that bridge otherwise disconnected clusters provide access to non-redundant information[16], and individuals who occupy the gaps between clusters gain access to privileged information that may be unavailable to those who belong to a single group[17]. For all its advances in understanding social relations through a systems-level lens, network science has focused less on the individual minds that form the basic units of the networks it studies. Structure is treated as an object of study in its own right—something that emerges from aggregated behavior—rather than as something people actively represent and reason about. Network science tools were developed primarily to characterize structural properties of systems. Therefore, adapting them to address questions about internal mental representations requires additional computational and behavioral methodology.

Between these two well-established areas of research lies the question at the heart of this dissertation: how do people build and use predictive representations of network structure, and what happens to that capacity under the conditions of real social networks? Tackling this problem requires a level of analysis that neither field was built to provide—one that combines computational models of individual cognition with network-level data. Prior work has made this question tractable for static networks[18–20]. What remains open, and what this dissertation addresses, is what happens when the network itself is changing, and when the cognitive representations people encode must support not just perceptions of reality, but also reasoning about the potential consequences of social inference.

4. A computational framework for social network representation

Social networks describe a fundamentally abstract feature of our lives—social relations. While the individuals that comprise a network can be physically located, the relations between them are largely invisible, scaffolded onto physical acts whose social meanings are defined by convention (e.g., a hug versus a handshake). Conveniently, social networks can be brought into relief via mathematical representations, beginning with an adjacency matrix that indicates whether a relationship exists between any two individuals. Formulating a network this way opens the door to mathematical operations that help elucidate how social structure might be approximately instantiated in the mind.

Past work laid the groundwork for using algorithms such as the Successor Representation (SR) [21] to estimate connectivity between individuals, treating the problem as analogous to navigating from one physical location to another[18–20]. This was a substantial advance, as it formalized how multistep association mechanisms give rise to systematic biases in how people recall and infer ties in a network. The critical insight is that when two network members are observed together, the representation of one is updated to include not just the other, but also a discounted version of everyone that other person is connected to. This allows the model to capture the intuition that friends-of-friends are more likely to be connected than strangers, and to predict inferences about ties that have never been directly observed. However, the SR is limiting in the type of inference it can support. It represents the network as a series of probabilistic steps from one person to the next, where the probability of moving from person A to person B (i.e., transition probability) is determined in part by how many connections A has. The more connections that A has, the more ways the total probability can be split, so the chance of transitioning to any specific connection is lower. This means individuals with more connections have a lower probability of traversing to any one of them, which is counterintuitive to what popularity actually affords in social life. Consider two people in the same network: one with three friends, and one with fifteen. The SR assigns a higher transition probability to each of the three friends than to any one of the fifteen, which means the well-connected person appears, by this metric, harder to reach than the isolated one. But that is precisely the wrong answer for many social inference problems. A popular person is not less reachable; they are more so.

Katz centrality offers a point of view that aligns with the importance of popularity. Originally developed as a sociometric status index[22], it accounts not just for degree (number of direct connections), but for the idea that being connected to well-connected others matters too. By counting all paths that directly and indirectly connect two people—and down-weighting longer paths so that a friend-of-a-friend counts for less than a direct friend, who in turn counts for less than a mutual friend-of-friends—Katz communicability produces a matrix in which every entry is a number summarizing the total accumulated connection between two specific people in the network. Rather than quantifying the likelihood of reaching a given individual via a random walk through the network, Katz captures the total number of paths connecting to that individual, across all possible routes, weighting shorter paths more heavily than longer ones. The result is a matrix in which every value reflects the accumulated density of pathways in the network between two individuals, integrating both direct ties and the full cascade of indirect ones. This means that a highly connected person accumulates more pathways to and from others across the network. Their centrality is preserved rather than diluted, and the inferences that depend on it become recoverable in a way the SR cannot support. This dissertation is the first to test whether a representational format akin to Katz communicability holds up under the demands of real social life, where inference must be strategic, and the network itself may be changing.

5. The program of research

Navigating a social network requires knowing not just who is connected to whom, but how information travels through the network as a whole. **Chapter 2** establishes that people represent this global structure and use it strategically. Using gossip as a behavioral probe (via a task that requires broadcasting information while preventing it from reaching a specific target), I show that people flexibly combine sensitivity to structural distance and popularity to achieve their goal, even as the target changes and cognitive demands increase. The representation that best accounts for this flexibility is Katz communicability, which integrates all direct and indirect paths between network members and preserves the influence of popular individuals rather than diluting it. Katz outperforms the SR precisely because it reflects how centrality actually functions in social life—as accumulated influence, not just local reachability. What this study cannot address, however, is whether the same capacity operates when the network is not a fixed object to be learned, but an evolving structure the person is embedded in.

Chapter 3 takes that question into a real, emerging social network. By tracking the friendships of a cohort of undergraduates across an academic year, I ask whether people can predict the future state of the relationships around them. They can, and systematically so. Errors in judgments about others' relationships predict the actual state of those relationships nearly half a year later, and individual differences in how well people encode a Katz-like representation of the network determine whether those errors align with future outcomes. This establishes something Chapter 2 could not. The representational capacity is not just a response to network structure, it is a generative input into how that structure evolves. But the study treats each relationship as a shared unit, by assuming both people in a pair hold the same view of their tie.

Chapter 4 examines the concept of relational asymmetry, two people holding different views of the same relationship. Relational asymmetry is pervasive in real social networks, and volatile. While the overall proportion of asymmetric ties stays relatively stable, specific dyads are continuously transitioning between relationship states. I show that these transitions are not random. Asymmetric pairs move toward each other over time, resolving their discrepancy at rates well above chance. The broader community also tracks asymmetry within dyads. People's judgments about information flow and social integration are shaped by the degree of asymmetry in the relationships around them, revealing that people represent not just who is connected to whom, but how securely. Moreover, people use that signal to make adaptive social inferences about a network that is still taking shape.

Together, these studies reveal that people's sensitivity to the structural properties of their social networks supports a range of consequential social inferences. Scaled from controlled laboratory networks to a naturally emerging community of nearly 200 individuals, this sensitivity shapes strategic behavior, predicts how the network will evolve, and tracks the internal dynamics of the ties themselves. Taken together, they build a case that the relationship between social cognition and network structure is not one-directional. The representations people carry are shaped by the networks they are embedded in, and the inferences those representations support feed back into how those networks develop. In each case, what people represent about their social world is part of what makes that world take the shape it does.

References

1. Homans, G. C. Social behavior: Its elementary forms. (1974).
2. Dunbar, R. I. Coevolution of neocortical size, group size and language in humans. *Behav. Brain Sci.* 16, 681–694 (1993).
3. Killworth, P. & Bernard, H. Informant accuracy in social network data. *Hum. Organ.* 35, 269–286 (1976).
4. Krackhardt, D. Cognitive social structures. *Soc. Netw.* 9, 109–134 (1987).
5. Brashears, M. E. Humans use Compression Heuristics to Improve the Recall of Social Networks. *Sci. Rep.* 3, 1513 (2013).
6. Watts, D. J. & Strogatz, S. H. Collective dynamics of 'small-world'networks. *nature* 393, 440–442 (1998).
7. Freeman, L. C., Romney, A. K. & Freeman, S. C. Cognitive structure and informant accuracy. *Am. Anthropol.* 89, 310–325 (1987).
8. Freeman, L. C. Filling in the Blanks: A Theory of Cognitive Categories and the Structure of Social Affiliation. *Soc. Psychol. Q.* 55, 118 (1992).
9. Kumbasar, E., Rommey, A. K. & Batchelder, W. H. Systematic biases in social perception. *Am. J. Sociol.* 100, 477–505 (1994).
10. Basyouni, R. & Parkinson, C. Mapping the social landscape: tracking patterns of interpersonal relationships. *Trends Cogn. Sci.* 26, 204–221 (2022).
11. Brands, R. A. Cognitive social structures in social network research: A review. *J. Organ. Behav.* 34, (2013).
12. Smith, E. B., Brands, R. A., Brashears, M. E. & Kleinbaum, A. M. Social Networks and Cognition. *Annu. Rev. Sociol.* 46, 1–16 (2020).
13. Kossinets, G. & Watts, D. J. Empirical Analysis of an Evolving Social Network. *Science* 311, 88–90 (2006).
14. Borgatti, S. P., Mehra, A., Brass, D. J. & Labianca, G. Network Analysis in the Social Sciences. *Science* 323, 892–895 (2009).
15. Jackson, M. O. Networks and economic behavior. *Annu Rev Econ* 1, 489–511 (2009).
16. Granovetter, M. S. The strength of weak ties. *Am. J. Sociol.* 78, 1360–1380 (1973).
17. Burt, R. S. The social capital of structural holes. *New Econ. Sociol. Dev. Emerg. Field* 148, 122 (2002).
18. FeldmanHall, O., Son, J.-Y. & Bhandari, A. Abstract Cognitive Maps for Complex Social Systems. *Curr. Dir. Psychol. Sci.* 34, 349–356 (2025).
19. Son, J.-Y., Bhandari, A. & FeldmanHall, O. Cognitive maps of social features enable flexible inference in social networks. *Proc. Natl. Acad. Sci.* 118, e2021699118 (2021).
20. Son, J.-Y., Bhandari, A. & FeldmanHall, O. Abstract cognitive maps of social network structure aid adaptive inference. *Proc. Natl. Acad. Sci.* 120, e2310801120 (2023).
21. Dayan, P. Improving Generalization for Temporal Difference Learning: The Successor Representation. *Neural Comput.* 5, 613–624 (1993).
22. Katz, L. A new status index derived from sociometric analysis. *Psychometrika* 18, 39–43 (1953).

Chapter 2

Knowledge of information cascades through social networks facilitates strategic gossip

Alice Xia, Yi Yang Teoh, Matthew R. Nassar, Apoorva Bhandari, & Oriel FeldmanHall

Published in *Nature Human Behaviour* (2025)

1. Introduction

Many of our daily decisions are made against the backdrop of large, complex, and dense social networks, where each person has many ties to others. These relationships enable information to spread quickly through a community. While networks are highly efficient conduits of information, deciding who to communicate with, especially when the topic is about someone else, presents a challenge. Imagine wanting to share sensitive information about a mutual acquaintance with a friend. This is a form of gossip, since it is communication involving a sender, a receiver, and a target, where the target is not aware they are being discussed[1]. To avoid reputational damage or a breach of trust, it is critical to accurately assess how likely it is that the acquaintance—the target of gossip in this case—will discover you were talking about them. Decisions to gossip are therefore often strategic in nature[2,3].

Research in the lab, as well as our day-to-day experiences, illustrate that evaluations of others frequently travels through a network without reaching the target[4]. This implies a remarkable ability to assess the long-range consequences of gossip decisions even before they are made. Indeed, people appear to gossip in such a way that the target rarely finds out[1,5–7], which is critical for avoiding negative social consequences (e.g., retaliation)[8–10]. There is good evidence that individuals predominantly gossip within their close social networks, such as with trusted colleagues, demonstrating a tendency to carefully control its circulation to prevent it from reaching unintended recipients[11]. Even though gossip accounts for 10-20% of daily conversations[7] and an astounding 90% of people gossip in the workplace[5], the mechanisms that enable strategic gossiping remain unknown. Here, we interrogate how people achieve strategic gossiping.

What type of social knowledge might be useful for gossiping? If the goal is to share gossip while avoiding the target of gossip, one obvious and critical factor to consider is how often the individual you are gossiping with also tends to gossip with others. People evaluate this type of trustworthiness to reduce the risk of gossip reaching unintended audiences[1,12], and once it becomes evident that someone is unlikely to spread gossip, it is less risky to share sensitive information with that person[13]. However, the notion of trust becomes far more consequential when considered in the context of a group[14]. Research on organizational gossip, for instance, emphasizes that sharing sensitive information with someone who is untrustworthy but isolated within the community may have fewer repercussions than sharing with someone highly connected or influential[15]. Put simply, the number of friends a person has—that is, their popularity—should also inform whether confidential information is shared or withheld.

Similarly, information is more likely to travel between those who are closely connected in a social network, such as individuals with many mutual friends, compared to those who are distant and separated by many intermediary connections[16]. Research illustrates that knowledge about the proximity of potential spreaders to the target can shape how people choose to spread gossip[11,17], suggesting that the distance between the target and those who gossip is yet another determinant of how information propagates. While neither popularity nor network distances are directly observable, both are fundamental structural network properties that can determine information flow[18–21]. Thus, if one wanted to strategically gossip, a useful mental representation of the network should include these latent network properties. This places popularity and distance as critical information sources when building mental representations of how information might flow through a network.

We first examine what information people use to build network representations to aid strategic gossiping. In this case, we operationalize strategic gossiping as balancing the competing demands of spreading gossip as widely as possible while avoiding the target of gossip. Across three controlled experiments using artificial social networks, participants were tasked with learning about the relationships that comprise novel social networks. We find that people rely on knowledge about learned friendships to infer structural properties of the network that are not readily observable during learning—popularity and distance from target—and use this information to gossip strategically. In a fourth experiment leveraging a real-world social network, we observe that people use the same network properties to predict the likelihood of personal news passing between two individuals. We deploy computational modeling to test the format of mental representation that can account for these behavioral patterns. We observe that a mental representation that encodes the communicability (i.e., strength of association) between pairs of individuals in the network—capturing that information flows through a cascade-like process—best accounts for strategic decisions to gossip.

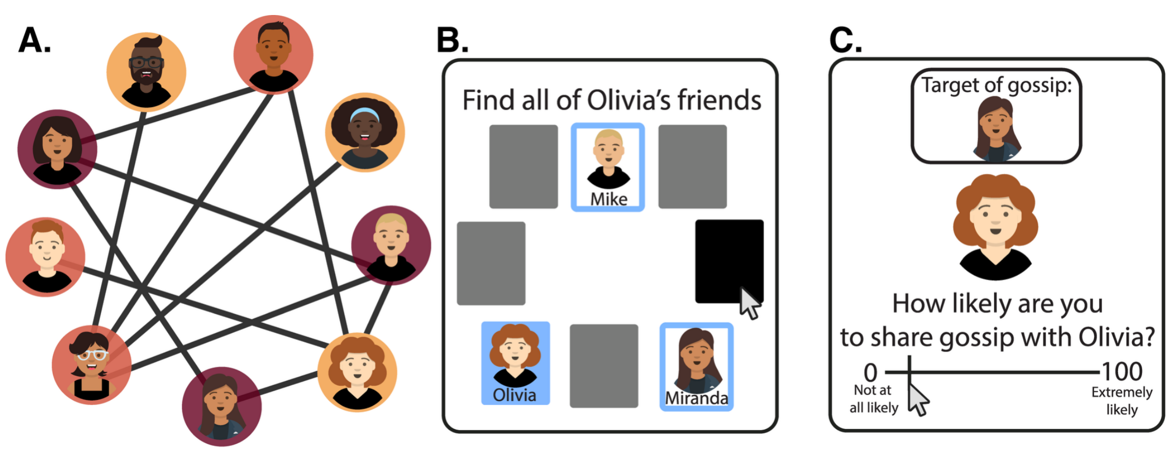

Figure 1 | Network and task structure used in online experiments. (A) The social network used in Experiments 1-3 *was composed of nine members who varied in popularity (i.e., number of friends); each line represents mutual friendship between two members. The color of each circle reflects the network member's transmissibility (Experiments 2 and 3 only), or general tendency to gossip, which is operationalized as the* *probability of sharing gossip with a friend: maroon = 100%; red = 70%; orange = 40%.* ***(B) Learning the network through 'flash cards' (Experiments 1 and 2 only).*** *Participants learned about friendships by selecting a gray square and then receiving feedback about whether that individual is a friend (blue outlined box) or not (black box) of a reference individual, e.g., Olivia (blue background box).* ***(C) Gossip Task.*** *After learning about the network, participants rated how likely they would be to share gossip with a Receiver in the network (Olivia), given the goal of spreading gossip widely while avoiding a given Target (Miranda). The Target was fixed within a block and Receiver changed on every trial.*

2. Results

a. Experiment 1: People spread gossip via popular members far from the Target

How do people leverage knowledge of latent social network structure in service of the competing goals of widely disseminating gossip, while minimizing the risk of the gossip reaching its subject? Adopting this dueling goal enabled us to repeatedly probe an individual's knowledge of the network's structure, even in the absence of being directly embedded within the network and therefore lacking personal motivation to share or withhold gossip. In Experiment 1, participants (N=201) first learned about a novel nine-person network, composed of 12 social ties (Fig. 1A-B). We assessed learning performance using the Jaccard index of guessed friendships relative to true friendships in the network. By the final learning block, group-level accuracy reached 74%, indicating that participants learned connections in the network. Participants were subsequently tasked with reporting their willingness to share gossip with an individual (the Receiver), given a specific Target of gossip, and the complex goal at hand (Fig. 1C). The network was designed to include a practical range of distance and popularity combinations, ensuring it was complex enough to test our hypotheses while remaining small enough for participants to learn its structure within a reasonable time. We then examined how two topological social network properties—popularity (how many friends the Receiver has) and distance to Target (shortest path between the Receiver and Target, given the network structure)—influence likelihood of gossiping with the Receiver.

We observe that people are more likely to share gossip with individuals who are more popular and farther away from the Target ($\beta_{distance*popularity}$ = 0.15, 95% CI [0.12, 0.17], z = 11.11, p < 0.001, Odds Ratio (OR) = 1.16; Fig. 2A; Supplementary Table 1). A simple slopes analysis shows that when Receivers are direct friends of the Target (distance = 1), participants are significantly less likely to share gossip with those Receivers, especially if they are also popular (simple slope: $\beta_{popularity@distance1}$= -0.25, 95% CI [-0.29, -0.20], z = -13.94, p < 0.001, OR = 0.78). However, farther distances from the Target weakens the popularity effect, as confirmed by a contrast analysis comparing the predicted effects at the closest and farthest distances from the Target (contrast: $\Delta[b_{popularity@distance1}$, $b_{popularity@distance4}]$ = -0.44, 95% CI [-0.54, -0.35], z = -11.11, p < 0.001, OR = 0.64; Supplementary Table 5).

Although likelihood of gossip-sharing appears to increase linearly with distance ($\beta_{distance}$ = 0.75, 95% CI [0.65, 0.84], z = 15.49, p < 0.001, OR = 2.12; Supplementary Table 1), when we assess whether the effect of distance might plateau (i.e., using quadratic terms in the model to test non-linearity, Supplementary Table 6), we find that the rate of gossiping decelerates or stabilizes at greater distances. This suggests that at farther distances, gossip may lose its relevance, and be less likely to be shared. As in any social network, distance and popularity are inherently linked—more popular individuals tend to be closer to the average network member. Likewise, in our network, the most popular individuals were always closer to a given Target. To ensure that our results are not simply an artifact of our network configuration, we conducted an additional analysis with popularity levels binned into low (popularity = 1 and 2) and high (popularity = 3 and 4) categories. Indeed, likelihood of gossiping increases with distance and is still enhanced with more popular individuals, i.e., those who three or four friends, relative to those who have only one or two friends ($\beta_{distance*popularity_{high}}$ = 0.40, 95% CI [0.33, 0.47], z = 11.23, p < 0.001, OR = 1.49; Supplementary Table 7).

To examine whether explicit memory of network connections influenced gossip decisions, we included d-prime, a measure of memory accuracy for pairwise friendships. If decisions were primarily guided by explicit, veridical knowledge of the network, we would expect a strong link between memory accuracy and gossip-sharing behavior. However, there was no evidence to suggest that memory of pairwise friendships predicted decisions to gossip (β_{dprime} = -0.01, 95% CI [-0.06, 0.04], z = -0.39, p = 0.693, OR = 0.99; Supplementary Table 1), or that memory accuracy modulated the use of distance ($\beta_{dprime*distance}$ = -0.02, 95% CI [-0.07, 0.04], z = -0.59, p = 0.553, OR = 0.98) or popularity ($\beta_{dprime*popularity}$ = 0.00, 95% CI [-0.01, 0.01], z = 0.38, p = 0.703, OR = 1.00). The fact that exhibiting specific network knowledge does not change the relationship between distance and popularity on gossip decisions, suggests that strategic behavior may emerge from sensitivity to the network structure rather than deliberate recall of detailed relationship knowledge. Overall, the combined use of distance and popularity reveals how people achieve the dual objectives of disseminating gossip widely while attenuating the risk of it reaching the Target. People share gossip *more* with popular individuals, since they serve as effective vehicles for dissemination, but do so only when those popular individuals are distant from the Target of gossip. On the other hand, when in the Target's proximity, people attenuate risk by being reticent, especially with the Target's most popular friends.

b. Experiment 2: Transmissibility further enhances gossip dissemination

While knowledge about friendships between individuals is fundamental for predicting gossip spread, people often also have access to other types of relevant information, such as an individual's general tendency to gossip—or their transmissibility—which serves as a proxy for trustworthiness. Therefore, in Experiment 2, we gave participants additional information about each network member's transmissibility to test how multiple pieces of goal-relevant information are leveraged to gossip strategically. We asked a new cohort of participants (N=198) to first learn about nine people's transmissibility through trial-and-error. Then, in a second task, participants were presented with these same nine people and learned that they were in fact part of the same social network, using the same flash cards task as the one described above (Fig 1B). Essentially, participants acquired these two pieces of information (transmissibility and friendships) in a piecemeal manner. This manipulation allowed us to examine whether people integrate multiple sources of social information to achieve the goal of strategic gossiping more effectively (see Supplementary Fig. 1 for simulation of expected effects of transmissibility).

First, network learning performance reached 68% in the final block of the task, suggesting that participants were able to learn a significant majority of the true friendships. As with Experiment 1, we observe that likelihood of sharing gossip increases with distance, and this effect is moderated by popularity ($\beta_{distance*popularity}$ = 0.10, 95% CI [0.08, 0.12], z = 8.61, p < 0.001, OR = 1.11; Supplementary Table 2), replicating the finding that these network properties provide useful scaffolding for achieving complex communication goals. Participants again showed a preference for avoiding gossip with popular individuals who are closest to the Target (simple slope: $\beta_{popularity@distance1}$= -0.21, 95% CI [-0.26, -0.17], z = -12.01, p < 0.001, OR = 0.81), which diminishes with distance from the Target (contrast: $\Delta[b_{popularity@distance1}, b_{popularity@distance4}]$ = -0.30, 95% CI [-0.38, -0.22], z = -8.61, p < 0.001, OR = 0.74; Supplementary Table 5). As before, this interaction remains significant when binning popularity into high and low levels ($\beta_{distance*popularity_{high}}$ = 0.25, 95% CI [0.18, 0.31], z = 7.80, p < 0.001, OR = 1.28; Supplementary Table 7).

Moreover, we find that greater transmissibility amplifies willingness to share gossip given a Receiver's popularity and distance to the Target ($\beta_{distance*popularity*transmissibility}$ = 0.07, 95% CI [0.01, 0.13], z = 2.29, p = 0.022, OR = 1.08; Fig. 2B; Supplementary Table 2). A contrast of simple slopes reveals that participants are less likely to share gossip with popular Receivers who are friends with the Target (simple slopes: $b_{popularity\&transmissibility100}$= -0.20, 95% CI [-0.27, -0.13], z = -7.92, p < 0.001, OR = 0.82; $b_{popularity\&transmissibility40}$= -0.23, 95% CI [-0.29, -0.16], z = -9.61, p < 0.001, OR = 0.79; Supplementary Table 8), regardless of the Receiver's transmissibility (contrast: $\Delta[b_{popularity\&transmissibility100}, b_{popularity\&transmissibility40}]$: 0.03, 95% CI [-0.06, 0.11], z = 0.75, p = 0.454, OR = 1.03). At the farthest distances from the Target, however, participants are more likely to gossip with popular Receivers, but only if those Receivers are highly transmissible and (simple slopes: $b_{popularity\&transmissibility100}$= 0.16, 95% CI [0.06, 0.26], z = 4.34, p < 0.001, OR = 1.17; $b_{popularity\&transmissibility40}$= 0.01, 95% CI [-0.08, 0.11], z = 0.22, p = 0.829, OR = 1.01). Participants thus incorporate information about each network member's transmissibility when considering how to gossip, even though they were never explicitly instructed to use this information—information that was learned in an entirely separate task.

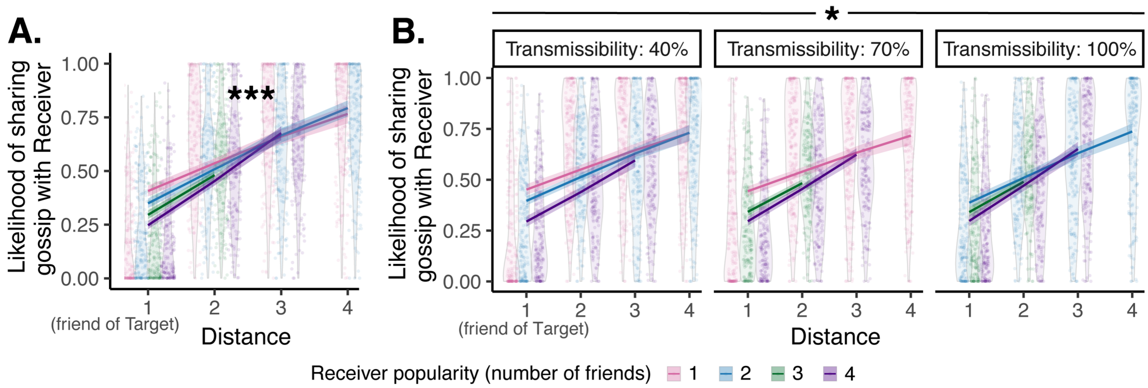

Figure 2* | *Likelihood of gossiping in Exp. 1 and 2. ***(A)*** *After learning about friendships in an artificial social network in Experiment 1, participants were able use distance (x-axis) and popularity (color) to make adaptive inferences about information flow, sharing gossip (y-axis) less with popular Receivers close to the Target, and sharing more with those same popular Receivers when they are farther away from the Target (Distance x Receiver Popularity interaction, ***p < .001).* ***(B)*** *Additional information about each network member's transmissibility (frequency of gossiping with a friend) was provided in Experiment 2. Likelihood of gossiping depends on the combined effect of distance and popularity, and transmissibility: transmissibility amplifies the likelihood of gossiping with popular individuals, but only at far distances (Distance × Receiver Popularity × Transmissibility interaction, *p < .05). Simple slopes analyses testing these effects are described in text. Lines represent regression model predictions, with 95% confidence intervals. Participant-level raw data shown as points. Distribution of data shown as violins.*

c. Experiment 3: Complex learning environment limits use of transmissibility.

While Experiment 2 revealed that people rely on multiple pieces of goal-relevant social information for strategic gossiping, in the real world, people typically receive multidimensional information about individuals simultaneously and not in a piecemeal fashion[22,23]. For example, we may observe someone betray a friend by telling a secret to another friend, thus learning about friendship *and* transmissibility in a single episode. However, given the known constraints on information processing, attending to multiple sources of information is likely to come at the cost of having a precise mental model of the network[24,25]. To test this possibility, in Experiment 3, we created an even more challenging learning environment by presenting information about friendships and transmissibility together, akin to how learning unfolds in the wild. Now, unlike the first two experiments, participants (N=169) were required to infer that two individuals are friends *if* they shared gossip with each other (Fig. 3A; Methods). Similarly, participants had to infer the transmissibility of each person based on how often an individual shared gossip relative to how often it was withheld.

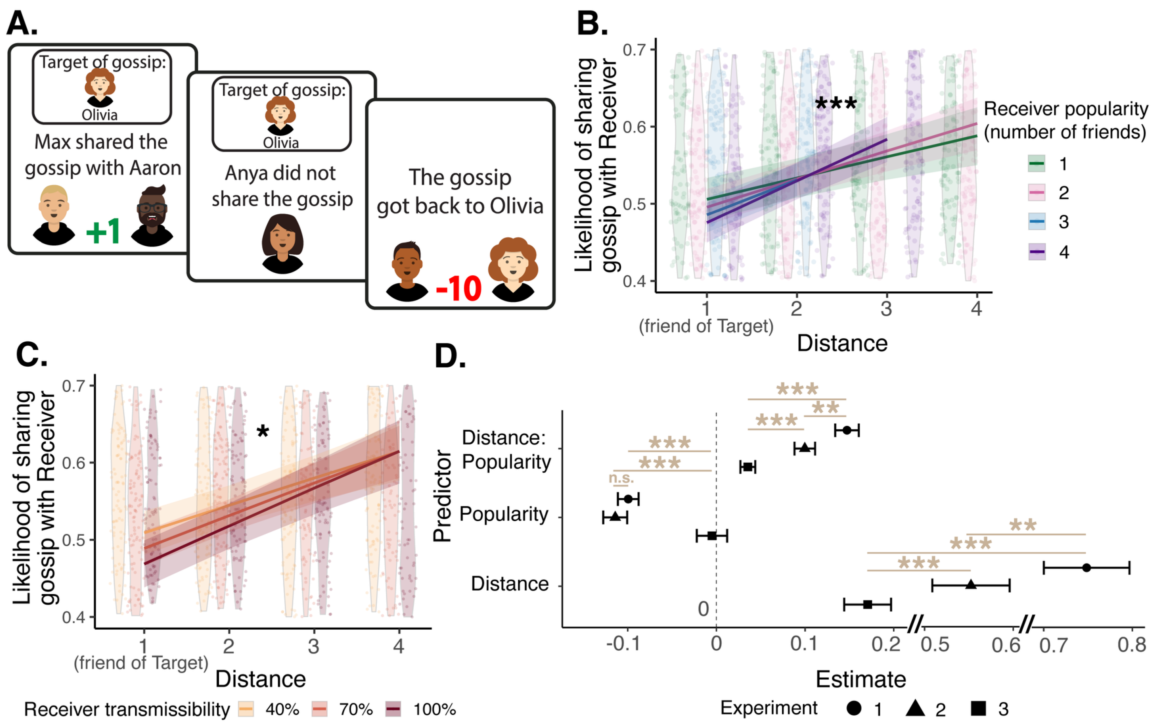

Figure 3 | (A) Experiment 3 learning task. *Participants simultaneously learned about friendships and network members' transmissibility by passively observing gossiping behavior in pairs. Transmissibility was reflected in the number of times a Receiver (Max) shared gossip with a friend (Aaron). The outcome was positive (+1) when gossip was shared with anyone who was not the current Target (Olivia), and neutral if gossip was not shared. If the friend was the Target, then sharing gossip with that person was considered a negative outcome (-10). Note: Points were provided only to help participants learn the outcomes associated with different types of gossip behavior; they did not contribute to monetary rewards.* ***(B) Likelihood of gossiping in Experiment 3:*** *Distance (x-axis) and popularity (color) are used to share gossip (y-axis) less with popular individuals close to the Target, and more with popular individuals who are far from the Target (Distance x Receiver Popularity interaction, ***p < .001).* ***(C)*** *The significant conditional relationship between distance (x-axis) and transmissibility (color) shows that gossip is shared (y-axis) less often with more transmissible individuals close to the Target, but this effect diminishes with distance (Distance x Transmissibility interaction, *p < .05). Lines represent regression model predictions, with 95% confidence intervals. Participant-level raw data shown as points. Distribution of data shown as violins.* ***(D) How does distance and popularity change with increasing learning demands?*** *Z-tests compare beta coefficients (x-axis) across experiments for each variable (y-axis). Different shapes correspond to different experiments. Each shape represents the beta coefficient estimate for a given effect, with standard error bars. As the learning environment becomes more challenging (Experiments 1 through 3), the effects of distance and popularity weaken (i.e., beta coefficients approach 0 on the x-axis), although participants' ability to use these two latent network properties to strategically gossip largely persists (except in the case of popularity in Experiment 3). All coefficients* *significantly differ from each other (***p < .001, **p < .01), except for the popularity coefficient between Experiments 1 and 2.*

In this more challenging learning environment, participants had far more difficulty integrating multidimensional information to gossip strategically. Although we observed an effect of distance, popularity, and transmissibility in a similar direction as the previous two experiments, once people were required to incorporate all three pieces of knowledge at the same time, the three-way interaction no longer robustly predicted decisions to gossip ($\beta_{distance*popularity*transmissibility}$ = 0.05, 95% CI [-0.02, 0.12], z = 1.36, p = 0.175, OR = 1.05). However, when considering only the conditional relationship between transmissibility and distance, as well as between transmissibility and popularity on their own (i.e., removing the three-way interaction term from the regression), we again observe a persistent relationship between distance and popularity ($\beta_{distance*popularity}$ = 0.04, 95% CI [0.02, 0.05], z = 4.22, p < 0.001, OR = 1.04; Fig. 3B; Supplementary Table 3), such that the effect of popularity on gossip decisions shifts as distance from the Target increases (contrast: $\Delta[b_{popularity_{distance1}}, b_{popularity_{distance4}}]$ = -0.11, 95% CI [-0.17, -0.05], z = -4.22, p < 0.001, OR = 0.89; Supplementary Table 5). The interaction between distance and popularity again holds when comparing high and low levels of popularity ($\beta_{distance*popularity_{high}}$ = 0.09, 95% CI [0.04, 0.13], z = 4.06, p < 0.001, OR = 1.09; Supplementary Table 7).

Even when learning in this more challenging environment, transmissibility fine-tunes decisions to share gossip with Receivers, depending on distance from the Target ($\beta_{distance*transmissibility}$ = 0.09, 95% CI [0.02, 0.16], z = 2.43, p = 0.015, OR = 1.09; Fig. 3C; Supplementary Table 3). Gossip is shared less with more transmissible Receivers who are close to the Target (simple slope: $b_{transmissibility@distance1}$= -0.27, 95% CI [-0.55, 0.01], z = -2.32, p = 0.021, OR = 0.76; Supplementary Table 9). This effect diminishes with distance (simple slope: $b_{transmissibility@distance4}$= 0.004, 95% CI [-0.32, 0.33], z = 0.03, p = 0.976, OR = 1.00), suggesting that transmissibility is most consequential for deciding how to gossip with the Target's immediate friends. Beyond that, the influence of transmissibility is minimal, highlighting the selective use of this person-specific information. While both popularity and transmissibility information influence gossip-sharing decisions, the challenging learning environment in Experiment 3 may have hampered people's ability to fully integrate these pieces of information in order to gossip strategically.

The goal of disseminating gossip widely while ensuring the Target never hears about it presents a challenge, as these two factors are in direct opposition to one another. On one hand, people disseminate information widely by sharing gossip with popular individuals, therefore treating those individuals as leaky conduits for information flow. On the other hand, to prevent gossip from reaching the Target, people are less likely to gossip with highly transmissible individuals who are friends with the Target. Regardless of an individual's transmissibility, we find that in all three controlled experiments, participants consistently extract the same topological network properties—distance and popularity—from their learning environments in order to gossip strategically.

However, these results also illustrate that the amount of social information and learning format impact participants' ability to maximally rely on these structural properties when building their network representations. Given the limitations of information processing, attending to multiple sources of information is likely to come at a cost of having a precise mental model of the network. To better understand how information load and learning format affect people's ability to use distance and popularity, we used z-tests to compare the coefficient estimates common across the regression models in the three experiments. Result reveal significant differences in how people use distance, and the distance and popularity interaction, depending on the learning environment (Fig 3D; Supplementary Table 10). In other words, information load shapes the representations people build. While increased information load reduces reliance on latent network structure (in Experiments 2 and 3), the ability to leverage these properties for adaptive gossiping remains largely intact, even when we control for the Target's popularity.

d. Experiment 4: Real-world network topology shapes inferred information flow

Unlike controlled experiments, however, tracing information spread among individuals in the real world is far more complex, not least because there are potentially hundreds of dynamically evolving connections that must be learned and tracked. Therefore, to demonstrate that knowledge of latent properties of social networks is functionally adaptive for real-world social navigation, we conducted an even stronger test of distance and popularity by leveraging the unique experience of college freshmen who are embedded in a new social environment, where knowledge about the relationships around them is intricately tied to personally relevant social stakes, and where the structure of the social network is itself rapidly changing. In Experiment 4, we first queried freshmen undergraduates (N=187) across three university dormitories about their friendships within the first two months of their first semester (see Methods), enabling us to build a ground-truth social network of mutual friendships between pairs of individuals (N=180) in this new and emerging network (Fig. 4A). We then asked a subset of participants (N=100) to judge the likelihood that someone in their network (Receiver) would hear news shared by someone else (Source), thus mimicking the challenge posed to our participants in Experiments 1-3, but without the explicit goal of disseminating information widely or restricting its reach (due to ethical considerations; Fig. 4B).

Importantly, we gave each participant a tailored set of stimuli consisting of network members in their immediate vicinity (see Methods), which allowed us to measure how knowledge about the topology of one's own social network shapes inferences of information flow between individuals. In short, this task allowed us to measure the network properties that individuals in the real world use to predict information flow, as a proxy of their knowledge of the relationships in their own network.

We find that participants' use of latent network properties is remarkably consistent with results from our controlled experiments. Similar to before, participants rely primarily on distance and popularity information to make inferences about whether news would travel between two individuals ($\beta_{distance*popularity}$ = -0.01, 95% CI [-0.02, -0.01], z = -3.74, p < 0.001, OR = 0.99; Fig. 4C; Supplementary Table 4). Specifically, when the Source and Receiver are closer to each other in network, participants believe it is likely that the Receiver will hear shared news, especially if the Source is popular (simple slope: $\beta_{popularity@-1SDmeanDistance}$ = 0.01, 95% CI [-0.001, 0.02], z = 2.09, p = 0.037, OR = 1.01; Supplementary Table 5). However, at the farther distances, the Source's popularity no longer informs perceptions about whether news will reach the Receiver ($\beta_{popularity@+1SDmeanDistance}$ = -0.01, 95% CI [-0.02, 0.004], z = -1.68, p = 0.094, OR = 0.99), suggesting that beliefs about information flow in a real-world social network are constrained by social proximity.

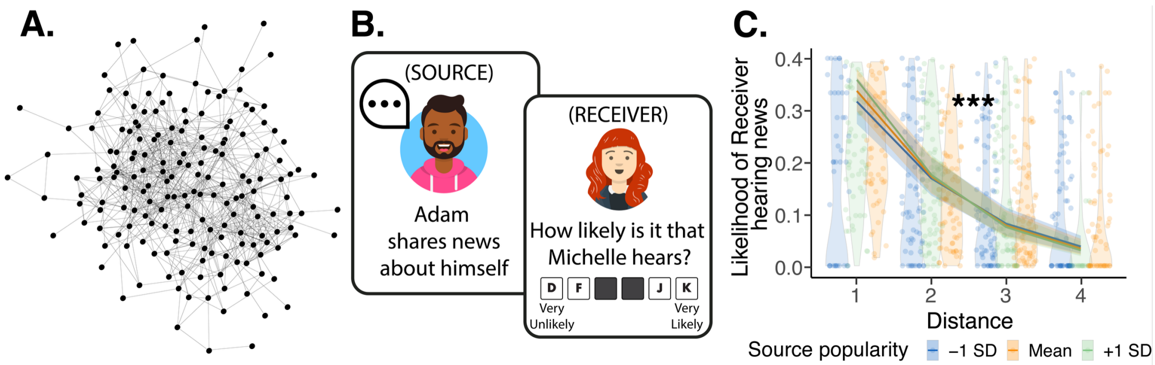

Figure 4* | *Predictions about information flow in a large, real-world social network (Exp. 4).

