A causal framework for the drivers of animal social network structure

Ben Kawam[1, 2, 3, †], Julia Ostner[1, 2, 4], Richard McElreath[3], Oliver Schülke[1, 2, 4, ‡], Daniel Redhead[3, 5, 6, ‡]

Abstract. A major goal of behavioural ecology is to explain how phenotypic and ecological factors shape the social relationships that animals form with one another. This inferential task is notoriously challenging. The social networks of interest are generally not observed, but must be approximated from behavioural samples. Moreover, these data are highly dependent: the observed network edges correlate with one another, due to biological and sampling processes. Failing to account for the resulting uncertainty and biases can lead to dysfunctional statistical procedures, and thus to incorrect results. Here, we argue that these problems should be understood—and addressed—as problems of causal inference. For this purpose, we introduce a Bayesian causal modelling framework that explicitly defines the links between the target interaction network, its causes, and the data. We illustrate the mechanics of our framework with simulation studies and an empirical example. First, we encode causal effects of individual-, dyad-, and group-level features on social interactions using Directed Acyclic Graphs and Structural Causal Models. These quantities are the objects of inquiry, our *estimands*. Second, we develop *estimators* for these effects—namely, Bayesian multilevel extensions of the Social Relations Model. Third, we recover the structural parameters of interest, map statistical estimates to the underlying causal structures, and compute causal *estimates* from the joint posterior distribution. Throughout the manuscript, we develop models layer by layer, thereby illustrating an iterative workflow for causal inference in social networks. We conclude by summarising this workflow as a set of seven steps, and provide practical recommendations.

Keywords: Bayesian modelling, causal inference, Directed Acyclic Graph (DAG), generative modelling, latent network, socioecology, social behaviour, Structural Causal Model (SCM).

[1] Department of Behavioural Ecology, University of Göttingen, Kellnerweg 6, 37077 Göttingen, Germany.
[2] Primate Social Evolution Group, German Primate Center Göttingen, Kellnerweg 4, 37077 Göttingen, Germany.
[3] Department of Human Behaviour, Ecology and Culture, Max Planck Institute for Evolutionary Anthropology, Deutscher Platz 6, 04103 Leipzig, Germany.
[4] Leibniz ScienceCampus Primate Cognition, German Primate Center, Leibniz Institute for Primate Research, Kellnerweg 4, 37077 Göttingen.
[5] Department of Sociology, University of Groningen, Grote Rozenstraat 31, 9712 TG Groningen, The Netherlands.
[6] Inter-University Center for Social Science Theory and Methodology, University of Groningen, Groningen, The Netherlands.
[†] Corresponding author: *Ben Kawam*, bkawam@dpz.eu.
[‡] These authors contributed equally to this work.

Table of Contents

1 Introduction **3**
- 1.1 Common problems in animal social network analysis 4
- 1.2 Causal inference in animal behavioural ecology 6

2 Mechanics of the framework **8**
- 2.1 Simulation study 1: random structuring features 9
- 2.2 Simulation study 2: individual-level features 19
- 2.3 Simulation study 3: dyad-level features 23

3 Empirical example **27**
- 3.1 Simulation study 4: kinship in female macaques 28
- 3.2 Empirical data 32

4 General discussion **38**
- 4.1 Future research avenues 40

S1 Supplementary Materials **43**
- S1.1 Simulation study 1: random structuring features 43
- S1.2 Simulation study 2: individual-level features 48
- S1.3 Simulation Study 2': categorical individual-level features 51
- S1.4 Simulation study 3: dyad-level features 55
- S1.5 Simulation study 3: alternative parameterisation 57
- S1.6 Simulation Study 3': categorical dyad-level features 59
- S1.7 Simulation Study 4: kinship in female macaques 64
- S1.8 Empirical Study 70
- S1.9 Simulation Study 4': group-level effects 76
- S1.10 Empirical Study' 80

References **81**

1. Introduction

A major goal of behavioural ecology is to explain how ecological and evolutionary processes affect social structure (Whitehead, 2008). Behavioural ecologists observe natural variation in social behaviour, and ask: *"why?"* (Davies et al., 2012; Tinbergen, 1963). Why do certain individuals cooperate by supporting each other against conspecifics, or by spending a substantial amount of time grooming? Why do other individuals confront each other in agonistic fights, sometimes at the risk of their lives? Ecology and evolution offer theoretical models to explain why individuals behave in certain ways (Fu et al., 2012; Hinde, 1976b; Nowak, 2006; Seyfarth, 1977; Smaldino, 2023). These models generate theoretical predictions about how the features of individuals (*e.g.*, age, social status), of dyads (*e.g.*, genetic relatedness, homophily) or of groups (*e.g.*, group size) can affect the social relationships they form with one another.

Social structure
Pattern of social interactions and relationships among the members of a population.

Network Science provides valuable analytical tools to bridge theoretical predictions with empirical data (Newman, 2018). To make inferences about the determinants of social structure, it is useful to first operationalise the system as a social network (Croft et al., 2008; Krause et al., 2015; Newman, 2018; M. J. Silk, 2023). Typically, a social network is composed of nodes that represent individual animals, and edges, which represent the social interactions or relationships between them. We find it important to distinguish three levels of abstraction for any social network under investigation (Figure 1). These levels differ regarding what the edges represent. In the highest level of abstraction, the network edges correspond to the *theoretical construct* that we wish to study (first level in Figure 1).[1] Often, that is the social relationship between two individuals, or an aspect of it like affiliation, dominance, agonistic support, tolerance, or friendship (Hinde, 1976a, 1976b, 1983; Kummer, 1978; J. B. Silk, 2002). These constructs cannot be observed directly (De Moor et al., 2024; Smaldino, 2023). They are abstractions, often assumed to be composed of—or expressed as—a diverse range of behavioural interactions, and can only be approximated.

Social network
Graph where nodes represent individuals, and edges represent either (i) an aspect of the social relationship, (ii) the true interaction rates for a given behaviour, or (iii) the measured interaction rates for a given behaviour, between individuals.

In the second level of abstraction, the network edges commonly correspond to one type of quantifiable social interaction, which is used as a proxy for the theoretical construct of interest (Carter et al., 2015; Whitehead, 2008). We refer to this level as the *interaction network* (Figure 1). Researchers choose which behavioural proxy to use based on knowledge of their study system (*"How do I best approximate what I want to capture?"*). For instance, allogrooming (*e.g.*, J. B. Silk et al., 2006), or spatial proximity (*e.g.*, Gerber et al., 2022), may capture affiliative relationships in different species. Researchers also determine which proxy to use based on practical constraints (*"What can I measure?"*). Certain behavioural interactions may be reliable indicators of affiliation, but occur too rarely to be of practical use (*e.g.*, "bridging" behaviour in macaques, see Kalbitz et al., 2017). How the theoretical construct of interest relates to the interaction network is of critical importance for future research (see De Bacco et al., 2023; Redhead, McElreath, et al., 2023; Young et al., 2020). It is, however,

[1] We number the layers *starting* with theory and ending with data, thereby flipping the order of Hinde's 1976b conceptual model.

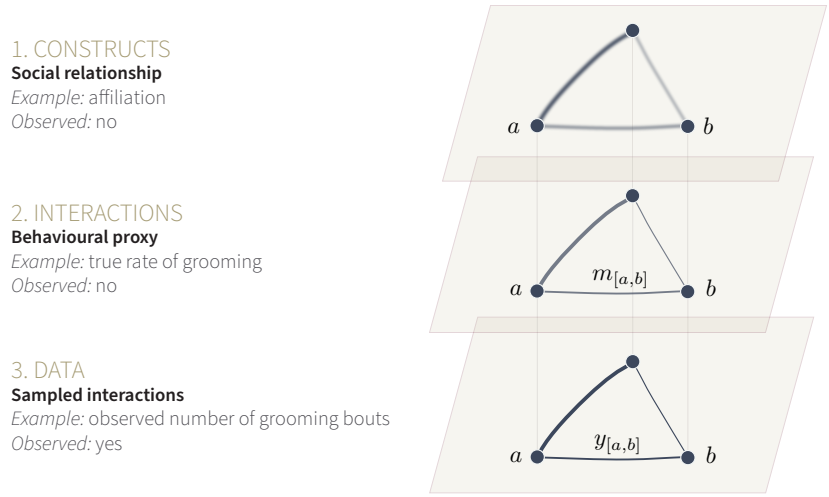

Figure 1 – *Three levels of abstraction for an animal social network.* On the three levels, the dark dots, or *nodes*, represent individual animals, two of which are named a and b. We show a network of only three individuals for simplicity of representation. The lines connecting individuals, or *edges*, depict the relationships or interactions among them. On all three levels, lines of different widths suggest variation in the strength of relationships, or in the number of interactions.

beyond the scope of the present manuscript. That is, we take for granted that the interaction network is a reasonable proxy for the theoretical construct of interest.

Third, it is typically impossible to observe every interaction event within the study system. Instead, researchers sample the study population using standardised behavioural protocols (Altmann, 1974). The data generated through the use of these protocols—*e.g.*, 1% of all grooming bouts—are then used to approximate the network of all interactions (Whitehead, 2008). These *data*, which we also further refer to as *sample* or *observations*, correspond to the third level of abstraction (Figure 1).

In the following sections, we use these concepts to describe common issues in animal social network analysis, and argue that they should be understood as problems of causal inference. We then introduce a novel, relatively simple, framework that explicitly addresses these issues.

1.1. Common problems in animal social network analysis

Establishing the connection between observations (third level) and unobserved interaction network (second level) is a major inferential challenge in animal social network analysis. A popular approach to approximate the interaction network is to *aggregate* the observed interactions at the individual- or dyad-level. These approaches include the use of Simple Ratio Index (Farine & Whitehead, 2015) and composite "Dyadic Sociality Index" (J. B. Silk et al., 2013), to quantify edge weight. The edge weights can be aggregated at the node level to calculate, for instance, node degree

or strength (Ellis et al., 2019). These metrics are then commonly used as predictor or outcome variables in multivariable regression procedures (Farine & Whitehead, 2015).

However, using such network summary metrics in statistical models amounts to treating the unobserved interaction network (second level) as being observed. In practice, the observed interactions are *noisy samples* of the interaction networks (Hart et al., 2023; Mielke & Samuni, 2021). Suppose we are interested in studying the rate at which two individuals groom. In one case, we observe these two individuals for 1 hour, and record 1 grooming bout (*i.e.* a grooming "event") between them. In a second case, we observe the same two individuals groom 100 times over 100 hours of observation. In both cases, the observed rate (*i.e.*, the Simple Ratio Index) is equal to 1 grooming bout per hour. Yet, we are more *uncertain* about the real grooming rate in the first case, where the individuals have been observed for a shorter period. The data-aggregating procedures mentioned above would typically treat them as equivalent, and this invalidates probability calculations—different amounts of evidence must be reflected in inferential uncertainty. To deal with the uncertainty inherent to aggregating observations (hereafter referred to as *problem I*), behavioural ecologists traditionally attempt to collect a large amount of data to produce an accurate approximation of the unobserved interaction network (Krause et al., 2015; Whitehead, 2008). But whether a given amount of data is "sufficient" for an accurate approximation of the network is hard to determine.