(A) *A social network consisting of 187 first-year undergraduate students.* ***(B)*** *A subset of individuals in the network (N=100) completed a task in which they rated the likelihood that someone in their network (e.g., Receiver: Michelle) would hear news about someone else (e.g., Source: Adam).* ***(C) Likelihood of hearing news.*** *Participants believe that the Receiver is more likely to hear news from the Source (y-axis) when those two individuals are closer to each other in the network (x-axis) and when the Source is more popular than the average individual in the network (green) (Distance x Receiver Popularity interaction, ***p < .001), which parallel our findings in Experiments 1-3 (datapoints at mean popularity +1 SD and distance = 5 are sparse and therefore removed for visual clarity). Lines represent regression model predictions, with 95% confidence intervals. Participant-level raw data shown as points. Distribution of data shown as violins.*

While it has been shown that people encode closeness and social status from an egocentric perspective in a social network[26,27], the way in which behavior unfolds based on one's perceptions of their network from an allocentric perspective has been under-explored. Here, we reveal that when tasked with inferring the dissemination of news between people in the network, individuals rely on a flexible combination of latent network properties to solve this challenge—a task particularly daunting in the real world, given the number of possible connections that must be learned. Despite differences between the demands of inferring information flow in the wild and the dueling goals of the Gossip Task used in our controlled laboratory experiments, leveraging knowledge about distance and popularity appears to be a general strategy to predict information trajectories through parts of a social network that one may have never experienced

Across four experiments, our results highlight the fundamental role that distance and popularity play in shaping perceptions of how information—gossip or otherwise—traverses through the network. Remarkably, these two latent network properties consistently emerge from different formats of noisy learning and in spite of imperfect memory. As soon as people are exposed to friendships within a group of individuals, they appear to stitch those discrete pieces of information to build a representation of an entire network, and importantly, that representation allows them to navigate its structure without necessarily relying on explicit knowledge of individual ties. Regardless of whether people are tasked with a specific communication goal, distance and popularity may be inherently prioritized as two pieces of topological knowledge ingrained in the process of forming a network representation.

e. Leaky cascade model accounts for strategic gossiping

While we show that people use topological properties of the network (distance and popularity), in combination with individual-level features of network members (transmissibility), to disseminate gossip while avoiding the target of gossip, it is unclear what type of mental representation of the network people are using to solve this challenging problem. In other words, what format of representation affords people the ability to make rapid and flexible inferences about the long-range consequences of spreading gossip, based on others' connections within the network?

Given knowledge of the discrete ties between individuals (i.e., a social graph), and the pathways that exist in the network, one may use an exhaustive planning process to flexibly work out how gossip could spread between any two individuals. However, such planning would be extremely resource-intensive and slow to execute[28–30]. On the other hand, through experience, people could rely on slowly and painstakingly learning the likelihood of gossip propagating along different pathways using a model-free learning mechanism, which would enable more rapid but largely inflexible assessments[31]. A third possibility is that, in order to trade off flexibility and speed, people instead combine elements of both to build a representation that encodes the cached likelihood of information traveling between any pair of individuals. We focus on two candidate models that embody distinct hypotheses about the nature of this information-caching process and the format of representation that supports strategic gossiping.

The first model is based on the concept of the *successor representation (SR)* from the field of reinforcement learning, which stores knowledge useful for solving sequential decision-making problems by encoding the probability of where an agent may end up in a graph given a starting point[32–35]. The SR not only captures mental abstractions across other domains, such as spatial navigation[33,36]and reward learning[34], but it has also been shown to predict people's inferences about friendships[37,38]. This type of 'social map' can be built with a multistep abstraction mechanism applied to knowledge of pairwise ties, providing predictions about the likelihood of information reaching each individual in a social network[35,39,40]. By integrating over piecemeal knowledge about pairs and chaining together paths that information might traverse, the SR would provide a parsimonious account of the fuzzy mental representations of social networks people build to support flexible social navigation[37,41].

However, in the case of strategic gossip, a potential weakness of the SR is that it models the dynamics of information spreading as a series of walks through the network, from one time step (Time 1) to the next (Time 2; Fig. 5A left). Information can only travel through the network along one chain of relationships at a time, akin to a person walking through a neighborhood. Yet, a notable feature of gossip is that it can spread to multiple people simultaneously and propagate through the network via multiple paths (Fig. 5A right). Critically, our behavioral results suggest that people are keenly attuned to this feature of gossip: the popularity of individuals, which can be re-conceptualized as the number of network paths that pass through them, factors heavily into whether people decide to share gossip.

We therefore consider a second model that captures such dynamics—Katz criterion—which measures the relative influence of each individual in the network based on every other individual's connections (see Methods). Just like the SR, the Katz criterion can be derived from knowledge of pairwise ties in a network. However, in contrast to the SR's assumptions of serial walks, the Katz criterion instead assumes that information flows through a leaky cascade[42]: from a starting node, information propagates by making copies of itself and then simultaneously cascading through all connected paths in parallel, thereby distributing gossip quickly to close contacts, with its influence diminishing as it moves to more remote nodes in the network. Consequently, if information flow through networks is more akin to leaky cascades than serial walks, then a mental representation that captures these dynamics—especially if both distance and popularity information can be reflected in a single value—may be more effective for making strategic decisions to gossip. While the Katz model is theoretically more suitable for capturing the dynamics of gossip flowing in a network, whether Katz actually captures human mental representations is an open empirical question.

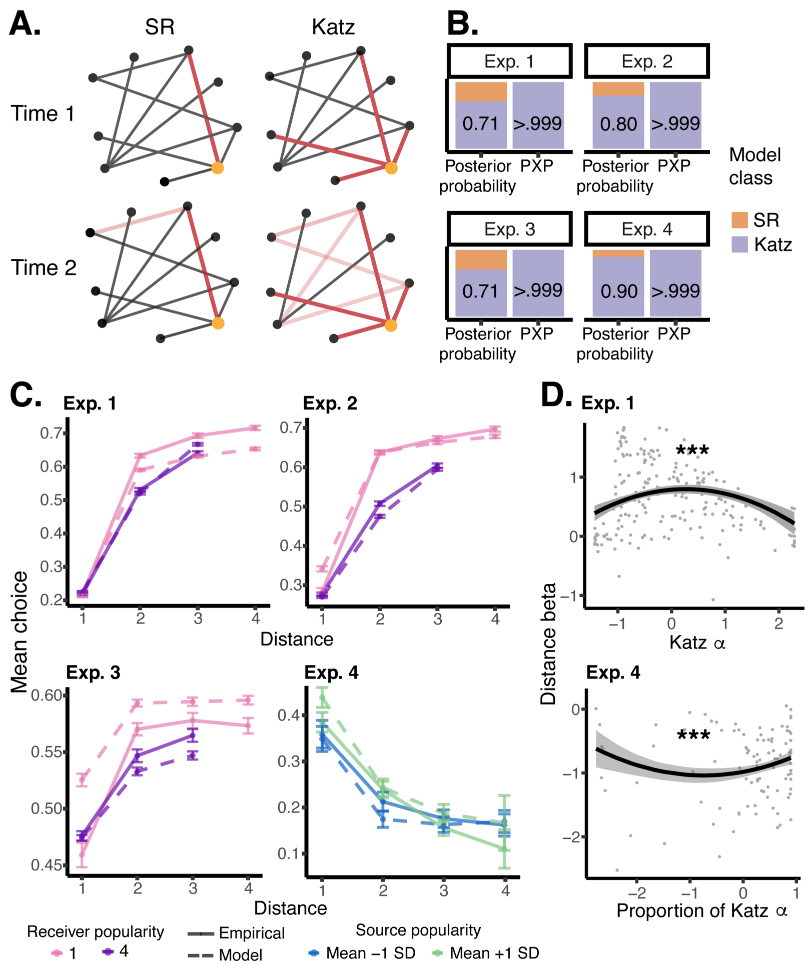

Figure 5* | *Computational modeling results*. *(A) *Toy illustration of the successor representation (SR) model (left) vs. the Katz model (right) at the first time point of information dissemination (Time 1) and the second (Time 2). Yellow dots represent the origin node.* ***(B)*** *Bayesian Model Selection demonstrates that the Katz model consistently exhibits higher posterior probabilities and achieves PXPs greater than 0.99 across all experiments, indicating a superior fit compared to the SR model.* ***(C)*** *Across all four experiments, posterior predictive checks from the Katz model recapitulate behavioral data (dashed lines represent the model's estimated likelihood of sharing gossip, and solid lines represent empirical likelihood judgments (y-axis); data are averaged over minimum and maximum popularity levels).* ***(D)*** *Curvilinear relationship between Katz alpha (α) parameter and beta coefficients of the distance term from our regression models suggest that intermediate degrees of abstraction of network structure may be optimal for leveraging distance to evaluate information flow (Katz α x Distance beta interaction, ***$p < .001$). Lines represent regression model predictions, with 95% confidence intervals. Points are participant-level regression beta/ α parameter estimates.*

Which mechanistic model provides a better explanation for how people succeed in disseminating gossip widely without the target of gossip finding out? While we observed individual differences in model fits for specific variants of both the SR and Katz models (see

Supplementary Fig. 2), protected exceedance probabilities (PXP) indicate that the majority of participants across all four experiments is best fit by the Katz model (Fig. 5B; Methods). Posterior predictive checks reveal that this cascade model closely aligns with patterns observed in the empirical data, capturing the network properties most critical for gossip-sharing decisions and beliefs about information flow (Fig. 5C). Mean within-subject correlations between empirical and predicted data were significantly positive, as confirmed by one-sample *t*-tests, indicating that the best-fitting Katz model variant in each experiment captures individual judgments about information flow (Exp. 1: mean $r = 0.75$, 95% CI [0.69, 0.81], $t(200) = 24.98$, $p < 0.001$; Exp. 2: $r = 0.74$, 95% CI [0.68, 0.80], $t(197) = 26.16$, $p < 0.001$; Exp. 3: $r = 0.56$, 95% CI [0.49, 0.62], $t(168) = 16.39$, $p < 0.001$; Exp. 4: $r = 0.66$, 95% CI [0.60, 0.73], $t(97) = 18.95$, $p < 0.001$), though note that in Experiment 3, the model tended to overestimate the probability that participants would gossip with low-popularity Receivers. Overall, these findings suggest that people rely on a mental representation consistent with the Katz criterion to anticipate how information propagates through their social network.

Reliance on the Katz criterion implies that individuals integrate both direct and indirect relationships to form a cohesive representation of the network. In our model, the degree of this integration—or abstraction over the network structure—is captured by the Katz alpha parameter (α), which reflects the extent to which individuals consider longer paths when evaluating information flow. Higher α is associated with greater abstraction of the network; consequently, distant connections should be perceived as more easily reachable. We examined this possibility by using individuals' estimated Katz α values to predict the distance beta coefficients extracted from our regression. In Experiment 1, we observed that α was curvilinearly related to distance, showing an inverted U-shaped relationship (β_{α^2} = -2.20, 95% CI [-3.00, -1.40], $t(196) = -5.28$, $p < 0.00r = -0.35$; Fig. 5D top; Supplementary Table 14). This result suggests that individuals who engage in moderate abstraction (reflected by intermediate values of α) may be best equipped to use information about distance for strategic gossiping, and thus able to balance dissemination of gossip with minimizing the risk of it reaching the Target. Conversely, individuals who abstract to a greater degree (higher α) may perceive the network as overly compressed, and those who abstract less (lower α) may rely primarily on direct connections—in both cases, the effect of using longer paths for gossip dissemination is diminished. This inverted U-shaped relationship was replicated in Experiment 2 (β_{α^2} = -3.81, 95% CI [-4.90, -2.70], $t(191) = -6.95$, $p < 0.001$, $r = -0.45$), while in Experiment 3, it was not significant (β_{α^2} = -0.34, 95% CI [-0.95, 0.27], $t(162) = -1.10$, $p = 0.271$, $r = -0.08$,). In Experiment 4, we observe an opposite (U-shaped) pattern (β_{α^2} = 1.07, 95% CI [0.45, 1.70], $t(92) = 3.44$, $p < 0.001$, $r = 0.33$; Fig. 5D bottom). Although the pattern is flipped (perhaps because participants were not required to use distance information in a strategic manner), the underlying conceptual relationship remains the same: minimal and excessive abstraction both result in reduced sensitivity to distance. This hints that there may be an optimal degree of abstraction that supports inferences about information flow. Together, these findings suggest that even though the model was not explicitly parameterized to encode topology, the information cached in the Katz model reflects sensitivity to key network properties for inferring information flow. The model's multistep integration process, driven by the alpha parameter, implicitly encodes this sensitivity, offering a computationally useful account of how individuals navigate social networks.

To model the effects of transmissibility in Experiments 2 and 3, we included a weight in our model's choice rule (w_T; see Methods), which quantifies the extent to which participants incorporate individual-level transmissibility information into their gossip-sharing decisions. We find a significant positive correlation between w_T and the beta coefficients of the three-way interaction between transmissibility, distance, and popularity in Experiment 2 ($r = 0.15$, 95% CI [0.00, 0.28], $t(196) = 2.06$, $p = 0.041$; Supplementary Fig. 3), suggesting that w_T effectively tracks individual differences in the degree to which transmissibility is integrated into gossip-sharing decisions. In contrast, in Experiment 3, where transmissibility information had to be learned alongside network structure, the effects of transmissibility were overall less robust (see Supplementary Table 3), and we did not observe any significant relationships with w_T.

3. Discussion

Gossip fundamentally shapes many aspects of our social lives, from transmitting cultural norms[43] to uncovering information that may otherwise be difficult to obtain first-hand[44]. It is used for a litany of social purposes, from maintaining positions in a social hierarchy[45–47] to promoting and sustaining cooperation[48–52], but in particular, gossip can be helpful for making more adaptive decisions (e.g., #MeToo whisper networks about sexual predators in Hollywood). Failing to gossip strategically can carry significant social costs, including loss of trust, social exclusion, and reputational damage[8,53]. People therefore often actively conceal gossip—which can shape how it spreads[1,54]—supporting the idea that gossip can be an indirect strategy for regulating social behavior. The need for strategic gossiping arises because individuals often do not know whether the target will learn the gossip's content, its source, or both. The consequences of gossip reaching the target stem from the uncertainty about being identified as the source—a consequence that carries considerable reputational risk. In short, sharing gossiping widely while avoiding the target of gossip is a common social phenomenon, and may even be a default form of how evaluations of others are communicated within a social group[4]. Despite this, we have little understanding of how strategic gossiping is achieved, even though it is the bedrock of so many social interactions. To reveal the socio-cognitive mechanisms that support strategic gossiping, we examined the type of social information people leverage to construct mental models of a social network, and then tested different computational models that might underlie the construction and use of such network representations.

We found that after learning who is friends with whom in a social network, people's gossip decisions are shaped by properties of the network that are not directly observable, but rather, must be inferred by stitching bits of information together that build out a holistic representation of the network—one that reveals its latent topological structures. With that knowledge in hand, people gossip more with popular individuals, but only when those individuals are located farther away from the Target of gossip, i.e., not direct friends of the Target (Experiment 1). By tracking a network member's popularity and taking advantage of connections that are far-reaching, people use distance and popularity to predict how information might propagate through the network. When presented with an additional piece of information about each individual's general tendency to gossip, i.e., transmissibility, people flexibly use this knowledge to gossip more with those who are likely to pass on the information and who are popular and not immediately connected to the Target (Experiment 2). When people learn about friendship and transmissibility information simultaneously, similar to how we learn about others in the real world, the effect of popularity and distance on gossip decisions is significantly attenuated (Experiment 3). While there were no explicit and direct costs of gossip reaching the target in our studies, participants still exhibited sensitivity to network structure, suggesting that our observed effects would be even stronger in real-world settings. Even when there are no personal stakes to sharing gossip, specific network properties (distance and popularity) nevertheless shape perceptions of information flow and gossip behavior. We conceptually replicate these findings in a noisy, large-scale social network. People indicate that news shared by someone in the network would be more likely to reach those who are closer and more popular. This underscores the fact that network structure, which is not readily observable, provides critical information for inferring how information flows, even in real-world contexts (Experiment 4). Although we did not directly investigate who people chose to gossip with in this naturalistic setting, we examined a perhaps more fundamental ability to track how information flows in one's own personal network, which serves as a likely precondition for effective gossiping.

The belief that information will be amplified as a function of social proximity is consistent with network science principles that predict greater efficiency in information transfer along paths of least resistance[18,55,56]. Distance not only serves as the structural foundation of networks but is also useful for building mental representations of relational structure. In particular, it can act as a dial that tunes the influence of other properties, such as popularity, depending on the specific goal at hand. Although our main findings suggest that gossip-sharing increases linearly with distance, we also find that at far enough distances, this trend stabilizes, a finding that could be further explored in future studies by examining much larger networks. Our results also hint at why popular individuals are often considered to have greater access to information[57–60]: people strategically use them as key points for spreading gossip or news because of their capacity to propagate information. By choosing well-connected individuals to share information, people tap into the extensive social reach of those individuals, ensuring that the gossip flows through the network effectively. However, considering the nature of the goal we gave our participants, being a conduit for information spread can be a double-edged sword. Gossip disseminated through popular individuals will be better broadcasted through the network, but this increases the risk of gossip getting back to the Target. People's behavior aligned with this critical function of popularity and, more impressively, popularity was leveraged in a context-dependent manner, such that people choose to either disseminate or withhold gossip based on the specific individual they encountered. Furthermore, we found selective integration of transmissibility information when deciding how to gossip, suggesting that people's network representations also include perceptions of individual-level attributes. These findings reflect the tactical use of relevant individual-level features to propagate gossip to parts of the network that otherwise may not have otherwise been exposed to the information.

We note, of course, that gossiping is not always strategic in the way we have defined it here, nor did we probe situations in which there were outcomes of consequence (breach of trust or damage to reputational concerns)—largely for ethical reasons of tampering with relationships in a real-world network. There are many situations in which people only gossip with a few individuals in a targeted manner or actively refrain from gossiping altogether, for example, when there is strong interdependence between potential gossipers in a cooperative context[54]. Strategic use of gossip can also serve a variety of different goals, including promoting generosity towards individuals who may be future interaction partners[17]. The specific goal we test here is therefore only a sliver of possible functions of both strategic and non-strategic gossip. To build a comprehensive account of gossiping, future work would need to examine how people flexibly draw on their knowledge of network dynamics in different social environments where the stakes of these decisions vary.

What representational format allows people to compute how information swiftly propagates through a social network? We tested two classes of computational models to address competing hypotheses about how people infer information moves through communities. We pitted the SR—which builds up a cognitive map of relational information among network members through random walks[37,41], such that the network is seen as a series of probabilistic transitions from one person to the next—against an alternative psychological account of how information spreads: the Katz criterion, which embodies a set of propagation dynamics that is distinct from the SR. The walk-based dynamics of the SR requires individuals to pay attention to local connectivity and the immediate transitions between nodes, caching knowledge about the network's structure as it pertains to sequential information flow. While the SR can be flexibly adapted to different goals, it may be less efficient for capturing how information diffuses through multiple paths in the network at once. Katz criterion, on the other hand, reflects simultaneous flow and cascade-like dynamics, thus capturing how information can spread through many people simultaneously. Such a representational format aligns with how information can diffuse rapidly and broadly across a network, mirroring real-world phenomena[61].

Here, we find that to solve challenges associated with network information flow—such as gossip spreading widely while not getting into the wrong hands—people rely on a mental model of their network that emphasizes information cascades. This type of representation allows for fairly rapid evaluations of how likely it is for gossip to spread from one corner of the network to another, paralleling the amplification and cascade of initially arbitrary choices within animal swarms[62], information diffusion through social media networks[63], and the swift adoption of customs and beliefs, even when based on inaccurate information[64]. The fact that our participants had relatively poor memory of specific pairs of friends, but were still able to exhibit strategic gossiping behaviors, aligns with field work illustrating that individuals correctly identify a link between two fellow villagers about one-third of the time[65] and yet can still predict information transmission through the network[61,66]. Since cascade processes are naturally shaped by network structure[67], such reliance on a network's topology—rather than explicit memory of each friendship—likely drives the construction of a mental model that prioritizes information cascades. This is similar to the propagation dynamics seen in real world case studies, such as when early adopters of a new policy act as seeds that determine how quickly others follow suit[61]. We speculate that learning network topology allows us to build representations that mimic systems-level diffusion dynamics in service of individual-level social goals.

An important remaining consideration is whether representations such as the Katz criterion can be incrementally learned and effectively adapt to changes in network structure. Unlike the SR, which readily supports trial-and-error learning through updating methods such as temporal difference learning, it remains unclear how people might learn the Katz criterion in the same manner[68]. The Katz criterion may be particularly advantageous in stable networks where frequent updates are unnecessary. However, in dynamic environments—such as real-world social networks that evolve over time—the challenge lies in how a Katz-like representation balances ease of updating with the ability to efficiently integrate information across multiple paths to adapt to structural changes. Understanding how individuals balance these competing demands could offer valuable insight into *how* people incrementally construct mental models for decision-making in complex social spaces, a potential avenue for future research.

While the role of network topology in shaping the diffusion of social phenomena has been extensively studied in network science[19–21,69], the mechanism by which humans track information flow or leverage topological knowledge in the service of strategic decision-making remain poorly understood. Here, we offer a concrete account of how individuals efficiently solve the problem of disseminating information both widely and in a targeted way. To build a comprehensive account of gossiping, future work should examine how people flexibly draw on these mechanisms in different social environments where the stakes of such interpersonal decisions vary.

4. Methods

The study protocol was approved by Brown University's Institutional Review Board, and informed consent was obtained from all participants prior to each experiment.

Participants.

Participants in our three online experiments were recruited via Prolific and location was restricted to the U.S. to minimize potential confounding factors related to language and culture. In Experiment 1, we recruited 202 participants (108 female; 60 non-White or mixed race; mean age = 38.20, *SD* ± 13.79). In Experiment 2, we recruited 199 participants (97 female; 64 non-White or mixed race; mean age = 39.20, *SD* ± 13.47). In Experiment 3, we recruited 220 participants (113 female; 64 non-White or mixed race; mean age = 37.44, SD ± 12.75). To ensure that participants in Experiment 3 were paying attention as they learned the network through passive observation, we included catch trials that required participants to press a key whenever they saw an upside-down image of a face. Data were excluded from participants who missed more than 20% of catch trials during learning (N=38), whose data were not recorded properly (N=12). Across all three experiments, data were also excluded from participants whose responses in the Gossip Task never varied (N=1 in each experiment). Our final sample size was therefore 201 participants in Experiment 1, 198 participants in Experiment 2, and 169 participants in Experiment 3.

In Experiment 4, we asked 187 students from the freshmen cohort at a university in the U.S. (100 female; 140 non-White or mixed race; mean age = 18.06, SD ± 0.42) to report their friendships, thus allowing us to construct a ground-truth social network. Within this sample, a subset of participants (N=100) completed the Information Flow Task in-lab. Two participants were removed from the model fitting procedure (described below) because they completed the Information Flow Task without completing the original friendship survey. Therefore, we were only able to appropriate model the data of 98 participants.

Procedural details.

No statistical methods were used to pre-determine sample sizes, but our sample sizes are similar to those reported in previous publications[37,38] . Data collection and analysis were not performed blind to the conditions of the experiments. Online participants were paid at a rate of $8 per hour and they could earn up to a $5 bonus based on their performance during the study. Participants in the on-campus experiment were paid $15 per hour.

Artificial social network.

We created an artificial social network consisting of nine nodes representing the network members and 12 edges representing mutual friendships (Fig. 1A). This configuration allowed us to manipulate and orthogonalize network properties such as popularity (number of friends) and distance (shortest path between two members), which are latent properties of the network that can either increase or decrease the spread of information. We used the same network configuration across all three in-lab experiments. For each participant in each experiment, a photograph of a face from the Chicago Face Database[70] was randomly assigned to a node in the network. The photos were evenly split between male and female faces, distributed as even as possible across four different racial/ethnic (Asian, Black, White, Latinx) categories.

Real-world social network.

Undergraduate students at a university in the U.S. completed an online questionnaire in which they reported their friendship status with all other participants (yes or no response). The network was constructed using these self-reported friendships, where a node represents an individual in the network and an edge a relationship between two individuals (Fig. 4A). To mimic the connections in our artificial social network, we used only mutual friendships (i.e., where both individuals mutually reported a relationship) to construct the ground truth network (N=180).

Network Learning Task.

In Experiment 1, participants only had access to information about who was friends with whom in the network. In a self-paced task, participants viewed one network member's photo while the others remained hidden by a gray square (Fig. 1B). On each trial, they clicked the gray squares one at a time to learn whether that individual is friends with the network member initially shown. The trial ended when all the friends of a given network member were revealed. This procedure was repeated for five rounds for each member in the network. In each round, the network member whose friendships participants learned about was randomly selected. In Experiment 2, we provided additional information about each network member's transmissibility, which was defined by how frequently an individual shared gossip with a friend (40%, 70%, or 100% of the time, represented by orange, red, and purple circles, respectively, in Fig. 1A). For example, the most transmissible individuals (100%) were those who always gossiped with a friend. Participants learned this individual-level information separately from learning about friendships (which was the same as in Experiment 1). In each block, participants learned about a randomly selected network member through trial and error, across 10 trials. In Experiment 3, participants passively observed how frequently the Target found out they were the topic of gossip (Fig. 3A). This had to be learned simultaneously alongside friendship information—akin to how we learn in the real world where multidimensional information is experienced as one coherent episode. Participants were instructed to remember how often each person shares gossip, and who is friends with whom. At the start of each block, participants were told that gossip emerged about a network member (the Target), and for all the trials within that block, they observed what all the other network members (the Receivers) did with that piece of gossip. Both the Target and the sequence of gossip exchanges were randomly selected for each participant. Again, as in Experiment 2, the three levels of transmissibility determined the frequency with which each Receiver shared gossip. To ensure that participants were able to learn the value associated with outcomes of different types of gossiping behavior, we provided a points structure on each trial to reinforce behaviors that align with the goal of the experiment (-10 when gossip reaches the Target, +1 if gossip is shared with a friend who is not the Target, and no points when gossip is not shared at all). These outcomes roughly aligned with the goal that participants were asked to achieve in the Gossip Task, described below. Participants were informed that the points corresponded only to positive or negative outcomes and did not translate into monetary rewards; they were solely intended to facilitate learning (see Supplementary Methods for task instructions).

Gossip Task (Experiments 1-3).

After learning the network, participants in all three experiments completed a gossip decision task in which they were hypothetically embedded within the network (Fig. 1C). Trials were organized into blocks with a constant Target, who was randomly selected network member. Within each block, the order of Receivers presented to participants was randomized. On each trial, participants decide how likely it is they will share gossip with a network member (Receiver) on a scale ranging from 0 to 100. The goal is to spread gossip about the Target as widely as possible through the network, while also preventing gossip from getting back to the Target. To fulfill the first part of the goal, people might share gossip with popular network members who are able to effectively broadcast information, but that increases the risk of gossip getting back to the Target. Therefore, to behave strategically, participants should also consider where an individual sits in the network, relative to the Target.

We used mixed-effects beta regression models to examine the effect of network properties (distance and popularity) and individual-level information (transmissibility) on participants' likelihood ratings. Ratings should scale upwards with the number of Receivers that a piece of gossip can reach without the Target hearing about it, which reflects optimal behavior.

Information Flow Task (Experiment 4)

To assess the type of network knowledge that members of the real-world social network use to predict whether information would travel between two individuals, we first asked each participant to provide information about their own friendships through an online questionnaire. Using that data, we generated a unique set of stimuli, one for each participant, by randomly sub-sampling 25 network members from the participant's immediate neighborhood, ensuring that these members varied in their path distance from the participant (e.g., direct friends, friends-of-friends, etc.). At the start of each block in the Information Flow Task, we presented one individual from the stimulus set (the Source) who has hypothetically shared some news. On each trial within the block, participants used a four-point scale ranging from 'very likely' to 'very unlikely' to rate the likelihood that the other individuals (the Receivers) would hear about the news. Each individual was presented as the Source once. As in the Gossip Task, each block consisted of the same Source, randomly selected from the participant's stimulus set, and the order of Receivers within each block was randomized.

Behavioral analysis.

We used mixed-effects beta regressions to predict participants' ratings of how likely they would be to share gossip with the source. This type of regression is suitable for modeling response variables that are continuous and bounded between 0 and 1. Because beta regression models the mean of a continuous, bounded outcome using a link function (e.g., logit), it does not assume normally distributed residuals or equal variances. All regression models were implemented using the *glmmTMB* (v.1.1.10) package in R[71](v. 4.3.1). Statistical inference was based on two-tailed hypothesis tests with a significance threshold of $p < 0.05$. The main predictors in each model for Experiments 1-3 were the shortest path distance between a given Receiver and Target on each trial, and the popularity (number of friends) of the Receiver. For Experiments 2 and 3 only, we included a term for transmissibility. In all three regression models, we included memory accuracy for friendships in the network (as measured by d-prime), the Target's popularity, and the interaction between d-prime and popularity, as well as the interaction between d-prime and distance as control variables. In Experiment 4, we binarized the response variable (0 = 'Unlikely' or 'Very Unlikely'; 1 = 'Likely' or 'Very Likely') to enhance statistical power, especially since we were not focused on participants' confidence. We used a mixed-effects logistic regression with the same terms as in the beta regressions, but without a term for memory accuracy. Across all experiments, we included random intercepts for each main effect for each participant, and separately, random slopes for each interaction term, to control for potential correlations between the main effects and the interactions (see Supplementary Tables 1-4 for more information).

Computational models.

We built two classes of computational models to test the cognitive mechanisms underlying strategic gossiping. Each class consisted of two main components—a representation of the social network and a choice function. In each model, the network representation took the form of an n x n matrix given n nodes, representing n individuals in the network. Computational models were built using custom code in MATLAB 2024b. Below we describe the two classes of models (see Supplementary Tables 12-13 for full details about each model variant).

Successor representation (SR) model

The SR model consisted of a network representation given by:

$$M = (I - \gamma T)^{-1},$$

where M is a matrix of successor values that represent the likelihood of gossip reaching some network member, given its starting point; I is the identity matrix; and T is the transition matrix, which is the adjacency matrix (an *n x n* matrix containing a 1 where a friendship between two individuals exists, and 0 otherwise, where *n* is the number of network members) normalized by its row sums. This normalization ensures that each row of the transition matrix represents the probabilities of transitioning from each node to its neighboring nodes in the network. The SR gamma (γ) specifies how much participants discount the probability of gossip reaching network members farther away from the starting point. γ is bounded between 0 and 0.99, with a higher γ reflecting less discounting with distance.

In Experiment 4 only, participants were embedded in the social network. While we were able to construct the true adjacency matrix using self-reported friendships across all members of the network, we assume that individuals have greater access to parts of the network that are closer to them. Therefore, each participant's observed adjacency matrix would be shaped by their own egocentric vantage point in the network. We constructed these participant-specific adjacency matrices by starting with the true adjacency matrix (ground truth network) and assuming that the probability (p^d) that the participant knows about social ties in this true matrix exponentially *decreases* with participants' path distance *d* from the pair. This reflects our assumption that participants have less information about network members who are farther away. We estimated *p* for each participant, which was bounded between 0 and 1, to capture individual differences in discounting more distant connections. Since there was no reward or penalty associated with decisions, we did not compute a reward vector on each trial. Value was simply the successor value in matrix M from the Receiver to the Target.

On each trial, the value of sharing gossip with a Receiver i, given the Target j, is computed as the dot product of the vector of SR values corresponding to the Receiver, and the reward function (R), which is a vector of dimensions *1 x n*, with +1 for each node except for the Target *j,* for which we specified as a free parameter r_j:

$$V(i,j)_t = \overrightarrow{M(i)} \cdot \vec{R}$$

Katz representation model.

We additionally tested another class of models based on Katz criterion, which reflects the total number of paths between a pair of network members discounted by path length. The network representation in the Katz model is computed as:

$K = (I - \alpha A)^{-1}$,

where each cell of K represents the communicability strength between each pair of individuals in the network, i.e., the sum of the weighted number of all paths, of all lengths, between two specific nodes[72]. I is the identity matrix, while A is the true adjacency matrix for Experiments 1-3. The bounds on α ($0 < \alpha < \frac{1}{\lambda}$, where λ is the absolute value of the largest eigenvalue of A) ensure that the resulting matrix is invertible and converges to its inverse.

In Experiment 4, we constructed the adjacency matrix (A) based on data specific to each participant. This results in different α values for each participant, reflecting their unique connectivity within the real-world social network. Direct estimation of α as a free parameter is therefore impractical. Consequently, we estimate α_{prop}, which is a scaling factor multiplied by each participant's α, allowing us to capture individual-level sensitivity to Katz criterion scores. In Experiments 1-3, value was computed in the same way as for the SR model. Similarly, for Experiment 4, value was the Katz criterion between two individuals in the network.

The main difference between the Katz model and the SR is the use of the adjacency matrix instead of a transition matrix. While the SR matrix M summarizes the results of taking random walks discounted by the number of steps between the start and end points, the Katz matrix K summarizes the results of a cascade in which information copies itself and propagates along multiple channels simultaneously.

Choice rule.

The logistic function was used as a choice rule to translate value on each trial t from each model into a probabilistic decision. In the most complex variants of each model, we included parameters that reflect different sources of noise: k, which determines sensitivity to changes in value (slope of the logistic curve), and v_0, which determines the value at which gossip would be shared with a 50% probability (midpoint or inflection point of the curve):

$$p(share\ gossip)_t = \frac{1}{1 + e^{-(k * v_t - v_0)}}$$

To test the possibility that our empirical observations emerge from less complex computations, we altered this choice rule by setting k to 1 (effectively assuming that sensitivity to changes in value is constant across participants), setting v_0 to 0 (where probability of sharing gossip is 50%), or both. This resulted in 4 variants of each class of models for Experiments 1-3. Because the task in Experiment 4 simply required participants to report their beliefs about whether information would travel between two individuals in the network, we omitted the penalty parameter (*r*), thus resulting in 3 variants for each class of models (see Supplementary Tables 12-13).

To model transmissibility in Experiments 2 and 3, we included a weight (w_T) in the choice function that captured differences in participants' sensitivity to each network member's transmissibility:

$$p(share\ gossip)_{trial} = \frac{1}{1 + e^{-(k*(w_T*T+1)*v_{trial} - v_0)}}$$

T represents the ground-truth transmissibility of a given node (either 40%, 70%, or 100%) and w_T reflects participants' individual sensitivity to this information when making gossip-sharing decisions. A constant (1) is included to ensure that the choice probability remains above zero when w_T is small, thereby maintaining a baseline influence of the trial value. This adjustment allows the transmissibility term to scale the trial value multiplicatively, paralleling the interaction terms in our regressions, while preserving baseline contributions from Katz or the SR.

Model fitting and model comparison.

Model fitting was performed using the *fmincon* nonlinear optimization function in MATLAB to estimate free parameters for each participant 100 times, with random parameter initializations for each iteration. We kept only the parameter values associated with the minimum mean squared error (MSE). Since the data in Experiment 4 is binary, we used negative log-likelihood instead of MSE as the loss function. To compare the fits of the different computational models, we leveraged a two-stage approach. As a primary comparison, we first focused on models that differed in their format of representation (Katz or SR). To do this, we first computed the Schwarz weights of each model variant we fitted[73] and summed these weights for models that used the same format of representation (Katz or SR). This generated the relative model evidence for each model class, which we then used to compute the protected exceedance probability (PXP)[74] over model classes to infer whether participants' representations were more Katz-like or SR-like. We then conducted comparisons among different variants of the Katz models by computing PXPs over evidence for each specific variant, providing further insights into the relative performance of these variants (see Supplementary Fig. 2). Parameter recovery analyses confirmed that key parameters in these variants were reliably identifiable (see Supplementary Fig. 5). Model confusion analyses further confirmed that data simulated from the Katz and SR models could be recovered by the correct model across all four experiments (see Supplementary Fig. 6). To assess how well the model captured participants' data, we simulated data using parameters from the best-fitting model for each experiment and calculated the within-subject correlation between empirical and observed data across levels of distance, popularity, and transmissibility (Exp. 2 and 3).

References

1. Cruz, T. D. D., Nieper, A. S., Testori, M., Martinescu, E. & Beersma, B. An Integrative Definition and Framework to Study Gossip. *Group Organ. Manag.* **46**, 252–285 (2021).
2. Feinberg, M., Willer, R. & Schultz, M. Gossip and Ostracism Promote Cooperation in Groups. *Psychol. Sci.* **25**, 656–664 (2013).
3. Sommerfeld, R. D., Krambeck, H.-J. & Milinski, M. Multiple gossip statements and their effect on reputation and trustworthiness. *Proc Royal Soc B Biological Sci* **275**, 2529–2536 (2008).
4. Blumberg, H. H. Communication of interpersonal evaluations. *J Pers Soc Psychol* **23**, 157–162 (1972).
5. Grosser, T. J., Lopez-Kidwell, V., Labianca, G. (Joe) & Ellwardt, L. Hearing it through the grapevine Positive and negative workplace gossip. *Organ. Dyn.* **41**, 52–61 (2012).
6. Ellwardt, L., Labianca, G. (Joe) & Wittek, R. Who are the objects of positive and negative gossip at work? A social network perspective on workplace gossip. *Soc. Netw.* **34**, 193–205 (2012).
7. Robbins, M. L. & Karan, A. Who Gossips and How in Everyday Life? *Soc. Psychol. Pers. Sci.* **11**, 185–195 (2020).
8. Molho, C., Tybur, J. M., Lange, P. A. M. V. & Balliet, D. Direct and indirect punishment of norm violations in daily life. *Nat. Commun.* **11**, 3432 (2020).
9. Wu, J., Balliet, D. & Lange, P. A. M. V. Gossip Versus Punishment: The Efficiency of Reputation to Promote and Maintain Cooperation. *Sci. Rep.* **6**, 23919 (2016).
10. Wu, J., Balliet, D. & Lange, P. A. M. V. Reputation, Gossip, and Human Cooperation. *Soc. Pers. Psychol. Compass* **10**, 350–364 (2016).
11. Ellwardt, L., Steglich, C. & Wittek, R. The co-evolution of gossip and friendship in workplace social networks. *Soc Networks* **34**, 623–633 (2012).
12. Hess, N. H. & Hagen, E. H. The Oxford Handbook of Gossip and Reputation. 274–302 (2019) doi:10.1093/oxfordhb/9780190494087.013.15.
13. Tassiello, V., Lombardi, S. & Costabile, M. Are we truly wicked when gossiping at work? The role of valence, interpersonal closeness and social awareness. *J. Bus. Res.* **84**, 141–149 (2018).
14. Burt, R. S. & Knez, M. Kinds of Third-Party Effects on Trust. *Ration. Soc.* **7**, 255–292 (1995).
15. Grosser, T. J., Lopez-Kidwell, V. & Labianca, G. A Social Network Analysis of Positive and Negative Gossip in Organizational Life. *Group Organ Manage* **35**, 177–212 (2010).
16. Carrasco, J. A., Miller, E. J. & Wellman, B. How Far and with Whom Do People Socialize? *Transport Res Rec* **2076**, 114–122 (2008).
17. Wu, J., Balliet, D. & Lange, P. A. M. V. Reputation management: Why and how gossip enhances generosity. *Evol. Hum. Behav.* **37**, 193–201 (2016).
18. Newman, M. E. J. The Structure and Function of Complex Networks. *SIAM Rev.* **45**, 167–256 (2003).
19. Kleinberg, J. M. Navigation in a small world. *Nature* **406**, 845–845 (2000).
20. Barabási, A.-L. & Albert, R. Emergence of Scaling in Random Networks. *Science* **286**, 509–512 (1999).

21. Watts, D. J. & Strogatz, S. H. Collective dynamics of 'small-world' networks. *Nature* **393**, 440–442 (1998).
22. Freeman, J. B. & Ambady, N. A Dynamic Interactive Theory of Person Construal. *Psychol. Rev.* **118**, 247–279 (2011).
23. Ambady, N. & Rosenthal, R. Thin slices of expressive behavior as predictors of interpersonal consequences: A meta-analysis. *Psychol. Bull.* **111**, 256–274 (1992).
24. Rouse, W. B. & Morris, N. M. On Looking Into the Black Box: Prospects and Limits in the Search for Mental Models. *Psychol. Bull.* **100**, 349–363 (1986).
25. Sweller, J. Cognitive Load During Problem Solving: Effects on Learning. *Cogn. Sci.* **12**, 257–285 (1988).
26. Parkinson, C., Kleinbaum, A. M. & Wheatley, T. Spontaneous neural encoding of social network position. *Nat Hum Behav* **1**, 0072 (2017).
27. Courtney, A. L. & Meyer, M. L. Self-Other Representation in the Social Brain Reflects Social Connection. *J Neurosci Official J Soc Neurosci* **40**, 5616–5627 (2020).
28. Doll, B. B., Duncan, K. D., Simon, D. A., Shohamy, D. & Daw, N. D. Model-based choices involve prospective neural activity. *Nat. Neurosci.* **18**, 767–772 (2015).
29. Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P. & Dolan, R. J. Model-Based Influences on Humans' Choices and Striatal Prediction Errors. *Neuron* **69**, 1204–1215 (2011).
30. Huys, Q. J. M. *et al.* Bonsai Trees in Your Head: How the Pavlovian System Sculpts Goal-Directed Choices by Pruning Decision Trees. *Plos Comput Biol* **8**, e1002410 (2012).
31. Drummond, N. & Niv, Y. Model-based decision making and model-free learning. *Curr. Biol.* **30**, R860–R865 (2020).
32. Dayan, P. Improving Generalization for Temporal Difference Learning: The Successor Representation. *Neural Comput* **5**, 613–624 (1993).
33. Momennejad, I. *et al.* The successor representation in human reinforcement learning. *Nat Hum Behav* **1**, 680–692 (2017).
34. Geerts, J. P., Gershman, S. J., Burgess, N. & Stachenfeld, K. L. A Probabilistic Successor Representation for Context-Dependent Learning. *Psychol. Rev.* (2023) doi:10.1037/rev0000414.
35. Stachenfeld, K. L., Botvinick, M. M. & Gershman, S. J. The hippocampus as a predictive map. *Nat Neurosci* **20**, 1643–1653 (2017).
36. Garvert, M. M., Dolan, R. J. & Behrens, T. E. A map of abstract relational knowledge in the human hippocampal–entorhinal cortex. *Elife* **6**, e17086 (2017).
37. Son, J.-Y., Bhandari, A. & FeldmanHall, O. Cognitive maps of social features enable flexible inference in social networks. *Proc National Acad Sci* **118**, e2021699118 (2021).
38. Son, J.-Y., Vives, M.-L., Bhandari, A. & FeldmanHall, O. Replay shapes abstract cognitive maps for efficient social navigation. *Nat. Hum. Behav.* **8**, 2156–2167 (2024).
39. Momennejad, I. Learning Structures: Predictive Representations, Replay, and Generalization. *Curr Opin Behav Sci* **32**, 155–166 (2020).
40. Behrens, T. E. J. *et al.* What Is a Cognitive Map? Organizing Knowledge for Flexible Behavior. *Neuron* **100**, 490–509 (2018).
41. Son, J.-Y., Vives, M.-L., Bhandari, A. & FeldmanHall, O. Replay shapes abstract cognitive maps for efficient social navigation. *bioRxiv* 2023.12.19.572418 (2023) doi:10.1101/2023.12.19.572418.

42. Zamora-López, G. & Gilson, M. An integrative dynamical perspective for graph theory and the analysis of complex networks. *Chaos: Interdiscip. J. Nonlinear Sci.* **34**, 041501 (2024).
43. Baumeister, R. F., Zhang, L. & Vohs, K. D. Gossip as Cultural Learning. *Rev Gen Psychol* **8**, 111–121 (2003).
44. Jolly, E. & Chang, L. J. Gossip drives vicarious learning and facilitates social connection. *Curr Biol* **31**, 2539-2549.e6 (2021).
45. Thornborrow, J. & Morris, D. Gossip as strategy: The management of talk about others on reality TV show 'Big Brother.' *J. Socioling.* **8**, 246–271 (2004).
46. McAndrew, F. T., Bell, E. K. & Garcia, C. M. Who Do We Tell and Whom Do We Tell On? Gossip as a Strategy for Status Enhancement1. *J. Appl. Soc. Psychol.* **37**, 1562–1577 (2007).
47. Besnier, N. Information withholding as a manipulative and collusive strategy in Nukulaelae gossip*. *Lang. Soc.* **18**, 315–341 (1989).
48. Wu, J. *et al.* Honesty and dishonesty in gossip strategies: a fitness interdependence analysis. *Philos. Trans. R. Soc. B* **376**, 20200300 (2021).
49. Cruz, T. D. D., Beersma, B., Dijkstra, M. T. M. & Bechtoldt, M. N. The Bright and Dark Side of Gossip for Cooperation in Groups. *Front. Psychol.* **10**, 1374 (2019).
50. Pan, X., Hsiao, V., Nau, D. S. & Gelfand, M. J. Explaining the evolution of gossip. *Proc. Natl. Acad. Sci.* **121**, e2214160121 (2024).
51. Fonseca, M. A. & Peters, K. Will any gossip do? Gossip does not need to be perfectly accurate to promote trust. *Games Econ. Behav.* **107**, 253–281 (2018).
52. Dunbar, R. I. M. Gossip in Evolutionary Perspective. *Rev. Gen. Psychol.* **8**, 100–110 (2004).
53. Balliet, D., Molho, C., Columbus, S. & Cruz, T. D. D. Prosocial and punishment behaviors in everyday life. *Curr. Opin. Psychol.* **43**, 278–283 (2022).
54. Giardini, F. & Wittek, R. P. M. Silence Is Golden. Six Reasons Inhibiting the Spread of Third-Party Gossip. *Front. Psychol.* **10**, 1120 (2019).
55. Granovetter, M. S. The Strength of Weak Ties. *Am J Sociol* **78**, 1360–1380 (1973).
56. Milgram, S. The small-world problem. *Psychology Today* **1**, 61–67 (1967).
57. Borgatti, S. P. Centrality and network flow. *Soc. Netw.* **27**, 55–71 (2005).
58. Borgatti, S. P., Mehra, A., Brass, D. J. & Labianca, G. Network Analysis in the Social Sciences. *Science* **323**, 892–895 (2009).
59. Burt, R. S. The Network Structure Of Social Capital. *Res. Organ. Behav.* **22**, 345–423 (2000).
60. Ellwardt, L., Steglich, C. & Wittek, R. The co-evolution of gossip and friendship in workplace social networks. *Soc. Netw.* **34**, 623–633 (2012).
61. Banerjee, A., Chandrasekhar, A. G., Duflo, E. & Jackson, M. O. The Diffusion of Microfinance. *Science* **341**, 1236498 (2013).
62. Couzin, I. D. Collective cognition in animal groups. *Trends Cogn. Sci.* **13**, 36–43 (2009).
63. Chen, X. *et al.* Prediction of retweet cascade size over time. *Proc. 21st ACM Int. Conf. Inf. Knowl. Manag.* 2335–2338 (2012) doi:10.1145/2396761.2398634.
64. Bikhchandani, S., Hirshleifer, D. & Welch, I. A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades. *J. Political Econ.* **100**, 992–1026 (1992).

65. Breza, E., Chandrasekhar, A. G. & Tahbaz-Salehi, A. SEEING THE FOREST FOR THE TREES? AN INVESTIGATION OF NETWORK KNOWLEDGE. *NBER Working Paper Series* (2018).
66. Bloch, F., Jackson, M. O. & Tebaldi, P. Centrality measures in networks. *Soc. Choice Welf.* **61**, 413–453 (2023).
67. Delvenne, J.-C., Lambiotte, R. & Rocha, L. E. C. Diffusion on networked systems is a question of time or structure. *Nat. Commun.* **6**, 7366 (2015).
68. Katz, L. A new status index derived from sociometric analysis. *Psychometrika* **18**, 39–43 (1953).
69. Newman, M. E. J. & Park, J. Why social networks are different from other types of networks. *Phys. Rev. E* **68**, 036122 (2003).
70. Ma, D. S., Correll, J. & Wittenbrink, B. The Chicago face database: A free stimulus set of faces and norming data. *Behav. Res. Methods* **47**, 1122–1135 (2015).
71. Brooks, M., E. *et al.* glmmTMB Balances Speed and Flexibility Among Packages for Zero-inflated Generalized Linear Mixed Modeling. *R J.* **9**, 378 (2017).
72. Benzi, M. & Klymko, C. Total communicability as a centrality measure. *J. Complex Netw.* **1**, 124–149 (2013).
73. Wagenmakers, E.-J. & Farrell, S. AIC model selection using Akaike weights. *Psychon. Bull. Rev.* **11**, 192–196 (2004).
74. Rigoux, L., Stephan, K. E., Friston, K. J. & Daunizeau, J. Bayesian model selection for group studies — Revisited. *NeuroImage* **84**, 971–985 (2014).

Supplement

<u>Supplementary Methods</u>. Task instructions.

Network Learning Task (Exp. 1-2).

You will be learning about a friend network from a university. Your job is to learn which people are friends with each other. You will first see a screen with the photos and names of the people in this network. This is to help familiarize you with each person before you start learning about their friendships. You can take as long as you want during this part.

To learn which people are friends, you will play a version of the classic card matching game. In this version, the photos of the people from the social network will be facedown and spread out in a circle on the screen. However, at the start of each round, one person's photo card will be face up. Your job to guess who is friends with the person who is face up. You can do this by clicking the facedown card one by one. Each time you click a card, you will get feedback about whether the two individuals are friends. A round is over once all the friends of a person have been identified. In the next round you will do the same for someone else in the network. You will have several chances to learn about the friends of each person.

PLEASE NOTE: You will need the friendship information you learn to do well on Part II and Part III of the study, so please pay close attention! You can earn a bonus up to $5 depending on your performance in those later parts of the study.

Network Learning Task (Exp. 3).

You will be told that gossip has emerged about someone in a social network. Then, you will see what different people in the network do with this gossip. Please keep in mind the following:

(1) Gossip is only shared between two people who are friends.

(2) It's good to keep spreading gossip through the network (+1).

(3) It's bad when gossip gets back to the person whom it is about (-10).

Your job is to remember how often each person shares gossip, and who is friends with whom. You will be tested on this knowledge later in this study. Depending on your performance, you may earn a bonus up to $5.

PLEASE NOTE: Whenever you see an upside-down face, please press the spacebar on your keyboard. Otherwise, you will NOT receive a cash bonus for your performance later in this study.

Transmissibility Learning (Exp. 2 only; administered prior to Network Learning Task).

You will be learning about a group of people who work at the same company. A piece of gossip has emerged about a colleague that everyone in this group knows. On each trial, you will guess whether someone in the group shared or didn't share the gossip. You will have 3 seconds to make your guess. You will see several different instances in which that person either shared or didn't share gossip, before moving on to the next person. How frequently someone shares gossip is an indication of their overall gossip-sharing tendencies.

If you think the person shown on the screen did NOT share the gossip, please press F on your keyboard. If you think the person shown on the screen shared the gossip, please press J on your keyboard. (You don't have to memorize this information – it will be shown on each trial.) After you make your guess on each trial, you will learn what that person actually did. As shown below, a green box indicates that that person shared the gossip and a red box indicates that they did NOT share the gossip.

PLEASE NOTE: Afterwards, you will be tested on your knowledge of each person's gossip-sharing tendencies, so pay close attention! You can earn a bonus up to $1 depending on your performance at test.

Gossip Task (Exp. 1-3).

This part of the study will be split into blocks. In each block, you will receive gossip about someone in the social network you just learned about. Your job is to indicate, on a scale of 0-100, how likely you are to share this gossip with other people in the network. For example, in Block 1, you may receive gossip about Mary. You will then be asked how likely it is that you will share gossip about Mary with Anna, Peter, Dan, etc.

Your goal is to PREVENT gossip from getting back to the person that the gossip is about, while also trying to SPREAD the gossip as widely as possible within the rest of the network.

You may be asked about the same person multiple times. Don't worry about staying consistent with what you said in the past. Just try your best and respond according to what you think in the moment. Please try to respond as quickly as possible. Click the "Continue" button each time after making your rating, otherwise your response will not be recorded.

PLEASE NOTE: If you make the same rating on all trials, you will NOT receive a cash bonus for your performance in this part of the study. Assuming you do this task correctly, you may earn a cash bonus up to $2.50.

Friendship Memory Task (Exp. 1-3).

You will now be asked to remember which individuals in the social networks are friends with each other. On each trial, you will see a single individual (the "target") at the top of the screen. At the bottom of the screen, you will see all other people in the network. Your job is to indicate who is friends with the target at the top of the screen using a rating scale. You will see the rating scale on each screen, so you do not have to memorize the following instructions. Please use the number keys on the keyboard to rate whether that person is friends with the target. You can type any whole number between 1 (VERY SURE they are NOT friends) and 10 (VERY SURE they ARE friends). All numbers in between indicate how sure you are. For example: 5 = NOT FRIENDS with the target, but you are not sure; 6 = FRIENDS with the target, but you are not sure."

Information Flow Task (Exp. 4).

At the start of each block, you will see that someone in your social network (Lucas) has shared some news. On each trial, you will see the person who shared news in the corner. In the center of the screen, you will see someone else in your network, outlined in black. Your job is to decide how likely it is that the person outlined in black (Amy) hears about the news (from Lucas or someone else in the network). You will use the D, F, J, K keys to rate the likelihood, on a scale ranging from Very Unlikely (D) to Very Likely (K).

<u>Supplementary Table 1</u>. Experiment 1 beta regression output.

Model was implemented using *glmmTMB* in R. Dependent variable (likelihood of sharing gossip with the Receiver) is a continuous measure bounded between 0 and 1. Distance is defined as the shortest path between the Target and Receiver on a given trial and ranges from 1 to 4. *Receiver_popularity* is the number of direct friendships that the Receiver has, and also ranges from 1 to 4. *Target_popularity* was included as a covariate to account for directional effects of popularity. *D-prime* was also included as a covariate to control for participants' memory accuracy for specific pairs of friends in the network. In addition for controlling for *d-prime*, we control for its interaction with *receiver_popularity* and *distance* since we might expect knowledge of specific friendships to modulate the use of these network properties during decision-making.

Predictors (mean centered)	***Estimates***	**95% *CI***	**z**	***p***
(Intercept)	-0.02	-0.11, 0.07	-0.44	0.664
distance	0.75	0.65, 0.84	15.49	**<0.001**
receiver popularity	-0.10	-0.12, -0.08	-8.52	**<0.001**
target popularity	-0.06	-0.07, -0.05	-10.72	**<0.001**
dprime	-0.01	-0.06, 0.04	-0.39	0.693
distance × receiver popularity	**0.15**	**0.12, 0.17**	**11.11**	**<0.001**
dprime × distance	-0.02	-0.07, 0.04	-0.59	0.553
dprime × receiver popularity	0.00	-0.01, 0.01	0.38	0.703
σ^2	0.48			
$\tau_{00\ sub}$	0.39			
$\tau_{11\ sub.distance}$	0.46			
$\tau_{11\ sub.receiver\_popularity}$	0.02			
$\tau_{11\ sub.distance:receiver\_popularity}$	0.03			
$\rho_{01\ sub.distance}$	0.45			
$\rho_{01\ sub.receiver\_popularity}$	0.01			
ICC	0.62			
N_{sub}	201			
Observations	43416			
Marginal R^2 / Conditional R^2	0.295 / 0.735			

<u>Supplementary Figure 1</u>. Simulation of the effects of transmissibility effects on gossip spread as a function of distance.

To better understand the role of transmissibility in gossip spread, we simulated how higher transmissibility, combined with distance from the Target, influences dissemination. Dissemination was defined as the number of nodes reached within one timestep starting from a randomly selected node. To account for slower and inconsistent spread from low-transmissibility network members (40%), gossip was initiated from all nodes except the Target and simulated over 10 timesteps. Each node influenced its direct neighbors based on its transmissibility value, with nodes reached within a timestep becoming new sources in the following step. If gossip reached the Target, it was deliberately excluded from further propagation, in order to simulate strategic behavior. At each timestep, the proportion of nodes reached was calculated by dividing the number of nodes reached by the total number of nodes in the network. The results show that choosing to gossip with highly transmissible individuals (yellow line) enhances gossip dissemination (y-axis), even at greater distances from the Target (x-axis)."

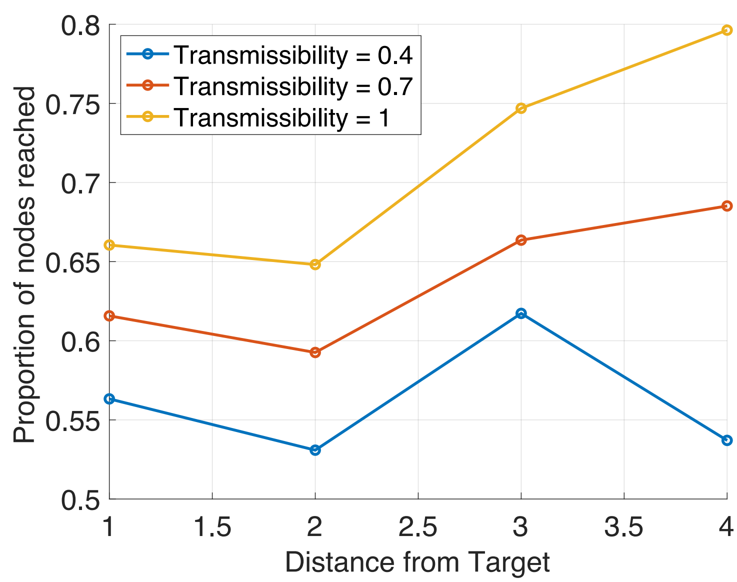

Supplementary Table 2. ***Experiment 2 beta regression output.***

Main effects in Experiment 2 were modeled the same way as in Experiment 1 with one exception. We included an additional term, *transmissibility*, which reflects the probability of a given Receiver sharing gossip with a friend, and this value was fixed for each Receiver at 40%, 70%, or 100%. All possible two- and three-way interactions among *distance*, *receiver_popularity*, and *transmissibility* were also modeled as covariates. The random effects structure was specified in the same way as in Experiment 1. (Note: *transmissibility* is abbreviated as *tr.*)

Predictors (mean centered)	***Estimates***	**95% *CI***	**z**	***p***
(Intercept)	-0.01	-0.08, 0.06	-0.23	0.822
distance	0.55	0.47, 0.64	12.66	**<0.001**
receiver popularity	-0.11	-0.14, -0.09	-8.46	**<0.001**
target popularity	-0.06	-0.08, -0.05	-11.60	**<0.001**
transmissibility	0.05	-0.04, 0.15	1.08	0.279
dprime	0.03	-0.01, 0.07	1.44	0.150
distance × receiver popularity	**0.10**	**0.08, 0.12**	**8.61**	**<0.001**
distance × transmissibility	0.09	0.01, 0.16	2.34	**0.019**
receiver popularity × transmissibility	0.11	0.02, 0.21	2.45	**0.014**
dprime × receiver popularity	0.00	-0.01, 0.02	0.25	0.807
dprime × distance	0.03	-0.01, 0.08	1.37	0.172
(distance × receiver popularity) × transmissibility	**0.07**	**0.01, 0.13**	**2.29**	**0.022**
	Random Effects			
σ^2	0.43			
$\tau_{00\ sub}$	0.27			
$\tau_{11\ sub.distance}$	0.36			
$\tau_{11\ sub.receiver\_popularity}$	0.03			
$\tau_{11\ sub.tr}$	0.38			
$\tau_{11\ sub.distance:receiver\_popularity}$	0.02			
$\tau_{11\ sub.distance:tr}$	0.10			
$\tau_{11\ sub.receiver\_popularity:tr}$	0.32			
$\tau_{11\ sub.distance:receiver\_popularity:tr}$	0.07			
$\rho_{01\ sub.distance}$	0.39			
$\rho_{01\ sub.receiver\_popularity}$	-0.02			
$\rho_{01\ sub.tr}$	0.13			
ICC	0.60			
N_{sub}	198			
Observations	42768			
Marginal R^2 / Conditional R^2	0.241 / 0.698			

<u>Supplementary Table 3.</u> Experiment 3 beta regression output.

Model included the same terms as that of the model in Experiment 2, but with the three-way interaction between *distance*, *receiver_popularity*, and *transmissibility* removed. The random effects structure was specified in the same way as in Experiment 1 and 2 above. (Note: *transmissibility* is abbreviated as *tr.*)

Predictors (mean centered)	***Estimates***	***95% CI***	***z***	***p***
(Intercept)	0.13	0.04, 0.21	2.94	**0.003**
distance	0.17	0.12, 0.22	6.46	**<0.001**
receiver popularity	0.00	-0.04, 0.03	-0.29	0.770
target popularity	0.01	0.00, 0.02	1.15	0.249
transmissibility	-0.18	-0.40, 0.04	-1.61	0.107
dprime	0.05	-0.05, 0.16	0.99	0.320
distance × receiver popularity	**0.04**	**0.02, 0.05**	**4.22**	**<0.001**
distance × transmissibility	**0.09**	**0.02, 0.16**	**2.43**	**0.015**
receiver popularity × transmissibility	0.07	-0.04, 0.18	1.23	0.218
receiver popularity × dprime	0.00	-0.05, 0.04	-0.22	0.824
distance × dprime	-0.03	-0.09, 0.04	-0.88	0.378
	Random Effects			
σ^2	0.29			
$\tau_{00\ sub}$	0.30			
$\tau_{11\ sub.distance}$	0.11			
$\tau_{11\ sub.receiver\_popularity}$	0.04			
$\tau_{11\ sub.tr}$	1.96			
$\tau_{11\ sub.distance:receiver\_popularity}$	0.01			
$\tau_{11\ sub.distance:tr}$	0.08			
$\tau_{11\ sub.receiver\_popularity:tr}$	0.43			
$\rho_{01\ sub.distance}$	0.19			
$\rho_{01\ sub.receiver\_popularity}$	-0.01			
$\rho_{01\ sub.tr}$	0.11			
ICC	0.66			
N_{sub}	169			
Observations	36084			
Marginal R^2 / Conditional R^2	0.032 / 0.668			

***Supplementary Table 4.* Experiment 4 logistic regression output.**

In Experiment 4, the response variable was modeled as a binary choice (yes/no). We therefore specified a logistic regression model, implemented using the *glmmTMB* package in R with a binomial error distribution. *Source_popularity* and *receiver_popularity* were operationalized as the number of friends that the individual reported in a separate questionnaire used to construct the ground-truth social network; popularity ranges from 1 to 28 (median = 11). Distance is the shortest path between the Source and Observer based on the true network structure. Distance ranges from 1 to 5 (median = 2). As in our controlled experiments, *receiver_popularity* was included to account for any directional effects of popularity.

Predictors (mean centered)	***Estimates***	***95% CI***	***z***	***p***
(Intercept)	*-1.88*	*-2.16, -1.60*	*-13.12*	***<0.001***
distance	-0.90	-1.02, -0.77	-13.97	**<0.001**
receiver popularity	0.01	0.00, 0.01	4.50	**<0.001**
source popularity	0.00	-0.01, 0.01	0.00	0.999
distance × source popularity	**-0.01**	**-0.02, -0.01**	**-3.74**	**<0.001**
	Random Effects			
σ^2	3.29			
$\tau_{00\ sub}$	2.01			
$\tau_{11\ sub.distance}$	0.36			
$\tau_{11\ sub.source\_popularity}$	0.00			
$\tau_{11\ sub.distance:source\_popularity}$	0.00			
$\rho_{01\ sub.distance}$	0.81			
$\rho_{01\ sub.source\_popularity}$	-0.08			
ICC	0.41			
N_{sub}	100			
Observations	60000			
Marginal R^2 / Conditional R^2	0.089 / 0.462			

***<u>Supplementary Table 5.</u>* Simple slopes of popularity across distance.**

To further investigate the interaction between distance and popularity observed across all four experiments, we conducted a simple slopes analysis. This analysis examines how the effect of *receiver_popularity* varies at the shortest (*distance* = 1) and farthest (*distance* = 4) distances from the Target. For Experiment 4, we estimated the trend of *source_popularity* at -/+ 1 SD from the mean distance. (Note: As in our regression models, all predictors were mean centered in this analysis. However, for ease of interpretation, distance values shown below align with those used in our paradigm, e.g., distance of 1 = friend of the Target of gossip.)

		Contrast	***Receiver popularity slope***	***95% CI***	***z***	***p***
Exp. 1	Distance 1	.	-0.25	-0.29, -0.20	-13.94	<0.001
	Distance 4	.	0.20	0.13, 0.26	6.75	<0.001
	.	(Distance 1) - Distance 4	-0.44	-0.54, -0.35	-11.11	<0.001
Exp. 2	Distance 1	.	-0.21	-0.26, -0.17	-12.01	<0.001
	Distance 4	.	0.09	0.02, 0.15	3.19	0.001
	.	(Distance 1) - Distance 4	-0.30	-0.38, -0.22	-8.61	<0.001
Exp. 3	Distance 1	.	-0.04	-0.09, 0.01	-2.11	0.035
	Distance 4	.	0.07	0.01, 0.12	2.77	0.006
	.	(Distance 1) - Distance 4	-0.11	-0.17, -0.05	-4.22	<0.001
Exp. 4			***Source popularity slope***			
	-1 SD mean distance	.	0.01	-0.001, 0.02	2.09	0.037
	+1 SD mean distance	.	-0.01	-0.02, 0.004	-1.68	0.094
	.	(-1 SD mean distance) - (+1 SD mean distance)	0.02	0.006, 0.03	3.74	<0.001

<u>Supplementary Table 6.</u>* *Quadratic effects of distance in each experiment.

We examined possible non-linear effects of *distance* by including a second-order polynomial term in the original regression models for each experiment (*distance*^{2}). In Experiments 1-3, the linear term remains strongly positive and significant, indicating that increasing distances reliably increases the likelihood of gossip-sharing, reflecting strategic decisions to share information with those farther away from the target. The significant negative quadratic term observed in all three experiments suggests this linear increase is tempered at greater distances, but the overall trend remains driven by the positive effect of the linear distance term. In Experiment 4, the linear term for distance was negative, reinforcing our original finding that the likelihood of news traveling between two individuals decreases linearly with distance. The significant positive quadratic term indicates a modest rebound at greater distances, possibly reflecting the belief that information can still flow through indirect connections in a network. (Note: *transmissibility* is abbreviated as *tr*).

	Predictors (mean centered)	***Estimate***	***95% CI***	***z***	***p***
Exp. 1	(Intercept)	-0.12	-0.04, -0.12	-2.79	0.005
	distance	**126.76**	**110.59, 142.93**	**15.37**	**<0.001**
	distance^2	**-82.43**	**-94.32, -70.55**	**-13.6**	**<0.001**
	receiver_popularity	-0.12	-0.15, -0.10	-10.35	<0.001
	target_popularity	-0.07	-0.08, -0.05	-11.76	<0.001
	dprime	-0.01	-0.05, 0.04	-0.31	0.761
	distance × receiver_popularity	-10.27	-13.03, -7.50	-7.28	<0.001
	distance^2 × receiver_popularity	-3.21	-7.30, 0.87	-1.54	0.123
	distance × dprime	-1.68	-10.33, 6.97	-0.38	0.703
	distance^2 × dprime	2.01	-4.33, 8.35	0.62	0.534
	receiver_popularity × dprime	0.00	-0.01, 0.02	0.63	0.527

Exp. 2	(Intercept)	-0.07	-0.15, 0.00	-2.03	0.042
	distance	**91.13**	**76.03, 106.24**	**11.82**	**<0.001**
	distance²	**-61.74**	**-71.71, -51.77**	**-12.14**	**<0.001**
	receiver_popularity	-0.13	-0.15, -0.10	-9.22	<0.001
	target_popularity	-0.08	-0.09, -0.07	-14.11	<0.001
	tr	0.04	-0.06, 0.14	0.80	0.424
	dprime	0.03	-0.01, 0.07	1.40	0.161
	distance × receiver_popularity	-8.84	-11.99, -5.70	-5.51	<0.001
	distance² × receiver_popularity	-3.30	-7.04, 0.44	-1.73	0.084
	distance × tr	1.54	-14.06, 17.14	0.19	0.847
	distance² × tr	37.97	23.79, 52.15	5.25	<0.001
	receiver_popularity × tr	0.09	-0.01, 0.18	1.75	0.080
	receiver_popularity × dprime	0.00	-0.01, 0.02	0.13	0.897
	distance × dprime	5.68	-2.61, 13.97	1.34	0.180
	distance² × dprime	-3.07	-8.48, 2.34	-1.11	0.267
	distance × receiver_popularity × tr	26.72	13.27, 40.17	3.90	<0.001
	distance² × receiver_popularity × tr	-4.90	-19.25, 9.45	-0.67	0.503
Exp. 3	(Intercept)	0.10	0.02, 0.19	2.45	0.014
	distance	**25.42**	**16.87, 33.96**	**5.83**	**<0.001**
	distance²	**-19.88**	**-26.12, -13.63**	**-6.24**	**<0.001**
	receiver_popularity	-0.01	-0.04, 0.03	-0.39	0.699
	target_popularity	0.01	-0.01, 0.02	0.99	0.325
	tr	-0.20	-0.42, 0.02	-1.81	0.071
	dprime	0.05	-0.05, 0.16	1.03	0.303
	distance × receiver_popularity	-2.83	-6.29, 0.63	-1.60	0.110
	distance² × receiver_popularity	0.55	-2.46, 3.56	0.36	0.722
	distance × tr	8.78	-1.39, 18.96	1.69	0.091
	distance² × tr	-10.53	-20.50, -0.56	-2.07	0.039
	receiver_popularity × tr	0.06	0.01, 0.10	2.20	0.028
	receiver_popularity × dprime	-0.00	-0.04, 0.04	0.00	0.999
	distance × dprime	-4.70	-14.85, 5.45	-0.91	0.364
	distance² × dprime	4.65	-2.69, 11.98	1.24	0.214

Exp. 4	(Intercept)	-1.84	-2.12, -1.57	-13.07	<0.001
	distance	**-159.79**	**-182.54, -137.05**	**-13.77**	**<0.001**
	distance^2	**29.94**	**18.74, 41.13**	**5.24**	**<0.001**
	receiver_popularity	0.01	0.01, 0.01	5.27	<0.001
	source_popularity	0.00	-0.01, 0.01	0.70	0.482
	distance × source_popularity	-0.99	-2.08, 0.10	-1.78	0.075
	distance^2 × source_popularity	-0.94	-1.89, 0.00	-1.95	0.051

***<u>Supplementary Table 7.</u>* Binned popularity regression models.**

To ensure that our reported interactions between *distance* and *receiver_popularity* in Experiments 1-3 are robust and not due to overextrapolation from combinations of popularity and distance that do not exist in our network structure, we tested the interaction with popularity binned into two categories (0 = popularity levels 1 and 2, and 1 = popularity levels 3 and 4). The models were specified in the same way as our original regressions, but with the *receiver_popularity* term replaced by *popularity_category*. (Note: *transmissibility* is abbreviated as *tr*).

	Predictors (mean centered)	***stimate***	***95% CI***	***z***	***p***
Exp. 1	(Intercept)	0.05	-0.04, 0.13	1.14	0.255
	distance	0.57	0.49, 0.64	14.74	<0.001
	popularity$_{high}$	-0.11	-0.15, -0.06	-4.49	<0.001
	target_popularity	-0.03	-0.05, -0.02	-6.40	<0.001
	dprime	-0.01	-0.06, 0.03	-0.56	0.573
	distance × popularity$_{high}$	**0.40**	**0.33, 0.47**	**11.23**	**<0.001**
	distance × dprime	-0.00	-0.03, 0.03	-0.01	0.988
	popularity$_{high}$ × dprime	0.01	-0.02, 0.03	0.63	0.531
Exp. 2	(Intercept)	0.07	-0.01, 0.15	1.75	0.079
	distance	0.46	0.39, 0.52	13.49	<0.001
	popularity$_{high}$	-0.14	-0.20, -0.09	-5.11	<0.001
	target_popularity	-0.04	-0.05, -0.03	-7.30	<0.001
	tr	-0.20	-0.34, -0.06	-2.73	0.006
	dprime	0.03	-0.01, 0.07	1.46	0.143
	distance × popularity$_{high}$	**0.25**	**0.18, 0.31**	**7.80**	**<0.001**
	distance × tr	0.03	-0.06, 0.12	0.63	0.528
	popularity$_{high}$ × tr	0.50	0.29, 0.71	4.60	<0.001
	popularity$_{high}$ × dprime	-0.00	-0.03, 0.03	-0.10	0.919
	distance × dprime	0.01	-0.02, 0.04	0.65	0.513

	distance × popularity$_{high}$ * tr	0.14	0.00, 0.28	1.89	0.058
Exp. 3	(Intercept)	0.11	0.02, 0.21	2.46	0.014
	distance	0.14	0.09, 0.18	5.51	<0.001
	popularity$_{high}$	0.02	-0.06, 0.09	0.42	0.672
	target_popularity	0.01	0.00, 0.02	2.11	0.035
	tr	-0.31	-0.56, -0.05	-2.34	0.019
	dprime	0.04	-0.07, 0.16	0.76	0.444
	distance × popularity$_{high}$	**0.09**	**0.04, 0.13**	**4.06**	**<0.001**
	distance × tr	0.11	0.03, 0.18	2.84	0.004
	popularity$_{high}$ × tr	0.23	-0.01, 0.48	1.89	0.058
	popularity$_{high}$ × dprime	0.02	-0.08, 0.12	0.38	0.707
	distance × dprime	-0.01	-0.07, 0.05	-0.42	0.678

***<u>Supplementary Table 8.</u>* Simple slopes of popularity across distance and transmissibility (Experiment 2)**.

To understand how the combined effect of *distance* and *receiver_popularity* changes across transmissibility (*tr*) in Experiment 2, we conducted a simple slopes analysis to estimate the trend of popularity at low (40%) and high (100%) *tr*, separately for the shortest (*distance* = 1) and farthest (*distance* = 4) distances from the Target. This allowed us to compare the degree to which the effect of popularity on gossip-sharing varies depending on both distance from the Target and transmissibility. (Note: As in our regression models, all predictors were mean centered in this analysis. However, for ease of interpretation, distance values shown below align with those used in our paradigm, e.g., distance of 1 = friend of the Target of gossip.)

Distance	***tr***	***Contrast***	***Receiver popularity slope***	***95% CI***	***z***	***p***
1	100%	.	-0.20	-0.27, -0.13	-7.92	<0.001
	40%	.	-0.23	-0.29, -0.16	-9.61	<0.001
4	100%	.	0.16	0.06, 0.26	4.34	<0.001
	40%	.	0.01	-0.08, 0.10	0.22	0.829
1	.	tr 100% - (tr 40%)	0.03	-0.06, 0.11	0.75	0.454
4	.	tr 100% - (tr 40%)	0.16	0.03, 0.28	3.27	0.001

<u>Supplementary Table 9.</u>* *Simple slopes of distance across transmissibility (Experiment 3).

To test how the effect of transmissibility (*tr*) changes across *distance* in Experiment 3, we conducted a simple slopes analysis to estimate the trend of *tr* at the shortest and farthest distances from the Target. (Note: As in our regression models, all predictors were mean centered in this analysis. However, for ease of interpretation, distance values shown below align with those used in our paradigm, e.g., distance of 1 = friend of the Target of gossip.)

Distance	***Contrast***	***Transmissibility slope***	***95% CI***	***z***	***p***
1	.	-0.27	-0.55, 0.01	-2.32	0.021
4	.	0.004	-0.32. 0.33	0.03	0.976
.	(Distance 1) - Distance4	-0.27	-0.54, -0.004	-2.43	0.015

***Supplementary Table 10*. Comparison of common coefficient estimates across experiments.**

To test how increasing information load and learning demands across our controlled experiments influence participants' use of latent network properties, we used z-tests to the effects of *distance*, *receiver_popularity*, and their interaction on gossip-sharing decisions across Experiments 1, 2, and 3. P-values indicate whether the differences in regression coefficients between experiments are statistically significant.

Comparison	***Predictor (mean centered)***	***95% CI***	***z***	***p***
Exp. 1 vs. Exp. 2	distance	0.07, 0.32	3.01	0.003
	receiver_popularity	-0.02, 0.05	0.82	0.410
	distance × receiver_popularity	0.01, 0.08	2.70	0.007
Exp. 1 vs. Exp. 3	distance	0.47, 0.69	10.50	<0.001
	receiver_popularity	-0.14, 0.05	-4.56	<0.001
	distance × receiver_popularity	0.08, 0.14	7.12	<0.001
Exp. 2 vs. Exp. 3	distance	0.28, 0.48	7.49	<0.001
	receiver_popularity	-0.15, -0.07	-5.00	<0.001
	distance × receiver_popularity	0.04, 0.09	4.49	<0.001

<u>Supplementary Table 11</u>. Comparison of Receiver and Target popularity effects (Exp. 1-3).

To provide insight into the relative salience of *receiver_popularity* versus *target_popularity* in our controlled experiments, we used z-tests to compare the effects of these predictors on gossip-sharing decisions. In Experiment 3, the conditional main effects of *receiver_popularity* and *target_popularity* are non-significant, resulting in no significant difference between the two.

Comparison	***Exp.***	***95% CI***	***z***	***p***
receiver_popularity vs. target_popularity estimate	1	-0.06, -0.02	-3.17	0.002
	2	-0.08, -0.02	-3.36	0.001
	3	0.05, 0.02	-0.65	0.515

Supplementary Table 12. Computational model variants.

Plus signs indicate the free parameters that were estimated for each participant in each model variant.

Model variants (Exp. 1)	γ	α	$penalty$	k	v_0	w_T
SR (gamma)	+					
SR (gamma-penalty)	+		+			
SR (gamma-penalty-slope)	+		+	+		
SR (gamma-penalty-slope-midpoint)	+		+	+	+	
Katz (alpha)		+				
Katz (alpha-penalty)		+	+			
Katz (alpha-penalty-slope)[1]		+	+	+		
Katz (alpha-penalty-slope-midpoint)		+	+	+	+	

Model variants (Exp. 2 & 3)	γ	α	$penalty$	k	v_0	w_T
SR (gamma-transmissibility)	+					+
SR (gamma-transmissibility-penalty)	+		+			+
SR (gamma-transmissibility-penalty-slope)	+		+	+		+
SR (gamma- transmissibility-penalty-slope-midpoint)	+		+	+	+	+
Katz (alpha-transmissibility)		+				+
Katz (alpha-transmissibility-penalty)		+	+			+
Katz (alpha-transmissibility-penalty-slope)		+	+	+		+
Katz (alpha-transmissibility-penalty-slope-midpoint) [2,3]		+	+	+	+	+

Model variants (Exp. 4)	γ	α_{prop}	p	k	v_0
SR (gamma-probKnow)	+		+		
SR (gamma-probKnow-slope)	+		+	+	
SR (gamma-probKnow-slope-midpoint)	+		+	+	+
Katz (alphaProp-probKnow)		+	+		
Katz (alphaProp-probKnow-slope)		+	+	+	
Katz (alphaProp-probKnow-slope-midpoint)[4]		+	+	+	+

[1]Best fitting model for Experiment 1, [2]Best fitting model for Experiment 2, [3]Best fitting model for Experiment 3, [4]Best fitting model for Experiment 4

***Supplementary Table 13.* Descriptions of model parameters.**

MODEL PARAMETER	PARAMETER DESCRIPTION	LOWER-BOUND	UPPER-BOUND
γ	Discount factor in the successor representation, reflecting the weight given to future states	0, Behavior is not influenced by transitions to future states beyond the next immediate step	0.99 – Behavior is influenced by transitions to states far into the future, with only a minor discount
α **(EXP. 1-3 ONLY)**	Attenuation factor in Katz centrality, indicating how much influence is decreased over longer paths in the network	0 – Influence from indirect connections is eliminated	0.3269 – Maximal influence of both direct and extended network connections; defined by the reciprocal of the largest eigenvalue of the adjacency matrix to ensure convergence
r **(EXP. 1-3 ONLY)**	Value associated with gossip reaching the Target	-50 – Gossip reaching the Target is maximally punishing	50 – Gossip reaching the Target is maximally rewarding
w_T **(EXP. 2-3 ONLY)**	Weight that scales the importance of each node's transmissibility (probability of sharing gossip with a friend)	-5 – Transmissibility has minimal influence on choice	5 – Transmissibility has maximal influence on choice
α_{prop} **(EXP. 4 ONLY DUE TO USING PARTICIPANT-SPECIFIC ADJACENCY MATRICES)**	Proportion of maximum α value	0 – Influence from indirect connections is eliminated	0.99 – Maximal influence of both direct and extended network connections

p **(EXP. 4 ONLY)**	Probability of having knowledge about friendships in the real-world social network, which determines participant-specific adjacency matrix based on participant's position in the network	0 – Access to knowledge about direct friends only	1 – Access to knowledge of entire ground-truth network
k	Slope of logistic function	-10 (-5 in Exp. 4) – As value increases, likelihood of sharing gossip/likelihood of hearing news *decreases*	10 (5 in Exp. 4) – As value increases, likelihood of sharing gossip/likelihood of hearing news *increases*
v_0	Midpoint of logistic function	0 – Any positive value of perceived association between two individuals leads to above-chance response	50 (10 in Exp. 4) – Value of perceived association between two individuals must reach 50 (or 10 in Exp. 4) to overcome a conservative bias against judgments of information flow

Supplementary Figure 2. Best-fitting model variants for each experiment.