A related issue is that of "*biases*" introduced by the sampling process (*problem II*). The sampling regime itself may make some individuals appear more or less sociable than they actually are. Farine (2017) provides an example where females are more gregarious than males in a hypothetical population of animals. However, females are sampled less regularly than the males (*e.g.*, because the males are more visible). As a result, the two sexes appear equally gregarious. Permutation methods have been proposed to solve these sorts of issues, sometimes framed as problems of "non-independence of social network data" (Farine, 2017; Farine & Carter, 2022; Farine & Whitehead, 2015). However, the utility of permutation methods for social network analysis has been challenged, as they carry several, fundamental, flaws (see Hart et al., 2022; McElreath, 2020; Ross et al., 2023; Weiss et al., 2021). Recent methodological work proposes to solve the issues of noise and biases introduced by the sampling process, by defining the true interaction rates as parameters in a Bayesian Model (Duboscq et al., 2023; Hart et al., 2023; Neumann & Fischer, 2023a; Redhead, McElreath, et al., 2023; Ross et al., 2023; Young et al., 2020). We will return to this idea in detail, further below.

Bayesian model
Statistical model where inference is based on the posterior probability distribution, which describes the plausibility of different parameter values for all parameters in the model. The posterior probability is computed by combining the observed data with prior probability distributions for the model parameters, using Bayes Theorem. Bayesian models are also sometimes called *probabilistic models*.

Social network data present yet another challenge—often framed as the "non independence in social network data" as well—but here, due to structuring features of the network. As opposed to the data dependencies caused to the sampling process, structural dependencies among edges are already present in the interaction network (second level). Consider, for instance, three edges from a hypothetical social network: the number of interactions given from an individual a to another individual b, the interactions from a to c; and those from a to d. The three edges will likely be more similar to each other than they are to the other edges in the network, because

Network structuring features
Variables shaping the network of social interaction. Features can be at the individual (*e.g.*, age), dyadic (*e.g.*, genetic relatedness), or supra-dyadic level (*e.g.*, predation pressure).

the actor, in the three cases, is a (Back & Kenny, 2010; Brask et al., 2023; Kenny & La Voie, 1984). Interactions from a to b, and those in the opposite direction, from b to a, often covary as well. In the presence of dyadic reciprocity, dyads with a high rate of interaction in one direction will have a high interaction rate in the opposite direction too (Puga-Gonzalez et al., 2018; Van Duijn et al., 2004). A third example regards triadic closure: it is often possible to learn about the value of the edge between $a - c$, given knowledge about the value of the edges $a - b$ and $b - c$. This can be the case because a, b, and c all belong to the same kin group, where they engage in more interactions than with non-kin. The line between the two types of dependencies we have introduced (*i.e.* due to sampling, and structural) is however not strict. An individual's features might affect how they behave—creating structural dependency among its edges—and, in turn, its behaviour can affect the sampling process, *e.g.*, if more gregarious individuals are sampled more regularly.

Triadic closure
In a triad composed of individuals a, b, and c, given that there is a connection between $a - b$, and one between $b - c$, there often exists a connection between $a - c$.

Ignoring structuring features of the network can have deleterious consequences, when the goal of a statistical analysis is to make inference about the effect of a given variable, like sex or age, on social behaviour. First of all, the effect of interest might be confounded (*problem III*). In this case, a statistical model describing the association between the predictor of interest and social behaviour will not recover the effect of interest, even with infinite sample size. Second, statistical models that ignore this structure can have low efficiency and low predictive accuracy (*problem IV*). That is, the models won't precisely recover generative parameters, and won't accurately predict unobserved data from the inferred interaction network (Feller & Gelman, 2015; McElreath, 2020). How to build statistical models in light of structuring features of the network has received ample attention in the field of (human) Social Network Analysis (Carrington et al., 2005; Ogburn et al., 2022; Rawlings et al., 2023; Shalizi & Thomas, 2011; Van Duijn et al., 2004; Wasserman & Faust, 1994). Statistical models in behavioural ecology are rarely built in light of the structuring features of the social network, and might greatly benefit from incorporating these approaches.

Confounder
A confounder of a "treatment" X and "outcome" Y is a variable that makes the observed statistical association between X and Y different than if X had been intervened upon—for instance, using a randomised control experiment. That is, $P(Y|do(X)) \neq P(Y|X)$, in *do*-calculus. Note that confounding is a causal, and not a statistical, concept.

In summary, inference from animal social networks implies several difficulties. Samples often look very different from the actual interaction network, because of noise (problem I) and potential biases (problem II) introduced by the sampling process. Structuring features of the social network further create associations between edges in the interaction network, and should be considered when building statistical models (problems III and IV). In the remaining sections of this manuscript, we argue that all of these issues are part of—and should be addressed as—one larger kind of problem: a problem of causal inference.

Causal inference
Refers to both a discipline and a process. Discipline: field studying causal relationships among variables. Process: inferring causal effects from data.

1.2. Causal inference in animal behavioural ecology

Behavioural ecologists ask mechanistic questions about behaviour, causal questions. However, like other scientists, they face a dilemma. Many are taught that the *only* way to infer causality from data is to run a randomised controlled experiment (Shipley, 2016). Arguing otherwise can be perceived as endorsing the statement that *"correlation implies causation"* (Pearl & Mackenzie, 2018; Rohrer, 2018, 2023). Yet, researchers often *cannot* run such experiments in their study system, for ethical or

Causal effect
A variable X has a causal effect on a variable Y, if a (hypothetical) intervention on X results in a change in Y: $P(Y) \neq P(Y|do(X))$, in *do*-calculus.

practical reasons. They must address causal questions using observational data, but under the idea that causal inference from non-experimental data is impossible.

As a result, many observational studies are causally ambiguous (Byrnes & Dee, 2024; Grosz et al., 2020; Rohrer, 2024). It is common to see scientific papers reporting observational studies with causal claims in their titles and abstracts—*e.g.*, X *drives* Y in species Z. These papers begin with a largely causal reading of the literature (*e.g.*, authors J found an *effect* of X on Y in species W). The Methods section, however, rarely contains transparent causal assumptions, even if statistical models are usually adjusted by "control" variables to avoid confounding biases—a clearly causal consideration. The Results section is most of the time free of causal language as well (Rohrer, 2023). For instance, a variable, X, is said to be simply *associated*, or *correlated*, with the outcome of interest, Y. But the discussions and conclusions often turn to explicitly causal vocabulary again: the authors discuss causal explanations for the observed pattern of association, as well as their implications.

In the end, statistical models from which causal evidence is assessed have generally not been logically linked to licensing causal assumptions (Pearl, 2001; Pearl & Bareinboim, 2014). For this reason, it can be hard to evaluate how strong the causal evidence actually is. It may be difficult to know what exactly is the quantity a statistical model is supposed to estimate (what is the causal estimand), why is a statistical model adjusted by certain variables and not others (under which assumptions is the effect identifiable); or how to interpret the statistical parameters in terms of the causal structure of the study system (Lundberg et al., 2021; Pearl et al., 2016; Westreich & Greenland, 2013).

Instead, inappropriate, non-causal approaches are often used to decide which "control" variables to include in a statistical model when the aim of the analysis is causal inference. A core problem is the use of predictive techniques that are given causal interpretation, for instance by selecting the predictor variables of a multivariable regression based on Information Criterion (*e.g.*, AIC) or Cross-Validation procedures (Arif & MacNeil, 2022). Another approach is to select covariates based on their p-values, for example by dropping non-significant variables. Researchers may also be advised to include all predictors that are assumed to affect the response variable. These approaches are *insufficient* to licence a statistical model designed to estimate a causal effect, because their logic has nothing to do with control of confounding, but other goals like forecasting or error control under random treatment assignment. In fact, well-meaning use of covariates may introduce biases—*e.g.*, collider bias or posttreatment bias (Byrnes & Dee, 2024; Greenland, 2022; McElreath, 2020; Pearl, 2009; Pearl et al., 2016; Pearl & Mackenzie, 2018; Rohrer, 2018).

We propose an alternative to the approaches mentioned above, by integrating tools from the field of formal Causal Inference with Bayesian statistical modelling. Causal inference in experiments depends upon the logic of random treatment assignment for modifying a causal model of the system under study. Causal inference in observational settings depends upon the same logic, in the absence of random treatment assignment. Both settings, experiments and observation, require in principle an explicit model of the data-generating process, a causal model. The causal model, com-

Licensing causal assumptions
Assumptions expressed in a formal language (*e.g.*, Directed Acyclic Graph, mathematical equation) describing causal relationships among variables. These assumptions are *licensing* in that they provide conditions under which the causal interpretation of statistical estimates is justified from first principles.

Estimand
Target quantity for a given data analysis, defined outside of a statistical model. For instance, the total causal effect of a variable X on another variable Y in population Z.

Identifiability
A causal estimand is identifiable if it can be theoretically computed using observed data. For instance, suppose that X causally affects Y, and that both X and Y are caused by an unobserved variable U. Then, the effect of X on Y is not identifiable, because of the confounder U.

Statistical parameter
Unobserved variable, whose value is estimated using a posterior probability distribution. Note that statistical parameters only describe associations, but contain no causal information.

bined with one or more questions, can transparently and objectively justify one or more statistical models.

In our case, what is needed is a model of the causal factors driving the observed social interaction network, as well as a transparent workflow that links it to hypotheses and data analysis. We describe this workflow in three parts. First, we show how to represent hypotheses about social networks as formal causal models. For the sake of generality and communication, we use Directed Acyclic Graphs, or DAGs (Pearl, 2001, 2009; Pearl et al., 2016; Pearl & Mackenzie, 2018), to describe the causal structures underlying animal social interaction networks. Second, we show how to explore the empirical implications of the DAG by building Structural Causal Models (SCM), which can be used for synthetic data simulation. Third, we combine these two types of causal models with matched multilevel extensions of the *Social Relations Model*, a generative statistical model (Kenny & La Voie, 1984; McElreath, 2020; Redhead, McElreath, et al., 2023; Ross et al., 2023; Snijders & Kenny, 1999; Van Duijn et al., 2004), and show that it can recover structural parameters from the simulations. This provides the grounds for the development and justification of an empirical workflow in which statistical models have been developed and tested transparently under specific causal assumptions.

Directed Acyclic Graphs
Graphical causal model showing which variables are assumed to affect each other in a given study system. DAGs only encode qualitative knowledge.

Structural Causal Model
Type of causal model where the functional relations among variables can be specified with quantitative knowledge. For instance, certain causal effects (corresponding to arrows in the DAG) can be defined as strong, weak, linear, non-linear, etc.

Our aim is not to provide a "one size fits all" model. Rather, the framework emphasises on the basic causal structure of animal social interaction networks, and on how they can be estimated with statistical models. Accordingly, empiricists interested in studying the drivers of social network structure can build upon our models, by tailoring them to their specific questions and study system.

Generative statistical model
Statistical model that can be used to simulate data, as they are implied by the model's assumptions.

2. Mechanics of the framework

We build the framework step by step, through four simulation studies that incrementally present its elements in detail.

Simulation study
Small, synthetic world created—in our case—to understand how a statistical model behaves when confronted with a *known* structural causal model.

In simulation 1, we introduce the basic elements of our causal model. They define the link between the observed network (third level of abstraction) and the interaction network (second level), and specify assumptions about the factors shaping the interaction network. We encode—*i.e.* translate to a formal language—these assumptions using two types of causal models, a DAG and a SCM. We start by assuming that the structuring features are random. We then introduce the Social Relations Model, and show it can recover the parameters of the SCM.

Next, we show how individual- (simulation 2) and dyad-level features (simulation 3) can affect the interaction network. We highlight that well-specified statistical models can accurately recover the causal effects encoded in the simulations. We also show how statistical models that are not adjusted by the structuring feature are affected by the causal effects, through their varying-effects structure, and discuss the interpretation of these effects. The aim of simulations 1-3 is not to be realistic. Instead, they reveal the internal *mechanics* of our framework.

Structural parameters
Parameters of a structural causal model. Note that structural parameters do contain causal information.