Each bar depicts the number of participants for whom the labeled model on the left is identified as the best-fitting model. Only models that were deemed the best fit for at least one participant are displayed. In each experiment, the majority of participants were best fit by a Katz model variant. Across all experiments, the majority of participants were best fit by a Katz model variant, highlighting the robustness of Katz in capturing individual behavior. Bayesian Model Selection results further confirm dominance of the Katz models. Protected exceedance probabilities (PXP) calculated from each variant's BIC, which account for model uncertainty and participant-level variability, consistently approach 1.0 for Katz variants across all experiments, suggesting that these models provide the best account of the data compared to the SR.

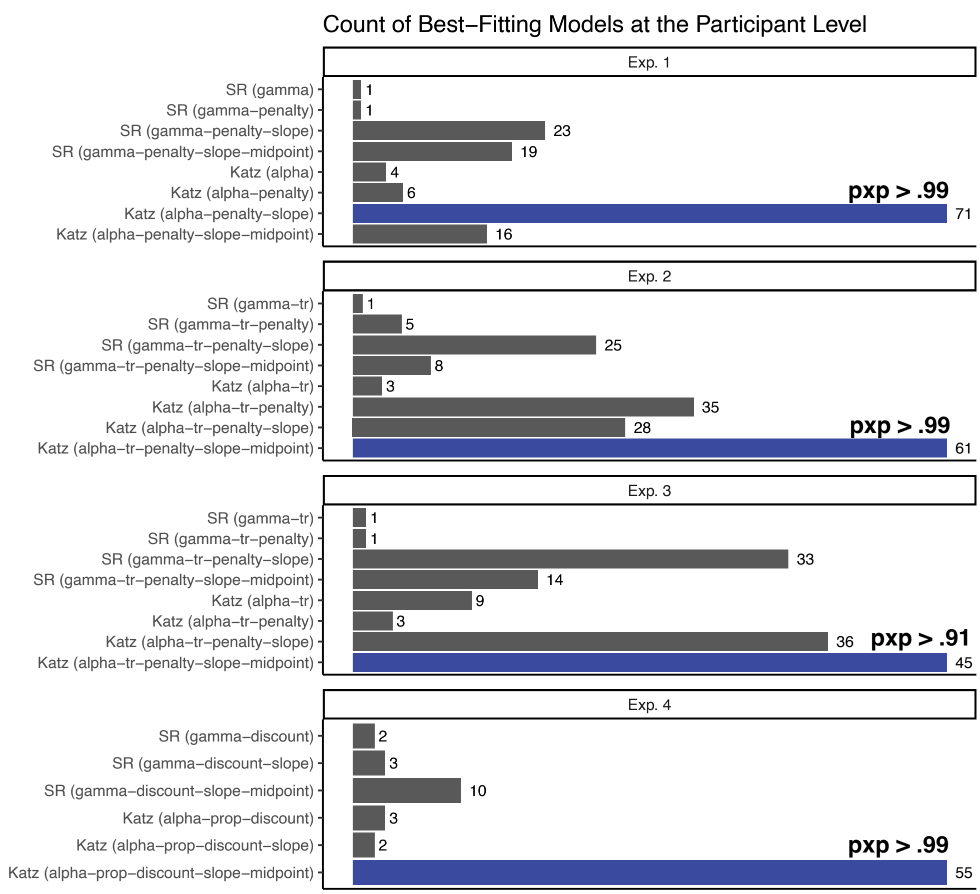

<u>Supplementary Table 14.</u> Outputs of linear regressions testing the relationship between Katz alpha parameter and regression model distance beta coefficients.

We were interested in whether the degree of network abstraction—captured by the Katz α parameter, which reflects sensitivity to longer paths in the network—relates to individuals' reliance on distance information. To test this, we extracted beta coefficients for the *distance* term from each regression model across experiments, which quantify the extent to which individuals incorporate distance into their evaluations of gossip or information flow. We predicted these coefficients by modeling individuals' Katz α estimates as a second-order polynomial term to examine potential nonlinear relationships between network abstraction and reliance on distance, while controlling for the other parameter estimates from our computational model.

	Predictor (mean centered)	***Estimate***	***95% CI***	***t***	***p***
Exp. 1	(Intercept)	0.74	0.68, 0.80	26.124	<0.001
	alpha	-0.73	-1.60, 0.10	-1.73	0.085
	alpha^2	**-2.20**	**-3.00, -1.40**	**-5.28**	**<0.001**
	slope	0.41	0.35, 0.47	13.24	<0.001
	penalty	-0.14	-0.20, -0.07	-4.31	<0.001
Exp. 2	(Intercept)	0.54	0.47, 0.62	15.82	<0.001
	alpha	-1.50	-2.80, -0.18	-2.24	0.026
	alpha^2	**-3.81**	**-4.90, -2.70**	**-6.95**	**<0.001**
	slope	0.13	0.01, 0.24	2.23	0.027
	penalty	0.17	0.09, 0.24	4.38	<0.001
	midpoint	0.14	0.06, 0.22	3.61	<0.001
	t_wt	0.03	-0.04, 0.10	0.81	0.421
	(Intercept)	0.16	0.13, 0.21	7.82	<0.001

Exp. 3	alpha	-1.30	-2.00, -0.58	-3.57	<0.001
	alpha²	**-0.34**	**-0.95, 0.27**	**-1.10**	**0.271**
	slope	0.07	0.00, 0.14	2.00	0.047
	penalty	0.13	0.09, 0.18	5.52	<0.001
	midpoint	0.01	-0.05, 0.06	0.29	0.776
	t_wt	-0.01	-0.05, 0.04	-0.31	0.757
Exp. 4	(Intercept)	-0.88	-0.94, -0.82	-29.73	<0.001
	alpha_prop	0.02	-0.73, 0.77	0.04	0.960
	alpha_prop²	**1.07**	**0.45, 1.70**	**3.44**	**<0.001**
	prob_know	-0.02	-0.08, 0.04	-0.52	0.608
	midpoint	-0.48	-0.54, -0.41	-15.34	<0.001
	slope	-0.05	-0.13, 0.03	-1.28	0.203

<u>Supplementary Figure 3.</u> Correlation between transmissibility weight parameter (w_T) and Transmissibility x Distance x Receiver Popularity interaction beta coefficients (Experiment 2).

We tested whether the transmissibility weight parameter (w_T), which quantifies how much individuals incorporate individual-level transmissibility into gossip-sharing decisions, relates to the three-way interaction between *transmissibility*, *distance*, and *receiver_popularity* in Experiment 2 (N=198 samples). A Pearson correlation revealed a significant positive relationship ($t(196) = 2.06$, $p = 0.041$, $r = 0.15$, 95% CI [0.006, 0.279] suggesting that w_T effectively tracks individual differences in the integration of transmissibility.

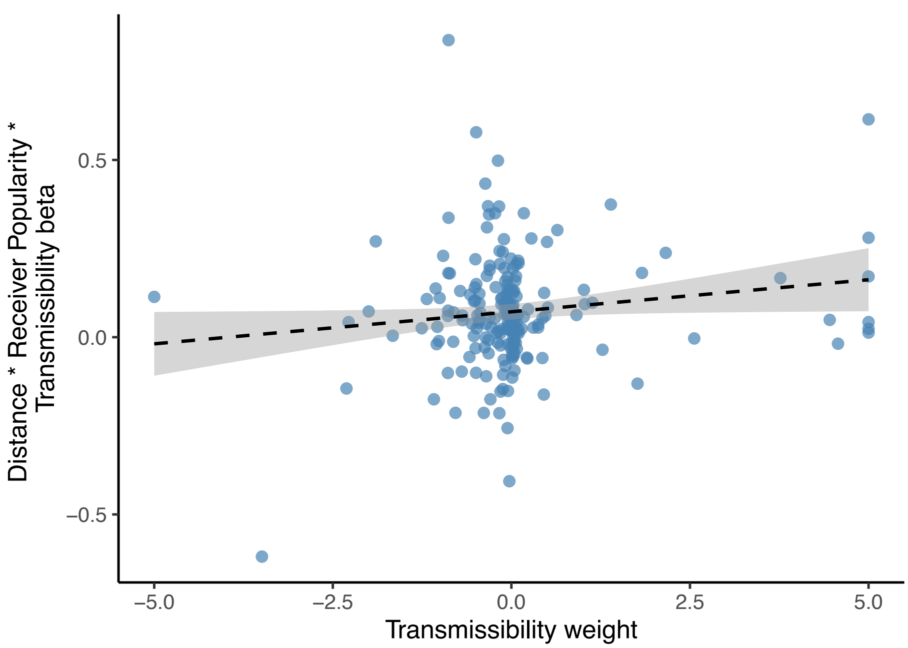

<u>Supplementary Figure 4.</u> Distribution of estimated parameter values from best-fitting Katz model variant in each experiment.

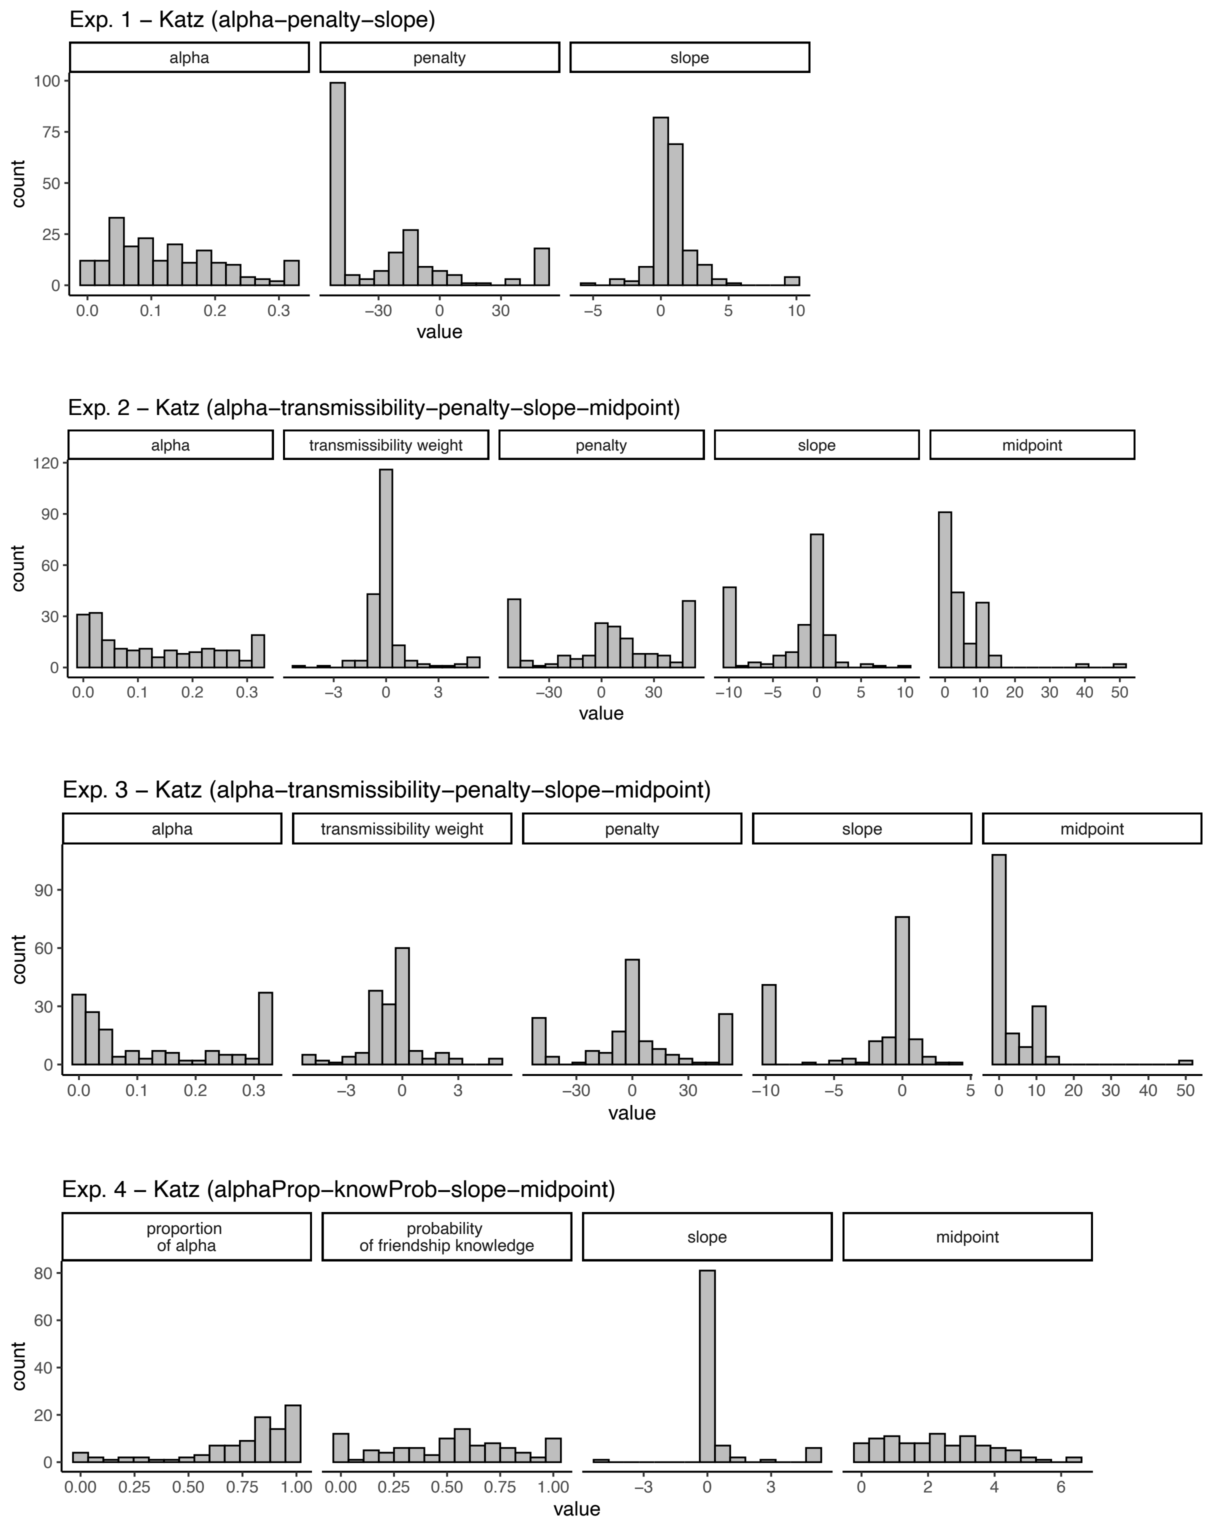

Supplementary Figure 5. ***Parameter recovery results for best-fitting Katz model variant in each experiment.***

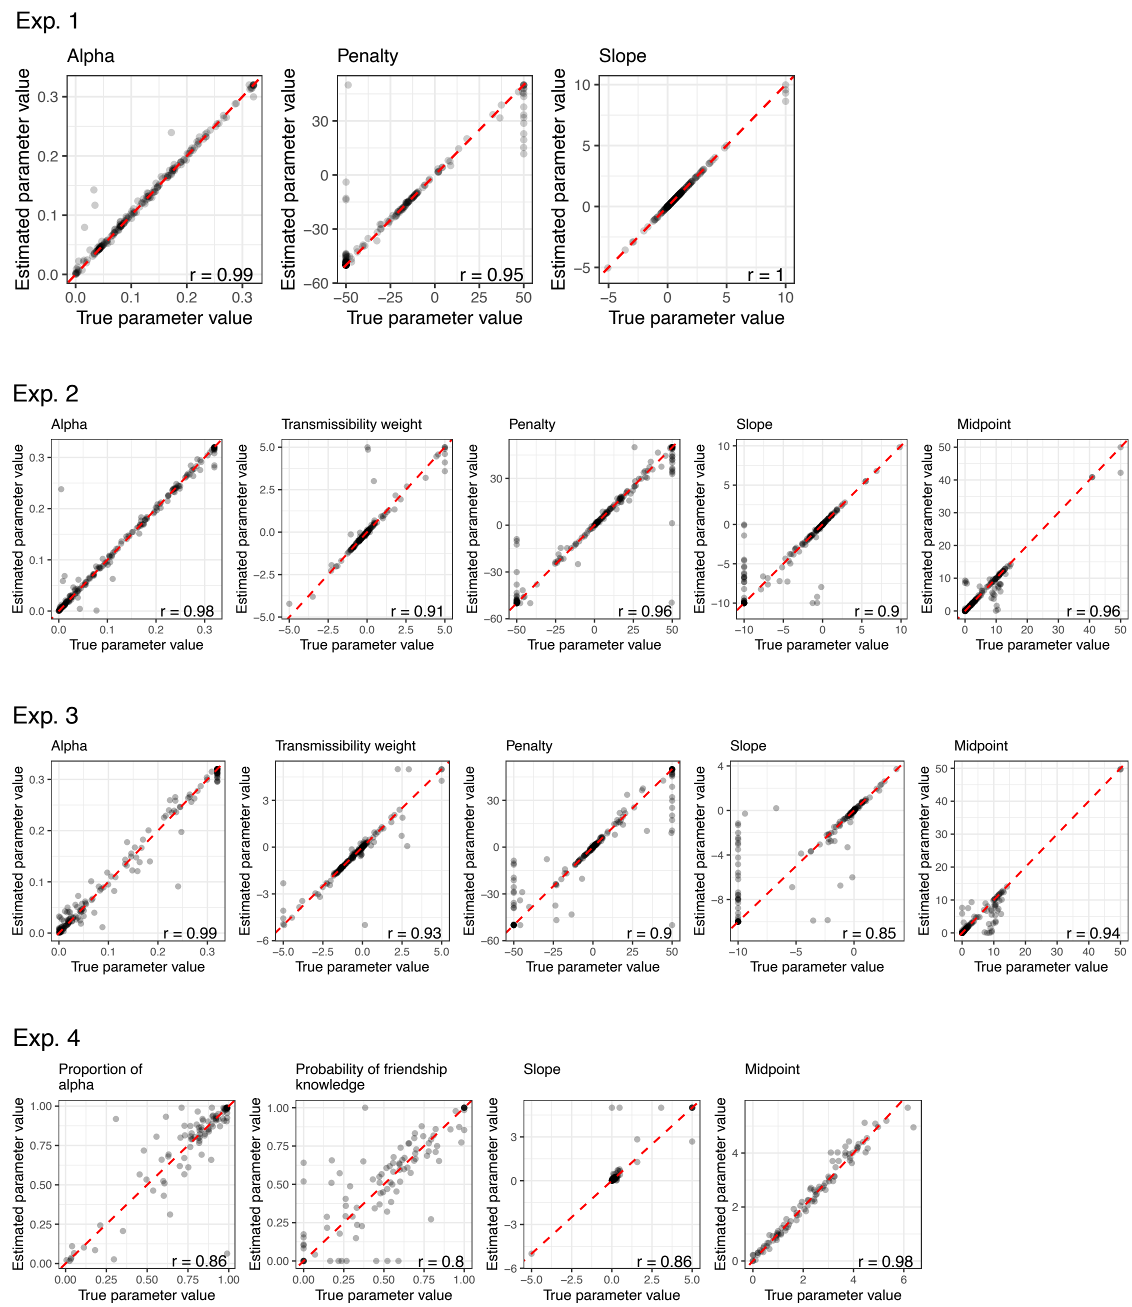

Supplementary Figure 6. ***Model confusion analysis.*** *Data simulated from each participant's best-fitting Katz and SR parameters were refit by both models; BIC identified the generating model with high accuracy across all four experiments (see below)), confirming that the two models produce distinguishable predictions.*

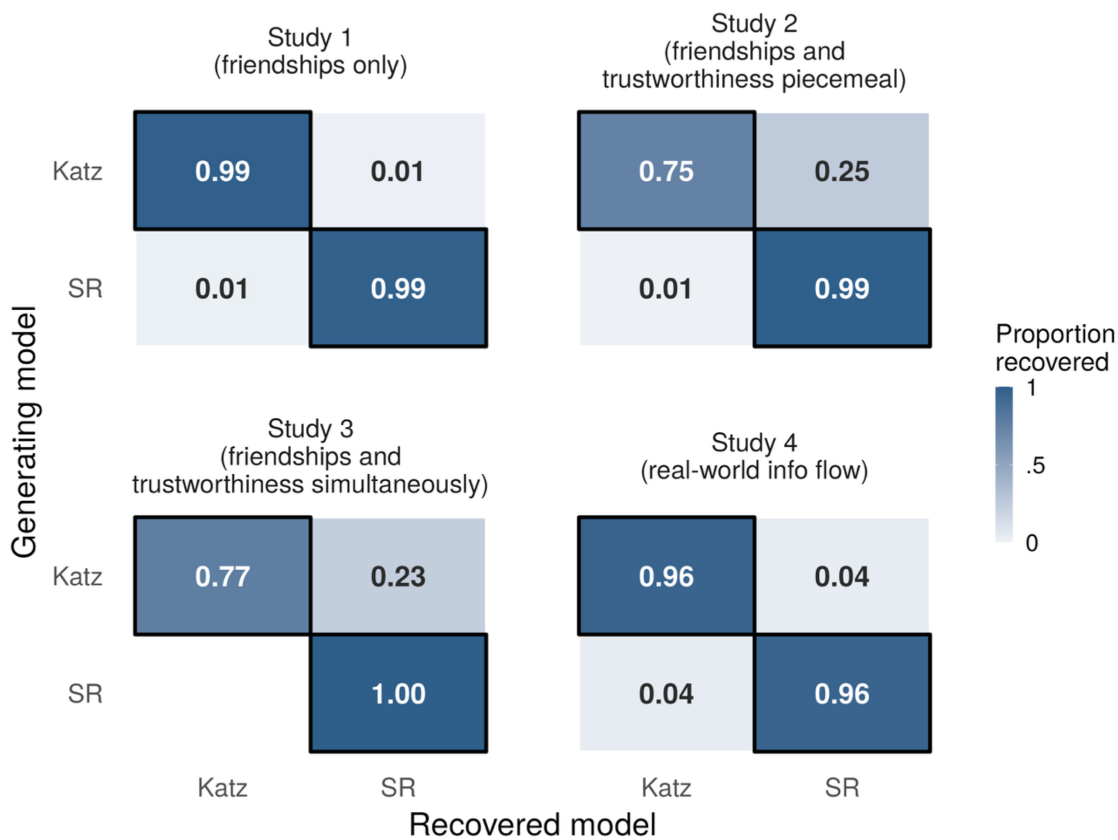

Chapter 3

Errors in social network knowledge predict how ties evolve in the future

Alice Xia, Yi Yang Teoh, Apoorva Bhandari, & Oriel FeldmanHall

Preprinted: https://osf.io/preprints/psyarxiv/cs9ux_v1 (2026)

1. Introduction

In Jane Austen's *Emma*, the heroine delights in her success at arranging a match between her former governess and a worthy gentleman, only for an ever-pragmatic family friend to dismiss her accomplishment as a "lucky guess." Emma retorts, "A lucky guess is never merely luck." What may appear as social fortunetelling in fact often reflects a deeper understanding of the webs of relationships surrounding an individual. The inferences people make about the social networks in which they are embedded are not haphazard—even the errors in recalling who is connected to whom in familiar groups are systematic (*1*). Across contexts, from classrooms to organizations, people believe relationships are more clustered, reciprocal, and balanced than they actually are (*2–4*). Rather than maintaining a veridical view of their social world, people construct error-ridden relational maps based on the topological contours of the network (*5–10*). Why do people make systematic mistakes when judging social relationships?

One enduring view, most prominently highlighted in episodic memory research, is that people compress experiences into simplified knowledge structures, otherwise known as schemas (*11–13*). Schemas capture the statistical regularities of the world, trading precision for generality, thus producing errors about specific details while preserving global accuracy. Because human memory capacity imposes similar constraints in the social domain, a parallel account has been proposed for how people learn and encode social relationships (*7, 14*). Social relational schemas feature frequently encountered social patterns (*15*), such as the tendency for an individual's friends to become friends with each other. Empirical work in this area has largely examined how such schemas either serve as priors that guide learning a novel network (*7, 16*), or as sources of systematic bias that produce predictable errors (*2, 17*). However, an underexplored possibility is that the errors themselves are not merely unimportant consequences of noisy and limited memory but are instead functionally useful for predicting how the network structure will evolve. Although it has been proposed that compressed representations of social networks could enable such kinds of predictions (*18*–*20*), this idea has never been empirically tested.

Errors in evaluating the existence of social ties may arise from a cognitive system that is fundamentally anticipatory. Distortions across different domains, from vision to memory, reveal how the brain predicts rather than passively records the world. In vision, top-down probabilistic inferences explain illusions such as the Kanizsa triangle, where the visual system "fills in" structure that is not explicitly observable (*21*–*23*). Biases in episodic memory follow the same underlying principle. For example, false memories can arise from recombining past experiences, in service of predicting what will most likely occur later (*24*–*25*). The constructive nature of recall has led to explicit frameworks proposing that episodic memory evolved precisely for simulating possible events (*26*). Recent proposals similarly argue that social knowledge—others' actions, traits, and mental states—is organized in a format that facilitates efficient predictions of social dynamics (*27*). By extension, knowledge of social relationships may also be represented in a manner that supports identifying parts of a network likely to undergo structural shifts.

An emerging line of work suggests that the ability to make inferences in complex social networks depends on mental representations that encode not only direct relationships but also the indirect connections that link individuals to distant others. Such multistep connectivity captures higher-order relational patterns and provides a basis for flexible social navigation, from identifying who bridges separate cliques to inferring how information might flow through the network (*9, 28*–*29*). A representation built on this computational principle mirrors the topology of the network and provides the mental scaffolding needed to predict which new ties are likely to form and which existing ties will dissolve. While recent work reveals the computational underpinnings of inferences about unobserved links within static networks (*9*), it remains unclear whether similar principles extend to forecasting the temporal evolution of social ties, known as link prediction.

The ability to anticipate which friendships will form and which will not dissolve should allow individuals to reap the social benefits of having early insight into their and shifting local social landscape. Correct predictions should thus be associated with advantages such as gains in influence or greater access to social capital. Identifying whether humans can track the development of connections in real-world (offline) social networks that naturally change over time is especially powerful for understanding how people navigate complex relational spaces to strategically sculpt the social worlds in which they are embedded. We therefore leverage a longitudinal approach that captures a social network from its very inception, and tracks relationships as they are made and broken over the course of an academic year. Longitudinal social network analysis critically enables us to measure how far into the future an individual can predict the status of other relationships in the community, with longer time horizons indicating that link prediction taps into structured representations of the network consistent with how it evolves.

Here, we provide the first behavioral evidence that systematic errors in third-party friendship judgments predict future structural changes in the network up to six months later. Specifically, we compare people's early judgments about the existence of others' relationships with ground-truth snapshots of the social network and show that these misperceptions accurately predict how those relationships will evolve. Computational modeling further reveals how this is achieved: participants' predictive accuracy is best explained by a representational format sensitive to both direct and indirect paths between network members. The ability to predict where new relationships emerge or dissolve, in turn, shapes how individuals' positions within the network evolve.

2. Results

a. Current errors in relationship judgments predict future tie formation and dissolution

We tracked evolving friendships among a cohort of incoming undergraduate students over the course of an academic year by asking individuals to report their friends from a roster of all network members (Fig. 1). Using mutually confirmed friendships (where both individuals independently report a connection in a first-person Friendship Survey), we constructed snapshots of the ground-truth friendship network across 6 waves starting in October and ending the following May, with between 186 and 173 students participating in each wave (see supplementary materials).

To assess participants' knowledge of others' relationships in the network, we recruited a subset of network members to complete a third-person survey (N = 100). In this Network Knowledge Task, each participant evaluated all possible pairs between a set of 30 individuals sampled within three degrees of separation from them in the network (see supplementary materials). For our primary analyses, judgments from the Network Knowledge Task in November were evaluated against the ground-truth network in December 2022—the first subsequent snapshot of the network made after the third-person survey (Fig. 1). This helped avoid treating judgments as errors when they may instead reflect socially updated knowledge about relationships already in flux, given that the October Friendship Survey was already underway when the Network Knowledge Task was administered.

In December, the network contained 591 mutual friendships out of a possible 16,110 ties. We tracked which new friendships emerged by the end of the school year in May, and which friendships dissolved over that same period. During this six-month window, 151 new ties emerged, and 181 ties dissolved. Among all pairs that were not friends in December and remained in the study through May (13,824 pairs), only 1.09% subsequently formed a mutual friendship, reflecting a low network-level base rate of tie formation despite an environment rich with opportunities for new social bonds.

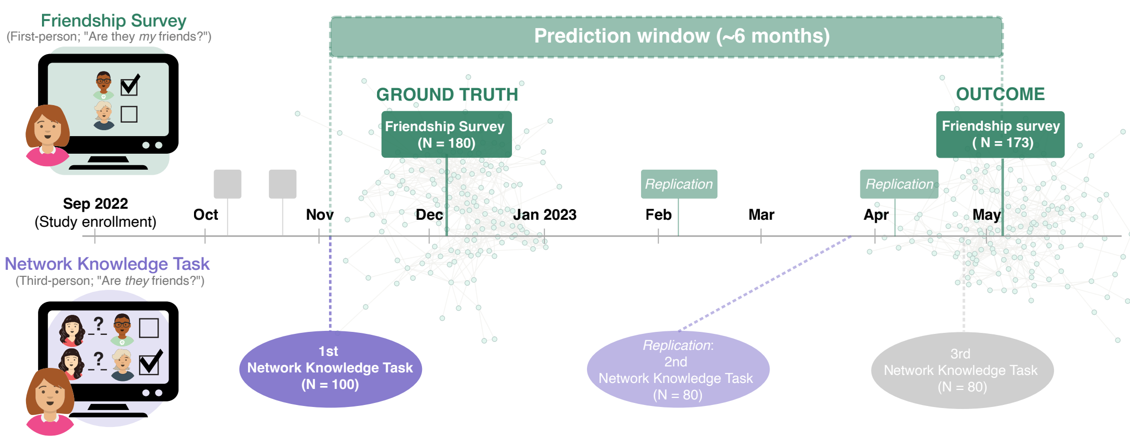

Figure 1* | *Study design and social network measurements. *Participants (N = 196) were enrolled in the study in September 2022. In November, a subset of participants (N = 100) completed a third-person survey, the Network Knowledge Task, in which they indicated whether pairs of other network members were friends. Errors in these judgments were defined relative to the ground-truth friendship network measured in December (N = 180), based on a first-person Friendship Survey. Primary analyses tracked tie formation and dissolution through May 2023 (N = 173), approximately 6 months after the Network Knowledge Task. Results were replicated with intermediate outcome time windows, as well as in a second wave of the Network Knowledge Task. Note: The Friendship Survey was administered 6 times throughout the academic year (rectangles), and the Network Knowledge Task was administered 3 times (ovals). Grey elements indicate data collected but not analyzed here.*

How accurately did participants report others' friendships in the network? Participants rated 11,877 unique pairs in total (73.7% of the 16,110 possible pairs in the December network), of which 571 were mutual friendships. On average, participants misclassified 13.6% of rated pairs (*SD* = 5.51%), judging 5.55% of pairs that were not friends in December as friends (*SD* = 5.81%) and missing 8.07% of existing friendships (*SD* = 1.98%).

These errors were not random—they systematically predicted *future* changes in the social network. When participants mistakenly perceived a friendship where none existed, that pair was nearly 3x more likely to subsequently form a friendship compared to when a pair was correctly judged to be not friends (6.76% vs. 2.38%; $\Delta_{\text{log-odds}} = 1.144$, $\beta_{\text{friendship judgment}} = 1.112$, $SE = 0.138$, $z = 8.040$, $p < .001$; Fig. 2A; table S1). The predictive effect of misjudging who is friends on later tie formation was robust across shorter prediction windows (table S2) and replicated in a second wave of the Network Knowledge Task (Fig. 1; table S3). Participants varied considerably in their ability to identify friendships that would later form. When we consider the top 10% of participants who accurately predicted future, 37.46% of pairs they judged as friends subsequently formed a friendship—approximately 14 times above the base rate of 2.65% among the pairs they rated. The top quartile accurately predicted tie formation for 26.10% of judged pairs on average.

Given that errors in judging others' friendships may simply reflect participants' direct interactions with the pairs they rated—for instance, having recently witnessed two people spending time together, or hearing that two people were growing closer—we aimed to rule out this alternative explanation by administering a Novel Tie Prediction Task (see supplementary materials). The same participants who completed the Network Knowledge Task now inferred the likelihood of a hypothetical transfer student becoming friends with an existing network member, given the transfer student's friendships with two other network members. Because the transfer student was fictional, participants could *only* draw on knowledge of existing network structure.

Predictions were evaluated by identifying real dyads in the network whose structural relationships to their mutual friendships matched those of each trial of the Novel Tie Prediction Task. The actual relationship outcomes of these structurally matched real pairs served as a proxy of ground truth, allowing us to assess whether participants' inferences tracked genuine network dynamics despite the fictional framing. Predictions forecasted which real, structurally similar pairs would go on to form a friendship ($\beta = 0.464$, $SE = 0.070$, $z = 6.582$, $p < .001$; fig. S1), confirming that accurate predictions reflect—at least in part—structural inference rather than firsthand knowledge of specific individuals.

Are network members also sensitive to which relationships will dissolve? Among existing friendships in December, 33.46% dissolved by May. When participants incorrectly judged two friends as not being friends, those existing friendships were more than 1.2x more likely to dissolve compared to when participants correctly identified them as friends (30.83% vs. 24.82%; $\Delta_{\text{log-odds}} = 0.550$; $\beta_{\text{friendship judgment}} = -0.553$, $SE = 0.103$, $z = -5.347$, $p < .001$; Fig. 2B; table S4). Participants' errors therefore also accurately tracked which friendships were at greater risk of ending. This pattern of results held across shorter prediction windows (table S5), as well as in the second Network Knowledge Task (table S6), indicating that participants' predictive errors were not specific to a single timepoint.

To understand whether predicting tie evolution six months into the future reflects a genuine ability to forecast the changed topology of the social network, we began by comparing participants' accuracy against complementary null models: we randomly shuffled each participant's judgments across dyads 1,000 times, preserving participants' overall response frequency but breaking any systematic link between judgments and outcomes. The predictive power of making errors about others' friendships on later tie formation and tie loss in the empirical data significantly exceeded chance levels (all estimates outside the 95% CI; Fig. 2C-D; see supplementary materials). We further confirmed that hit rates (proportion of dyads correctly identified as future ties) exceeded what would be expected from participants' own tendency to label pairs as 'Friends' or 'Not friends', ruling out response bias as an explanation (fig. S2A-B; results replicated in the second wave of the Network Knowledge Task; fig. S2C-F).

Real social networks tend toward triadic closure (*30–32*), and mental representations of social networks are weighted toward balanced triads (*3*, *7*, *14*), i.e., configurations in which a friend's friend is also one's friend, both of which may independently account for our findings. We therefore tested whether errors predicted future tie evolution beyond shared network neighborhood. Controlling for the number of mutual friends between two individuals (i.e., structural similarity), errors in relationship judgments still robustly predicted future tie formation ($\beta_{\text{friendship judgment}} = 0.545$, $SE = 0.125$, $z = 4.354$, $p < .001$; table S7), as well as tie dissolution, though the effect was attenuated ($\beta_{\text{friendship judgment}} = -0.352$, $SE = 0.175$, z = -2.010, $p = .044$; table S8).

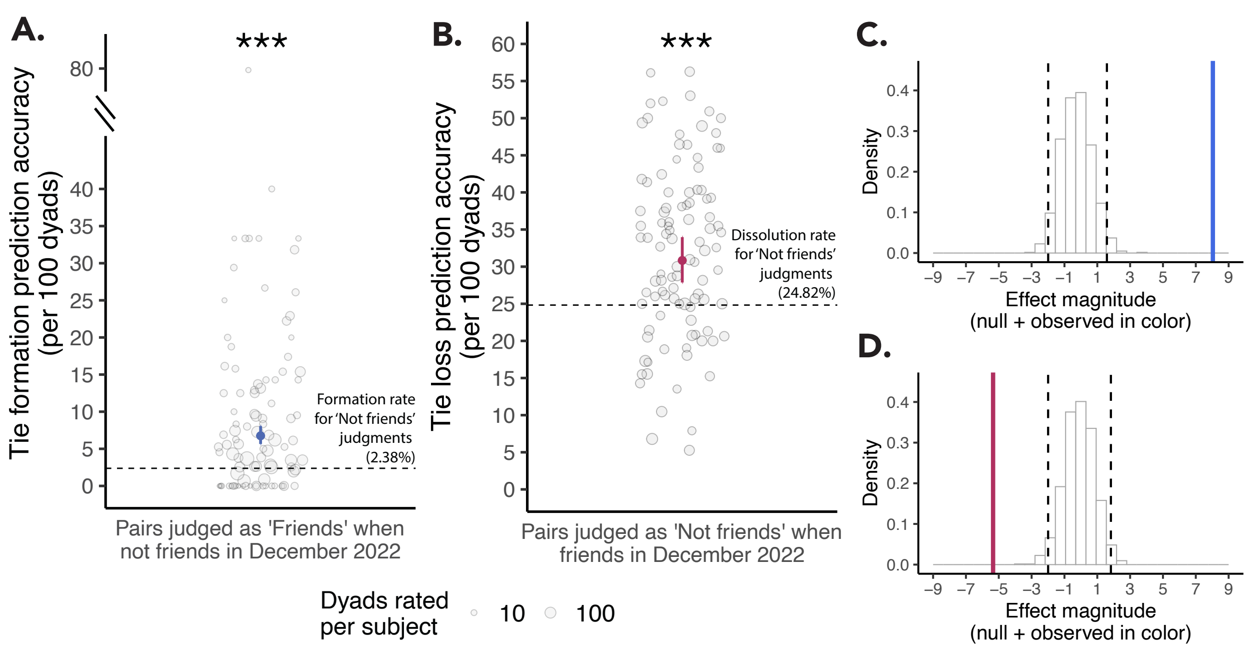

Figure 2 | Errors in friendship judgments predict future relationship change above chance. ***(A)*** *Dyads misjudged as 'Friends' (x-axis) in November 2022 are significantly more likely to form a friendship six months later. Y-axis shows future tie formation rate (number of future ties formed per 100 dyads) as a function of friendship judgment type. Each point represents one participant; point size reflects the number of dyads rated in each judgment category. The dashed line represents the rate of tie formation when participants judged a pair as 'Not Friends' (2.38%). Blue point reflects the mean; error bars are 95% CIs corrected for within-subject variance (33). Logistic mixed-effects model; N = 9,420 dyads, 100 raters.* ***(B)*** *Pairs misjudged as 'Not Friends' are more likely to dissolve their existing friendship by the following May. Dashed line represents the tie dissolution rate when participants judged a pair as 'Friends' (24.82%). Maroon point reflects the mean; error bars are 95% CIs corrected for within-subject variance (33). Logistic mixed-effects model; N = 523 dyads, 100 raters.* ***(C–D)*** *Observed effects exceed chance levels for both tie formation (C) and tie loss (D). Histograms show the null distribution of effect magnitudes generated by randomly shuffling each participant's judgments across dyads 1,000 times while preserving their overall response rate; colored vertical lines show the observed effects (formation: $z = 8.040$; loss: $z = -5.347$) and dashed lines indicate the 95% null interval (formation: [-1.991, 1.586]; loss: [-1.995, 1.829]). Observed effects fall outside the null distribution (both $p < .001$). *$p < .05$, **$p < .01$, ***$p < .001$.*

b. Representations of latent network structure support link prediction

Homophily, the tendency for people with similar backgrounds and interests to become friends (*34–35*), similarly distorts memories of ties (*2*, *18*, *36*) and could explain the same effects. We therefore tested whether errors predict future tie evolution over and above those expected by these tendencies. Additionally, a separate model that controlled for classic forms of homophily (specifically, shared gender, ethnicity, and dormitory floor), alongside the number of mutual friends, revealed that the effect of errors in social knowledge still remains highly significant ($\beta_{\text{friendship judgment}} = 0.451$, $SE = 0.127$, $z = 3.560$, $p < .001$; Fig. 3A left; table S9). In the tie loss domain, including measures of homophily attenuates the effect such that it is no longer significant ($\beta_{\text{friendship judgment}} = -0.321$, $SE = 0.178$, $z = -1.808$, $p = .071$; table S10), suggesting that the prediction of tie dissolution is more dependent on homophily, relative to the prediction of tie formation. Put simply, knowledge about someone's network position relative to others is what drives accurate predictions about which new friendships will form. However, knowing whether an existing friendship will dissolve likely requires additional, person-level information.

Given that people can predict who becomes friends and which friendships dissolve up to half a year later, we next asked what kind of mechanisms beyond triadic closure and homophily support such prediction. We reasoned that if two people are indirectly linked through multiple pathways, then the structural pathways between them create more opportunities for interaction, thus making a future friendship more likely. Prior work suggests that social inferences draw on representations that integrate both direct and indirect paths through a network (*9, 28–29*). Because the same structural dependencies that shape current relationships also constrain which ties are likely to emerge or terminate in the future, this kind of representation should support link prediction. To assess this possibility, we specified a computational model that formalizes multistep connectivity using Katz communicability (*37*), a measure of how strongly an individual is connected to others in the network through both direct and indirect paths (see supplementary materials).

Because people's knowledge of a network is necessarily partial, shaped by who they interact with and learn social information from, we constructed our model using directed friendship nominations rather than mutually confirmed ties, weighted by distance from the self. This approach is further motivated by our prior work showing that an outbound-based friendship nomination model systematically outperformed ones built on mutual ties when fitting participants' social judgments (*28*; see supplementary materials). Fitting this model to friendship judgments allowed us to extract estimates of each participant's inferred relational strength between any two individuals in the network, i.e., Katz communicability values (see schematic in Fig. 3B and supplementary materials). These estimates robustly predict future tie formation, even while accounting for homophily metrics ($\beta_{\text{Katz communicability}} = 0.329$, $SE = 0.026$, $z = 12.444$, $p < .001$); furthermore, they abolish the predictive power of explicit friendship judgments ($\beta_{\text{friendship judgment}} = 0.070$, $SE = 0.137$, $z = 0.509$, $p = .611$;

Fig. 3C left; table S11). This result suggests that perceiving latent structural connectivity between pairs of individuals is the knowledge that people rely on to correctly anticipate emerging friendships.

Participants' model-estimated Katz communicability values did not predict tie loss ($\beta_{\text{Katz communicability}} = 0.060$, $SE = 0.033$, $z = 1.820$, $p = .069$) after accounting for the dominant predictor of friendship dissolution: the number of mutual friends in December ($\beta_{\text{mutual friends}} = -1.466$, $SE = 0.098$, $z = -14.988$, $p < .001$), and homophily factors. However, explicit friendship judgments still significantly predicted which friendships would end ($\beta_{\text{friendship judgment}} = -0.446$, $SE = 0.191$, $z = -2.336$, $p = .019$; Fig. 3C right; table S12). Taken together, these results suggest that accurate predictions of when friendships may weaken or break may depend on aspects of relational strength that are not fully encoded in the network's multistep connectivity structure.

These findings raise the question of whether the predictive signal in participants' Katz communicability estimates merely reflects the network's own natural developmental trajectory. Social networks tend to evolve toward greater multistep connectivity over time, as new ties preferentially form between people who are already indirectly connected (*30–32*). If inferred connectivity strength simply tracks this structural trajectory, then its predictive value would be explainable by network dynamics alone. To test this, we included the objective Katz communicability of each dyad derived directly from the network structure in December 2022 as a covariate (see supplementary materials). The effect of participants' model-estimated Katz communicability values remained robustly predictive of future tie formation ($\beta_{\text{Katz communicability}} = 0.345$, $SE = 0.027$, $z = 13.003$, $p < .001$), even while the ground-truth connectivity of each dyad was also predictive ($\beta_{\text{ground-truth Katz communicability}} = 0.920$, $SE = 0.049$, $z = 18.602$, $p < .001$; table S13). Together, these results demonstrate that people's mental representations of who is indirectly connected to whom encode genuinely predictive information about the future of the network—information that goes beyond what the current structure of the network itself would anticipate.

The Katz model also estimates, for each participant, a free parameter γ that reflects the degree to which that individual encodes multistep paths in the network. If the ability to predict tie changes depends on this type of representation, then participants characterized by higher γ (greater connectivity) should also have more accurate predictions about the future of the network. This association would reveal that predictive accuracy stems from a representation attuned to the overall connectivity within the network. Since participants varied substantially in their ability to predict future tie evolution, and because each participant rated a different subset of pairs from their local network neighborhood, it was important to distinguish genuine predictive sensitivity from a general tendency to under- or over-report ties.

We quantified each participant's predictive accuracy using d′, a signal detection measure that compares correct identification of future tie change (hits) against incorrect judgments of relationships that did not change (false alarms). We computed d′ separately for tie formation and tie loss. Participants showed above-chance sensitivity to both friendship formation (mean d′ = 0.720, 95% CI [0.604, 0.835], $SD = 0.581$, $t(99) = 12.386$, $p < .001$) and dissolution (mean d′ = 0.268, 95% CI [0.140, 0.396], $SD = 0.645$, $t(99) = 4.162$, $p < .001$), though this was significantly more pronounced in the formation domain ($\beta_{\text{formation d}'} = 0.452$, $SE = 0.082$, $t(99) = 5.480$, $p < .001$; table S14). Consistent with our prediction, individuals who more strongly encoded multistep paths (greater connectivity, indexed by higher γ) in the network showed significantly greater sensitivity to tie formation relative to tie loss ($\beta_{\text{Katz } \gamma * \text{formation d}'} = 2.323$, $SE = 0.725$, $z = 3.205$, $p = .002$; Fig. 3D; table

S15). Decomposing this interaction revealed that the simple slope of γ was significantly positive for formation (β = 1.387, SE = 0.555, z = 2.501, p = .012) but not for loss (β = -0.936, SE = 0.555, z = -1.687, p = .092; table S16). In other words, individuals who integrate longer chains of connections are better at predicting the emergence of new ties, but not their termination. Mental representations of multistep connectivity therefore appear especially tuned to new ties in the network.

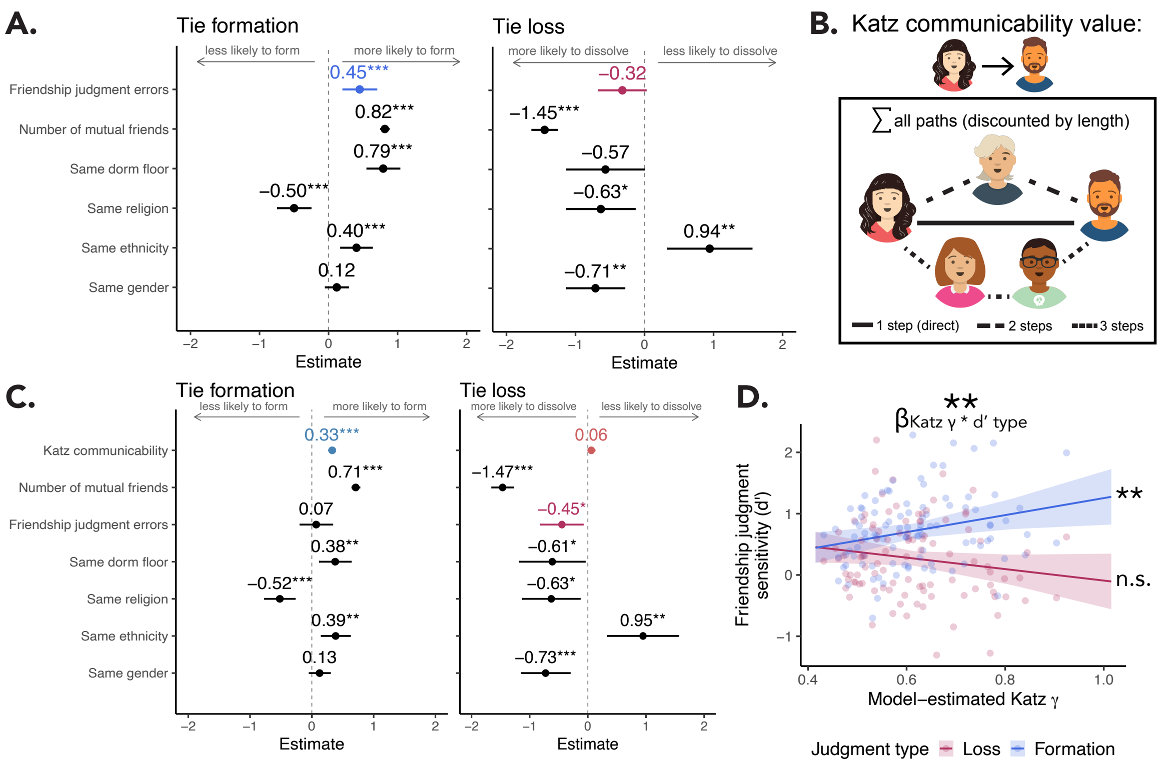

Figure 3 | Sensitivity to overall network connectivity underlies accurate tie formation predictions. ***(A)*** *Errors in friendship judgments predict future tie formation and dissolution, even when accounting for structural and demographic controls. Coefficient estimates and 95% confidence intervals from a mixed-effects model predicting future tie formation (left) and tie loss (right), including friendship judgments, number of mutual friends, and homophily controls (same gender, ethnicity, religion, and dormitory floor).* ***(B)*** *Schematic illustration of how Katz communicability is computed between two individuals, integrating both direct and indirect pathways of increasing length.* ***(C)*** *Including model-estimated pairwise Katz communicability values (i.e., inferred multistep connectivity between two network members) eliminates the predictive effect of errors in friendship judgments on tie formation (left) but not tie loss (right), though errors remain predictive in the loss domain (note: negative coefficient reflects lower dissolution probability for correct vs. erroneous judgments). Coefficient estimates and 95% confidence intervals from mixed-effects models additionally including Katz communicability alongside friendship judgments, number of mutual friends, and homophily controls.* ***(D)*** *Reliance on multistep paths in the network (higher model-estimated Katz γ) is associated with greater sensitivity to tie formation but not tie loss. Each point represents one participant's d′ for tie* *formation (blue) and tie loss (maroon) plotted against their model-estimated* Katz γ; lines show lines of best fit for each judgment type. *$p < .05$, **$p < .01$, ***$p < .001$.

c. Divergent social consequences of accurate link prediction

If the ability to predict future tie formation and loss enables social adaptation, then accurate link prediction should be associated with tangible social benefits. Even if not consciously aware of it, accurately anticipating how others' relationships evolve may give people the opportunity to pre-emptively adjust their own social behavior and, in principle, influence their own position in the network. Network position can be shaped by different forms of social importance, such as being embedded in a dense, well-connected core (reflected in measures such as eigenvector centrality), or serving as a bridge between otherwise separated groups (reflected in betweenness or brokerage; see supplementary materials). To understand whether individual differences in predictive accuracy influence these forms of social currency, we examined how both eigenvector centrality and betweenness change over time, while accounting for individuals' initial positions in the network.

Greater link prediction accuracy was more strongly associated with increases in eigenvector centrality for participants who started as being more central ($\beta_{\text{initial centrality * formation d}'} = 1.914$, $SE = 0.592$, $z = 3.232$, $p = .001$; Fig. 4A; table S17). Simple slopes analysis revealed that correctly predicting future ties significantly boosted future centrality for those who were already highly central and well-connected ($\beta = 0.513$, $SE = 0.175$, $z = 2.927$, $p = .003$) but conferred no benefit to those who were classified as medium ($\beta = 0.014$, $SE = 0.120$, $z = 0.117$, $p = .907$) or low centrality ($\beta = -0.448$, $SE = 0.204$, $z = -2.192$, $p = .028$; table S18). The social utility of accurate link prediction is therefore conditional on where an individual is initially located in the network. Anticipating tie formation amplifies existing social advantage, but only for those who initially occupy relatively influential positions in the network.

We observe a different pattern for betweenness. Even though betweenness was moderately correlated with eigenvector centrality at both timepoints (initial: r = .356; future: r = .342), its trajectory over time as a function of accurate link prediction was markedly distinct. Individuals who started off as brokers—reflected by greater betweenness—lose their brokerage position over time if they accurately predicted which new friendships would form ($\beta_{\text{initial betweenness * formation d}'}$ = -0.487, SE = 0.165, z = -2.958, p = .003; Fig. 4B; table S19). This effect was specifically driven by individuals who initially occupied the strongest bridging positions in the network (β = -0.050, SE = 0.017, z = -2.848, p = .004; table S20). In other words, accurate link prediction is thus associated with structural changes consistent with closing the open triads that high-betweenness individuals once connected.

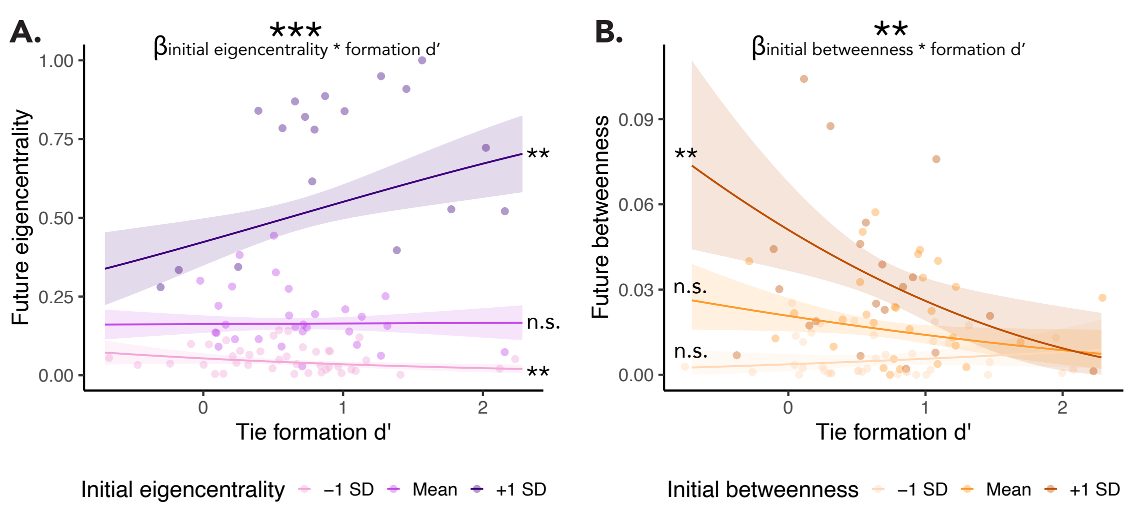

Figure 4 | Initial network position moderates the social benefits of accurate link prediction. (A) *Accurate tie formation prediction boosts future eigenvector centrality (eigencentrality), but only for individuals who were already highly central. Each point represents one participant's future eigencentrality (May 2023) plotted against tie formation d′, colored by initial eigencentrality in December 2022. Lines and shaded ribbons show model-predicted values and 95% confidence intervals at each level of initial eigencentrality.* ***(B)*** *Accurate tie formation prediction is associated with loss of brokerage, but only for individuals who began as strong brokers. Each point represents one participant's future betweenness (May 2023) plotted against tie formation d′, colored by initial betweenness in December 2022. Lines and shaded ribbons show model-predicted values and 95% confidence intervals* *at each level of initial betweenness. Betweenness values were square-root transformed for modeling and back-transformed for visualization. *p < .05, **p < .01, ***p < .001.*

3. Discussion

Social relational knowledge is notoriously prone to error. Here we show that these errors—rather than being symptoms of cognitive limitations (*7, 14*), simplified schemas (*2, 17*), or idiosyncratic perceptual tendencies (*38*)—serve as adaptive forecasts of how a social network will evolve. We provide the first behavioral evidence that 'errors' in judging others' friendships reliably anticipate tie formation and dissolution six months into the future. While we found at the group level that people are 3x better than chance at predicting future social ties and losses, those who are best at predicting these future social connections can anticipate who becomes friends nearly 14x better than chance. We provide behavioral and computational evidence that this predictive capacity is grounded in a mental model of multistep connectivity and demonstrate that this capacity influences an individual's future social standing. Collectively, our results provide striking evidence that the human brain does not just represent the social world as it is, but what it is expected to become in the future.

Prior work has demonstrated that the dynamics of social networks over time are strongly constrained by network topology, with triadic closure being one of the dominant patterns of tie formation (*30*). Foundational analyses of large-scale communication networks document how indirect connections between individuals progressively produce direct ties (*31*). Crucially, however, such work relies on behavioral traces, i.e., who contacts whom, rather than on individuals' *perceptions* of others' relationships. this prior work is therefore agnostic about whether the people embedded in these networks mentally represent the structural forces driving their evolution. Our findings fill this gap by not only providing evidence for the capacity to predict the formation and dissolution of future ties, but also by revealing that these predictions are afforded by certain representations of the network's topology. Accurate prediction does not simply rely on a bias toward triadic closures, but incorporates knowledge of multistep connectivity. Specifically, we found that participants' predictive errors are best explained by Katz communicability, integrating both direct ties and indirect pathways to estimate the likelihood of a relationship forming or dissolving.

Beyond its cognitive implications, the form of social acumen we observe here—ability to foresee where ties will emerge within the network—also reshapes individuals' future positions within the network. Our results reveal that predictive accuracy acts as a catalyst for social embeddedness, though its benefits are socially stratified. Predictive accuracy enhances the prominence of initially well-connected individuals, while diminishing the bridging role of those who start out as brokers. Correctly forecasting tie formation places already prominent individuals on a trajectory toward greater embeddedness and connection in the network core. In a similar manner, when individuals lose their brokerage positions, it reflects that the gaps they once spanned are closing. Rather than representing a decline in influence, this shift suggests that accurate forecasters are the first to navigate toward the closing of 'structural holes.' In this sense, both gaining centrality and losing brokerage reflect a singular strategic outcome: accurate link prediction allows people to align themselves with the network's emerging core.

By offering this cognitive account of link prediction in a large real-world social network, we also build on decades of computer science research that has sought to predict unobserved or missing ties from incomplete network information (*39–41*). A range of well-established approaches that focus on path-based metrics and, more recently, graphembedding methods, typically require full adjacency information to perform well (*42–44*). In contrast, we show that humans solve a version of this problem under far more constrained and noisy conditions, relying only on the sparse and locally observable cues available to them—specifically, the people others name as friends (outbound friendship nominations) and the graded layers of separation from themselves (distance from the self).

The ability to evaluate others' relationships in a way that anticipates future tie change is a striking feature of human social cognition. While multistep connectivity offers a parsimonious computational account of why present-day errors are tightly coupled with actual network evolution, a critical next step is to examine how behavioral signatures of link prediction map onto the neural geometry that supports forward-looking social inference. Additionally, unobserved offline dynamics may have influenced our measurements in ways we could not capture. Participants' representations were very likely shaped by unobserved experiences outside the laboratory, including direct interactions, conversations, and other cues about the relationships of others. Future work using more fine-grained experience-sampling approaches could more directly measure the social information people encounter and test how such exposure gives rise to accurate link prediction. Thus, although explicit judgments and inferred connectivity strength both outperformed ground-truth homophily, the real-world learning processes that support these representations remain to be identified.

Rather than treating distortions in social relational knowledge merely as outcomes of limited memory capacity or biases that arise from simplified social schemas, our findings show that they are the natural consequences of a mind that continuously generates expectations and predictions about the trajectory of one's social environment. On this view, a cognitive map of the social world is organized around structural cues that enable people to forecast relationship dynamics, not simply recall them. By relying on the current network structure to project which relationships are likely to form or break, people can better identify partnerships with more potential to thrive, or to invest in friendships that are less likely to fracture. Such predictive errors may therefore be a fundamental feature of how humans achieve social success in large, evolving communities, revealing that effective social navigation requires representing not only the present environment but also its most probable future.

4. Methods

Participants

In September, at the start of the academic year at a U.S. university, we recruited 198 first-year undergraduate students to take part in a longitudinal social network study. Two were excluded for not meeting eligibility criteria (wrong class year and not following task instructions), leaving 196 eligible participants (103 female, 89 male, 4 other gender identities; 146 identifying as non-white or mixed race; mean age = 18.06 years, *SD* = 0.42; see table below for demographic breakdown). Participants provided photos and demographic information during an initial enrollment session. Although we collected a broad range of demographic data, only variables theoretically relevant to homophily-based judgements were analyzed (i.e., gender, ethnicity, religion, and dormitory floor).

Overall timeline.

Over one academic year, from October 2022 to May 2023, we contacted participants six times to complete the Friendship Survey (first-person), in which they reported their friendships with other participants. We focus on the surveys administered in December 2022 (N = 180; 95 female, 81 male, 4 other gender; 136 non-white or mixed race; age *M* = 18.06 years, *SD* = 0.42) and May 2023 (N = 173 after one exclusion for failing to follow instructions; 93 female, 76 male, 4 other gender; 128 non-white or mixed race; age *M* = 18.06 years, *SD* = 0.43; see table below for demographic breakdown). December to May defines our primary prediction interval. The December Friendship Survey provides the ground-truth network against which the first wave of the Network Knowledge Task (administered from mid-November to mid-December) was evaluated. The first-person survey from May captures which ties formed or dissolved over the following six months. We did not use surveys preceding December because they were completed before the Network Knowledge Task. We use two intermediate first-person surveys collected in February and April to replicate the predictive effect of November-December friendship judgments on tie formation at earlier timepoints.

In November 2022, a subset of participants completed two behavioral tasks. First, the Network Knowledge Task measured their knowledge of ties in the broader social network by asking whether pairs of other network members were friends. Second, the Novel Tie Prediction Task measured their explicit predictions about new tie formation (same participants from first wave of the Network Knowledge Task). Our analyses test whether errors in the earlier Network Knowledge Task predicted how friendships subsequently evolved in the network. To minimize the possibility that apparent "errors" reflected relationship changes already underway, errors were defined relative to mutually confirmed friendships reported in December 2022.

All procedures were approved by the university's Institutional Review Board (protocol 1911002585). Participants provided informed consent, and for minors, consent was obtained from legal guardians alongside participant assent. Compensation was $10 per hour for online surveys and $15 per hour for in-lab behavioral tasks. Completion bonuses were also offered to increase retention over the academic year.

Friendship survey (first-person).

To construct ground-truth snapshots of the evolving friendship network, we asked participants to report their friendships with every other network member at six distinct time points throughout the academic year. In the first two rounds of the Friendship Survey, participants were provided with a roster of all other network members' names and photos and asked to indicate whether each person was a friend (yes/no). In subsequent versions of the survey, we measured social relationships using a 4-point scale ('Friend', 'Acquaintance', 'Familiar', or 'Do Not Recognize'). For all analyses, we focus on mutually reported friendships and therefore collapse these ratings into two categories: 'Friends' and 'Not Friends'. In our primary analyses, we focus on the ground-truth relationship status mutually confirmed by each dyad during the third wave of data collection (December 2022), which occurred after the first administration of the Network Knowledge Task and therefore allows us to compare participants' perceptions of others' relationships with those dyads' own most up-to-date reports.

<u>Network Knowledge Task (third-person).</u>

We assessed people's knowledge of the relationships in the social network by asking participants to indicate whether they believed two others are friends. This task was administered three times—in November 2022, mid-March 2023, and mid-April 2023—but all analyses reported here focus on the first wave because it is farthest in time from May 2023 and therefore serves as the strongest test for showing link prediction across time. Because inquiring about each possible pair of network members was infeasible due to timing constraints and managing retention of our participants, we opted for a sampling procedure that produced a custom set of stimuli for each participant. Specifically, we selected 5 individuals who are direct friends of the participant, 10 individuals who are two degrees removed, and 15 who are three degrees removed, based on the ground-truth network that was measured at the closest previous timepoint (October 2022). If there were not enough network members at a given distance level, additional members from the next distance level were included to complete the participant's stimulus set. Each pair was rated twice, once in each direction (i.e., ‘Is A friends with B?’ and ‘Is B friends with A?’), resulting in 870 total judgments per participant. Because customized stimulus sets overlapped across participants, dyads were evaluated repeatedly at the group level. Each unique pair was judged a median of 6 times (range: 2-38) across both directions.

4.1 Network and task demographics.

	Initial Recruitment (N =196)	**Network Knowledge Task (third-person survey; N = 100)**	**December 2022 Friendship Survey (first-person; N = 180)**	**May 2023 Friendship Survey (first-person; N = 176)**
Gender				
Female	103	54	95	96
Male	89	44	81	77
Other	4	2	4	3
Ethnicity				
Asian	54	27	51	49
Black	23	14	21	19
East Asian	5	3	4	4
Hispanic/Latinx	21	13	21	21
Middle Eastern	4	1	4	3
Mixed	18	8	16	15
Native American	1	0	1	1
South Asian	16	12	14	12
White	50	21	44	45
Other	4	1	4	4
Age				
Mean	18.1	18.1	18.1	18.1
SD	0.421	0.424	0.419	0.433
Min	17	17	17	17
Max	19	19	19	19

Statistical null models.

To rule out the possibility that accurate predictions in the Network Knowledge Task arose by chance, we tested two null models. In a permutation model, we preserved each participant's overall pattern of responses in the Network Knowledge Task but randomly shuffled their responses across dyads. This breaks any systematic relationship between participants' judgments and the ground-truth relationship outcomes (formation or loss) while preserving how frequently a participant tended to label a pair as friends or not friends. We repeated this procedure 1,000 times to produce a null distribution of model coefficients expected under chance.

Separately, we tested whether participants' predictions exceeded the level expected from their own response bias. For each participant, we generated 1,000 simulated datasets by sampling each judgment from a Bernoulli distribution parameterized by that participant's empirical rate of endorsing a pair as 'Friends' or 'Not Friends'. Each simulation preserves how often a participant tended to make each judgment, but removes any true association between those judgments and actual relationship outcomes. a null distribution of correctly identified future tie changes expected under pure response bias. This model produced a null distribution of correctly identified future tie changes expected under pure response bias.

Significance for both null models was assessed by comparing the empirical statistic to its corresponding null distribution: the model coefficient in the permutation model, or the rate of correctly identifying dyads that subsequently formed or dissolved a friendship (i.e., hit rate) in the response-bias model. Two-tailed p-values were computed as the proportion of simulated values more extreme than the empirical value.

Measures of network position.

Using the mutually confirmed friendship network as ground truth, we computed eigenvector centrality (eigencentrality) and betweenness for each participant at each timepoint using the R packages *tidygraph* v.1.3.1 (*1*) and *igraph* v.2.1.4 (*2*). Eigencentrality indexes how well-connected an individual's neighbors are, while betweenness centrality reflects how often an individual falls on the shortest path between two others. Both measures were normalized to the 0–1 range. For behavioral analyses, betweenness values were square root transformed to improve model stability. Participants who were not part of the largest connected component of the network graph at each timepoint were excluded prior to computing these centrality metrics.

Katz communicability model.

Katz communicability model. To formalize participants' representation of the social network, and to capture how they encode multistep connectivity, we modeled each participant's judgments using a representation derived from directed friendship nominations. Following existing work showing that judgments about others' relationships depend more strongly on the outbound direction of friendship nominations (i.e., who reports someone as a friend) rather than on mutually confirmed ties (*28*), we constructed a directed adjacency matrix A from the December 2022 self-reported friendships survey. Outbound nominations (M) provide a more feasible signal of relational information because they account for asymmetries in people's beliefs about their own relationships that are lost when ties are assumed to be mutual. Moreover, we discounted the relational information based on the minimum outbound-distance (d) from each pair of network members to the participant to capture people's reduced access to information about more distant network members, estimating a free-parameter (w) that reflects variation in such access. Consistent with this, Teoh et al. (*28*) found that models built on these distance-weighted outbound nominations systematically outperformed those based on mutual friendships, indicating that people rely on partial knowledge of others' relationships when inferring how people are connected.

$$A_{ij} = \begin{cases} \omega^{d_i}, & M_{ij} = 1 \wedge d_i \leq d_j \\ \omega^{d_j}, & M_{ij} = 1 \wedge d_i > d_j \\ 0, & otherwise \end{cases}$$

In our matrix, $A_{ij} = 1$ indicates that the participant is perfectly confident that network members i and j are friends, and less and less confident as $A_{ij} \to 0$. When $A_{ij} = 0$, participants are perfectly confident that members i and j are not friends. Katz communicability quantifies the extent to which two individuals are connected through direct and indirect pathways. For two nodes, i and j, Katz communicability is defined as:

$$KC_{ij} = \sum_{k=0}^{\infty} \left(\frac{\gamma}{\lambda(A)}\right)^k \times (A^k)_{ij},$$

where $(A^k)_{ij}$ is the number of directed walks of length k from i to j, qualified by participants' ability to acquire information about each of the k steps, and γ discounts the contribution of longer paths. The free parameter γ determines the degree of multistep integration and is bounded between 0 and 1. $\lambda(A)$ represents the maximum eigenvalue of the A and scales γ with reference to the participant-specific adjacency matrix of the network to ensure stable Katz series expansion across participants. Model fitting was performed with custom optimization scripts implemented in Python v3.9.18, also following procedures established in prior work (*28*–*29*).

Objective Katz communicability. To test whether the predictive signal in participants' inferred Katz communicability values reflects information beyond what the network's own structure would predict, we computed objective Katz communicability directly from the ground-truth network in December 2022. Unlike our participant-level model, which is fit to each participant's friendship judgments and yields a participant-specific γ, the objective measure uses the undirected adjacency matrix A, constructed from mutually confirmed friendships in the first-person friendship survey in December 2022. We draw on prior work to obtain the objective multistep connectivity between members of the ground-truth network (*47–48*), which estimates that the integration of multistep paths approximates 0.85 of the largest eigenvalue of adjacency matrix of the network (A). Thus, the objective Katz communicability value for each pair in the network was then computed as:

$$KC_{ij}^{\text{objective}} = \sum_{k=0}^{\infty} \left(\frac{0.85}{\lambda(A)}\right)^{k} (A^{k})_{ij}$$

Software.

All computational modeling was conducted in Python v3.9.18. using *numpy* v1.26.4, *pandas* v2.2.1, *numba* v0.59.0, *scipy* v1.12.0, *arviz* v0.16.0, and spyder v5.5.1 in *Anaconda* v2–2.4.0. All statistical analyses were conducted in RStudio v2023.12.1+402 (*6*) with R v4.5.1 (*7*). Logistic mixed-effects models were fitted to binary outcomes and Gaussian mixed-effects models to continuous outcomes using *glmmTMB* v1.1.11 (*8*), with maximal random effects where possible. Ordered beta regressions for outcomes bounded between 0 and 1 were also fitted using *glmmTMB*. Linear mixed-effects models were fitted using *lme4* v1.1–37 (*9*) with *lmerTest* v3.1–3 (*10*) for Satterthwaite degrees of freedom. Simple slopes were extracted using the *marginaleffects* package v0.31.0 (*11*). Model predictions and participant data were visualized using *ggplot2* v3.5.2 (*12*) and *ggeffects* v2.2.1 (*13*). Stimuli presentation and response recording were supported by Qualtrics in the Friendship Survey and Network Knowledge Task, and jsPsych v7.1 in the Novel Tie Prediction Task.

References

1. Smith, E. B., Brands, R. A., Brashears, M. E., & Kleinbaum, A. M. (2020). Social Networks and Cognition. Annual Review of Sociology, 46(1), 1–16. https://doi.org/10.1146/annurev-soc-121919–054736
2. Freeman, L. C. (1992). Filling in the Blanks: A Theory of Cognitive Categories and the Structure of Social Affiliation. Social Psychology Quarterly, 55(2), 118–127. https://doi.org/10.2307/2786941
3. Krackhardt, D., & Kilduff, M. (1999). Whether close or far: Social distance effects on perceived balance in friendship networks. Journal of Personality and Social Psychology, 76(5), 770–782. https://doi.org/10.1037/0022–3514.76.5.770
4. Brewer, D. D., & Webster, C. M. (2000). Forgetting of friends and its effects on measuring friendship networks. Social Networks, 21(4), 361–373. https://doi.org/10.1016/S0378–8733(99)00018–0
5. Killworth, P. D., & Bernard, H. R. (1976). Informant Accuracy in Social Network Data. Human Organization, 35(3), 269–286. https://doi.org/10.17730/humo.35.3.10215j2m359266n2
6. Bernard, H. R., Killworth, P. D., & Sailer, L. (1982). Informant accuracy in social-network data V. An experimental attempt to predict actual communication from recall data. Social Science Research, 11(1), 30–66. https://doi.org/10.1016/0049–089X(82)90006–0
7. Brashears, M. E. (2013). Humans use Compression Heuristics to Improve the Recall of Social Networks. Scientific Reports, 3(1), 1513. https://doi.org/10.1038/srep01513
8. Freeman, L. C., Romney, A. K., & Freeman, S. C. (1987). Cognitive Structure and Informant Accuracy. American Anthropologist, 89(2), 310–325. https://doi.org/10.1525/aa.1987.89.2.02a00020
9. Son, J.-Y., Bhandari, A., & FeldmanHall, O. (2023). Abstract cognitive maps of social network structure aid adaptive inference. Proceedings of the National Academy of Sciences, 120(47), e2310801120. https://doi.org/10.1073/pnas.2310801120
10. Aslarus, I. C., Son, J.-Y., Xia, A., & FeldmanHall, O. (2025). Early insight into social network structure predicts climbing the social ladder. Science Advances, 11(25), eads2133. https://doi.org/10.1126/sciadv.ads2133
11. Bartlett, F. C. (1932). Remembering: A study in experimental and social psychology (pp. xix, 317). Cambridge University Press.
12. Ghosh, V. E., & Gilboa, A. (2014). What is a memory schema? A historical perspective on current neuroscience literature. Neuropsychologia, 53, 104–114. https://doi.org/10.1016/j.neuropsychologia.2013.11.010
13. Preston, A. R., & Eichenbaum, H. (2013). Interplay of Hippocampus and Prefrontal Cortex in Memory. Current Biology, 23(17), R764–R773. https://doi.org/10.1016/j.cub.2013.05.041
14. Brashears, M. E., & Quintane, E. (2015). The microstructures of network recall: How social networks are encoded and represented in human memory. Social Networks, 41, 113–126. https://doi.org/10.1016/j.socnet.2014.11.003

15. Hu, Y.-F., Vives, M.-L., Bhandari, A., & FeldmanHall, O. (2025). Using Statistical Learning to Navigate Signed Social Networks (2fpt5_v1). OSF Preprints. https://doi.org/10.31219/osf.io/2fpt5_v1
16. 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
17. Krackhardt, D. (1987). Cognitive social structures. Social Networks, 9(2), 109–134. https://doi.org/10.1016/0378–8733(87)90009–8
18. Janicik, G., & Larrick, R. (2005). Social Network Schemas and the Learning of Incomplete Networks. Journal of Personality and Social Psychology, 88, 348–364. https://doi.org/10.1037/0022–3514.88.2.348
19. 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
20. FeldmanHall, O., Son, J.-Y., & Bhandari, A. (2025). Abstract Cognitive Maps for Complex Social Systems. Current Directions in Psychological Science, 34(6), 349–356. https://doi.org/10.1177/09637214251342742
21. Clark, A. (2013). Expecting the World: Perception, Prediction, and the Origins of Human Knowledge. The Journal of Philosophy, 110(9), 469–496.
22. Kanizsa, G. (1976). Subjective Contours. Scientific American, 234(4), 48–53.
23. Rao, R. P. N., & Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. Nature Neuroscience, 2(1), 79–87. https://doi.org/10.1038/4580
24. Newman, E. J., & Lindsay, D. S. (2009). False memories: What the hell are they for? Applied Cognitive Psychology, 23(8), 1105–1121. https://doi.org/10.1002/acp.1613
25. Schacter, D. L. (2012). Adaptive constructive processes and the future of memory. The American Psychologist, 67(8), 10.1037/a0029869. https://doi.org/10.1037/a0029869
26. Addis, D. R., Wong, A. T., & Schacter, D. L. (2007). Remembering the past and imagining the future: Common and distinct neural substrates during event construction and elaboration. Neuropsychologia, 45(7), 1363–1377. https://doi.org/10.1016/j.neuropsychologia.2006.10.016
27. Thornton, M. A., & Tamir, D. I. (2024). Neural representations of situations and mental states are composed of sums of representations of the actions they afford. Nature Communications, 15(1), 620. https://doi.org/10.1038/s41467–024–44870–7
28. Teoh, Y. Y., Son, J.-Y., Xia, A., Bhandari, A., & FeldmanHall, O. (2025). Medial temporal lobe encodes cognitive maps of real-world social networks (p. 2025.03.18.643925). bioRxiv. https://doi.org/10.1101/2025.03.18.643925
29. Xia, A., Teoh, Y. Y., Nassar, M. R., Bhandari, A., & FeldmanHall, O. (2025). Knowledge of information cascades through social networks facilitates strategic gossip. Nature Human Behaviour, 9(10), 2169–2182. https://doi.org/10.1038/s41562–025–02241–2
30. Granovetter, M. S. (1973). The strength of weak ties. American Journal of Sociology, 78(6), 1360–1380.
31. Kossinets, G., & Watts, D. J. (2006). Empirical Analysis of an Evolving Social Network. Science, 311(5757), 88–90. https://doi.org/10.1126/science.1116869

32. Mosleh, M., Eckles, D., & Rand, D. G. (2025). Tendencies toward triadic closure: Field experimental evidence. Proceedings of the National Academy of Sciences, 122(27), e2404590122. https://doi.org/10.1073/pnas.2404590122
33. Franz, V. H., & Loftus, G. R. (2012). Standard errors and confidence intervals in within-subjects designs: Generalizing Loftus and Masson (1994) and avoiding the biases of alternative accounts. Psychonomic Bulletin & Review, 19(3), 395–404. https://doi.org/10.3758/s13423–012–0230–1
34. McPherson, M., Smith-Lovin, L., & Cook, J. M. (2001). Birds of a Feather: Homophily in Social Networks. Annual Review of Sociology, 27, 415–444.
35. Lewis, K. (2015). How Networks Form: Homophily, Opportunity, and Balance (pp. 1–14). https://doi.org/10.1002/9781118900772.etrds0164
36. 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, 118(39), e2021699118. https://doi.org/10.1073/pnas.2021699118
37. Katz, L. (1953). A New Status Index Derived from Sociometric Analysis. Psychometrika, 18(1), 39–43. https://doi.org/10.1007/BF02289026
38. Casciaro, T. (1998). Seeing things clearly: Social structure, personality, and accuracy in social network perception. Social Networks, 20(4), 331–351. https://doi.org/10.1016/s0378–8733(98)00008–2
39. Liben-Nowell, D., & Kleinberg, J. (2003). The link prediction problem for social networks. Proceedings of the Twelfth International Conference on Information and Knowledge Management, 556–559. https://doi.org/10.1145/956863.956972
40. Lichtenwalter, R. N., Lussier, J. T., & Chawla, N. V. (2010). New perspectives and methods in link prediction. Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '10, 243–252. https://doi.org/10.1145/1835804.1835837
41. Liu, Z., Zhang, Q.-M., Lü, L., & Zhou, T. (2011). Link prediction in complex networks: A local naïve Bayes model. Europhysics Letters, 96(4), 48007. https://doi.org/10.1209/0295–5075/96/48007
42. Kipf, T. N., & Welling, M. (2017, February 6). Semi-Supervised Classification with Graph Convolutional Networks. International Conference on Learning Representations. https://openreview.net/forum?id=SJU4ayYgl
43. Perozzi, B., Al-Rfou, R., & Skiena, S. (2014). DeepWalk: Online learning of social representations. Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 701–710. https://doi.org/10.1145/2623330.2623732
44. Grover, A., & Leskovec, J. (2016). node2vec: Scalable Feature Learning for Networks. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '16, 855–864. https://doi.org/10.1145/2939672.2939754
45. Pedersen, T. (2025). tidygraph: A Tidy API for Graph Manipulation. R package version 1.3.1.9000. https://github.com/thomasp85/tidygraph
46. Csárdi, G., Nepusz, T., Müller, K., Horvát, S., Traag, V., Zanini, F., & Noom, D. (2024). igraph for R: R interface of the igraph library for graph theory and network analysis. Zenodo. https://doi.org/10.5281/zenodo.10681749

47. Benzi, M., & Klymko, C. (2013). Total communicability as a centrality measure. Journal of Complex Networks, 1(2), 124–149. https://doi.org/10.1093/comnet/cnt007
48. Garvert, M. M., Dolan, R. J., & Behrens, T. E. (2017). A map of abstract relational knowledge in the human hippocampal–entorhinal cortex. eLife, 6, e17086. https://doi.org/10.7554/eLife.17086
49. Posit team (2025). RStudio: Integrated Development Environment for R. Posit Software, PBC, Boston, MA. http://www.posit.co/
50. R Core Team (2025). R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria. https://www.R-project.org/
51. Brooks, M., Kristensen, K., van Benthem, K., Magnusson, A., Berg, C., Nielsen, A., Skaug, H., Mächler, M., & Bolker, B. (2017). glmmTMB Balances Speed and Flexibility Among Packages for Zero-inflated Generalized Linear Mixed Modeling. The R Journal. https://digitalcommons.unl.edu/r-journal/675
52. Bates, D., Maechler, M., Bolker, B., Walker, S., Jagan, M., & Ly, A. (2003). lme4: Linear Mixed-Effects Models using "Eigen" and S4. https://doi.org/10.32614/CRAN.package.lme4
53. Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. (2017). lmerTest package: tests in linear mixed effects models. Journal of Statistical Software, 82, 1–26.
54. Arel-Bundock, V., Greifer, N., & Heiss, A. (2024). How to Interpret Statistical Models Using marginaleffects for R and Python. Journal of Statistical Software, 111, 1–32. https://doi.org/10.18637/jss.v111.i09
55. Wickham, H. (2011). ggplot2. WIREs Computational Statistics, 3(2), 180–185. https://doi.org/10.1002/wics.147
56. Lüdecke, D. (2018). ggeffects: Tidy data frames of marginal effects from regression models. Journal of Open Source Software, 3(26), 772.