In simulation 4, we turn to a more realistic scenario. We build an estimator for the

effect of genetic relatedness on affiliation in females of a population of Assamese macaques, where this effect is confounded by dominance rank. We first validate our statistical model using synthetic data, and then fit it to empirical observations. Throughout the simulations, we show how our framework naturally addresses the issues mentioned above (problems I to IV).

2.1. Simulation study 1: random structuring features

2.1.1. Introduction to Directed Acyclic Graphs

DAGs are a simple and powerful tool to describe how variables in a system causally affect one another (Pearl, 2009; Pearl et al., 2016). The mathematical and applied literature on DAGs is large. We give here only a necessary conceptual introduction. DAGs are simple, but they are sufficient to describe the fundamental obstacles to, and derive solutions for, causal inference—including the design of experiments and observational studies.

DAGs are *graphs*. They are composed of nodes—here, representing variables—and directed edges, encoding causal relationships among these variables. For instance, the following DAG:

$$X \rightarrow Y$$

can be read as: "X influences Y". There is no universal definition of what "influences" means in science, but in a DAG it means that an intervention on X changes the distribution of Y. In contrast, an intervention on Y would not change X, because of the direction of the arrow. It is important to note that DAGs are qualitative: they do not say anything about the functional form of the causal effects. They simply posit the presence and absence of causal relationships among a set of variables.

A *path* is a structure formed by two or more adjacent nodes (*i.e.* nodes directly connected by an arrow). Paths connecting nodes can be causal, as in:

$$X \rightarrow Z \rightarrow Y,$$

Where, X transmits a causal effect to Y, through a *mediator*, Z. This specific structure (*i.e.*, a directed path containing three variables) is called a *chain*. Paths can also be non-causal, as in:

$$X \leftarrow Z \rightarrow Y.$$

This particular path, where a variable Z is a common cause of two other variables X and Y, is called a *fork*.

Although causal effects (consequences of interventions) only flow through causal paths, statistical association can arise through non-causal paths too. This is a fundamental reason that statistical association does not imply causal relationships. For instance, X and Y are associated in a fork, because of their common cause Z. Their association arises from a common cause, not the influence of one on the other.

A graph is *acyclic* if it is impossible to start at any node, and return to it by following only directed paths, or in other words: a node never influences itself. If one were to model feedback loops, they would need to represent time explicitly on a DAG (*e.g.*, $X_t \rightarrow Y_t \rightarrow X_{t+1}$; see Runge et al., 2023). What makes DAGs *directed* is that causal arrows can only be single-headed. Double-headed arrows can however be used as shortcuts, to indicate that an unspecified open path connects two variables. For instance, if an unobserved variable U affects both X and Y, it can be represented as:

$$X \longleftrightarrow Y, \text{ which is equivalent to: } X \leftarrow U \rightarrow Y.$$

For a general introduction to causal graphs, see Pearl et al. (2016) and McElreath (2020).

2.1.2. Directed Acyclic Graph for simulation study 1

In Table 1, we introduce the core elements of our causal framework. It links observed social interactions y in a dyad to the unobserved true rate of interaction m in that dyad, which is caused by features of the individuals, γ and ρ, and of the dyad τ. Each individual is assigned a general tendency to perform a particular behaviour, across all dyads. We call this *giving*, and represent it with the parameter γ ("gamma"). Each individual is also assigned a general tendency to receive a behaviour, across dyads. We call this parameter *receiving*, and represents it with ρ ("rho"). In addition to these general tendencies to give and receive, each individual can have a specific *tie* with each other individual. We denote these ties with τ ("tau"). Ties, like m and y are directional—their values in the two directions (from a to b, and from b to a) are not necessarily the same.

We illustrate how these parameters affect one another in Figure 2a. This DAG embodies assumptions about *which factors shape each observed edge* $y_{[a,b]}$. It is the simplest data-generating process in our framework: a system where γ, ρ, and τ share no common influences. More formally, we can say that they are only affected by exogenous, independent variables. The most common way to represent this scenario on a DAG is to draw no arrow pointing into these variables. Thus, the propensity of an individual, a, to give interactions to others ($\gamma_{[a]}$), its propensity to receive interactions from others ($\rho_{[a]}$), and the propensity it has to give interactions specifically to any other individual, b, in its group ($\tau_{[a,b]}$), are all independent. These three parameters affect the true rate of interactions from a to b: $m_{[a,b]}$. Then, $m_{[a,b]}$ affects the observed number of interactions, $y_{[a,b]}$. This slightly contrasts with Figure 1, where the link between y and m was not explicitly causal.[2]

Exogenous variable
Variables whose causes are not explicitly modeled.

We now have a graphical representation of our causal model for simulation 1. However, what it exactly means for a variable to affect the "tendency" of an individual to give, or receive, interactions, as well as the exact nature of the above-mentioned

[2] The arrow from m to y might seem counter-intuitive. Recall that, as we have defined it earlier, a variable has a causal effect on another if a (hypothetical) intervention on the former results in a change in the latter. As such, because a hypothetical intervention on the true rate m would, indeed, result in a change in the observed y, we draw an arrow from m to y.

Table 1 – *Core elements of the causal framework.* The variables in this table allow researchers to encode assumptions about the effects of individual- and dyad-level features on social interactions (see section S1.9.1 for group-level effects). The causal relationships among them are defined in Directed Acyclic Graphs (*e.g.*, Figure 2) and in Structural Causal Models (*e.g.*, equations 1.1.1-1.1.5). We do not define the statistical parameters G, R, and T in this table, because their meanings differ across models.

Parmeter	*Indexed by*	*Definition/Function*	*Example*	*Notes*
y	Directed dyad, *i.e.*, one value per dyad, and per direction: $y_{[a,b]}$ (direction 1) $y_{[b,a]}$ (direction 2)	Number of *observed* social interactions given from an individual, a, to another individual, b, over a sampling period.	N. of observed grooming bouts from a to b over 12 hours of combined focal follows of a or b.	y has the same meaning as in Figure 1, except that here, the interactions are directed.
m	Directed dyad: $m_{[a,b]}$ (direction 1) $m_{[b,a]}$ (direction 2)	True rate/number of social interactions given from an individual, a, to another individual, b, over a sampling period.	Average number of grooming bouts given by a to b per period of 12 hours.	m is unobserved (see Figure 1-2). It has the same meaning as in Figure 1, except that here, the interactions are directed.
γ (*giving*)	Individual: $\gamma_{[a]}$	Encodes assumptions about the effect of individual-level features on the general tendency of an individual to *give* social interactions to others in the group.	Given that $X_{[a]}$ represents the age of an individual a, $X_{[a]} \rightarrow \gamma_{[a]}$ can be read as: "*age affects how many social interactions an individual, a, tends to give to others*".	γ can, for example, encode the assumption that older individuals spend less time engaging in social interactions (*e.g.*, grooming), because they may be constrained by their low foraging efficiency.
ρ (*receiving*)	Individual: $\rho_{[a]}$	Encodes assumptions about the effect of individual-level features on the general tendency of an individual to *receive* social interactions from others in the group.	Given that $X_{[a]}$ represents the age of an individual a, $X_{[a]} \rightarrow \rho_{[a]}$ can be read as: "*age affects how many social interactions an individual, a, tends to receive from others*".	ρ may, for instance, describe that individuals of a certain age are more attractive interaction partners.
τ (*tie*)	Directed dyad: $\tau_{[a,b]}$ (direction 1) $\tau_{[b,a]}$ (direction 2)	Encodes assumptions about the effect of dyad-level features on the tendency of an individual, a, to give social interactions with a specific other individual, b.	Given that $X_{\|a,b\|}$ is the genetic relatedness between a and b, $X_{\|a,b\|} \rightarrow \tau_{[a,b]}$ encodes the following assumption: "The relatedness between a and b affects how many interactions a gives to b".	We index *symmetric* dyadic variables, like genetic relatedness, using vertical bars ($X_{\|a,b\|} = X_{\|b,a\|}$). We index *directed* dyadic variables using square brackets ($\tau_{[a,b]}$ is not necessarily equal to $\tau_{[b,a]}$).

"exogenous variation", depends on the mathematical function that we assign to it. This is why, in the next section, we turn the qualitative assumptions of the DAG into quantitative assumptions.

2.1.3. Structural Causal Model

Structural Causal Models (SCM) are a type of causal model where the functional relationship among variables is defined by a mathematical function (Pearl, 2009). It means that they contain assumptions about whether the effects among variables—corresponding to arrows on a causal graph—are positive, negative, strong, linear, etc. SCMs are closely linked to causal graphs, such that there exists a causal graph for each SCM. As SCMs encode quantitative causal assumptions, they can be used as *generative models*, to simulate observations.

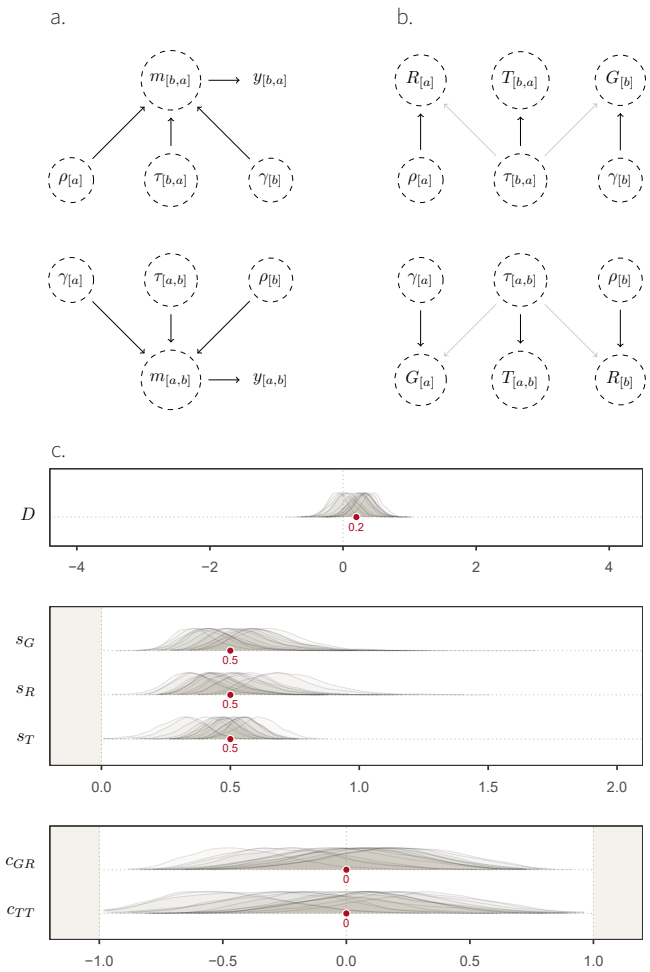

Figure 2 – *Simulation study 1.* **a.** DAG showing the core elements of our causal framework, and how they causally affect one another, for two individuals, a and b. No arrows enter the structuring parameters γ, ρ and τ. This means that exogenous, independent, noise is the only factor affecting: (i) how many social interactions an individual, a, generally gives to others ($\gamma_{[a]}$); (ii) how many interactions a generally receives from others ($\rho_{[a]}$); and (iii) how many interactions a specifically gives to another individual, b ($\tau_{[a,b]}$). These three structural parameters affect the true rate of interactions given from a to b ($m_{[a,b]}$), which causes the number of observed interactions ($y_{[a,b]}$). We show the two directions—*i.e.* from a to b, and from b to a, and represent unobserved variables with a dashed circle. **b.** Mapping between structural parameters (Greek letters) and statistical parameters (Latin letters) of the non-adjusted Social Relations Model (equations 1.2.1-1.2.4). The transparent arrows between the structural dyadic parameter τ and the statistical individual parameters G and R indicate that these effects are possible (*e.g.* simulation 3), but do not exist here. **c.** Marginal posterior distributions, over a range of parameter values (*x-axis*) for the fixed-effects (*y-axis*) of posterior model 1. Their respective target values are shown as red dots. These fixed effects—except for D which does not appear on the DAG—capture the patterns of (co-)variation of G, R, and T, shown in panel **b**.