Supplement

Supplementary Table 1. Effects of friendship judgment errors on future tie formation in the first wave of the Friendship Judgment Task (Nov 2022).

	Tie formation		
Predictors	*Log-Odds*	*CI*	*p*
(Intercept)	-6.03	-6.48 – -5.58	**<0.001**
Friendship judgment	1.11	0.84 – 1.38	**<0.001**
Random Effects			
σ^2	3.29		
$\tau_{00\ \mathrm{rater}}$	0.04		
$\tau_{00\ \mathrm{to}}$	2.83		
$\tau_{00\ \mathrm{from}}$	2.80		
$\tau_{11\ \mathrm{rater.friend\_guess1}}$	0.92		
$\rho_{01\ \mathrm{rater}}$	-0.60		
ICC	0.63		
N_{rater}	100		
N_{to}	162		
N_{from}	162		
Observations	66156		
Marginal R^2 / Conditional R^2	0.008 / 0.638		

Supplementary Table 2. Robustness of friendship judgment effect on tie formation in intermediate time windows.

	Tie formation		
Predictors	*Log-Odds*	*CI*	*p*
(Intercept)	-3.80	-3.89 – -3.71	**<0.001**
Friendship judgment	1.04	0.76 – 1.32	**<0.001**
timeApr	-0.33	-0.41 – -0.24	**<0.001**
timeFeb	-0.36	-0.45 – -0.27	**<0.001**
Friendship judgment:April	0.06	-0.16 – 0.29	0.582
Friendship judgment:Feb	0.11	-0.12 – 0.34	0.338
Random Effects			
σ^2	3.29		
$\tau_{00\ \text{rater}}$	0.12		
$\tau_{11\ \text{rater.friend\_guess1}}$	1.02		
$\tau_{11\ \text{rater.timeApr}}$	0.00		
$\tau_{11\ \text{rater.timeFeb}}$	0.01		
ρ_{01}	0.00		
	0.00		
	0.00		
ICC	0.05		
N_{rater}	100		
Observations	179046		
Marginal R^2 / Conditional R^2	0.026 / 0.077		

Supplementary Table 3. Replication of friendship judgment effects on tie formation in second wave of the Friendship Judgment Task (Mar 2023).

	Tie formation		
Predictors	*Log-Odds*	*CI*	*p*
(Intercept)	-7.81	-8.55 – -7.07	**<0.001**
Friendship judgment	1.39	1.15 – 1.62	**<0.001**
Random Effects			
σ^2	3.29		
$\tau_{00\ \mathrm{rater}}$	0.01		
$\tau_{00\ \mathrm{to}}$	5.49		
$\tau_{00\ \mathrm{from}}$	5.53		
$\tau_{11\ \mathrm{rater.friend\_guess1}}$	0.59		
$\rho_{01\ \mathrm{rater}}$	0.00		
ICC	0.77		
N_{rater}	80		
N_{to}	152		
N_{from}	152		
Observations	48030		
Marginal R^2 / Conditional R^2	0.019 / 0.776		

Supplementary Results 1. Novel Tie Prediction Task.

Procedure. Participants were told that a hypothetical transfer student had become friends with two existing network members (a "reference pair") and were asked to judge whether the transfer student was also friends with a third network member (a "probe"; Supplementary Fig. 3A). Critically, because the transfer student is fictional, participants could not rely on observed interactions or any ground-truth social information. Instead, inferences depend on estimating the likelihood of friendship based on the relative network positions of the reference pair and the probe. To evaluate predictive validity, we identified real dyads whose structural configuration mirrored each reference pair-probe combination—specifically, cases where two individuals (analogous to the reference pair) shared a mutual friend (the proxy, analogous to the transfer student), and a third person (analogous to the probe) who was not friends with the proxy in December.

Novel tie prediction responses reflect true network structure. We first confirmed that task responses reflected knowledge of true network structure. Participants were more likely to infer friendship between the transfer student and the probe when the probe was friends with a larger proportion of the reference pair's mutual friends ($\beta = 1.581$, $SE = 0.187$, $z = 8.467$, $p < .001$; Fig. Supplementary 3B), accounting for the number of mutual friends the pair had.

Novel tie prediction responses predict future tie formation. Having established that task responses reflect genuine structural knowledge, we tested whether they predict real future tie formation. When participants judged the transfer student and probe as more likely to be friends, the analogous real-world probe-proxy pair was also more likely to become friends months later ($\beta = 0.464$, $SE = 0.070$, $z = 6.582$, $p < .001$; Fig. Supplementary 3C). Because these judgements are made about a fictional individual, they cannot reflect privileged knowledge of real-world interactions, providing strong evidence that participants extract structural patterns from the network and use them to make genuine predictions about how relationships in their real-world social environment will change.

Novel tie prediction responses do not predict tie loss. We next tested whether predictions also forecast the dissolution of existing friendships—that is, whether judging the transfer student and probe as *unlikely* to be friends predicted that the analogous real-world probe-proxy pair would lose their friendship by May. Although predictions forecasted tie loss in the expected direction, the effect was not significant (β = -0.077, SE = 0.083, z = -0.931, p = .352), suggesting that the structural patterns participants extract from the network are specifically informative about friendship formation rather than relationship change more broadly.

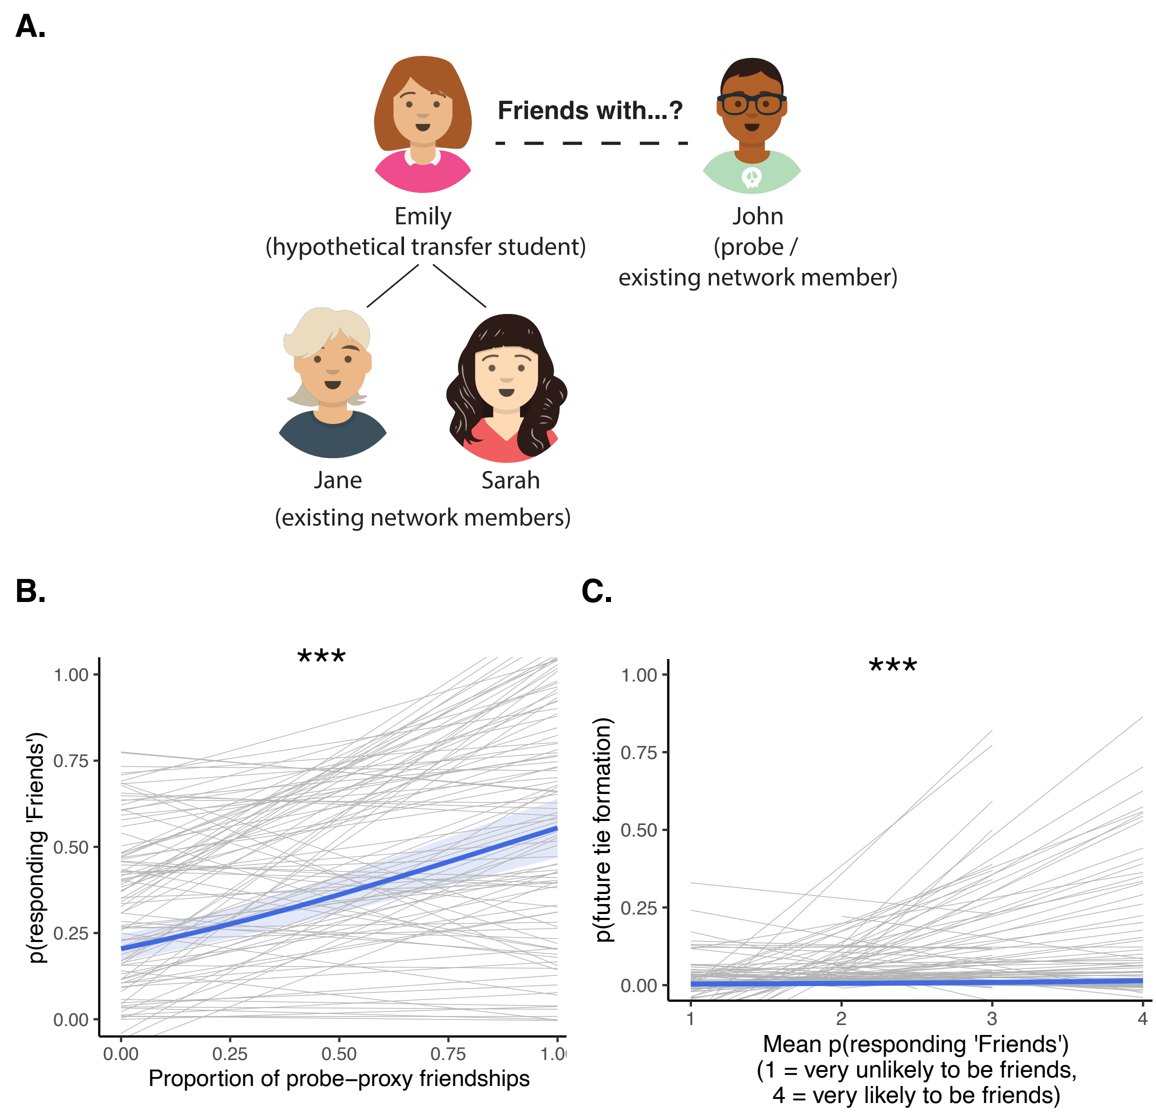

Supplementary Figure 3. Novel tie prediction task. (A) Task schematic. Participants judged whether a hypothetical transfer student, who is already friends with two existing network members, would also be friends with a third member (the probe). **(B)** Probability of friendship inference as a function of the proportion of real probe-proxy friendships in the network. Individual lines reflect per-subject linear fits; the bold line reflects the fixed-effects prediction from a mixed-effects logistic regression controlling for the number of proxy friends presented ($\beta = 1.58$, $z = 8.47$, $p < .001$). **(C)** Probability of future tie formation as a function of mean friendship inference rating in the novel tie prediction task. Points reflect binned means ± SE; the line reflects fixed-effects predictions from a mixed-effects logistic regression ($\beta = 0.46$, $z = 6.58$, $p < .001$).

Supplementary Table 4. Effects of friendship judgment errors on future tie loss in the primary (first wave) of the Friendship Judgment Task (Nov 2022).

	Tie loss		
Predictors	*Log-Odds*	*CI*	*p*
(Intercept)	-1.74	-2.42 – -1.07	**<0.001**
Friendship judgment	-0.55	-0.76 – -0.35	**<0.001**
Random Effects			
σ^2	3.29		
$\tau_{00\ \text{rater}}$	0.00		
$\tau_{00\ \text{to}}$	7.57		
$\tau_{00\ \text{from}}$	7.52		
$\tau_{11\ \text{rater.friend\_guess1}}$	0.21		
$\rho_{01\ \text{rater}}$	0.00		
ICC	0.82		
N_{rater}	100		
N_{to}	151		
N_{from}	151		
Observations	8146		
Marginal R^2 / Conditional R^2	0.003 / 0.822		

Supplementary Table 5. Robustness of friendship judgment effect on tie loss in intermediate time windows.

	Tie loss		
Predictors	*Log-Odds*	*CI*	*p*
(Intercept)	-1.27	-1.70 – -0.83	**<0.001**
Friendship judgment	-0.36	-0.57 – -0.15	**0.001**
timeFeb	-0.79	-0.91 – -0.66	**<0.001**
timeApr	-0.01	-0.12 – 0.10	0.805
Friendship judgment:Feb	-0.34	-0.58 – -0.10	**0.005**
Friendship judgment:Apr	-0.21	-0.43 – 0.01	0.056
Random Effects			
σ^2	3.29		
$\tau_{00\ \text{rater}}$	0.02		
$\tau_{00\ \text{to}}$	3.23		
$\tau_{00\ \text{from}}$	3.19		
$\tau_{11\ \text{rater.friend\_guess1}}$	0.43		
$\tau_{11\ \text{rater.timeFeb}}$	0.13		
$\tau_{11\ \text{rater.timeApr}}$	0.07		
$\rho_{01\ \text{rater.friend\_guess1}}$	0.00		
$\rho_{01\ \text{rater.timeFeb}}$	0.00		
$\rho_{01\ \text{rater.timeApr}}$	0.00		
ICC	0.67		
N_{rater}	100		
N_{to}	147		
N_{from}	147		
Observations	22194		
Marginal R^2 / Conditional R^2	0.021 / 0.675		

Supplementary Table 6. Replication of friendship judgment effects on tie loss in second wave of the Friendship Judgment Task (Mar 2023).

	Tie loss		
Predictors	*Log-Odds*	*CI*	*p*
(Intercept)	-10.28	-13.45 – -7.12	**<0.001**
Friendship judgment	-0.45	-0.68 – -0.23	**<0.001**
Random Effects			
σ^2	3.29		
$\tau_{00\ \text{rater}}$	0.00		
$\tau_{00\ \text{to}}$	47.09		
$\tau_{00\ \text{from}}$	47.06		
$\tau_{11\ \text{rater.friend\_guess1}}$	0.00		
$\rho_{01\ \text{rater}}$	0.00		
ICC	0.97		
N_{rater}	80		
N_{to}	143		
N_{from}	143		
Observations	6758		
Marginal R^2 / Conditional R^2	0.001 / 0.966		

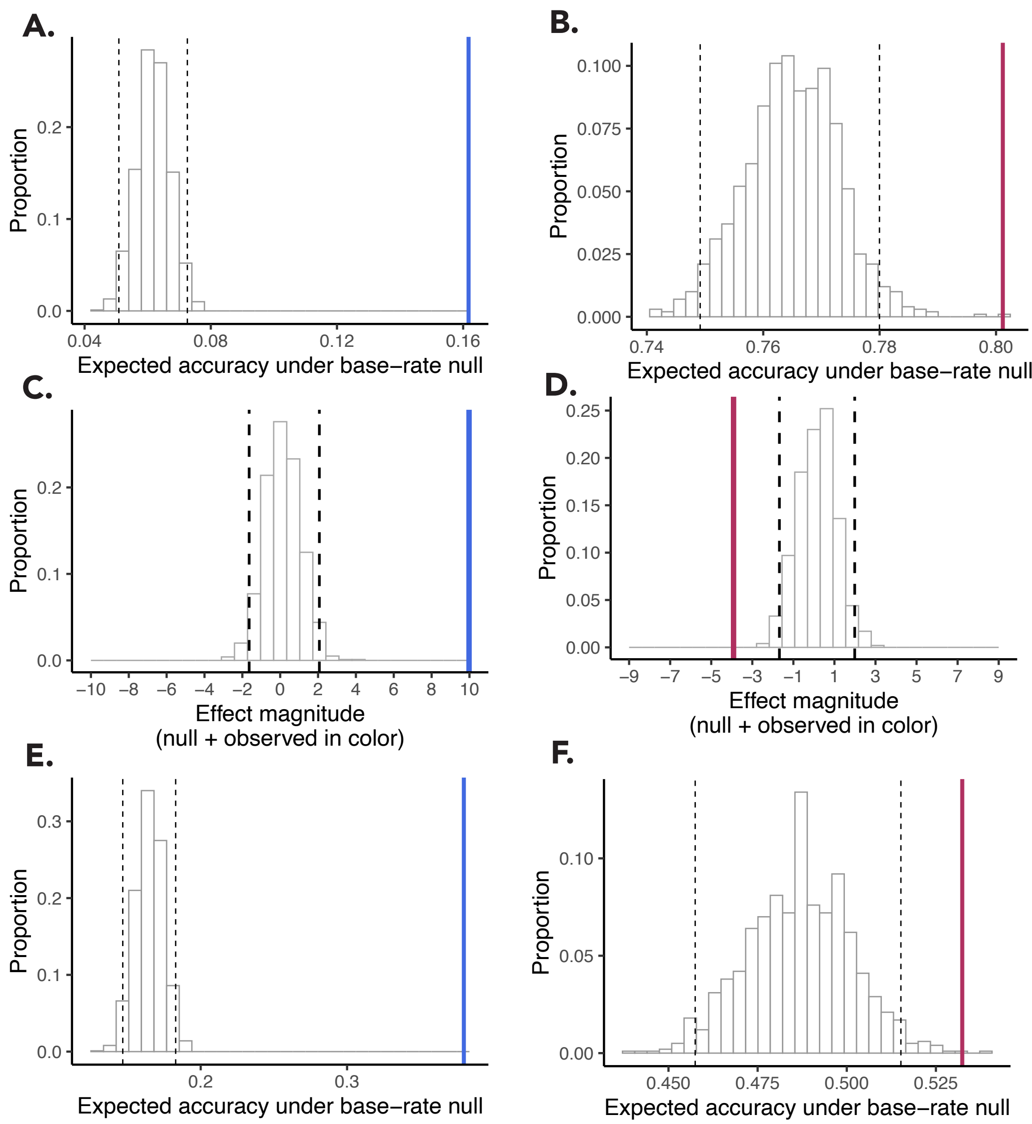

Supplementary Figure 2. Results of null model comparisons. **(A)** Base-rate null for tie formation from the 1st wave of the Friendship Judgment Task (observed hit rate = 0.162; $p < .001$). **(B)** Base-rate null for tie loss from the 1st wave of the Friendship Judgment Task (observed accuracy = 0.801; $p = .001$). **(C)** Permutation null for tie formation from the 2nd wave of the Friendship Judgment Task (observed $z = 9.978$; $p < .001$). **(D)** Permutation null for tie loss from the 2nd wave of the Network Knowledge Task (observed $z = -3.913$; $p < .001$). **(E)** Base-rate null for tie formation from the 2nd wave of the Network Knowledge Task (observed hit rate = 0.380; $p < .001$). **(F)** Base-rate null for tie loss from the 2nd wave of the Network Knowledge Task (observed accuracy = 0.532; $p = .001$). In all panels, the colored vertical line shows the observed value and dashed lines denote the 95% CI of the null distribution. For permutation null panels (C–D), the null was generated by randomly shuffling each participant's judgments 1,000 times. For base-rate null panels (A–B, E–F), the null was generated by simulating predictions based solely on each rater's tendency to judge a pair as 'Friends' or 'Not Friends'. In all cases, observed values fell outside the null distribution (all $p \leq .001$).

Supplementary Table 7. Effect of friendship judgment on tie formation, controlling for number of mutual friends.

	Tie loss		
Predictors	*Log-Odds*	*CI*	*p*
(Intercept)	-10.28	-13.45 – -7.12	**<0.001**
Friendship judgment	-0.45	-0.68 – -0.23	**<0.001**
Random Effects			
σ^2	3.29		
$\tau_{00\ \mathrm{rater}}$	0.00		
$\tau_{00\ \mathrm{to}}$	47.09		
$\tau_{00\ \mathrm{from}}$	47.06		
$\tau_{11\ \mathrm{rater.friend\_guess1}}$	0.00		
$\rho_{01\ \mathrm{rater}}$	0.00		
ICC	0.97		
N_{rater}	80		
N_{to}	143		
N_{from}	143		
Observations	6758		
Marginal R^2 / Conditional R^2	0.001 / 0.966		

Supplementary Table 8. Effect of friendship judgment on tie loss, controlling for number of mutual friends.

	Tie loss		
Predictors	*Log-Odds*	*CI*	*p*
(Intercept)	-1.61	-3.76 – 0.55	0.144
Number of mutual friends in Dec	-1.35	-1.52 – -1.17	**<0.001**
Friendship judgment	-0.35	-0.70 – -0.01	**0.044**
Random Effects			
σ^2	3.29		
$\tau_{00\ rater}$	0.00		
$\tau_{00\ to}$	33.38		
$\tau_{00\ from}$	57.14		
$\tau_{11\ rater.n\_mf\_ego3}$	0.00		
$\tau_{11\ rater.friend\_guess1}$	0.00		
$\rho_{01\ rater.n\_mf\_ego3}$	0.00		
$\rho_{01\ rater.friend\_guess1}$	0.00		
ICC	0.96		
N_{rater}	100		
N_{to}	123		
N_{from}	127		
Observations	4073		
Marginal R^2 / Conditional R^2	0.128 / 0.969		

Supplementary Table 9. Effect of friendship judgment on tie formation, controlling for number of mutual friends and homophily between rated pairs.

	Tie formation		
Predictors	*Log-Odds*	*CI*	*p*
(Intercept)	-7.86	-8.62 – -7.11	**<0.001**
Shared gender	0.12	-0.05 – 0.29	0.175
Shared ethnicity	0.40	0.17 – 0.64	**0.001**
Shared religion	-0.50	-0.74 – -0.25	**<0.001**
Shared dorm floor	0.79	0.55 – 1.03	**<0.001**
Number of mutual friends in Dec	0.82	0.75 – 0.88	**<0.001**
Friendship judgment	0.45	0.20 – 0.70	**<0.001**
Random Effects			
σ^2	3.29		
$\tau_{00\ \mathrm{rater}}$	0.00		
$\tau_{00\ \mathrm{from}}$	4.63		
$\tau_{00\ \mathrm{to}}$	4.77		
$\tau_{11\ \mathrm{rater.n\_mf\_ego3}}$	0.00		
$\tau_{11\ \mathrm{rater.friend\_guess1}}$	0.00		
$\rho_{01\ \mathrm{rater.n\_mf\_ego3}}$	0.00		
$\rho_{01\ \mathrm{rater.friend\_guess1}}$	0.00		
ICC	0.74		
N_{rater}	100		
N_{from}	161		
N_{to}	161		
Observations	33078		
Marginal R^2 / Conditional R^2	0.074 / 0.760		

Supplementary Table 10. Effect of friendship judgment on tie loss, controlling for number of mutual friends and homophily between rated pairs.

	Tie loss		
Predictors	*Log-Odds*	*CI*	*p*
(Intercept)	-1.16	-3.51 – 1.20	0.336
Shared gender	-0.71	-1.14 – -0.29	**0.001**
Shared ethnicity	0.94	0.33 – 1.56	**0.003**
Shared religion	-0.63	-1.13 – -0.13	**0.013**
Shared dorm floor	-0.57	-1.14 – 0.00	0.051
Number of mutual friends in Dec	-1.45	-1.64 – -1.26	**<0.001**
Friendship judgment	-0.32	-0.67 – 0.03	0.071
Random Effects			
σ^2	3.29		
$\tau_{00\ rater}$	0.00		
$\tau_{00\ from}$	67.28		
$\tau_{00\ to}$	39.18		
$\tau_{11\ rater.n\_mf\_ego3}$	0.00		
$\tau_{11\ rater.friend\_guess1}$	0.00		
$\rho_{01\ rater.n\_mf\_ego3}$	0.00		
$\rho_{01\ rater.friend\_guess1}$	0.00		
ICC	0.97		
N_{rater}	100		
N_{from}	127		
N_{to}	123		
Observations	4073		
Marginal R^2 / Conditional R^2	0.129 / 0.974		

Supplementary Table 11. Effect of model-estimated pairwise Katz communicability values on tie formation.

	Tie formation		
Predictors	*Log-Odds*	*CI*	*p*
(Intercept)	-8.14	-8.92 – -7.35	**<0.001**
Shared ethnicity	0.39	0.14 – 0.63	**0.002**
Shared gender	0.13	-0.05 – 0.30	0.163
Shared religion	-0.52	-0.77 – -0.27	**<0.001**
Shared dorm floor	0.38	0.12 – 0.64	**0.004**
Katz communicability	0.33	0.28 – 0.38	**<0.001**
Number of mutual friends in Dec	0.71	0.64 – 0.78	**<0.001**
Friendship judgment	0.07	-0.20 – 0.34	0.611
Random Effects			
σ^2	3.29		
$\tau_{00\ \mathrm{rater}}$	0.00		
$\tau_{00\ \mathrm{from}}$	4.85		
$\tau_{00\ \mathrm{to}}$	5.06		
$\tau_{11\ \mathrm{rater.friend\_guess1}}$	0.00		
$\tau_{11\ \mathrm{rater.mean\_kc\_scaled}}$	0.00		
$\rho_{01\ \mathrm{rater.friend\_guess1}}$	0.00		
$\rho_{01\ \mathrm{rater.mean\_kc\_scaled}}$	0.00		
ICC	0.75		
N_{rater}	100		
N_{from}	161		
N_{to}	161		
Observations	33078		
Marginal R^2 / Conditional R^2	0.084 / 0.772		

Supplementary Table 12. Effect of model-estimated pairwise Katz communicability values on tie loss.

	Tie loss		
Predictors	*Log-Odds*	*CI*	*p*
(Intercept)	-1.28	-3.66 – 1.10	0.293
Shared ethnicity	0.95	0.33 – 1.56	**0.002**
Shared gender	-0.73	-1.15 – -0.30	**0.001**
Shared religion	-0.63	-1.13 – -0.13	**0.014**
Shared dorm floor	-0.61	-1.19 – -0.04	**0.037**
Katz communicability	0.06	-0.00 – 0.12	0.069
Number of mutual friends in Dec	-1.47	-1.66 – -1.27	**<0.001**
Friendship judgment	-0.45	-0.82 – -0.07	**0.019**
Random Effects			
σ^2	3.29		
$\tau_{00\ rater}$	0.00		
$\tau_{00\ to}$	39.73		
$\tau_{00\ from}$	68.25		
$\tau_{11\ rater.friend\_guess1}$	0.00		
$\tau_{11\ rater.mean\_kc\_scaled}$	0.00		
$\rho_{01\ rater.friend\_guess1}$	0.00		
$\rho_{01\ rater.mean\_kc\_scaled}$	0.00		
ICC	0.97		
N_{rater}	100		
N_{to}	123		
N_{from}	127		
Observations	4073		
Marginal R^2 / Conditional R^2	0.131 / 0.974		

Supplementary Table 13. Effect of Katz communicability values estimated from ground-truth network structure in Dec 2022 on tie formation.

	Tie formation		
Predictors	*Log-Odds*	*CI*	*p*
(Intercept)	-7.56	-8.32 – -6.80	**<0.001**
Shared ethnicity	0.42	0.18 – 0.66	**0.001**
Shared gender	0.12	-0.06 – 0.29	0.183
Shared religion	-0.39	-0.64 – -0.15	**0.002**
Shared dorm floor	0.54	0.29 – 0.79	**<0.001**
Katz communicability	0.35	0.29 – 0.40	**<0.001**
Friendship judgment	0.09	-0.18 – 0.35	0.508
Ground-truth Katz communicability	0.92	0.82 – 1.02	**<0.001**
Random Effects			
σ^2	3.29		
$\tau_{00\ rater}$	0.00		
$\tau_{00\ from}$	4.62		
$\tau_{00\ to}$	4.86		
$\tau_{11\ rater.friend\_guess1}$	0.00		
$\tau_{11\ rater.mean\_kc\_scaled}$	0.00		
$\rho_{01\ rater.friend\_guess1}$	0.00		
$\rho_{01\ rater.mean\_kc\_scaled}$	0.00		
ICC	0.74		
N_{rater}	100		
N_{from}	161		
N_{to}	161		
Observations	33078		
Marginal R^2 / Conditional R^2	0.104 / 0.769		

Supplementary Table 14. Linear mixed model predicting d′ from judgment type (tie formation vs. tie loss)

	d'		
Predictors	*Estimate*	*CI*	*p*
(Intercept)	0.27	0.15 – 0.39	**<0.001**
Formation vs. Loss	0.45	0.29 – 0.61	**<0.001**
Random Effects			
σ^2	0.34		
$\tau_{00\ \mathrm{rater}}$	0.04		
ICC	0.10		
N_{rater}	100		
Observations	200		
Marginal R^2 / Conditional R^2	0.120 / 0.207		

Supplementary Table 15a. Effects of γ (estimated from Katz centrality model) on predictive accuracy (d′) across tie formation, tie loss, and current friendships in the Friendship Judgment Task.

	d′		
Predictors	*Estimate*	*CI*	*p*
(Intercept)	0.27	0.15 – 0.39	**<0.001**
Katz gamma	-0.94	-2.03 – 0.16	0.093
Formation	0.45	0.30 – 0.61	**<0.001**
Katz gamma * Formation vs. Loss	2.32	0.89 – 3.75	**0.002**
Random Effects			
σ^2	0.31		
$\tau_{00\ rater}$	0.05		
ICC	0.15		
N_{rater}	100		
Observations	200		
Marginal R^2 / Conditional R^2	0.157 / 0.281		

Supplementary Table 15b. Marginal effects of Katz model-estimated γ on d′ by judgment type.

Judgment type	Estimate	SE	z	p	2.50%	97.50%
Formation	1.387	0.555	2.501	0.012	0.3	2.474
Loss	-0.936	0.555	-1.687	0.092	-2.023	0.152

Supplementary Table 16a. Effects of tie formation prediction accuracy (d′) and initial eigencentrality on future eigencentrality.

	Future eigencentrality (May 2023)		
Predictors	*Estimate*	*CI*	*p*
Intercept	-3.11	-3.56 – -2.67	**<0.001**
Initial eigencentrality (Dec 2022)	6.74	5.15 – 8.33	**<0.001**
d′formation	-0.45	-0.85 – -0.05	**0.028**
d′loss	0.19	-0.21 – 0.59	0.358
d′current	0.18	-0.26 – 0.62	0.433
Initial eigencentrality * d′formation	1.91	0.75 – 3.07	**0.001**
Initial eigencentrality * d′loss	-0.81	-1.90 – 0.28	0.146
Initial eigencentrality * d′current	-1.35	-2.72 – 0.03	0.055
Observations	88		
R^2	0.864		

Supplementary Table 16b. Simple slopes of d′ formation on future eigencentrality by initial eigencentrality.

Initial eigencentrality	Estimate	SE	z	p	2.50%	97.50%
Low (-1 SD)	-0.448	0.204	-2.192	0.028	-0.848	-0.047
Mean	0.014	0.12	0.117	0.907	-0.221	0.249
High (+1 SD)	0.513	0.175	2.927	0.003	0.17	0.857

Supplementary Table 17a. Effects of tie formation prediction accuracy (d′) and initial betweenness on future betweenness.

	Future eigencentrality (May 2023)		
Predictors	*Estimate*	*CI*	*p*
Intercept	-3.11	-3.56 – -2.67	**<0.001**
Initial eigencentrality (Dec 2022)	6.74	5.15 – 8.33	**<0.001**
d′formation	-0.45	-0.85 – -0.05	**0.028**
d′loss	0.19	-0.21 – 0.59	0.358
d′current	0.18	-0.26 – 0.62	0.433
Initial eigencentrality * d′formation	1.91	0.75 – 3.07	**0.001**
Initial eigencentrality * d′loss	-0.81	-1.90 – 0.28	0.146
Initial eigencentrality * d′current	-1.35	-2.72 – 0.03	0.055
Observations	88		
R^2	0.864		

Supplementary Table 17b. Simple slopes of d′ formation on future betweenness by initial betweenness.

Initial betweenness	**Estimate**	**SE**	**z**	**p**	**2.50%**	**97.50%**
Low (-1 SD)	0.023	0.016	1.435	0.151	-0.009	0.056
Mean	-0.012	0.011	-1.063	0.288	-0.034	0.01
High (+1 SD)	-0.048	0.017	-2.848	0.004	-0.081	-0.015

Chapter 4

Asymmetric social ties are dynamic, detectable, and consequential for inference

Alice Xia*, Haoxue Fan*, Matthew R. Nassar, & Oriel FeldmanHall

Preprinted: https://osf.io/preprints/psyarxiv/8bfsp_v1 (2026)

1. Introduction

Two people holding fundamentally different views of the same relationship is one of the most basic yet understudied features of social life. One person may perceive a close friendship where the other sees mere acquaintanceship—a phenomenon we refer to as relational asymmetry. The origins of such asymmetries are multifaceted, possibly reflecting disparities in social standing[1], or differences in the rate at which individuals update their beliefs about a relationship. This dynamic updating may contribute to the prevalence of relational asymmetry in social network, which comprises up to 50% of friendship nominations by some estimates[2]. This high degree of asymmetry is surprising given that unreciprocated ties produce uncertainty about the nature of the relationship. Interpersonal uncertainty typically motivates people to resolve it, placing pressure on the dyad to move toward either mutual closeness or mutual disengagement[3,4]. When two people hold discrepant views of their relationship, this mismatch should trigger information seeking and direct interaction that push the relationship toward convergence[6-9].

The puzzle deepens when we consider a second empirical regularity: asymmetric ties are not only common but appear stable at the network level, with a roughly constant proportion of unreciprocated ties observed across successive snapshots[9–11]. Social Network Science (SNS) has largely treated this stability as a structural property, assuming that asymmetry in a dyad reflects directional relations such as hierarchy or authority—as when unreciprocated citations signal expertise[12,13] or one-sided advice nominations reveal organizational rank[14]. As a result, research has predominantly focused on mutual ties[15–18] and research on directed ties have proceeded largely in parallel[19–21], with little attention to how asymmetries persist despite expectations of reciprocity. An alternative account is that the apparent stability of relational asymmetry reflects not fixed structure but dynamic equilibrium—an ongoing process in which individual dyads continuously cycle through varying degrees of asymmetry as some relationships resolve toward reciprocity while new mismatches emerge elsewhere. From this perspective, systems-level appearance of stability is in fact an aggregate signature of constant, dyad-level change.

While it is likely that any interpersonal uncertainty in a given dyad can be felt by those in the relationship, an open question is whether others in the community also sense relational asymmetry in a particular dyad. People routinely observe and reason about the ties around them, inferring, for example, which relationships cluster together and how information might flow between individuals[22–25]. An assumption is that these social inferences are made under the belief that relationships are stable and mutually recognized. Relational asymmetry challenges this assumption. When two individuals' views about the same relationship are misaligned, outside observers may infer this ambiguity about the nature of that connection and its role in the broader network—for example, whether the tie supports a robust pathway by which information can flow. In this sense, relational asymmetry may itself be an important social cue that the community leverages to make inferences. If the broader network is sensitive to patterns of asymmetry across ties, then these asymmetric signals could be used to infer which relationships are more stable than others, shaping higher-order judgments including those about whether information will quickly spread or become bottlenecked between two people. In this way, relational asymmetry may influence not only how individual relationships evolve but also how observers perceive the structure of the network.

Here, we track how relational asymmetry evolves over an academic year within a large emerging real-world social network (N=187)—from milder disagreements, such as one person considering the other a friend while the other views them merely as an acquaintance, to starker ones, such as one person considering the other a friend while the other reports not recognizing them at all (although this is rare). We leverage these inconsistencies in relationship status to test two key questions: How does asymmetry within a dyad evolve over time? Does the broader community have knowledge of this dyadic asymmetry, and does that knowledge shape how people navigate their social network? Consistent with prior work, we observe that the overall proportion of asymmetric ties remains stable across a one-year period. Beneath this stability, however, individual pairs frequently transition between relationship states. The degree of disagreement within a dyad predicts resolution towards symmetry at the next time point in a nonlinear manner. We further show that relational asymmetry is not merely a naturally occurring, descriptive property of dyads but also a social signal that the network detects and uses to make inferences that are consequential for social success. The more misaligned two individuals are, the more salient their asymmetric status is to others in the network. This perceived relational asymmetry is then used by others in the network to infer how likely information is to flow between two individuals and how new ties may fit into existing groups. Additionally, greater relational misalignment surrounding an individual reduces their ability to rely on network structure to make these types of inferences.

2. Results

a. Relational asymmetry is pervasive and stable in a converging network

Over the course of an academic year, we mapped out an emerging social network comprising first-year undergraduate students (N = 187) using relationships surveys that captured the granularity of social tie strength. Participants rated every other network member as either Friend, Acquaintance, Familiar or Stranger (Fig. 1A). For full details of participants, design, and analytic methods, see Methods, p. 162). We operationalized relational asymmetry as pairs in which two individuals provided different ratings of each other (red vs. blue in Fig. 1B). Ratings were converted to numeric values, from 1 (Friend) to 4 (Stranger), and the difference between two values served as a continuous measure of asymmetry magnitude. Graded relationship ratings were collected at four timepoints through the year, allowing us to assess relational asymmetry over time (Fig. 1C).