Practically, we used SCMs for data simulations that "obey" the causal relations represented on the DAGs, and implemented them in R, version 4.2.1 (R Core Team, 2022). For each arrow on the DAGs, we defined a function, with arbitrary parameter values. We did so with two main goals. First, each SCM serves as a simple *example* of data-generating process encoded by its DAG. These simulations can be more tangible than the abstract, non-parametric, knowledge encoded in causal graphs. Second, and perhaps most importantly, SCMs can be used to derive and validate *statistical models*. We will focus on the former goal in this section, before turning to the latter, in the next section.

Below, we provide a mathematical description of the SCM for simulation study 1—hereafter *SCM 1*. We start by defining f_γ, the function determining γ for each individual :

The function determining gamma is defined as follows:

...is drawn from the following distribution:

$$f_\gamma : \quad \gamma_{[a]} \sim \text{Normal}(0, 0.5) \tag{1.1.1}$$

The value of gamma, for each individual, a...

Mean

Standard deviation

Where $a \in \{1, ..., N\}$. This notation with the $\in$ symbol implies that a represents an arbitrary individual index that can take any value from 1 to N, where N is the number of unique individuals in the sample. Accordingly, equation 1.1.1 should be interpreted as follows: $\gamma_{[1]}, \gamma_{[2]}, \ldots, \gamma_{[N]}$ are distributed as a normal (or Gaussian) distribution with mean 0 and standard deviation 0.5.[3]

$$f_\rho : \quad \rho_{[a]} \sim \text{Normal}(0, 0.5) \tag{1.1.2}$$

$$f_\tau : \quad \tau_{[a,b]} \sim \text{Normal}(0, 0.5). \tag{1.1.3}$$

b, like a, is an arbitrary individual index (where $b \in \{1, ..., N\}$ and $a \neq b$). Thus, $\tau_{[a,b]}$ represents the value of τ for any combination of a and b:

$$\tau_{[a,b]} \in \{\tau_{[1,2]}, \tau_{[2,1]}, \ldots, \tau_{[N-1,N]}, \tau_{[N,N-1]}\}.$$

In equations 1.1.1-1.1.3, we have defined what it meant for γ, ρ, and τ to be only affected by exogenous, independent, variables. We can now turn to define how these structural parameters affect m, the true interaction rate:

$$f_m : \quad m_{[a,b]} = \exp(0.2 + \gamma_{[a]} + \rho_{[b]} + \tau_{[a,b]}). \tag{1.1.4}$$

The intercept value 0.2 is the logarithm of how many interactions individuals give to, and receive from, others in the network on average. Here, this average rate is $\exp(0.2) \simeq 1.2$ interactions. The exponential function of f_m, in combination with

[3]Note that this equation is equivalent to:

$$f_\gamma : \quad \gamma_{[a]} = U^\gamma_{[a]} \quad \text{where} \quad U^\gamma_{[a]} \sim \text{Normal}(0, 0.5),$$

Which is more in line with the usual SCM nomenclature (see Pearl, 2009).

the additive (or "linear") combination in terms, ensures that the interaction rate m is always positive. It also has the following consequence: the effects of γ, ρ and τ on m are now ***multiplicative***. It means, for instance, that the effect of $\tau_{[a,b]}$ on $m_{[a,b]}$ is larger if $\gamma_{[a]}$ and $\rho_{[b]}$ have large values themselves. We will come back to this issue, because it is a fundamental and unavoidable aspect of modeling behavioural count data: causal effects are not simple and additive on the outcome scale, even though we model them as additive on the link scale.

The rate of behaviour $m_{[a,b]}$ is connected to the observable count $y_{[a,b]}$ through a stochastic process. The simplest choice is a Poisson process:

$$f_y : \quad y_{[a,b]} \sim \text{Poisson}(m_{[a,b]}). \tag{1.1.5}$$

Equation 1.1.5 defines the relationship between the true interaction rate (second level of abstraction) and the observed number of interactions (third level), represented in Figure 1. The number of observed interactions $y_{[a,b]}$ is drawn from a Poisson distribution, characterised by one parameter: its average rate of interaction, $m_{[a,b]}$, specific to each directed dyad.

In summary, SCM 1 mathematically defines *one* simple data-generating process compatible with the DAG in Figure 2a. The generated data is represented in Figure S1.1.1. Note that the number of individuals in the network, the value of the intercept in f_m, or the distributions we have chosen (normal, Poisson), are flexible. We might have, for instance, decided to model binary social events (*e.g.*, has a been seen grooming b during a focal-animal sampling protocol?) instead of unbounded counts, which would make a Bernoulli distribution rather than Poisson appropriate. The key idea in any case is that the target of inference, the rate m and its components, is not observed but must be modelled as filtered by a stochastic sampling process. This stands in contrast to aggregated indices of sociality which produce point estimates of rates, without properly characterizing uncertainty, before decomposing them as functions of explanatory variables.

In the next section, using the simple simulation defined above, we build a statistical model that aims to *uncover*—or, in this case, because we are dealing with synthetic data, *recover*—the parameters of the data-generating process.

2.1.4. The Social Relations Model

A non-parametric causal model, like a DAG, defines a qualitative hypothesis about how observations arise. A SCM makes this hypothesis into a data-generating algorithm that can be used for planning, understanding, and validating projects. Using data, whether simulated or real, requires a statistical model. Here, we explain how a statistical network model relates to a DAG and SCM, and how it can be developed and validated together with the causal models.

Statistical models and causal models are similar but distinct. Causal models contain directionality—they define the consequences of interventions—while statistical models do not. Statistical models contain additional elements that are useful for estimation. Importantly, they can include hierarchical distributions that function

to regularise estimates. These elements exist in both Bayesian and non-Bayesian approaches, because they produce better, more efficient estimates (Feller & Gelman, 2015).

Regularised estimates
Statistical estimates that tend to capture features of the target population (so called *regular* features), and not features of the specific sample (*irregular* features).

The statistical models that we use in our framework are extensions of the *"Social Relations Model"* (Kenny & La Voie, 1984), initially developed in the social sciences. Compared to Kenny & La Voie's initial model, ours have a multilevel varying-effect structure, can include covariates (Snijders & Kenny, 1999), and have non-Gaussian likelihoods (Van Duijn et al., 2004). They are further implemented in a Bayesian framework (see McElreath, 2020, Chapter 14) and can incorporate flexible additive effects, such as a stochastic blockmodel, *i.e.* a model capable of estimating causal effects of observable categorical variables, like sex, on network structure (see Redhead, McElreath, et al., 2023; Ross et al., 2023). We wrote the models presented in this manuscript in the *Stan* probabilistic programming language (Carpenter et al., 2017), and ran them through *R*, using *CmdStanR* (Češnovar et al., 2021)[4].

Multilevel model
Also sometimes called "hierarchical" or "mixed-effect" models, multilevel models learn about the value of certain clusters using what they have learnt in other clusters, when these clusters are part of a so-called "varying-effect" (or "random effect") structure. For instance, to estimate an individual $\boldsymbol{a}$'s propensity to give social interactions to others, a multilevel model learns, naturally, from the interactions given from $\boldsymbol{a}$ to others, but also from the other individuals' (*e.g.* $\boldsymbol{b}$, $\boldsymbol{c}$, $\boldsymbol{d}$) propensities to give interactions to others. It does so through a process called *partial-pooling*.

In practice, we use statistical models as *estimators*. Given (i) an *estimand*, and causal assumptions encoded in (ii) a DAG and (iii) a SCM, we build a statistical model whose goal is to *estimate* the estimand. In this section, we outline how such a model can learn, by fitting it to known data: the data generated through the simulation above (SCM 1). We feed the statistical model with the number of observed interactions $y_{[a,b]}$ for each directed dyad in the network. Given these data, the model has to recover the unobserved rates $m_{[a,b]}$, as well as other structural parameters. In this first simulation, we do not focus on one particular estimand, but instead describe how the statistical model generally learns about the social network structure.

Below, we mathematically define the statistical model of simulation study 1 (hereafter statistical model 1). We start by describing the observed data y—here, one data point per directed dyad—using a conservative choice of probability distribution for unbounded count data: the Poisson distribution (Blitzstein & Hwang, 2019; McElreath, 2020).

$$y_{[a,b]} \sim \text{Poisson}(m_{[a,b]}) \tag{1.2.1}$$

This line can be read as such: *"$y_{[a,b]}$ is described using a Poisson distribution, whose average rate is $m_{[a,b]}$"*. Here, $y_{[a,b]}$ is known, and $m_{[a,b]}$ is not: it has to be estimated. Recall that as opposed to equation 1.1.5, we are, here, describing a *statistical* model, which *learns* the value of unobserved variables, like m. For this reason, the model will return a posterior probability distribution for each parameter $m_{[a,b]}$. Statistical models and SCMs can be visually differentiated in this manuscript, because each structural equation is defined by a function f. $m_{[a,b]}$ is further defined in terms of several parameters:

$$m_{[a,b]} = \exp(D + G_{[a]} + R_{[b]} + T_{[a,b]}) \tag{1.2.2}$$

[4]The *STRAND* (Redhead, McElreath, et al., 2023; Ross et al., 2023) and *Rethinking* (McElreath, 2020) R packages can also express the Social Relations Model varying effects. Alternatively, comparable multilevel structures could be encoded in *BISoN* (Hart et al., 2023), *bamoso* (Neumann & Fischer, 2023a), or *brms* (Bürkner, 2017).

As in the SCM, the exp (exponential) function ensures that m remains positive; it corresponds to the so-called inverse-link function. D, G, R, and T are all unobserved.[5] Like m, they have to be estimated, and will be assigned a posterior probability distribution.

D, the intercept, is fixed across all dyads. It is equivalent to the intercept of 0.2 in SCM 1, which it should recover. In this simulation study, G, R, and T can be thought of as the statistical counterpart of the structural parameters γ, ρ, and τ (Figure 2b). $G_{[a]}$ captures the general propensity (*i.e.* the average rate, on the link scale), of an individual, a, to give social interactions to other individuals in the network. It is only affected by $\gamma_{[a]}$, and hence, $G_{[1]}$ should, for instance, recover $\gamma_{[1]}$. $R_{[b]}$, on the other hand, captures the propensity of an individual, b, to generally receive social interactions from others. It is, analogously, affected by $\rho_{[b]}$ only. Finally, $T_{[a,b]}$ represents the residual tendency of a to give interactions to b, conditional on $G_{[a]}$ and $R_{[b]}$. Here, it should recover the dyad-specific tendency effect $\tau_{[a,b]}$. In later simulation studies (*e.g.*, in simulation study 3), we will see cases where the statistical and structural parameters are not equivalent.

To efficiently estimate the components of the rate $m_{[a,b]}$, it is useful to employ hierarchical distributions that relate the components to one another and allow them to share information through those statistical relationships. In practice, this means that $G_{[a]}$ and $R_{[a]}$ are described as varying-effects in a Multivariate Normal distribution.

The average individual giving and receiving rates...

Variance in the dist. of G

Covariance between G-R

$$\begin{pmatrix} G_{[a]} \\ R_{[a]} \end{pmatrix} \sim \text{MVNormal}\left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} s_G^2 & c_{GR}s_Gs_R \\ c_{GR}s_Gs_R & s_R^2 \end{pmatrix}\right] \tag{1.2.3}$$

...are described using a multivariate normal dist., whose parameters are:

Mean

Variance in the dist. of R

The distribution, in this case, has two dimensions: G and R. Its mean of $\{0, 0\}$ implies that $G_{[a]}$ and $R_{[a]}$ are *deviations* from the average interaction rate in the network, D. s_G and s_R describe the amount of variation in $G_{[a]}$ and $R_{[a]}$, respectively. If, for instance, individuals vary a lot in how many interactions they give to others, but do not vary much in how many interactions they receive from others, then s_G will be large, and s_R will be small. The pattern of covariation between these two dimensions is then captured by c_{GR}. Suppose a network where individuals who give a lot of interactions also receive a lot of interactions. This situation would result in a positive estimate for c_{GR}.

[5] The individual-level parameter $G_{[a]}$ has received several names, across disciplines and contexts: *e.g.* generalised giving (McElreath, 2020), actor effect (Back & Kenny, 2010), productivity or sociability (Van Duijn et al., 2004), expansiveness (Rawlings et al., 2023), focal/sender/nominator effect (Ross et al., 2023), giver effect (Redhead, Dalla Ragione, et al., 2023). $R_{[a]}$ has been called generalised receiving (McElreath, 2020), partner effect (Back & Kenny, 2010), attractiveness or popularity (Van Duijn et al., 2004), target/receiver/nominee effect (Ross et al., 2023), alter effect (Redhead, Dalla Ragione, et al., 2023).

Similarly the directed ties in the network are modeled hierarchically, using parameters for structure within and among dyads:

$$\begin{pmatrix} T_{[a,b]} \\ T_{[b,a]} \end{pmatrix} \sim \text{MVNormal}\left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} s_T^2 & c_{TT}s_T^2 \\ c_{TT}s_T^2 & s_T^2 \end{pmatrix}\right]. \tag{1.2.4}$$

Equation 1.2.4 is similar to equation 1.2.3, with one exception. There is only one variance parameter in the variance-covariance matrix: s_T^2. Recall that a and b are only arbitrary labels. Thus, there is no reason why the variation in $T_{[a,b]}$ should be different from that of $T_{[b,a]}$. c_{TT} hence describes the association between the number of interactions that a gives to b, and those that b gives to a, conditional on their respective averages $G_{[a]}$ and $R_{[b]}$. We will see further below how this parameter may be interpreted biologically, in light of causal models. We provide details about the prior probability distributions, as well as the exact parametrisation of our statistical model, in the Supplementary section S1.1.2-S1.1.3.

In this section, we have defined the basic architecture of the Social Relations Model (which we sometimes refer to as *non-adjusted* Social Relations Model). A feature of this model we want to highlight is that it describes the true rate of interactions m as a *statistical parameter*, for which a probability distribution is computed using Bayesian updating. This represents a principled solution to deal with the uncertainty and noise inherent to sampling social interactions (problem I). It notably contrasts with data-aggregation methods that we mentioned above, which treat data y (third level of abstraction) as the *known*, true interaction rates (second level). Moreover, the varying effects of the Social Relations Model explicitly describe several patterns of (co)variations—*i.e.* structural "dependencies"— present in social networks, like those mentioned in the introduction (problem IV). This results in more accurate and efficient estimators (Feller & Gelman, 2015; McElreath, 2020).

2.1.5. Posterior model

In this section, we describe *posterior model 1*: the model fit obtained by feeding the data generated by SCM 1 into statistical model 1. We ran 15 iterations of the SCM, and fitted the data from each iteration into the statistical model. Thus, we obtained 15 marginal posterior distributions per "fixed-effect", or population level parameters (Figure 2c). We show each parameter's posterior distributions with its *target value*, *i.e.* the value it should recover, given that we know the generative model. The pattern one would expect from an estimator that accurately recovers structural parameters is the following: over several iterations, the high-density regions of the marginal distributions should overlap with the target value. Accordingly, we see that our estimator works, as the posterior distributions overlap well with the target values. This kind of model checking can be formalized as *simulation based calibration* (Talts et al., 2020), but even a few simulations are often sufficient to spot problems in conceptualising and programming the SCM, the statistical model, or both. The general idea is: until we are prepared to interpret estimates from synthetic data, we are unprepared to interpret estimates from real data.

Marginal posterior distribution
Posterior probability of a parameter regardless of (*i.e.* unconditional on) the value of the other parameters.

The mapping between the structural and statistical parameters is, in the case of SCM 1, rather straightforward. Starting with D, we expect the statistical parameter to re-

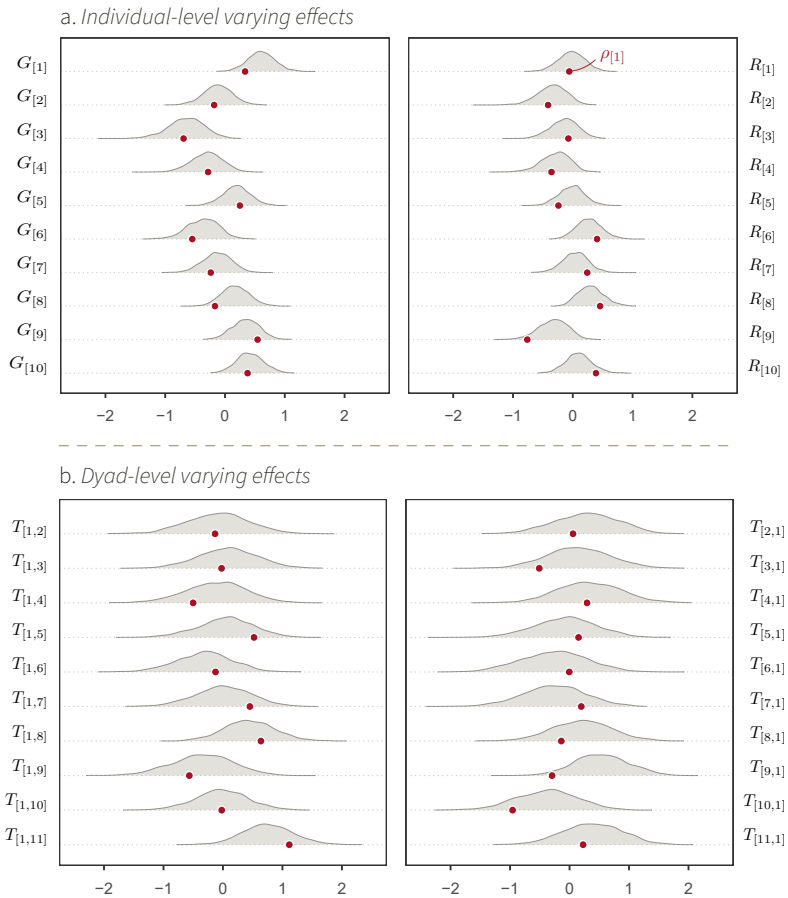

Figure 3 – *Marginal posterior distributions of a set of statistical model 1's varying-effects.* The probability density of these parameters (*y-axis*) is shown for a range of parameter values (*x-axis*), for one iteration of the simulation. The red dots indicate the parameters' target values. For instance, $\gamma_{[1]}$ is the target value of $G_{[1]}$, whose posterior probability distribution is represented as a density curve on top of the red dot.

cover the intercept value of 0.2 (equation 1.1.4). G, R, and T are only caused by, respectively, γ, ρ, and τ (Figure 2b and Figure 3). Furthermore, s_G, s_R, and s_T, should capture the variation in γ, ρ, and τ, respectively, and thus, be equal to 0.5 (see equations 1.1.1-1.1.3).

The correlation coefficient between G and R, c_{GR}, does not have a counterpart in the SCM. However, looking at the DAG in Figure 2b, we can see that there is no path connecting $G_{[a]}$ and $R_{[a]}$ (or $G_{[b]}$ and $R_{[b]}$). Hence, there should be no association between them ($c_{GR} = 0$). The same reasoning applies to c_{TT}: there is no path between $T_{[a,b]}$ and $T_{[b,a]}$, and thus, we expect $c_{TT} = 0$. The Multivariate Normal

posterior distribution describing $G_{[a]}$ and $R_{[a]}$, as well as the posterior distributions for m, are shown in supplementary Figure S3-S4.

2.1.6. Conclusions

In simulation 1, we have modelled a simple data-generating process, which describes the factors shaping the observed edges of a social network. We assumed that the general tendencies of individuals to interact with conspecifics, as well as their propensity to engage in social interactions with specific partners, were random. We started by describing the abstract causal structure of this system, using a DAG. Next, we implemented this DAG as a SCM: a type of data simulation, where we specified parametric features of the causal system. We then fed the synthetic data into a statistical model, and showed that the statistical model could recover the structural parameters of the SCM. In the following section, we will turn to the following question: how can we use the models introduced in this section, to estimate the causal effect of individual-level features on their general propensity to engage in social interactions?

2.2. Simulation study 2: individual-level features

Here, we build upon the tools presented in the previous section to add features of individuals that may be targets of causal inference.

2.2.1. Causal model

Suppose that we wish to study the effect of an individual-level phenotypic variable $X_{[a]}$ (*e.g.*, age) on the overall tendency of individuals to give and receive social interactions (*e.g.*, grooming), in a specific social system (*e.g.*, a group of monkeys). We might, for instance, hypothesise that increasing age causes individual monkeys to generally disengage from their social group, and thus, to give and receive fewer grooming interactions (Sadoughi et al., 2024; Siracusa et al., 2022). Such a system may be represented as a DAG where $X_{[a]}$ has a direct effect on $\gamma_{[a]}$ and $\rho_{[a]}$ (see Figure 4a).

We wish to build an estimator that can estimate this effect. To do so, we specify quantitative causal assumptions using SCM 2. This model is identical to SCM 1, except with regard to the variable $X_{[a]}$ and its effects:

$$f_X : \quad X_{[a]} \sim \text{Normal}(0, 1). \tag{2.1.1}$$

$X_{[a]}$, an individual-level feature, is here defined as a standardised Gaussian distribution, with mean 0 and SD 1. Then,

$$f_\gamma : \quad \gamma_{[a]} \sim \text{Normal}(-0.7^* \cdot X_{[a]}, 0.5) \tag{2.1.2}$$

$$f_\rho : \quad \rho_{[a]} \sim \text{Normal}(-0.7^\ddagger \cdot X_{[a]}, 0.5). \tag{2.1.3}$$

These two equations encode the key aspects of simulation study 2: the causal effects of $X_{[a]}$ on $\gamma_{[a]}$ and $\rho_{[a]}$—and thus, on the rate $m_{[a,b]}$. These effects are, respectively, our two *estimands*; we mark them with $*$ and ‡, such that they can be tracked more