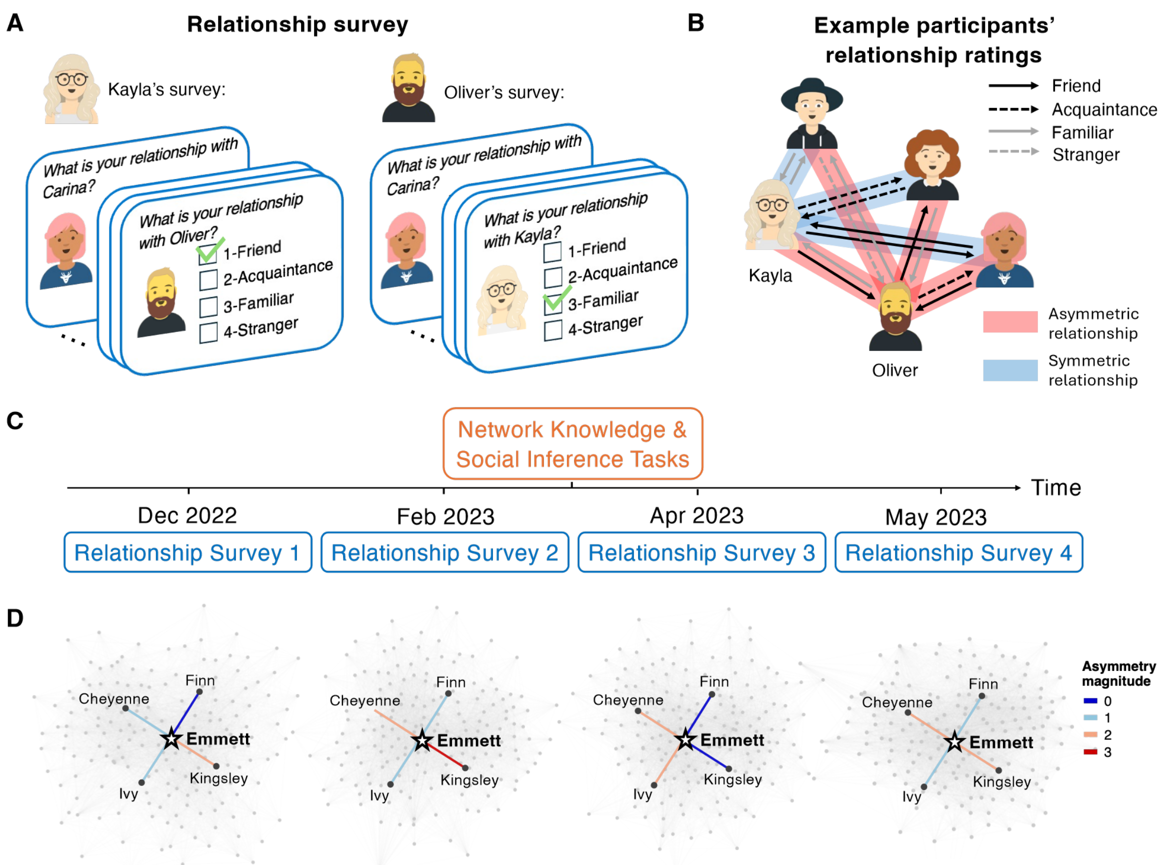

Fig. 1. Social network measurements. (A) Relationship survey structure. *Each participant viewed the names and faces of all others in the network and reported their relationship to them. Because ratings are bidirectional, data from this survey allows us to quantify relational asymmetry for every dyad.* ***(B) Example participants' relationship ratings.*** *Asymmetric relationships (red) are indicated by different arrow types between two people, while symmetric relationships (blue) are indicated by the same arrow type. Each participant is involved in multiple relationships and differ in the number of asymmetric (red) and symmetric (blue) relationships that they possess (e.g., Kayla overall has more symmetric relationships than Oliver in this example).* ***(C) Social network measurements and tasks timeline.*** *From December 2022 to May 2023, four relationship surveys with graded relationship ratings were administered. Participants also completed the Network*

Knowledge Task and two other social inference tasks immediately after completing Relationship Survey 2, to ensure to the best of our ability that responses in those tasks would be based on knowledge about the most up-to-date state of the network. ***(D) Temporal evolution of network and example participant's asymmetric ties.*** *Each figure depicts the full social network at one time point (Relationship Surveys 1-4, left to right). Highlighted edges show a subset of example participant's (Emmett) asymmetric relationships across survey waves, colored by asymmetry magnitude, illustrating that dyad-level relational asymmetry is not static.*

To compute a baseline measure of asymmetry, we collapsed snapshots of the network over time, observing that asymmetric ties account for approximately a quarter of all possible relationships (23.75%; Fig. 2A). However, since the majority of dyads in the network are mutual strangers (62.24%)—a symmetric relationship by definition—this estimate substantially understates the prevalence of asymmetric relationships in which one person is at least familiar with the other. When restricting our analysis to non-stranger dyads, we found that more than half of ties were asymmetric (62.89%). The distribution of these asymmetric ties followed a structured pattern: those marked by greater degrees of asymmetry were progressively less common (average proportion of symmetric relationships: 76.25%; relationships with asymmetry magnitude = 1: 21.12%; relationships with asymmetry magnitude = 2: 2.26%; relationships with asymmetry magnitude = 3: 0.37%), suggesting that extreme discrepancies in perceptions of the same relationship were relatively rare (Spearman correlation between relational asymmetry and proportion: $\rho = -0.756$, $p = .011$). While the most frequently observed asymmetric pairs fell under the Familiar-Stranger ($M = 14.70\%$, $SD = 1.74\%$, range: 12.51-16.34%) and Friend-Acquaintance ($M = 2.67\%$, $SD = 0.02\%$, range: 2.64-2.69%) categories, even the most extreme asymmetric pair types (Friend-Stranger) persisted at a steady rate across all time points ($M = 0.37\%$, $SD = 0.06\%$, range: 0.32-0.44%). While it appears common for network members to hold mismatched beliefs about the nature of a relationship in general, the relative scarcity of these highly-asymmetry dyads suggests that relationships characterized by such extreme opposing viewpoints may be difficult to develop and maintain, given the social uncertainty and tension they create.

We next examined whether the overall prevalence of relational asymmetry remains constant over time. Because the relationship states of dyads are continuously evolving, it is unclear whether the observed aggregate asymmetry reflects a stable property of the network or temporary fluctuations (e.g., start of the semester, where there is typically more turnover as individuals are entering a new social environment). To test this, we benchmarked the empirical distribution of relationships against a theoretical stationary distribution, derived from the transition dynamics of the network [26]. The theoretical stationary distribution represents the configuration that the network would reach if the dynamics observed at a given timepoint were projected infinitely (see Methods), providing a principled baseline for assessing long-term stability.

The empirical network closely aligned with the stationary distribution in terms of the relative frequency of relationship types (i.e., rank order of proportions: $\rho = 0.976$, $p < .001$ in the first snapshot; $\rho = 1$ in later snapshots), whereas numerical differences remained (Fig. 2B, grey vs. colored dots). While the relative ordering of relationship types was consistent, the actual proportions had not yet fully converged to the long-run equilibrium. To quantify the gap between the empirical distribution and stationary distribution, we calculated the total variation distance (TVD) between these two distributions (see Methods). TVD provides a direct measure of how much the distributions differ in absolute terms; a TVD of 0 indicates that the network's composition is identical to the theoretical stationary distribution. We found that the TVD steadily decreased across data collection points ($TVD_{Survey1} = 0.198$; $TVD_{Survey2} = 0.174$; $TVD_{Survey3} = 0.142$; $TVD_{Survey4} = 0.132$). Together, these results show that, despite continuous churn among individual dyads, the overall distribution of relationship types—including asymmetric ones—remains stable, which is consistent with a dynamic equilibrium in which relationships constantly shift but the network exhibits persistent patterns at the systems-level.

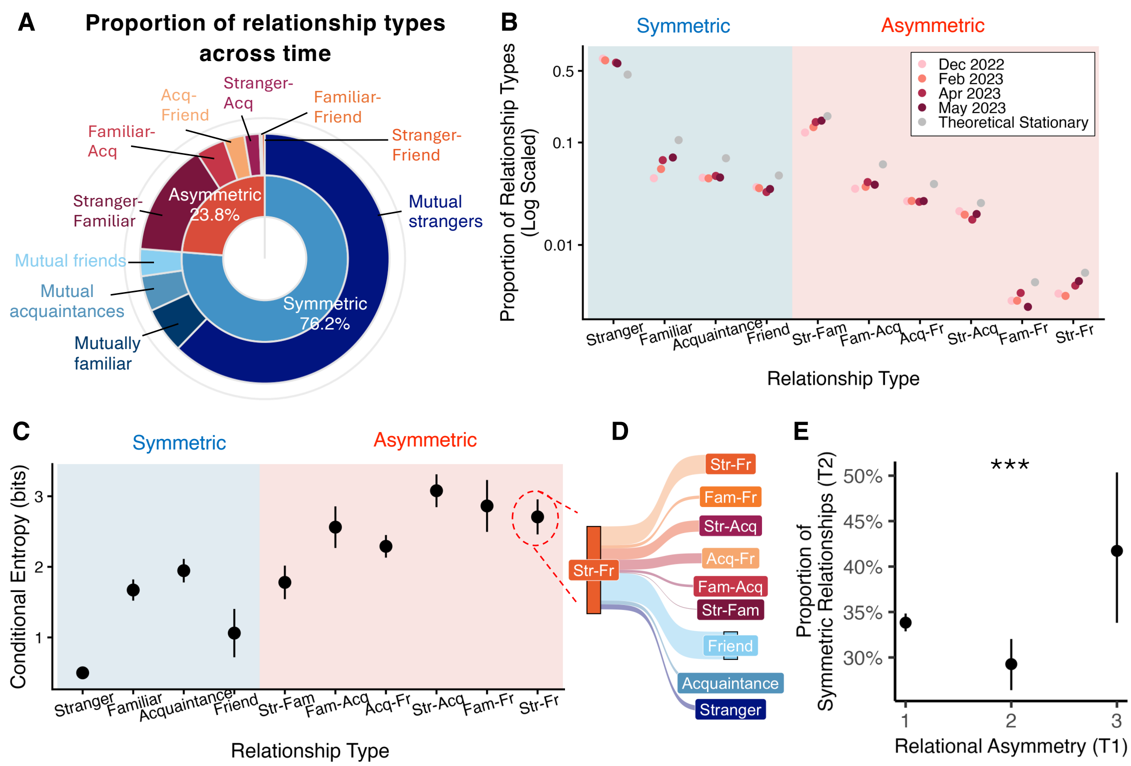

Fig. 2. Proportion of asymmetric versus symmetric ties and relationship dynamics over time.(A) Average proportion of different relationship types across time. *Around a quarter of all relationships in the network are asymmetric (red inner circle), and among asymmetric relationship types, 'Stranger-Familiar' pairs are most common.* ***(B) Relationship type occupancy over time***. *Proportion of different relationship types remains stable over time (colored dots) and approximates the expected long-run distribution derived from the network's own transition dynamics (grey dots).* ***(C) Conditional entropy of different relationship types across time.*** *Overall, conditional entropy of asymmetric relationships (red background) is higher than symmetric relationships (blue background). Conditional entropy is calculated as the Shannon entropy of the distribution of relationship types at the subsequent timepoint (T2) conditioned on the relationship type at the previous timepoint (T1), and quantifies the unpredictability of a dyad-level relationship across time.* ***(D) Distribution of relationship types at T2 that evolved from Stranger-Friend pairs at T1.*** *The thickness of each band denotes the transition probability, with thicker bands reflecting greater transition probability into a specific relationship type. While some Stranger-Friend pairs remain in their current state at T2 (22% probability), a considerable proportion transition into other relationship types, including symmetric ones (i.e., mutual Friend, Acquaintance, and Stranger).* ***(E) The proportion of symmetric relationships at T2 as a function of relational asymmetry at T1.*** *Asymmetric dyads at T1 are more likely to resolve to symmetry at T2 than to persist as asymmetric. More extreme relational asymmetry—whether low or high—is associated with increased likelihood of symmetry at T2. The error bar indicates standard error.* ***$p < .001$.

b. Asymmetric ties are volatile and resolve toward symmetry

Even as the network converges towards a stationary configuration, individual dyads continued to move between relationship states. For example, the observation that mutual stranger pairs exceeded stationary predictions (Fig. 2B) implied a consistent outflow from this relationship state into other states (confirmed by detailed balance analysis reported in *SI Appendix*, Note S1). To characterize how dyads move between distinct relationship types, we first modelled relationship transitions and found that these patterns are statistically consistent over time (see *SI Appendix*, Note S1 and Fig. S1). While transitions capture average flow, it does not account for the uncertainty inherent in these movements, nor how that uncertainty scales with relational asymmetry.

To quantify transition uncertainty, we computed conditional Shannon's entropy—a metric that captures the total dispersion of potential future states—as an index of a relationship's unpredictability [27]. Asymmetric relationships exhibit significantly higher conditional entropy than symmetric ones (Wilcoxon rank-sum = 1, $p = .019$; Fig. 2C), with conditional entropy scaling positively with the magnitude of relational asymmetry (Spearman correlation $\rho = 0.871$, $p = .001$). High conditional entropy reflects a diffuse transition profile, including both a reduced probability of persisting in the current state and a more uniform distribution of the relationship types that it transitions into, as illustrated by the Friend-Stranger pairs (Fig. 2D).

Despite this volatility, dyadic changes were not purely stochastic. Conditional entropy of all relationship types remained well below the theoretical maximum four bits, i.e., a uniform distribution over 16 possible states, indicating that structured movement is governed by specific biases. Asymmetric dyads transitioned into symmetric pairs at a rate nearly three times that at which symmetric dyads moved into asymmetry (33.50% vs.

12.01%), indicating that resolution is structured rather than due to random fluctuation (two-sample proportion test: $\chi^2(1) = 2468.13$, $p < .001$). Notably, the probability of achieving symmetry was nonlinearly related to the magnitude of relational asymmetry, with a quadratic relationship providing a reasonable fit to the data ($\beta_{\text{relational asymmetry}^2} = 7.314$, $SE = 2.139$, 95% CI [3.12, 11.51], $z = 3.434$, $p < .001$; Fig. 2E): pairs with the least—and most—discrepant views of the same relationship were both most likely to achieve greater alignment at the next timepoint. The robustness of this nonlinear effect was confirmed via permutation test ($p < .002$, n = 500 permutations), in which dyad-level asymmetry magnitude at the subsequent timepoint was randomly shuffled within time point, preserving the marginal distribution of asymmetry magnitudes at each wave but breaking the association between asymmetry at the initial time point and whether asymmetry within a dyad was resolved. The observed coefficient fell well outside the null distribution (M = 0.036, SD = 2.127, 95% CI [-3.973, 4.070]). We replicated this nonlinear pattern in a separate real-world social network consisting of a different cohort of undergraduate students (N = 179), where we again observed a quadratic relationship between relational asymmetry magnitude and subsequent asymmetry resolution (see *SI Appendix*, Result S1). Overall, these dyad-level dynamics facilitate the coexistence of network-level stability and dyad-level volatility, making relational asymmetry a persistent, yet evolving, property of the social structure.

c. The network can perceive relational asymmetry at the dyad level

If relational asymmetry is a meaningful feature of social structure that contains predictive signal, its effects should extend beyond the dyad itself and become visible in how other network members perceive that tie. Previous work has shown that people are attuned to the structural properties of the environment [22,24,25], yet whether this sensitivity can capture the granularity of specific asymmetric ties, remains unknown. To test this, we designed the Network Knowledge Task, in which participants judged the friendship status of dyads within three degrees of separation from them (orange boxes in Fig. 1C and Fig. 3A; see Methods). Here, we report results from the Network Knowledge Task following Relationship Survey 2, and replicate our findings using the data collected following Relationship Survey 3 (see *Appendix SI*, Result S2).We quantified group-level disagreement using the standard deviation of all ratings for each unique pair, and compared it against the difference between the ratings reported by the individuals in the pair (i.e., ground-truth). Group-level disagreement increased with greater relational asymmetry ($\beta_{\text{relational asymmetry}}$ = 0.015, SE = 0.002, 95% CI [0.012, 0.019], $p < .001$), suggesting that the broader community was less certain about the nature of two people's relationship when the pair themselves were also very uncertain. This effect held when using entropy as an alternative measure of group-level uncertainty (see *SI Appendix*, Result S3), suggesting that when the nature of a relationship is uncertain, this uncertainty is detectable by the community.

Does it matter if community members have different vantage points from which to evaluate relationships? Network members' judgments may be anchored by their social proximity to one of the individuals in the perceived dyad, which could result in the observed lack of consensus for asymmetric relationships. For example, imagine that Anya and Julie report an asymmetric tie (Anya reports being familiar with Julie, but Julie considers Anya to be a friend). In this case, network members who are closer to Julie should be more likely to share Julie's view—that there is a friendship between Julie and Anya. To test this possibility, we first categorized individuals within an asymmetric pair as either being 'more invested' or 'less invested' based on how socially close one rated the other (e.g., Friend =

more invested vs. Recognize = less invested), and then assessed whether social distance between the individuals in the pair and other network members (i.e., observers who rated the pairs) is associated with group-level disagreement (see Methods).

We first examined whether distance to the person in the dyad who was more invested contributed to group level ratings of the pair (see Fig. 3C for schematic). We found that network members were more likely to judge the relationship to be a friendship if they were closer to the more invested member of the dyad (β = 0.016, *SE* = 0.008, 95% CI [0.000, 0.031], p = .043). This main effect was modulated by the pair's relational strength (β = -0.024, *SE* = 0.008, 95% CI [-0.040, -0.009], p = .002; Fig. 3D). For ties that were nearing mutual friendship (i.e., Friend-Acquaintance), network members closer to the individual who was more invested were more likely to infer a friendship, relative to those closer to the less invested individual (post-hoc comparison for Friend-Acquaintance pairs: Δ = -0.056, SE = 0.015, $p < .001$). Such differences in proximity-based perceptions diminished when the dyad was not relationally close (post-hoc comparison for Familiar-Stranger pairs: Δ = 0.006, SE = 0.010, p = .586). These patterns suggest that group-level disagreement may be driven by a collection of heterogeneous, distance-dependent perceptions rather than uniform perceptions of uncertainty within the network.

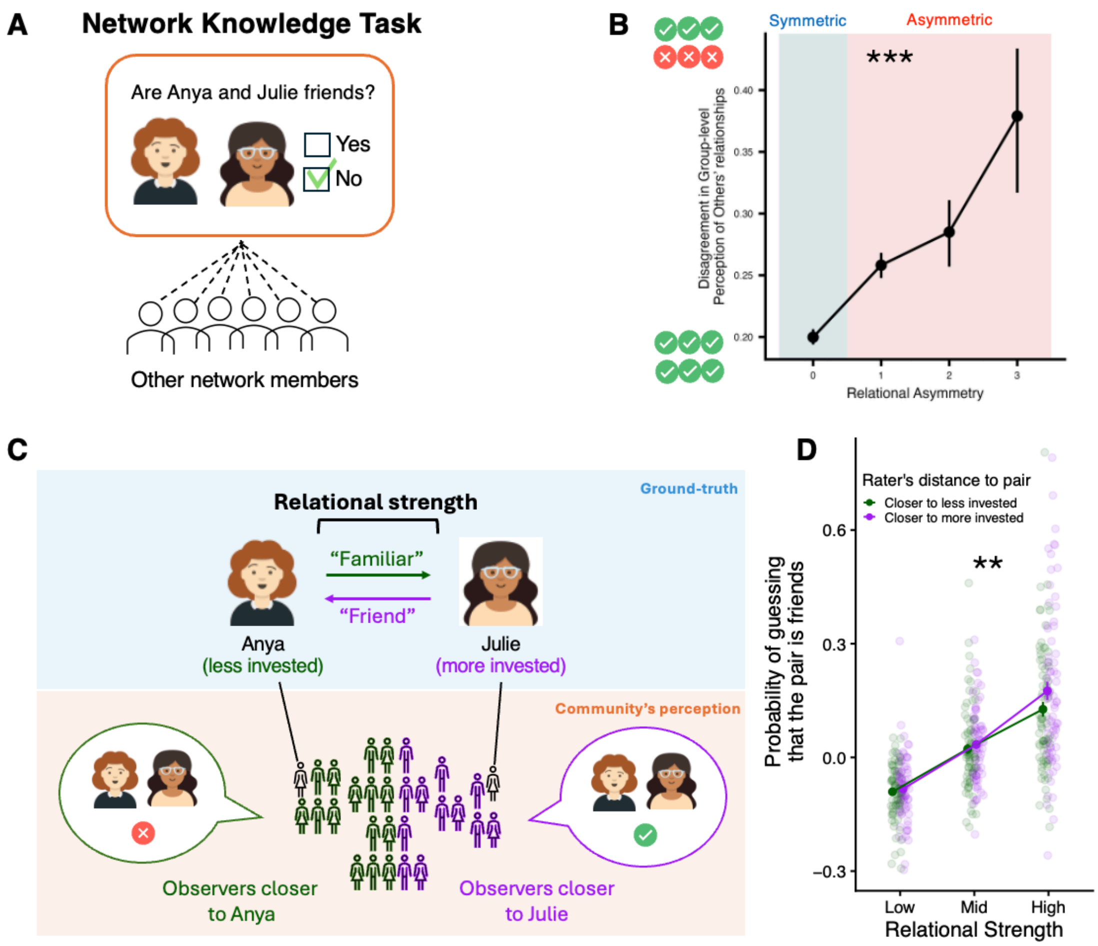

Fig. 3. Individuals in the broader community track dyad-level relational asymmetry. (A) Network Knowledge Task structure. *To measure how individuals in the broader community perceive the relationships of others in the network, we asked each participant to rate a subset of dyads in the network and report whether they believe each dyad were friends.* ***(B) Group-level disagreement increases with greater dyad-level relational asymmetry.*** *Distribution of community members' ratings is the most disperse (quantified by the standard deviation of all ratings for a given dyad), for dyads characterized by high relational asymmetry. An equal split between those believing and not believing a friendship exists reflects maximum group-level disagreement. Data shown here are restricted to pairs rated by at least two network members.* ***(C) Schematic illustrating the more "ainvested" side (Anya, who rates the other as Friend) and the less invested side (Julie, who rates the other as Familiar) of an asymmetric pair.*** *Observer proximity to each side is defined by shortest path geodesic distance in the network.* ***(D) Friendship judgments as a function of observer proximity to each side of the pair and the pair's relational strength.*** *Observers who sit closer in the network to the more invested side are more likely to judge an asymmetric pair as friends than observers closer to the less invested side, but only when the pair's average relational strength is moderate or high (x-axis). Relational strength within a pair reflects the mean of the two directional ratings within the pair. Both distance and relational strength are defined using ratings from the preceding relationship survey. Friendship judgments (y-axis) are shown after subtracting each observer's mean response across all rated pairs to account for overall response tendency. Error bars indicate standard error.* *$^{**}p < .01$; $^{***}p < .001$.*

d. People leverage knowledge of relational asymmetry to make adaptive social inferences

If relational asymmetry is a signal that is detected by the community, it stands to reason that such accurate perception is achieved for a social purpose. One possibility is that inferring relational asymmetry within the network provides a signal that can be used to generalize to novel social situations, such as estimating the likelihood of a new relationship between unknown individuals. For example, if a new student forms ties with two existing individuals in the network, an observer may rely on knowledge about the relational structure of nearby network members to infer how likely that new student is to also form a tie with other peers (Fig. 4A). These inferences should depend, in part, on the relationship strength among the existing members (operationalized as the mean of the existing network members' ground-truth relationship ratings, where lower values indicate greater closeness). As expected, greater relational strength within a triad predicted greater inferred likelihood of tie formation with an existing member ($\beta_{\text{triad's relational strength}} = 0.338$, $SE = 0.025$, 95% CI [0.290, 0.386], $p < .001$; Fig. 4B). But, relational asymmetry reduced this likelihood above and beyond relational strength ($\beta_{\text{triad's asymmetry}} = -0.241$, $SE = 0.061$, 95% CI [-0.350, -0.121], $p < .001$; Fig. 4C), suggesting that asymmetry provides an independent cue that cannot be merely reduced to confidence about how close a given tie is.

The degree of relational asymmetry within a dyad should also determine how readily information flows between the two individuals. A mutual friendship, for instance, is a more reliable channel than a tie in which only one party considers the other a friend [28]. We leveraged a task in which participants judged the likelihood that a piece of information will successfully spread from an individual who shares news (Source) to a different individual who might hear about it (Receiver; Fig. 4D). The likelihood of inferring information flow tracked relational strength within dyads, with mutual Friend and mutual Stranger representing the upper and lower bound of transmission probability, respectively ($\beta_{\text{relational strength}} = 0.137$, $SE = 0.002$, 95% CI [0.133, 0.142], $p < .001$; Fig. 4E). To isolate the specific effect of relational asymmetry, we restricted our analysis to dyads in which at least one individual reported the other as a friend. This provides a meaningful baseline against which we can assess how inferred likelihood of information changes with asymmetry. Within this subset, mutual friendships were judged to be more effective channels for information flow compared to one-sided, asymmetric ties ($\beta = 0.155$, $SE = 0.022$, 95% CI [0.111, 0.199], $p < .001$; Fig. 4D), revealing that participants were sensitive to the presence of asymmetry and treating relational mismatch as meaningfully distinct signals for how information travels through the network. We replicated these findings using data collected following Relationship Survey 3 (see *Appendix SI*, Result S4).

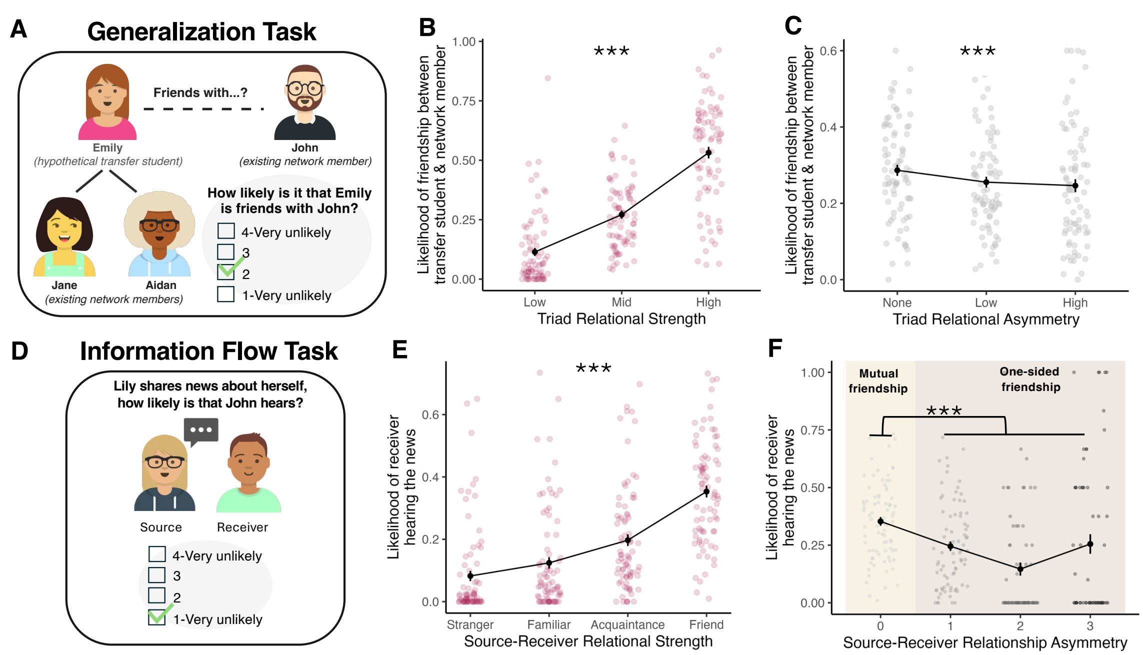

Fig. 4. Influence of relational strength and asymmetry on social inferences. (A) Generalization Task structure. *To test whether participants generalize knowledge about existing network structure to infer new ties, participants were first told that a hypothetical transfer student (Emily) is friends with two existing network members (Jane and Aidan). They were then asked to judge the likelihood that the transfer student would also be friends with a third, existing network member (John).* ***(B) Relational strength within a triad increases inferred likelihood about unseen***

ties. *When the average relationship strength within the triad of existing network members is stronger (lower on x-axis), participants are more likely to infer a new tie between transfer student and network member.* ***(C) Average relational asymmetry predicts inferences about unseen ties.*** *If a triad of existing network members on a given trial is characterized by greater average asymmetry, participants are less likely to infer a friendship between the transfer student and existing network member. "Low", "Mid", and "High" denote terciles of asymmetry within the triad (≤ 33rd percentile, 33rd-67th percentile, >67th percentile). Error bars indicate standard error.* ***(D) Information Flow Task structure.*** *Participants viewed a subset of dyads in the network and reported the likelihood that a given individual (the Receiver) would hear news shared by another individual in the network (the Source) on a 4-point Likert scale.* ***(E) Dyad-level relational strength positively influences people's judgment about information flow.*** *When the Source-Receiver relationship is symmetric, the rated likelihood of receiving information decreases as the relationship type moves farther away from mutual friendship.* ***(F) Average perceived likelihood of receiving news as a function of Source-Receiver relational asymmetry.*** *Participants believe that information is more likely to flow between mutual Friend pairs than between individuals in an asymmetric, one-sided friendship. Likelihood ratings were binarized and grouped for illustration purposes.* $^{**}p < .01$; $^{***}p < .001$.

e. Relational alignment shapes sensitivity to network structure

Although people rely on relational structure to make social inferences, they do not all do so equally well. One possible source of this variability is the degree to which an individual's own relationships are mutual—that is, whether their relationship ratings of others in the network are reciprocated in kind. We operationalize this as the correlation between one's own relationship ratings and the ratings one receives from others (higher values indicate greater relational alignment). Unlike relational asymmetry, which characterizes the degree of disagreement within a specific dyad, relational alignment is an individual-level measure that captures the overall consistency of bidirectional relationship ratings across an individual's entire set of social ties. People who have more misaligned relationships may also have noisier internal representations of others' social ties, making it harder to detect relational patterns in others, which can subsequently impact their ability to leverage true network structure to make social inferences.

Individuals in our cohort exhibited substantial variation in relational alignment (r = 0-0.906; M = 0.758, SD = 0.112; Fig. 5A and *SI Appendix*, Fig. S3), with the proportion of asymmetric relationships ranging from only 7% to 50% across participants. Relational alignment was included as a predictor in models of both social inference tasks, revealing that the less an individual's exhibits relational alignment in their own life, the less able they were to infer whether a new student will become friends with existing cliques ($\beta_{\text{relational strength:alignment}}$ = -0.069, SE = 0.024, 95% CI [-0.115, -0.022], *p* = .004; Fig. 5B). Judgments about information flow were also attenuated by the degree of an individuals' relational alignment ($\beta_{\text{relational strength:alignment}}$ = -0.017, *SE* = 0.007, 95% CI [-0.031, -0.002], *p* = .022; Fig. 5C). These findings were replicated using the second wave of data collection (see *SI Appendix*, Result S4). That is, people characterized by greater misalignment in their own relationships were similarly less able to use others' relationship structures to reason about information transmission.

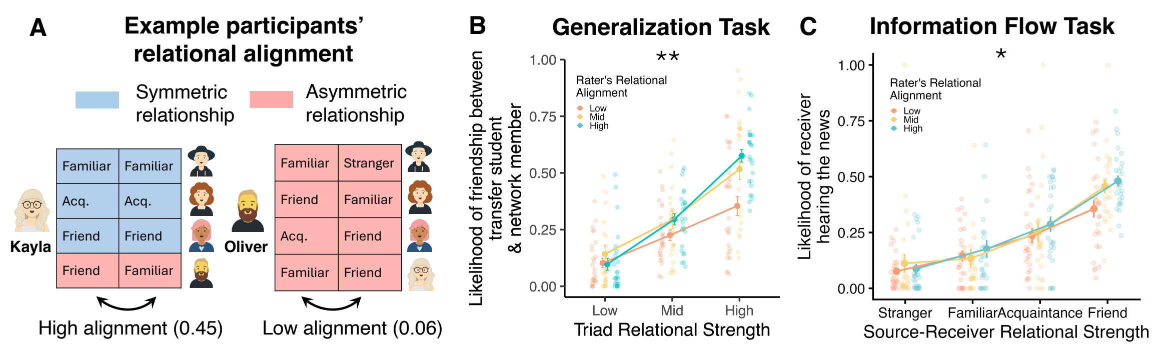

Fig. 5. Influence of participant-level relational alignment on social inferences. (A) Example participants' relational alignment. *Relational alignment is operationalized as the correlation between one participant's ratings of other members in the network and those members' ratings of the participant. This metric reflects the overall alignment between the participant's view of their social world and the perspectives of their peers, accounting for each individual's response tendency. Shown here, example participant Kayla has more symmetric relationships (blue) as well as greater local alignment, compared to Oliver, who has lower relational alignment because he has more asymmetric relationships and is inconsistent in being the more/less invested side (red).* ***(B) Lower participant-level relational alignment diminishes ability to generalize structural knowledge to unseen ties.*** *The effect of average relationship strength on new relationship inference is attenuated among participants with lower relational alignment. Color indicates relational alignment, divided into percentiles (≤ 33rd percentile; 33rd-67th percentile; >67th percentile). Shaded bands indicate standard error.* ***(C) Lower participant-level relational alignment reduces ability to judge information probability based on Source-Receiver relational strength.*** *People with lower relational alignment are less able to use graded relationship types to judge the likelihood of information flow. Color indicates rater-level alignment, divided into terciles (≤ 33rd percentile; 33rd-67th percentile; >67th percentile). Error bars indicate standard error. *: p < .05; **: p < .01.*

3. Discussion

Relational asymmetry is a persistent feature of social life, evident in how frequently it occurs across various real-world social networks. Even as the network converges toward a stationary distribution of relationship states, individual pairs continue to transition between different degrees of asymmetry, and in predictable directions. Relational asymmetry therefore has dynamics of its own that are lost when distilled to macro-level descriptions of network structure—dynamics that are adaptively leveraged by the community to make useful social inferences for who will become friends and whether information will move fluidly across the network. The fact that relational asymmetry is not only experienced between the two individuals in the relationship, but is also felt by the surrounding community, suggests that being able to perceive granular social dynamics play an important role for adaptive social navigation. Evidence of these network dynamics have been obscured by cross-sectional approaches[29,30] and by models that reduce relationships to binary ties[31,32], both of which overlook the patterned movement unfolding just beneath macro-level stability.

Relational asymmetry is not a static mismatch between two individuals, but belief states under continuous evolution. Dyads that began in asymmetric states tend to move toward symmetry over time, among which relationships with high relational asymmetry are especially unstable and are characterized by high transition entropy and low probability of maintaining their current state. This transient nature of asymmetric ties aligns with classic models of friendship formation, which posits that unreciprocated bonds are inherently unstable[20]. While this work suggests that one-sided friendships require an explicit act of 'acceptance' or 'rejection' to transition into a stable state (i.e., mutual Friend or mutual non-friend), we treated the development of asymmetric relationship as a continuous updating process driven by interpersonal uncertainty, which could result in symmetric relationships of varying degrees of relational strength. We observed that the trajectory from asymmetry to later resolution is nonlinear. Symmetry is especially likely to emerge when disagreements were either small enough to be easily reconciled or large enough to become difficult to ignore, suggesting that extreme mismatches in beliefs about a relationship may be especially likely to trigger resolution. Given the relative scarcity of extreme asymmetric relationships in the network, this curved relationship may have been masked in previous work studying networks with binarized[20]. The U-shaped trajectory of asymmetric resolution raises the possibility that different degrees of asymmetry may be associated with distinct updating processes. For example, extreme asymmetry could signal high relational uncertainty, which could prompt social exploration and result in new connections[33]. This pattern is consistent with the idea that people are active updaters who revise their judgments about where they stand with others in the face of interpersonal uncertainty.

Individuals do not merely experience their own ties; they actively detect and perceive the network structure around them to make adaptive social inferences about the community at large. These inferences are grounded by the nature of the underlying relationships. When existing relationships are strong (i.e., of high relational strength), people report higher likelihood of new tie formation and higher probability of information flow. This adaptive inferential process also manifests as being sensitive to the asymmetry in existing relationships, translating into lower estimated possibility of new friendship formation and information spread. This reliance on relational strength and asymmetry is supported by past work demonstrating that both weak and asymmetric ties are notably less effective for information diffusion[28,34]. Here, we suggest that human social reasoning is finely tuned to the structural constraints of the environment and the functions that relationships of different qualities can afford.

People were not uniformly sensitivity to network structure, however. When people's own relational environment was more uncertain, they were less sensitive to the relational strength of others, and struggled to infer how information would spread and to generalize network structure to ties they had not observed. Relational asymmetry may therefore create a feedback loop, in which relational misalignment undermines the quality of network representations, and representations that are characterized by greater ambiguity leave people less able to reason about other relationships. These findings converge on a view of social network representations as containing information about relational strength, which impact downstream social reasoning, not simply maps of who is connected to whom.

Tracking a cohort of first-year undergraduate students provides ecological validity, but not without several boundary conditions. Social networks of this type are relatively closed and thus have strong constraints on who can meet and interact[35,36], making it uncertain how well the present findings generalize to other social environments both in real life and online. Additionally, the granularity of our relationship measures improves on prior work that relied on binarized friendship judgments or binary classifications of asymmetry[1,20], but we nevertheless treat transitions between relationship states as equivalent, which may not reflect how such updates actually unfold. Future work should examine whether some kinds of relational change are more meaningful than others, rather than assuming that transitions between different relationship states are qualitatively equivalent.

More broadly, these findings speak to a recurring tension in the study of social networks, between descriptions of network-level structure and accounts of the local processes that produce it. Network-level stability is often taken as evidence that a system has reached equilibrium[10], but our results show that apparent stability can coexist with predictable, directional change at the level of individual relationships. This dissociation is particularly consequential for social networks, where the dynamics driving structural change are not governed by physical laws but by people actively updating their beliefs about one another. Focusing on changes at the dyad level is therefore not merely a methodological choice, but an understudied window into the psychological processes that shape the composition of our social worlds.

4. Methods

Overview

We recruited 198 first-year undergraduate students across three dorms. Two participants were excluded due to not meeting eligibility criteria, resulting in a final sample size of 196 (103 female, 89 male, 4 other; 146 non-white or mixed race; age M = 18.1, SD = 0.42). They completed four relationship surveys across their whole first year college and the interval between two surveys are approximately equal across all surveys (Fig. 1C). We refer to the four relationship survey waves as Surveys 1-4 (Survey 1: N = 180; 95 female, 81 male, 4 other; 136 non-white or mixed race; age M = 18.1, SD = 0.42; 16,110 unique pairs; Survey 2: N = 176; 96 female, 77 male, 3 other; 133 non-white or mixed race; age M = 18.1, SD = 0.42; 15,400 unique pairs; Survey 3: N = 172; 92 female, 76 male, 4 other; 130 non-white or mixed race; age M = 18.1, SD = 0.42; 14,706 unique pairs; Survey 4: N = 173; 93 female, 76 male, 4 other; 128 non-white or mixed race; age M = 18.1, SD = 0.43; 14,878 unique pairs). For all analyses tracking changes across time, we included pairs that have filled out the questionnaire for at least two consecutive time points, resulting in 14,878 pairs for Survey 1→2; 14,028 pairs for Survey 2→3; and 13,695 pairs for Survey 3→4.

In addition to four relationship surveys, a subset of the participants also completed Network Knowledge Task and social inference tasks (Generalization and Information Flow) two times during the experiment (N = 80; 45 female, 35 male; 62 non-white or mixed race; age M = 18.1, SD = 0.45; 78 of the participants also completed the second). These tasks were administered immediately after the second relationship survey and the third relationship survey to make sure that the data used to probe participant's knowledge of the network was up to date. All procedures were approved by Brown University's Institutional Review Board, and informed consent was obtained from all participants. For 17-year-old minors, assent was obtained from the participants and informed consent were obtained from their legal guardians.

Relationship Survey

Participants reported their relationships with other people by selecting one of four options: Friend, Acquaintance, Recognize, Do Not Know. We refer to Recognize as Familiar and Do Not Know as Stranger to ensure that descriptions of certain relationship types are more intuitive (e.g., Friend-Stranger). For each rating, participants were shown the name and a photograph of the person to be rated. This was repeated for all the network members so that we were able to map out the bidirectional ratings for all the pairs that participate in the study.

Constructing transition matrices. To examine the temporal resolution of relationships, we modeled the movement of relationship between discrete states using a transition matrix. Given that each individual rated their partner on a four-point scale (Friend, Familiar, Acquaintance, Stranger), every dyad occupied one of 16 possible joint states (e.g., Friend-Stranger or mutual Friend).We constructed the matrix by tracking the trajectory of every dyad across consecutive snapshots, recording the kind of transition that the dyad went through (a Friend-Stranger pair can stay Friend-Stranger, or become mutual Friend or mutual Stranger, etc.). These observations were aggregated into a frequency matrix, where rows represent the initial relationship type and columns represent the state following the transition. We then transform the raw counts into a probability matrix by normalizing the matrix row-wise. In the transition matrix, each element the *i,j* element represents the probability of relationship that starts in type *i* ends in type *j*, i.e., the conditional probability of relationship type moving to *j* provided that it starts in state *i*. Since we are interested in capturing the overall network movement across time, we collapse data across all timepoints to derive the aggregated transition matrix *M* (see *SI Appendix*, Fig. S1).