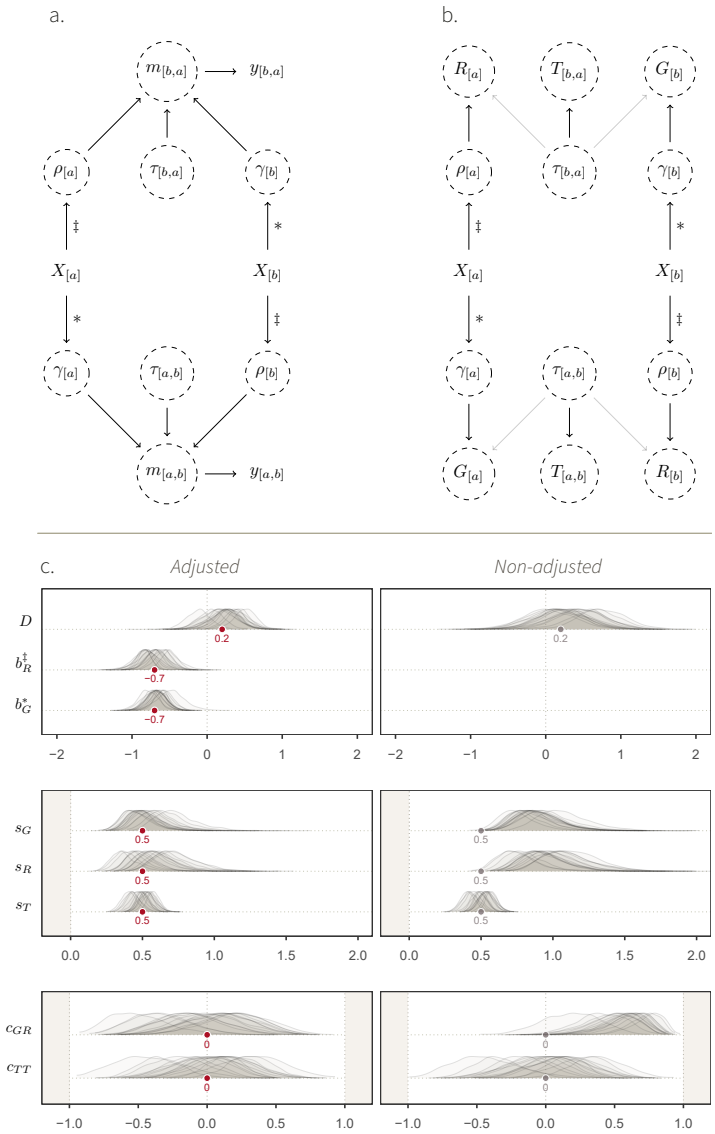

Figure 4 – *Simulation study 2.* **a.** DAG describing a causal system where an individual-level phenotypic trait $X_{[a]}$, like age, affects their overall tendency to give interactions, $\gamma_{[a]}$, and their general tendency to receive interactions, $\rho_{[a]}$. The $*$ and ‡ symbols mark our estimands: the causal effects we wish to estimate. The rest of the DAG is similar to Figure 2a. **b.** Mapping between structural parameters (SCM 2) and statistical parameters. As for Figure 2b, the transparent arrows indicate that these effects are possible (*e.g.*, simulation 3), but do not exist in SCM 2. **c.** Fixed-effect estimates, for two statistical models fitted to the data generated with SCM 2. *Left*: fixed effects of the social relations model adjusted by $X_{[a]}$ (statistical model 2). *Right*: fixed effects of the non-adjusted social relations model (statistical model 1). The target values of the well-adjusted model are shown in grey. Deviations from them allow us to understand how $X_{[a]}$ affects the varying effects of the model, when $X_{[a]}$ is not adjusted for.

easily. f_γ and f_ρ imply that $\gamma_{[a]}$ and $\rho_{[a]}$ are now sampled from distributions, whose means depend on the value of $X_{[a]}$. To illustrate how this effect works, imagine two individuals, a and b, with $X_{[a]} = 1$ (old individual) and $X_{[b]} = -2$ (young individual). Consequently, $\gamma_{[a]}$ will be drawn from a distribution with mean $-0.7 \cdot 1 = -0.7$, and $\gamma_{[b]}$, from a distribution with mean $-0.7 \cdot -2 = 1.4$. Thus, $\gamma_{[a]}$ will most likely be lower than $\gamma_{[b]}$, which is consistent with our hypothesis. Note that the effect of $X_{[a]}$ on $\gamma_{[a]}$ need not be the same as its effect on $\rho_{[a]}$.

Finally, f_τ, f_m and f_y, are the same as in SCM 1:

$$f_\tau : \quad \tau_{[a,b]} \sim \text{Normal}(0, 0.5) \tag{2.1.4}$$

$$f_m : \quad m_{[a,b]} = \exp(0.2 + \gamma_{[a]} + \rho_{[b]} + \tau_{[a,b]}) \tag{2.1.5}$$

$$f_y : \quad y_{[a,b]} \sim \text{Poisson}(m_{[a,b]}). \tag{2.1.6}$$

The resulting network of (synthetic) observations $y_{[a,b]}$ is shown in Figure S5.

2.2.2. Statistical model

Next, we describe *Statistical model 2*, an estimator for the causal effects of X on γ and ρ. This estimator maintains the basic architecture statistical model 1, but is slightly modified. Here, we define $\hat{G}_{[a]}$ and $\hat{R}_{[b]}$ as submodels of $m_{[a,b]}$, where $\hat{G}_{[a]}$ is stratified by the variables causing $\gamma_{[a]}$, and $\hat{R}_{[b]}$ is stratified by the variables causing $\rho_{[b]}$. Thus, both are stratified by X:

$$y_{[a,b]} \sim \text{Poisson}(m_{[a,b]}) \tag{2.2.1}$$

$$m_{[a,b]} = \exp(D + \hat{G}_{[a]} + \hat{R}_{[b]} + T_{[a,b]}) \tag{2.2.2}$$

$$\hat{G}_{[a]} = G_{[a]} + b^*_G \cdot X_{[a]} \tag{2.2.3}$$

$$\hat{R}_{[b]} = R_{[b]} + b^\ddagger_R \cdot X_{[b]}, \tag{2.2.4}$$

Where equations 2.2.2-2.2.4 are, together, mathematically equivalent to:

$$m_{[a,b]} = \exp(D + G_{[a]} + b^*_G \cdot X_{[a]} + R_{[b]} + b^\ddagger_R \cdot X_{[b]} + T_{[a,b]}).$$

We usually use the former notation, which is easier to read, and whose link with the non-adjusted Social Relations Model is more visible. Yet, these two sets of equations imply the same thing: b^*_G is a slope estimating the association between $X_{[a]}$ and $m_{[a,b]}$—whether directly, or through a sub-model $\hat{G}_{[a]}$—, conditional on the other parameters of the model. Similarly, $b^\ddagger_R$ estimates the conditional association between $X_{[b]}$ and $m_{[a,b]}$. On the DAG, b^*_G captures the path coefficient between $X_{[a]}$ and $m_{[a,b]}$; and $b^\ddagger_R$ the path between $X_{[b]}$ and $m_{[a,b]}$. Therefore, b^*_G and $b^\ddagger_R$ should respectively recover the effect of $X_{[a]}$ on $\gamma_{[a]}$ (marked by a $*$), and that of $X_{[b]}$ on $\rho_{[b]}$ (marked by a ‡), and take a value of -0.7.

Finally, $G_{[a]}$, $R_{[b]}$, and $T_{[a,b]}$ are varying effects. They are part of the exact same

multivariate adaptive priors described in statistical model 1:

$$\begin{pmatrix} G_{[a]} \\ R_{[a]} \end{pmatrix} \sim \text{MVNormal}\left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} s_G^2 & c_{GR}s_Gs_R \\ c_{GR}s_Gs_R & s_R^2 \end{pmatrix}\right] \tag{2.2.5}$$

$$\begin{pmatrix} T_{[a,b]} \\ T_{[b,a]} \end{pmatrix} \sim \text{MVNormal}\left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} s_T^2 & c_{TT}s_T^2 \\ c_{TT}s_T^2 & s_T^2 \end{pmatrix}\right]. \tag{2.2.6}$$

The full model with hyperpriors can be found in section S1.2.2. We also wish to describe how $X_{[a]}$ would impact the varying effects of the statistical model if $X_{[a]}$ was not observed. Thus, we also fitted the data generated with SCM 2 into statistical model 1.

2.2.3. Posterior model

Statistical model 2 successfully recovered the structural parameters of the simulation (Figure 4c, left panel). Most importantly, we accurately estimated our estimands: b_G^* recovered the true causal effects of $X_{[a]}$ on $\gamma_{[a]}$, and $b_R^\ddagger$, the effect of $X_{[a]}$ on $\rho_{[a]}$. This means that our estimator is capable of producing valid causal inference under the assumptions embodied by the DAG and the SCM.

On the right panel of Figure 4c, we observe important deviations between, on the one hand, the target values of the well-specified estimator, and, on the other hand, the posterior distributions of the non-adjusted statistical model. The posterior distributions of s_G and s_R—respectively quantifying the variation in $G_{[a]}$ and $R_{[a]}$—are now consistently higher than 0.5. This is because $G_{[a]}$ and $R_{[a]}$ are caused by $\gamma_{[a]}$ and $\rho_{[a]}$, which are themselves caused by *two* kinds of variables. They are affected by random noise, quantified by a Gaussian SD of 0.5, and by $X_{[a]}$ (see equation 2.1.2-2.1.3). Hence, s_G and s_R are larger than 0.5 if we don't stratify by $X_{[a]}$ (right panel), but are equal to 0.5 once we do condition on it (left panel).

Furthermore, the posterior distribution of c_{GR} is positive in the non-adjusted model (Figure 4; see Figure S6 for an additional visualisation). Recall that c_{GR} captures the association between $G_{[a]}$ and $R_{[a]}$. Looking at the DAG, we see that these parameters are connected by $X_{[a]}$. $X_{[a]}$ creates an open path,

$$G \leftarrow \gamma \leftarrow X \rightarrow \rho \rightarrow R,$$

Which lets the association flow between $G_{[a]}$ and $R_{[a]}$. Once we condition on $X_{[a]}$, however, we block this path, cancel the association between $G_{[a]}$ and $R_{[a]}$, and we return to a correlation coefficient $c_{GR} = 0$ (left panel). Finally, notice that the dyadic parameters s_T and c_{TT}, were unaffected by $X_{[a]}$. This can be deducted from the DAG, too, for there exist no open paths connecting $\gamma_{[a]}$ or $\rho_{[a]}$— carrying the effects of $X_{[a]}$—and $T_{[a,b]}$.

2.2.4. Conclusions

In simulation 2, we have modelled a data-generating process where an individual-level trait $X_{[a]}$ (*e.g.*, age) affected the general tendency of individuals to generally

give and receive social interactions. Following the same workflow as in simulation 1, we showed that a well-specified estimator—an adjusted version of the Social Relations Model—could accurately recover the causal effect of $X_{[a]}$ on social network structure. Next, we highlighted how the causal effect of $X_{[a]}$ impacted the varying-effects structure of the statistical model, when $X_{[a]}$ was not adjusted for. We refer the interested readers to the supplementary section S1.3.1-S1.3.4: there, we describe *simulation study 2'*, a variation of simulation study 2, where $X_{[a]}$ is a categorical variable (*e.g.*, sex). In the next section, we turn to the following question: how can we build a statistical model to estimate the effect of dyad-level features on the propensity of two individuals to socially interact with one another?

2.3. Simulation study 3: dyad-level features

Suppose that we wish to study how the genetic relatedness of individuals affects the way they interact with one another. We may hypothesise that individuals belonging to the same kin group engage in a greater number of affiliative interactions than non-kin. Furthermore, sampling effort may vary across individuals and dyads, for behavioural observation is conducted by following focal animals using standardised protocols. Individuals of some dyads might have been observed for several hours of behavioural sampling, while others might have been followed for very little time.

2.3.1. Causal model

We represent the causal structure of such a data-generating process on a DAG (Figure 5a), where we specify these causal relationships, as well as the link between individual- and dyad-level features. The combinations of two individuals' kin groups, $K_{[a]}$ and $K_{[b]}$, determines how genetically related they are to one another ($Re_{|a,b|}$), by definition. The sampling effort for each individual, $S_{[a]}$, in combination with the sampling effort for its partner, $S_{[b]}$, further determines the dyad-level sampling effort, $S_{|a,b|}$.[6] Note that there is no path between $Re_{|a,b|}$ and $m_{[a,b]}$ that passes through $S_{|a,b|}$. This means that $S_{|a,b|}$ does not *confound* the effect of $Re_{|a,b|}$. Yet, we still want to include it in our estimator (see next section), for doing so will increase the model's precision.

Model precision
Refers to how variable posterior estimates are from one another across data samples (*e.g.*, across iterations of a given SCM). More precise models are less variable across data samples.

Before building the estimator, we translate the assumptions of the DAG into *SCM 3*. We simulate 20 individual animals, belonging to 11 different kin groups: 10 individuals belong to one, large, kin group, and each of the 10 remaining individuals are alone in their kin group (Figure S9). For simplicity, we dichotomise genetic relatedness, such that:

$$f_{Re}: \quad Re_{|a,b|} = \begin{cases} 1 & \text{if } K_{[a]} = K_{[b]} \\ 0 & \text{if } K_{[a]} \neq K_{[b]} \end{cases} \tag{3.1.1}$$

[6] In practice, for any implementation of this DAG as a SCM, the arrows going from $S_{[a]}$ and $S_{[b]}$ to $S_{|a,b|}$ will represent *deterministic* causes. That is, they imply that $S_{|a,b|}$ is algebraically constructed from $S_{[a]}$ and $S_{[b]}$: $S_{|a,b|} = S_{[a]} + S_{[b]}$. This contrasts with the usual *probabilistic* causal relationships depicted in DAGs. The implications of using deterministic arrows in directed graphs are discussed in Berrie et al. (2022) and Koller and Friedman (2009).

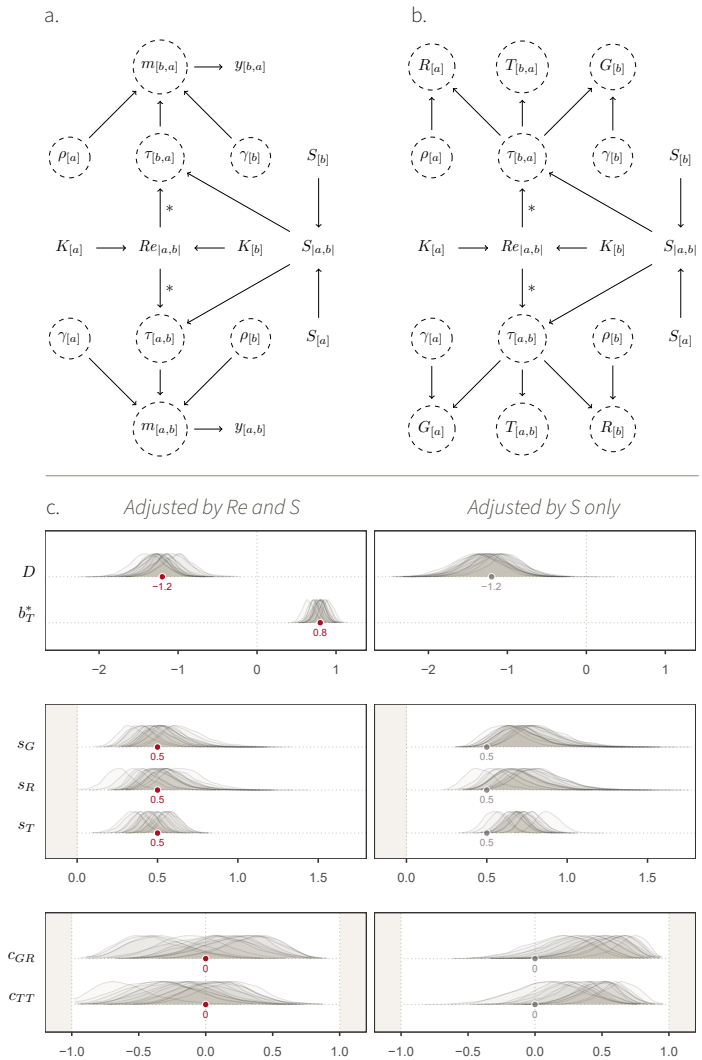

Figure 5 – *Simulation study 3.* **a.** DAG representing a causal system where the kin group $K_{[a]}$ of individuals affect how they preferentially interact with one another. Individuals belonging to the same kin group have a higher relatedness $Re_{|a,b|}$ than non-kin. $Re_{|a,b|}$, in turn, affects how the pair of individuals interact with one another, through $\tau_{[a,b]}$. This effect is our estimand, and we mark it $*$. The sampling effort of the two individuals, $S_{[a]}$ and $S_{[b]}$, determine $S_{|a,b|}$, which in turn affects the unit of $m_{[a,b]}$. The rest of the DAG is similar to Figure 2a. **b.** Mapping between structural parameters (SCM 3) statistical parameters. Compared to the simulation studies 1 and 2 where the only cause of $\tau_{[a,b]}$ was exogenous noise, $\tau_{[a,b]}$ is here affected by $Re_{|a,b|}$ and $S_{|a,b|}$, and it has an effect on $G_{[a]}$ and $R_{[b]}$. **c.** *Left*: fixed effects of the social relations model adjusted by $Re_{|a,b|}$ and $S_{|a,b|}$ (statistical model 3). *Right*: fixed effects of the social relations model adjusted by $S_{|a,b|}$ only. The target values of the well-adjusted model are shown in grey. Deviations from them allow us to understand how $Re_{|a,b|}$ affects varying effects, when $Re_{|a,b|}$ is not adjusted for.

We further assume that the variation in sampling effort is random:

$$f_{S_{[a]}}: \quad S_{[a]} \sim \text{Uniform}(1, 2.5). \tag{3.1.2}$$

The observation effort per individual takes a value between 1 and 2.5 time units—*e.g.*, the number of full days of focal-follows. We then combine the sampling effort of the two interacting individuals into their dyad-level sampling effort:

$$f_{S_{|a,b|}}: \quad S_{|a,b|} = S_{[a]} + S_{[b]}. \tag{3.1.3}$$

Then, we define how different variables affect γ, ρ, and τ:

$$f_{\gamma}: \quad \gamma_{[a]} \sim \text{Normal}(0, 0.5) \tag{3.1.4}$$

$$f_{\rho}: \quad \rho_{[a]} \sim \text{Normal}(0, 0.5) \tag{3.1.5}$$

$$f_{\tau}: \quad \tau_{[a,b]} \sim \text{Normal}(\log(S_{|a,b|}) + 0.8^{*} \cdot \tilde{Re}_{|a,b|}, 0.5), \tag{3.1.6}$$

Where the $\sim$ symbol on top of $\tilde{Re}_{|a,b|}$ means that we standardised it using a z-transformation.

$$f_{m}: \quad m_{[a,b]} = \exp(-1.2 + \gamma_{[a]} + \rho_{[b]} + \tau_{[a,b]}) \tag{3.1.7}$$

$$f_{y}: \quad y_{[a,b]} \sim \text{Poisson}(m_{[a,b]}). \tag{3.1.8}$$

Compared to the SCMs above, the intercept is set to -1.2, to make the interaction rates similar to those in the simulation studies above. In SCM 3, an unrelated average dyad observed for 3.5 time units, would have an expected value of $m_{[a,b]} = \exp(-1.2 + \log(3.5)) \simeq 1.1$ interactions, compared to $m_{[a,b]} = \exp(0.2) \simeq 1.2$ interactions in the previous simulation studies. We show the resulting network of synthetic observations, where the effect of genetic relatedness on social structure is visible by eye, in Figure S9.

To illustrate how the offset $\log(S_{|a,b|})$ scales the rates $m_{[a,b]}$—through its effect on $\tau_{[a,b]}$—, imagine two average, directed dyads $[1, 2]$ and $[12, 9]$, each composed of non-kin. Assume their respective rate of interaction per *one* time unit is identical, and is equal to: $\exp(-1.2 + \log(1)) = \exp(-1.2) = 0.3$. Let $S_{|1,2|} = 4$ and $S_{|12,9|} = 2$. Consequently, the dyads' interaction rates will be scaled by their respective values of $S_{|a,b|}$, such that: $m_{[1,2]} = \exp(-1.2 + \log(4)) = 1.2$ and $m_{[12,9]} = \exp(-1.2 + \log(2)) = 0.6$. Their meaning, thus, slightly differs from one another: $m_{[1,2]}$ is now the average number of interactions per *four* times units, and $m_{[12,9]}$ is the average number of interactions per *two* time units. In section S1.5, we show an alternative parameterisation of SCM 3 model where $S_{|a,b|}$ affects $y_{[a,b]}$ directly, instead of scaling $m_{[a,b]}$ through $\tau_{[a,b]}$.

2.3.2. Statistical model

Below, we describe *Statistical Model* 3, an estimator for the effect of genetic relatedness $Re_{|a,b|}$ on $m_{[a,b]}$.

$$y_{[a,b]} \sim \text{Poisson}(m_{[a,b]}) \tag{3.2.1}$$

$$m_{[a,b]} = \exp(D + G_{[a]} + R_{[b]} + \hat{T}_{[a,b]}). \tag{3.2.2}$$

Like in statistical model 1, $G_{[a]}$ and $R_{[a]}$ are part of a Multivariate adaptive prior distribution,

$$\begin{pmatrix} G_{[a]} \\ R_{[a]} \end{pmatrix} \sim \text{MVNormal}\left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} s_G^2 & c_{GR}s_Gs_R \\ c_{GR}s_Gs_R & s_R^2 \end{pmatrix}\right]. \tag{3.2.3}$$

$\hat{T}_{[a,b]}$ is, however, described by a submodel, stratified by the causes of $\tau_{[a,b]}$:

$$\hat{T}_{[a,b]} = \log(S_{|a,b|}) + T_{[a,b]} + b_T^* \cdot \tilde{Re}_{|a,b|} \tag{3.2.4}$$

$$\begin{pmatrix} T_{[a,b]} \\ T_{[b,a]} \end{pmatrix} \sim \text{MVNormal}\left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} s_T^2 & c_{TT}s_T^2 \\ c_{TT}s_T^2 & s_T^2 \end{pmatrix}\right]. \tag{3.2.5}$$

The submodel for $\hat{T}_{[a,b]}$ is offsetted by $\log(S_{|a,b|})$, which works like the offset of SCM 3. By making dyads comparable to one another, this term naturally deals with the issue of sampling biases that we introduced earlier (problem II). b_T^*, accordingly, captures the association between genetic relatedness and $m_{[a,b]}$, *conditional on* the variation in $S_{|a,b|}$ among dyads. This parameter should recover our estimand by taking a value of 0.8. The full specification of statistical model 3 can be found in section S1.4.2. In addition to statistical model 3, we also fitted the data generated by SCM 3 to a version of statistical model 3 that was not adjusted by $Re_{|a,b|}$. That is, equation 3.2.4 was replaced by:

$$\hat{T}_{[a,b]} = \log(S_{|a,b|}) + T_{[a,b]}.$$

2.3.3. Posterior model

Statistical model 3 successfully recovered the structural parameters of the generative model, SCM 3 (Figure 5c, left panel). Most notably, b_T^*, the statistical parameter whose aim was to estimate our causal estimand, accurately recovered the value of 0.8. This result validates our estimator, given the causal assumptions encoded in the DAG and SCM.