Stationary distribution of relationship types. The stationary state of a network refers to the stable configuration that the network would reach if its current transition dynamics were projected to infinity. If a network is at its stationary state, the distribution of relationship types becomes invariant after the transition. This property allows us to derive the stationary state by solving the following equation:

$$\tilde{\pi} = M^T \tilde{\pi},$$

where $\tilde{\pi}$ represents the stationary distribution of all relationship types. We analytically solve the equation by finding the eigenvector of the transpose of transition matrix M associated with eigenvalue = 1.

We compare the theoretical stationary distribution with the empirical distribution observed in our data to examine the relationship between empirical data and theoretical estimates. To capture the qualitative structural alignment, we compute the rank-order correlation between these two distributions. A larger correlation coefficient indicates that the relative ordering of relationship types remains stable across time. To quantify the difference between distributions, we calculated total variation distance (TVD; Levin & Peres, 2017):

$$TVD(P, Q) = \frac{1}{2}\sum_{i=1}^{k} |p_i - q_i|$$

Where P is the empirical relationship distribution (i.e., proportion of all possible relationship types observed in the network) and Q is the theoretical distribution. TVD is bounded between 0 and 1 and can be interpreted as the proportion of relationships that need to be redistributed to achieve a perfect match. Therefore, smaller TVD suggests that the empirical network relationship distribution more closely resemble the distribution achieved from the network at its stationary state.

Conditional entropy of relationship transitions. To quantify the degree of uncertainty associated with relational transitions, we calculated the conditional entropy H for each of the relationship states. While the transition matrix provides the probability of moving between relationship types, conditional entropy allows us to measure the volatility of these shifts. For a given initial state, the entropy is defined as

$$H(j|i) = -\sum_{j=1}^{k} p(j|i) log_2 p(j|i)$$

where $p(j|i)$ represents the probability of a dyad transitioning to relationship type j given that it originated as relationship i. An entropy value of zero bits would indicate a perfectly deterministic transition, where a relationship is fully predictable and always moves to the same relationship at the next timepoint. The theoretical maximum of four bits (i.e., $\log_2(16)$) would represent a uniform distribution, implying a high uncertainty transition process with the relationship being equally likely to transition into any of the 16 possible states.

Quantifying relational asymmetry. Asymmetric relationships are defined as pairs with mismatched ratings. For example, Friend-Acquaintance refers to pairs where A reports that B is a friend while B rates A as an acquaintance. To quantify the degree of relational asymmetry, we rank the options according to closeness and convert them to numerical value (Friend = 1, Acquaintance = 2, Familiar = 3, Stranger = 4). Relational asymmetry is therefore computed as the absolute difference between the bidirectional ratings, and greatest relational asymmetry exist in Friend-Stranger pairs (relational asymmetry = 3).

To examine whether relational asymmetry resolves over time, we used paired-sample t-tests to compare the degree of relational asymmetry across consecutive snapshots. Given that dyadic transitions were found to be temporally invariant and statistically equivalent with the aggregated transition matrix (see *SI Appendix*, Note S1), we collapsed observations across all timepoints to provide a high-powered, aggregated estimate of these dynamics. Additionally, we tested whether this asymmetry resolution is driven by a specific individual within the dyad. Specifically, we distinguished between the closer-rating partner (i.e., the actor initially providing a rating closer to friend) and the more distant-rating partner (i.e., the actor initially with a rating closer to stranger) within each asymmetry dyad. We then conducted separate paired-sample -tests for each role to examine how their respective ratings shifted between timepoints.

Quantifying individuals' relational alignment. To capture the relationship profile of each individual, we calculate the correlation between the participant's ratings for other network members and their partners' ratings for the participant. This metric captures the overall alignment between the belief of the individual of their social relationships and that of their partners, while accounting for individuals own response tendency. When alignment equals 1, participant's ratings of other network members perfectly correlate with their ratings of the participant. When alignment equals 0, participant's ratings fluctuate independent of other network members' ratings towards the participant.

Quantifying relational strength. Relational strength is computed as the average of the bidirectional ratings from the ground-truth relationship survey and reflects the overall social closeness of the pair. Among all relationship types, mutual Friend pairs possess the highest relational strength while mutual Stranger are of the lowest relational strength.

Network Knowledge Task

Participants reported their belief about whether pairs of other network members were friends (Fig. 3A). Given the sheer size of the network, we customized the stimulus for each participant and probed the relationships among a subset of 30 network members based on their network distance to the participant (measured using an undirected network at the preceding survey). This sample contains 5 randomly sampled friends, 10 friends-of-friends, and 15 friends-of-friends-of-friends. When the requirement cannot be met (e.g., some participant may have fewer than 5 friends), members at the next distance level was added to fill up the slot. Given this stimulus set composed of 30 network members, each participant makes in total of 870 friendship judgments, rating each pair twice.

Personalized stimulus set results in varied number of responses for each pair (see *SI Appendix*, Fig. S2). We constrained our analysis on pairs that have been rated at least two times. After exclusion, 6,003 pairs were analyzed, which are rated on average 4.88 times (SD = 4.35).

Quantifying group-level disagreement. In the Network Knowledge Task, each pair is rated twice by an observer (in both directions, i.e., A→B and B→A). We averaged each participant's ratings for a given pair and then quantified group-level disagreement as the standard deviation of these averaged ratings across participants. Given that individual guesses were constrained (only possible values were 0, 1, and 0.5), the standard deviation metric serves as a proxy for group-level uncertainty regarding the nature of a relationship.

To test whether group-level disagreement tracks relational asymmetry within a pair, we used a linear regression to model group-level disagreement as a function of relational asymmetry. Given the high correlation between the mean and the standard deviation of group ratings about a pair (r = 0.728, p < .001), we entered the mean group rating as a covariate. To account for the potential effect of number of raters on the standard deviation metric, we re-ran the model with the number of raters for a pair as an additional covariate. For robustness, we replicated our results using Shannon's entropy as an alternative, comparable group-level uncertainty measure (see *SI Appendix*, Result S3).

Quantifying relative social distance between the observer and individuals within a dyad. To determine the observer's proximity to each individual of a target dyad, we calculated the geodesic distance (i.e., shortest path) between the observer and each individual using the undirected graph constructed from mutual friendships from the preceding relationship survey. Observers were considered closer to the more invested side if their distance to that individual was shorter than to the less invested side individual, and vice versa for those closer to the less invested side (Fig. 3C). We excluded observations where the observer was equidistance to both individuals within the dyad since they could not provide any signal of interest for this particular analysis.

Social Inference Tasks

We administered two computerized tasks probing how people use existing network structure knowledge to make social inferences. These tasks were administered in the same session with network knowledge task (Fig. 1C).

Generalization Task. On each trial, participants were presented with the names and photographs of three existing network members and were told that a hypothetical transfer student—depicted with a name and stock photograph that varied across trials—was friends with two of them (Fig. 4A). Participants then rated the likelihood that the transfer student was also friends with the third network member on a four-point scale ranging from 'very unlikely' to 'very likely'. Responses were binarized for analysis, with 'likely' and 'very likely' coded as 1, and 'unlikely' and 'very unlikely' coded as 0. The stimuli were the same as that used for the participant in the Information Flow Task.

To test whether relational asymmetry among existing network members constrained generalization responses, we fit a mixed-effects logistic regression predicting binarized response from trial-level average asymmetry, computed as the mean absolute rating discrepancy across all three dyadic pairs within the triad of existing network members. We also included random intercepts and slopes by participant.

To test whether participant-level relational alignment moderated the use of relationship strength in generalization, we fit a second mixed-effects logistic regression predicting binarized response from the interaction between average probe-friends relationship strength and local alignment. Average probe-friend relationship strength was computed per trial as the mean of the two probe-friend pair ratings, where each pair rating was itself the average of the two directional ratings between the probe and each friend. This captures how strongly embedded the probe was within the transfer student's known social context on a given trial. Both predictors were entered as random slopes by participant.

Information Flow Task. Participants judged the likelihood of information flow from one network member to the other. On each trial, participants were presented with the name and photograph of one network member (Source). They then estimated the likelihood of another network member (Receiver) hearing about the information that the source has shared using a four-point Likert scale, ranging from 'very unlikely' to 'very likely' (Fig. 4D). The stimuli set used for each participant is personalized and contains a subsample of 25 network members used in the Network Knowledge task. Participants rated all the possible Source-Receiver pairing among the stimulus set, which results in 600 trials.

We binarized participants' responses for analysis, with 'likely' and 'very likely' coded as 1, and 'unlikely' and 'very unlikely' coded as 0. We averaged each participant's ratings for a given pair and used mixed-effects linear regression to examine the relationship between ground-truth pair and participants' perception of information flow. We first regressed the reported likelihood using the average relational strength of the Source-

Receiver pair. This model tested the veridicality of social judgments, i.e., whether observers are sensitive to the strength of a relationship when estimating the probability of information flow. We then tested whether relational asymmetry biases social inference by examining the reported likelihood of information flow between mutual Friend and pairs where at least one actor rates the other as Friend, i.e., one-sided friendship, controlling for relational strength. Finally, we examine the egocentric modulation of these likelihood judgments by including the rater's own relational alignment and its interaction with the Source-Receiver dyad's relational strength into the model. Random intercepts and slopes by participant were included in regressions.

References

1. Redhead, D. & Power, E. A. Social hierarchies and social networks in humans. *Philos. Trans. R. Soc. B* **377**, 20200440 (2022).
2. Almaatouq, A., Radaelli, L., Pentland, A. & Shmueli, E. Are you your friends' friend? Poor perception of friendship ties limits the ability to promote behavioral change. *PloS One* **11**, e0151588 (2016).
3. Berger, C. R. & Calabrese, R. J. Some explorations in initial interaction and beyond: Toward a developmental theory of interpersonal communication. *Hum. Commun. Res.* **1**, 99–112 (1974).
4. Knobloch, L. K. & Solomon, D. H. Information seeking beyond initial interaction: Negotiating relational uncertainty within close relationships. *Hum. Commun. Res.* **28**, 243–257 (2002).
5. Berger, C. R. & Calabrese, R. J. Some Explorations in Initial Interaction and Beyond: Toward a Developmental Theory of Interpersonal Communication. *Hum. Commun. Res.* **1**, 99–112 (1975).
6. Kellermann, K. & Reynolds, R. When Ignorance Is Bliss The Role of Motivation to Reduce Uncertainty in Uncertainty Reduction Theory. *Hum. Commun. Res.* **17**, 5–75 (1990).
7. Knobloch, L. K. Perceptions of turmoil within courtship: Associations with intimacy, relational uncertainty, and interference from partners. *J. Soc. Pers. Relatsh.* **24**, 363–384 (2007).
8. Knobloch, L. K. & Solomon, D. H. Relational Uncertainty and Relational Information Processing: Questions without Answers? *Commun. Res.* **32**, 349–388 (2005).
9. Escribano, D., Lapuente, F. J., Cuesta, J. A., Dunbar, R. I. M. & Sánchez, A. Stability of the personal relationship networks in a longitudinal study of middle school students. *Sci. Rep.* **13**, 14575 (2023).
10. González-Casado, M. A., Teixeira, A. S. & Sánchez, A. Evidence of equilibrium dynamics in human social networks evolving in time. *Commun. Phys.* **8**, 227 (2025).
11. Paul, S. & O'Malley, A. J. Hierarchical Longitudinal Models of Relationships in Social Networks. *J. R. Stat. Soc. Ser. C Appl. Stat.* **62**, 705–722 (2013).
12. Newman, M. E. J. Scientific collaboration networks. I. Network construction and fundamental results. *Phys. Rev. E* **64**, 016131 (2001).
13. Zhang, J., Ackerman, M. S. & Adamic, L. Expertise networks in online communities: structure and algorithms. in *Proceedings of the 16th international conference on World Wide Web* 221–230 (Association for Computing Machinery, New York, NY, USA, 2007). doi:10.1145/1242572.1242603.
14. Krackhardt, D. & Hanson, J. R. Informal networks. *Harv. Bus. Rev.* **71**, 104–111 (1993).
15. Borgatti, S. P., Mehra, A., Brass, D. J. & Labianca, G. Network Analysis in the Social Sciences. *Science* **323**, 892–895 (2009).
16. Granovetter, M. S. The Strength of Weak Ties. *Am J Sociol* **78**, 1360–1380 (1973).
17. Marsden, P. V. Network Data and Measurement. *Annu. Rev. Sociol.* **16**, 435–463 (1990).
18. Wasserman, S. & Faust, K. Social network analysis: Methods and applications. (1994).

19. An, W. & Mcconnell, W. R. The origins of asymmetric ties in friendship networks: From status differential to self-perceived centrality. *Netw. Sci.* **3**, 269–292 (2015).
20. Hallinan, M. T. The process of friendship formation. *Soc. Netw.* **1**, 193–210 (1978).
21. Holland, P. W. & Leinhardt, S. An Exponential Family of Probability Distributions for Directed Graphs. (1981).
22. Aslarus, I. C., Son, J.-Y., Xia, A. & FeldmanHall, O. Early insight into social network structure predicts climbing the social ladder. *Sci. Adv.* **11**, eads2133 (2025).
23. Basyouni, R. & Parkinson, C. Mapping the social landscape: tracking patterns of interpersonal relationships. *Trends Cogn. Sci.* **26**, 204–221 (2022).
24. Teoh, Y. Y., Son, J.-Y., Xia, A., Bhandari, A. & FeldmanHall, O. Medial temporal lobe encodes cognitive maps of real-world social networks. 2025.03.18.643925 Preprint at https://doi.org/10.1101/2025.03.18.643925 (2025).
25. Xia, A., Teoh, Y. Y., Nassar, M. R., Bhandari, A. & FeldmanHall, O. Knowledge of information cascades through social networks facilitates strategic gossip. *Nat. Hum. Behav.* **9**, 2169–2182 (2025).
26. Snijders, T. A. Statistical models for social networks. *Annu. Rev. Sociol.* **37**, 131–153 (2011).
27. Shannon, C. E. A mathematical theory of communication. *Bell Syst. Tech. J.* **27**, 379–423 (1948).
28. Friedkin, N. E. Information flow through strong and weak ties in intraorganizational social networks. *Soc. Netw.* **3**, 273–285 (1982).
29. Smith, E. B., Brands, R. A., Brashears, M. E. & Kleinbaum, A. M. Social Networks and Cognition. *Annu. Rev. Sociol.* **46**, 1–16 (2020).
30. Snijders, T. A., Van de Bunt, G. G. & Steglich, C. E. Introduction to stochastic actor-based models for network dynamics. *Soc. Netw.* **32**, 44–60 (2010).
31. Marsden, P. V. & Campbell, K. E. Measuring tie strength. *Soc. Forces* **63**, 482–501 (1984).
32. Xiang, R., Neville, J. & Rogati, M. Modeling relationship strength in online social networks. in *Proceedings of the 19th international conference on World wide web* 981–990 (Association for Computing Machinery, New York, NY, USA, 2010). doi:10.1145/1772690.1772790.
33. Tsang, S., Barrentine, K., Chadha, S., Oishi, S. & Wood, A. Social exploration: How and why people seek new connections. *Psychol. Rev.* **132**, 656–679 (2025).
34. Larson, J. M. The weakness of weak ties for novel information diffusion. *Appl. Netw. Sci.* **2**, 14 (2017).
35. Boda, Z., Elmer, T., Vörös, A. & Stadtfeld, C. Short-term and long-term effects of a social network intervention on friendships among university students. *Sci. Rep.* **10**, 2889 (2020).
36. Newcomb, T. M. The acquaintance process as a prototype of human interaction. (1961).
37. Levin, D. A. & Peres, Y. *Markov Chains and Mixing Times*. vol. 107 (American Mathematical Soc., 2017).

Supplement
Supplementary Methods

Detailed balance analysis

Overall relationship category transitions are consistent with a network approaching stationarity—i.e., general reduction of mutual strangers and increased proportion of other relationship categories. However, this aggregate movement does not reveal directional changes between relationship states. For example, is the decrease in the proportion of mutual strangers driven by strangers becoming mutual friends, or strangers becoming Familiar-Stranger pairs? To depict the directional flow between relationship states, we leveraged a detailed balance analysis, which is widely used in statistical physics to examine whether a system has reached equilibrium [1]. Specifically, for each relationship type (16 total), we apply the following equation:

$$P(i)P(j|i) = P(j)P(i|j),$$

where $P(i)$ is the proportion of relationship type i, and $P(j|i)$ is the probability of transitioning from i to j. If the equation holds, it suggests that the flow between these two relationship types is equivalent, canceling out each other at the aggregate level. If the left-hand side is larger than the right-hand side, it suggests a predominant outflow from relationship type i to relationship type j. Because each observed transition is a single sample from the underlying dynamics governing network evolution, we estimated confidence intervals (CIs) using the binomial variance approximation [2]:

$$95\%\ CI \approx p \pm 1.96 \times \sqrt{\frac{p \times (1-p)}{N}}$$

where $p = \frac{x}{N}$ represents the probability of interest, i.e., $P(i)$ or $P(j|i)$. N refers to the total number of observations, which is the total number of unique pairs that constitute $P(i)$ and pairs of relationship type i prior to transition for $P(j|i)$. Detailed balance was considered satisfied when the CIs for both sides of the equation overlapped.

A small proportion of transitions exhibited imbalanced flow (0.020, 0.027, and 0.08 of all possible transitions across the three time intervals; see Table S2). Among these, we observed a consistent directional pattern: outflow from mutual strangers toward categories such as Familiar-Stranger and Familiar. It is worth noting that violations of detailed balance do not contradict the stationarity finding; rather, it identifies the specific transitions driving aggregate change, revealing that network evolution is distributed across many transition pathways rather than dominated by any single transition.

Computing statistical equivalence between transition matrices

To test whether transitions between snapshots of the network are stable, we adopt the methods used in past work [2] and examine whether the transition matrices M^T are statistically equivalent to the average transition matrix $\overline{M}$, constructed by pooling all the transitions between consecutive snapshots. If each M^T is statistically equivalent to $\overline{M}$, it can be treated as a resample of the same underlying process, indicating that turnover dynamics remain constant over time. For each transition (e.g., between Survey 1 and 2), we constructed a simulated transition matrix $\widehat{M^T}$ by sampling the same number of pairs (e.g., 15,931 pairs for the transition from Survey 1 to 2) from the pooled transitions across all timepoints. We repeated this simulation 100 times, resulting in a distribution of simulated transition matrices. Since the simulated matrices were constructed using the same data that the average transition matrix is based on, each of the simulated matrix can be seen as a random sample that statistically equivalent to the average transition matrix. We then quantified, for each time interval, the distance between the observed transition matrix and the average matrix, relative to the distance between the simulated matrices and the average matrix. The goal of this was to determine whether the observed transition matrix is statistically different from the average matrix. For each element of the average matrix $\bar{M}_{ij}$, we compute the difference between the average matrix and the simulated matrices, which we termed as the distribution of null distance:

$$\hat{D}_{ij} = \widehat{M^T}_{ij} - \bar{M}_{ij}$$

This null distance distribution represents the level of variation that is likely to be caused by sampling noise. We can then compare the null distance distribution to the observed distance:

$$D_{ij} = M_{ij}^T - \bar{M}_{ij}$$

If the observed distance is extreme compared to the null distribution, it indicates a deviation for the time-specific transition from the average transition matrix. To quantify this, we can compute z-score:

$$z_{ij} = \frac{D_{ij} - \mu(\hat{D}_{ij})}{\sigma(\hat{D}_{ij})}$$

We also calculated the proportion of times the absolute observed distance is smaller than the absolute value of the null distance, which serves as the p-value for testing whether the transition probability of moving from relationship type i to j between two snapshots is different from the average transition probability from i to j in $\bar{M}$. Given multiple comparisons across the whole transition matrix, we adopted Bonferroni correction with a critical value of $0.05/256$ (total number of cells in the transition matrix). The proportion of transitions falling below this threshold was 0.004, 0.012, and 0.012 for each time-specific transition matrix respectively (z-scores for deviating transitions are shown in Table S1). This is comparable to, or smaller than, the range reported by González-Casado et al. (2025) for a stable network (0.0096-0.0224), suggesting that the overall stochastic process governing the transition between timepoints is stable.

Result S1: Replication of the nonlinear asymmetry resolution effect in a separate social network

To test the generalizability of our network-level and dyad-level findings, we collected relationship survey data from an independent freshman cohort (N = 176) with three survey waves administered within the first semester. This dataset was designed to replicate the network dynamics and dyad-level transition analyses. The Network Knowledge Task was not administered due to time constraints, so group-level perception analyses were not replicated.

Participants rated their relationship with other network members by selecting one of five options: Friend, Acquaintance, Recognize, Do Not Know, and Do Not Get Along. As the rating 'Do Not Get Along' carries negative affective valence and likely reflects relational conflict rather than mere asymmetry, we exclude pairs where at least one actor chose 'Do Not Get Along' to maintain construct consistency across datasets.

First, we replicated our network-level findings, showing that pairs with high relational asymmetry are less common in the network (average proportion of symmetric relationships: 78.63%; relationships with asymmetry magnitude = 1: 18.81%; relationships with asymmetry magnitude = 2: 2.29%; relationships with asymmetry magnitude = 3: 0.27%; Spearman correlation between relational asymmetry and proportion: $\rho = -0.69$, $p = .027$). Second, we again showed that asymmetric dyads have higher conditional entropy overall (Wilcoxon rank-sum = 2, $p = .038$; Spearman correlation $\rho = 0.77$, $p = .009$).

Finally, we replicated our dyad-level findings. In this network dataset, asymmetric dyads are more likely to reach symmetry relative to symmetric dyads moving into asymmetry (35.3% vs. 13.0%; two-sample proportion test: $\chi^2(1) = 491.33$, $p < .001$). We also replicated the nonlinear relationship between asymmetry magnitude and asymmetry resolution ($\beta_{\text{relational asymmetry}^2}$ = 7.26, SE = 2.42, 95% CI [2.53, 12.01], p = .003), suggesting that extreme relational asymmetry, at both the low and high ends, is associated with a higher chance of reaching symmetry at the subsequent timepoint.

Result S2: Wave 2 replication—Network Knowledge Task

We collected a second wave of the Network Knowledge Task in late spring, using ratings from the preceding Relationship Survey 3 as the ground-truth relationship status. Here, we report results replicating our main analyses.

We find that group-level disagreement is positively associated with greater relational asymmetry ($\beta_{\text{relational asymmetry}}$ = 0.007, *SE* = 0.002, 95% CI [0.004, 0.011], $p < .001$). Focusing on asymmetric pairs, while we do not observe the main effect of the network member proximity to either the more or less invested side (β = 0.002, *SE* = 0.007, 95% CI [-0.012, 0.015], p = .816), we replicated the finding that network members' judgments are jointly influenced by whether the network member is closer to the invested member of the dyad and the pair's relational strength (β = -0.020, *SE* = 0.007, 95% CI [-0.033, -0.007], p = .003).

Result S3: Entropy of group-level friendship judgments tracks relational asymmetry

We computed entropy of group-level friendship judgments as an alternative measure of disagreement, complementing the standard deviation reported in the main text:

$$H = -\sum_{i=1}^{k} p(i) log_2 p(i)$$

Here, i indicates the judgment of each participant in the Network Knowledge Task, which takes the value of 1 if a participant judged the pair as 'Friends' in the A→B direction, 0 if a participant judged the pair as 'Not friends' in the A→B direction, and 0.5 if the judgments for both directions are different. Entropy is highly correlated with standard deviation ($r(6281) = .894$, $p < .001$) and is invariant with respect to the number of judgments made for each unique pair. Entropy is also positively associated with ground-truth relational asymmetry of pairs ($\beta = 0.016$, $SE = 0.004$, 95% CI [0.007, 0.024], $p < .001$), providing converging evidence that uncertainty in group-level judgments is associated with uncertainty in a dyad's relationship state.

Result S4: Wave 2 replication—Social inference tasks (Generalization and Information Flow)

In the second wave of the Generalization Task, administered after Relationship Survey 3, we again find that people are more likely to infer the existence of an unseen tie when the existing members of the triad are closer to each other ($\beta_{\text{triad's relational strength}} = 0.351$, *SE* = 0.025, 95% CI [0.300, 0.401], $p < .001$), while greater relational asymmetry within the triad reduces the likelihood of inferring a tie ($\beta_{\text{triad's asymmetry}} = -0.317$, *SE* = 0.066, 95% CI [-0. 447, -0.187], $p < .001$). Similarly, in the Information Flow Task, we find that the likelihood of information flow positively correlates with the ground-truth relational closeness of the pair ($\beta_{\text{relational strength}} = 0.105$, *SE* = 0.002, 95% CI [0.101, 0.109], $p < .001$; Figure 4C). For pairs where at least one individual rates the other person as friend, mutual friendship is again inferred by observers to support a higher probability of information flow compared to one-sided friendship ($\beta = 0.108$, *SE* = 0.023, 95% CI [0.063, 0.153], $p < .001$).

We also replicate our findings that the participants' relational alignment modulates their ability to predict tie formation and information flow. Results from the Generalization Task indicates a negative interaction between local alignment and closeness of the existing triad ($\beta_{\text{relational strength:alignment}} = -0.064$, *SE* = 0.026, 95% CI [-0.151, -0.013], $p = .013$), suggesting that people with less aligned relationships are less able to reason about potential new relationships based on existing network structure. Similarly, we find a negative interaction between relational alignment and relationship closeness when predicting the probability of the information flow ($\beta_{\text{relational strength:alignment}} = -0.017$, *SE* = 0.002, 95% CI [-0.021, -0.013], $p < .001$), suggesting that people who are in a less aligned local social environment have a diminished ability to use ground-truth relationship status when making predictions.

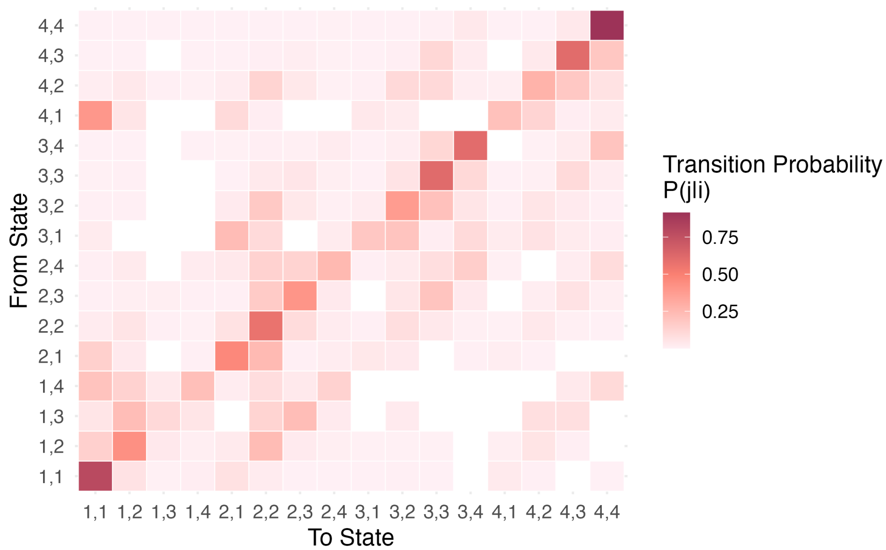

Fig. S1. Average transition matrix (pooled across all timepoints). The transition matrix is constructed using data collapsing across all transitions. From State indicates the relationship status before transition, and To States indicate the relationship type after transition obtained from the consecutive snapshot. Each relationship type is indicated by the ratings from two actors within the dyad. Rating scale: 1 = Friend, 2 = Acquaintance, 3 = Familiar, 4 = Stranger.

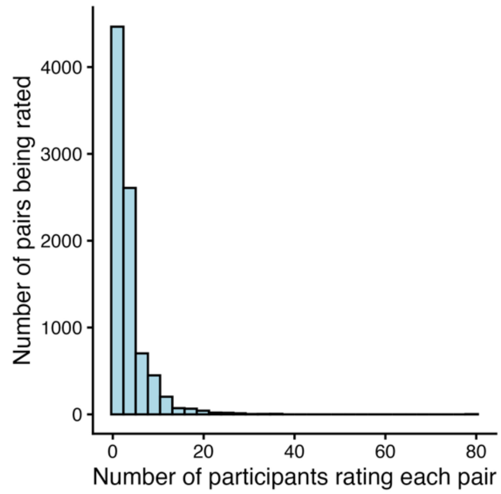

Fig. S2. Number of times that each pair get rated in the Network Knowledge Task.

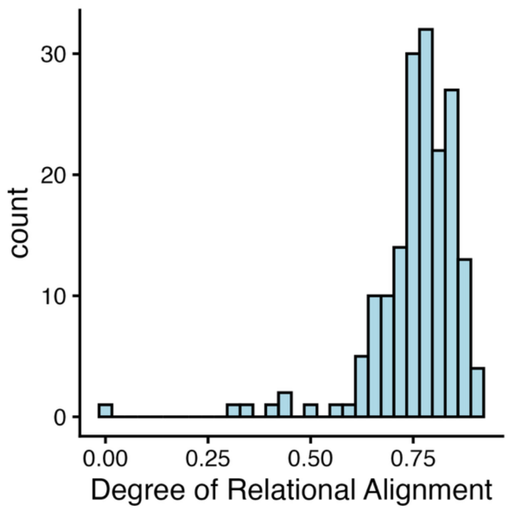

Fig. S3. Distribution of relational alignment among participants. Relational alignment is calculated as the correlation between individual's rating for other network members and the ratings from other network members to the individual, and is bounded between 0 and 1.

Table S1: Time-specific transitions that deviate from average transition matrix

From	To	Average Transition Probability	Transition Probability (Survey 1-2)	Transition Probability (Survey 2-3)	Transition Probability (Survey 3-4)	Z-score (Survey 1-2)	Z-score (Survey 2-3)	Z-score (Survey 3-4)
1,1	1,1	0.79	0.80	**0.74**	0.83	0.56	**-2.69**	2.36
4,2	1,2	0.04	0.01	**0.08**	0.02	-1.70	**3.37**	-1.15
1,2	2,2	0.24	0.24	**0.34**	**0.15**	0.25	**2.99**	**-2.84**
2,2	3,2	0.08	0.10	0.10	**0.05**	1.58	1.43	**-2.94**
2,3	3,2	0.05	**0.08**	0.02	0.04	**3.03**	-1.94	-0.58
4,2	4,2	0.28	0.23	0.23	**0.38**	-1.43	-1.31	**3.22**

Note: "From" column shows the relationship type before transition, and "To" column shows the relationship type after transition. Each relationship type is indicated by the ratings from two actors within the dyad. 1=Friend, 2=Acquaintance, 3=Familiar, 4=Stranger. Transitions that significantly deviate from the average transition probability are bolded, using a critical value of 0.05/256. Z-score indicates the degree of deviation from the average estimate and are bolded according to significant p-values.

Table S2: Transitions with imbalanced flow

From (Outflow)	To (Inflow)	Transition
4,4	2,1	Survey 1-2
4,2	3,3	Survey 1-2
4,4	3,3	Survey 1-2
4,4	3,4	Survey 1-2
4,4	4,3	Survey 1-2
1,2	1,1	Survey 2-3
2,2	2,1	Survey 2-3
4,4	3,2	Survey 2-3
4,4	3,3	Survey 2-3
4,4	3,4	Survey 2-3
4,4	4,2	Survey 2-3
4,4	4,3	Survey 2-3
4,4	2,1	Survey 3-4
4,4	4,3	Survey 3-4

Note: “From” column shows the relationship type that creates outflow, and “To” column shows the relationship type that receives inflow. Each relationship type is indicated by the ratings from two actors within the dyad. 1=Friend, 2=Acquaintance, 3=Familiar, 4=Stranger.

SI References:

1. Farkas, I., Derényi, I., Palla, G. & Vicsek, T. Equilibrium statistical mechanicsof network structures. in *Complex networks* 163–187 (Springer, 2004).
2. González-Casado, M. A., Teixeira, A. S. & Sánchez, A. Evidence of equilibrium dynamics in human social networks evolving in time. *Commun Phys* **8**, 227 (2025).

Chapter 5

General Discussion

1. Summary of results

How do people navigate a social world they can never fully observe? The three empirical studies presented in this dissertation converge on the answer that people encode flexible representations of multistep connectivity through a social network rather than veridical memory of direct ties. Representations of this nature facilitate a range of important social functions: gossiping strategically, forecasting who will become friends in the future, and detecting the internal tensions within dyads they are not personally part of. What changes from study to study is not the underlying cognitive machinery, but the conditions under which it must operate.

Chapter 2 established that people quickly acquire knowledge about the structure of a network, in service of a complex communication goal. Using a task that requires broadcasting information while preventing it from reaching a specific target, I showed that people's decisions were not governed by explicit memory of who is friends with whom. Memory accuracy for individual ties was, in fact, a poor predictor of strategic behavior. What predicted behavior was the participant's sensitivity to the combination of path-based distance and popularity, which is consistent with Katz communicability—a model that integrates all paths through a network and preserves the influence of well-connected individuals rather than diluting it. This dissociation between explicit relational memory and adaptive inference is one of the central empirical results of this dissertation. People appear to readily extract latent structural information about a set of ties, to determine how influence might accumulate across paths, without necessarily storing details about individual relationships.

In **Chapter 3**, I demonstrated how inferential processes relating to a small artificial network, constructed and test in the lab, generalize to the real world. Leveraging a longitudinal design within a cohort of incoming undergraduates, I asked whether the errors people make in judging others' relationships are informative about how those relationships will evolve, many months later. I found that misjudged ties—pairs that participants identified as friends when they were not yet friends in December—were nearly three times more likely to become friendships by the following May, an effect that was replicated across different time windows and was not reducible to triadic closure, homophily, or response bias. Computational modeling revealed that the predictive signal in participants' errors was best explained by their inferred Katz communicability values, which reflect the latent, multi-connectivity strength between pairs. Individuals who encoded a richer representation of such paths (indexed by a higher γ parameter) were selectively better at predicting tie formation, and not tie dissolution, suggesting that this representational format is specifically tuned to the emerging structure of the network rather than its decay. The consequences of this capacity were tangible. Accurate link prediction boosted the social standing of initially well-connected individuals, while closing the structural gaps previously spanned by brokers. The ability to model the network's future is not merely a cognitive feat but a mechanism through which social capital can be gained.

Chapter 4 turned to a phenomenon that Chapter 3 had treated as a simplifying assumption, and that most network science research treats the same way: the reciprocity of ties. In the real world, two people can hold quite discrepant views of the same relationship—one considering the other a close friend while the other sees them merely as an acquaintance. I documented that relational asymmetry of this kind is pervasive in the network of students, comprising more than half of all non-stranger ties in the network, and that it operates under its own dynamics. Despite a stable distribution of different types of asymmetric ties at the network level (e.g., Friend-Familiar, Familiar-Acquaintance), individual dyads continuously cycle through distinct states of alignment and misalignment, moving toward symmetry at above-chance rates. The trajectory of resolution is nonlinear, such that individuals that are members of dyads characterized by the smallest and largest discrepancies are both especially likely to move towards each other. Crucially, the community tracks this. Group-level disagreement about a pair's relationship status scales with the ground-truth degree of asymmetry within that pair. People also use this signal adaptively. The average amount of asymmetry within a triad constrains how readily people believe a new person will join their friendship cluster, and asymmetry between two individuals attenuates perceived information flow between them. Above and beyond being a dynamic property of dyads, an individual's own relational alignment—the degree to which their social perceptions are reciprocated—modulates the capacity to use existing network structure to make social inferences, suggesting a feedback loop in which relational misalignment degrades the precision of one's model of their surrounding network.

2. Limitations and future directions

There are several limitations of the present work. Chapters 3 and 4 both rely on a single cohort of first-year undergraduate students in a relatively closed residential network. The network structure that emerges from this kind of environment differs in important ways from the networks people navigate across their lifespans. The social density of a university dormitory, the temporal marking of the academic year, and the fact that participants have limited prior history with one another all shape the dynamics documented here. Whether the same inferential and computational principles apply in established friendship networks, or professional networks organized around hierarchy rather than homophily remains an open question.

The temporal resolution of survey-based design is another limitation. The present studies cannot fully disentangle how much of the predictive signal in people's representations reflects passive readout of structural patterns versus active shaping of those patterns through social behavior. People's expectations about who will become friends may themselves influence whether those friendships form. Chapter 3 provides indirect evidence that this loop exists: accurate link prediction shapes network position over time, suggesting that representations do not merely anticipate the network's future but participate in producing it. Disentangling these contributions would require finer-grained, experience-sampling designs that capture daily social interactions alongside periodic snapshots of the network, which would make it possible to track how specific encounters relate to the mental representations that people build.

A related open question concerns how a Katz-like representation gets built in the first place. The Successor Representation is constructed upon principles of trial-and-error, incremental learning, since it can be acquired through temporal difference updates as an agent moves through an environment. Currently, Katz communicability lacks an equivalent account. It is a closed-form estimate of all paths through a network, which makes it well-suited as a representational format, but it is not obvious what prediction error would drive an agent toward it, or perhaps even more challenging, how it would be updated as the network changes. The finding that the Katz model consistently outperforms the SR behaviorally, across these studies, raises a question the experiments cannot answer: how do people arrive at a Katz-like representation from the kind of piecemeal, sequential social experience that real networks afford? Identifying the computational principles by which people encode something like Katz communicability from the different types of learning experiences that real networks afford, and over different time horizons, is the next foundational question this computational framework needs to address.

A separate question relates to what a Katz-like representation is capable of encoding. Chapter 4 showed that relational asymmetry weakens the functional precision of people's inferences, which depend on having structure-based network representations. Although the variant of the Katz model used in the real-world network studies rely on an asymmetric adjacency matrix as a modeling accommodation, it is not a theoretical claim about people's perceptions of the pervasive relational asymmetry in their social network. It remains unclear how asymmetry should be incorporated into the representational formalism itself. The current model treats directed ties as a proxy for partial knowledge, not as a signal about the internal state of the dyad. But the results of Chapter 4's suggest that asymmetry contains useful, predictive signals separable from overall tie strength. A natural next step would be to extend the Katz framework to incorporate dyadic asymmetry explicitly, perhaps by using a weighting term that reduces the communicability contribution of more uncertain ties, and test whether this extension improves the fit to participants' inferences in scenarios where asymmetry is relevant to the task at hand.

3. Conclusions

Social life requires navigating a structure that is too large to observe directly, too dynamic to treat as fixed, and too consequential to navigate by random guesswork. The work presented in this dissertation argues that people meet these challenges not by maintaining a faithful record of who is connected to whom, but by building compressed representations of network topology that preserve the properties most relevant for inferences that go beyond a few pairs of individuals.

The computational framework at the core of this argument—Katz communicability—predates the cognitive questions posed here by decades, having originated as a metric of total accumulated influence (i.e., eigenvector centrality), rather than the simpler notion of degree, or one's number of friends. What this dissertation contributes is evidence that the human mind spontaneously constructs something like Katz communicability. People do not represent social networks as raw adjacency matrices, weighted by who they remember to be connected. They build a mental model of a social network as a structure through which influence radiates, favoring the well-connected and discounting the distant, in a way that naturally captures cascade-like dynamics through which information and other actions can propagate.

The relationship between social cognition and social structure has most often been understood as unidirectional. Structure shapes the information people have access to, and that access shapes what they can know. The present work suggests the relation runs in both directions. The representations people carry are shaped by the networks they inhabit, but those representations also feed back into how the network develops. The sophisticated yet efficient cognitive machinery leading us to everyday social success is not merely a system for perceiving an external world—it is one of the very forces that shape it.