Below, we briefly explain how the effect of genetic relatedness, a *dyadic* variable, results in variation in *individual-* and *dyad*-level varying effects, when relatedness is not included in the statistical model (Figure 5c, right panel). In simulation study 1 and 2, the structural and statistical parameters were equivalent: $\gamma_{[a]}$ was the only cause of $G_{[a]}$, $\rho_{[a]}$ the only cause of $R_{[a]}$, and $\tau_{[a,b]}$ the only cause of $T_{[a,b]}$ (*e.g.*, Figure 4). As a result, $G_{[a]}$ fully recovered $\gamma_{[a]}$, $T_{[a,b]}$ recovered $\tau_{[a,b]}$, etc., such that the parameters capturing their patterns of (co)variation—s_G, s_R, c_{GR}—could be directly mapped onto the structural model. However, this is not the case here. We see that the estimates for s_G, s_R and c_{GR} substantially deviate from the grey dots—*i.e.*, the value they wound take if none of the effect of $\tau_{[a,b]}$ "leaked" onto $G_{[a]}$ and $R_{[b]}$.

Recall that $G_{[a]}$, $R_{[b]}$, and $T_{[a,b]}$ are *statistical* parameters, meaning they only measure (conditional) averages. They cannot "see" that the effect was encoded as a dyadic effect through $\tau_{[a,b]}$. Instead, $G_{[a]}$ and $R_{[b]}$ measure interindividual differences in

average interaction rates. Genetic relatedness, although it acts at the dyadic level, creates such differences. Individuals belonging to the large kin group give, and receive, more frequent social interactions, than those in the small kin groups (see Figure 5c, right panel, and Figure S9). This interindividual variation is visible as an increase in s_G and s_R, and in covariation between these two individual features, indicated by a positive value for c_{GR}. This pattern can also be explained with the DAG (Figure 5b): there is an open path between $G_{[a]}$ and $R_{[a]}$, passing through $Re_{|a,b|}$. The effect of $Re_{|a,b|}$ that can be attributed to interindividual variation is absorbed by $G_{[a]}$ and $R_{[a]}$, and the *residual* variation is absorbed by $T_{[a,b]}$. This residual variation does *not* correspond to a process in the SCM, and therefore, s_T and c_{TT} do not have a straightforward biological interpretation. Once we condition on $Re_{|a,b|}$, we block the paths between $G_{[a]}$ and $R_{[a]}$, and between $T_{[a,b]}$ and $T_{[b,a]}$. Hence, the correlations captured by c_{GR} and c_{TT} disappear (Figure 5c, left panel).

2.3.4. Conclusions

In simulation study 3, we modelled a data-generating process where two dyad-level variables affected the observed social interactions among individuals. Genetic relatedness, a biological variable, affected the rate at which individuals interacted with one another; it was our estimand. Sampling effort, a variable describing the sampling process, further impacted the scale of the interaction rates. We then built our estimator in combination with a DAG and a SCM, and successfully recovered the true causal effect. Next, we showed that dyad-level causal effects could impact individual- and dyad-level varying parameters, when not accounted for. Once again, we refer the interested readers to simulation study 3', a variation of simulation study 3, where we estimate the causal effect of a categorical dyad-level variable, like the combination of sexes, using a stochastic block model structure (section S1.6.1-S1.6.5). In the next section, building upon our previous models, we ask: how can we develop a statistical model to estimate the effect of a dyad-level variable if this effect is *confounded* (problem III)?

3. Empirical example

In this section, we showcase the use of our framework to address a causal question in a specific empirical system. Our aim is to estimate the effect of maternal relatedness on affiliative behaviour in the females of a wild population of Assamese macaques (*Macaca assamensis*). The causal assumptions that we will make about this social system are, although crude, reasonably realistic. Thus we will conclude the section by fitting our estimator to an empirical dataset, and we will show how to compute causal estimates on the outcome scale from the joint posterior probability distribution.

Outcome scale
Scale of the outcome variable. For instance, rate of behavioural events, for our Poisson models.

Joint posterior probability distribution
A multidimensional probability distribution describing the plausible values of a statistical model's parameters, after it has been updated with data.

3.1. Simulation study 4: kinship in female macaques

3.1.1. Verbal description

Kinship in female macaques, like in other Cercopithecinae species, is thought to be a major driver of social network structure (Chapais, 1983; Cheney & Seyfarth, 2007; Hinde, 1983). The kin group of a female macaque not only determines who she is genetically related to, it also affects her position in the group's dominance hierarchy. Females (non-genetically) inherit their rank from their mothers, such that individuals in the same kin group generally occupy adjacent ranks. Furthermore, kinship affects the formation of the social groups themselves. When groups of macaques permanently split, they usually divide along matrilines, and members of the same matrilines stay together in the newly formed groups (Chepko-Sade & Sade, 1979; De Moor et al., 2020; Widdig et al., 2006). This phenomenon may result in smaller social groups—often, newly formed ones—to contain fewer kin groups, and thus to have a higher average degree of genetic relatedness.

Cercopithecinae
Sub-family of the African and Asian monkeys (sometimes referred to as "old world" monkeys) which comprises the baboons, the vervet monkeys, and the macaques.

The genetic relatedness between two individuals, as well as their respective positions in the dominance hierarchy, might both affect the pattern of affiliative interactions they exchange with one another (Seyfarth, 1977). As suggested earlier, genetically related individuals might exchange more affiliative interactions because of a preference for their kin. Additionally, dominant females may be attractive affiliation partners. Individuals might, accordingly, preferentially target their dominants—rather than their subordinates—with affiliative interactions.

3.1.2. Directed Acyclic Graph

We represent this causal system as a DAG (Figure 6a), and assume that the rate of grooming bouts $m_{[a,b]}$ is a good proxy for affiliative relationships (J. B. Silk, 2002). The causal graph shares its basic structure with that of simulation study 3. We add an effect of kin group $K_{[a]}$ on individual rank $Ra_{[a]}$. $Ra_{[a]}$, in combination with $Ra_{[b]}$ then determine the *rank difference* between two individuals, $\Delta Ra_{|a,b|}$, which in turn affects $\tau_{[a,b]}$. Finally, we represent group size $GS_{\text{gr}_{|a,b|}}$ as a group-level feature (hence the "gr" index). The double-headed arrow between K and GS_{gr} means that they are both connected by an open path. It simply encodes that larger social groups tend to contain more kin groups while remaining agnostic about the complex, underlying, causal mechanism (*i.e.* group fission).

An important insight provided by the DAG is that our estimand appears to be confounded by dominance rank (Figure 6a). There are now two paths connecting $Re_{|a,b|}$ and $m_{[a,b]}$. There is, first, a "frontdoor" path, transmitting the causal effect of interest[7]:

$$Re \rightarrow \tau \rightarrow m,$$

[7] An alternative way to look at the DAG is to consider two directed paths going from kinship K to the outcome m. The first one passes through relatedness (*i.e.* our estimand), the other one through rank. In either case, we need to condition on rank to identify our estimand.

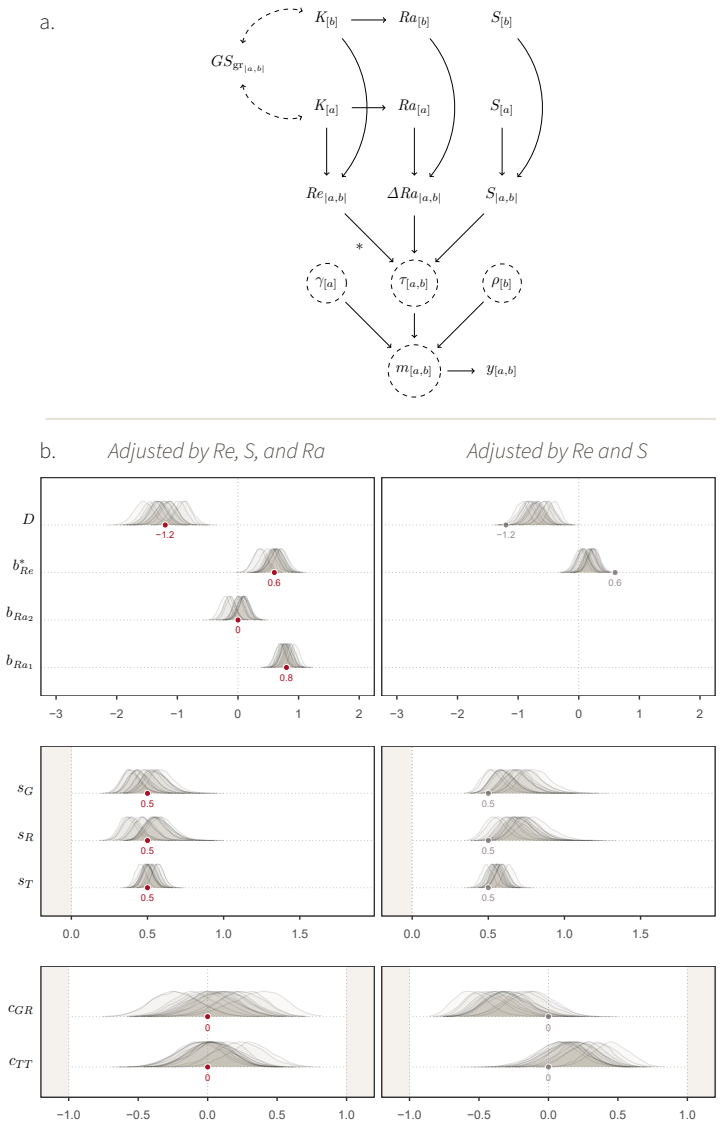

Figure 6 – *Simulation study 4.* **a.** DAG for the dual effect of kinship on grooming interactions in macaques. As in simulation study 3, the kin groups of two individuals, $K_{[a]}$ and $K_{[b]}$, determine their degree of genetic relatedness $Re_{|a,b|}$, which in turn affects $\tau_{[a,b]}$. This effect is our estimand, which we mark with $*$. $K_{[a]}$ and $K_{[b]}$ affect the dominance ranks $Ra_{[a]}$ and $Ra_{[b]}$ of the individuals, whose difference $\Delta Ra_{|a,b|}$ affect $\tau_{[a,b]}$. Dyads further vary in their respective sampling effort $S_{|a,b|}$. Finally, the size of social groups $GS_{\mathrm{gr}_{|a,b|}}$ covaries with the number of kin groups K within in. The rest of the DAG is similar to Figure 5a. **b.** *Left*: fixed effects of the social relations model adjusted by $Re_{|a,b|}$, $S_{|a,b|}$ and $\Delta Ra_{|a,b|}$ (statistical model 4). *Right*: fixed effects of the social relations model adjusted by $Re_{|a,b|}$ and $S_{|a,b|}$ only. The target values of the well-adjusted model are shown in grey. Deviations from them allow us to understand how $\Delta Ra_{|a,b|}$ counfounds our estimand, when $\Delta Ra_{|a,b|}$ is not adjusted for.

Second, there is a "backdoor" path transmitting a spurious association between Re and m[8]:

$$Re \leftarrow K \rightarrow Ra \rightarrow \tau \rightarrow m,$$

If we were to regress $m_{[a,b]}$ on $Re_{|a,b|}$ to estimate our estimand without stratifying by rank (as we did in simulation study 3), our effect would be biased by the backdoor path. To develop a better intuition for this issue, we first translate this DAG into SCM 4.

3.1.3. Structural Causal Model

In SCM 4, we create three groups of 10, 15, and 20 individuals, which vary in their maximum number of kin groups $N_{K_{\text{gr}_{|a,b|}}}$: 3, 5, 7, respectively. Each individual's kin group $K_{[a]}$ is drawn with replacement from the vector $\{1, \ldots, N_{K_{\text{gr}_{[a]}}}\}$. Doing so leads to a higher degree of average genetic relatedness in smaller social groups than in larger ones (double arrow in Figure 6). The kin groups are then assigned kin-group dominance rank $Ra_{K_{[a]}}$ (*i.e.* a rank for the whole matriline):

$$f_{Ra_K}: \quad Ra_{K_{[a]}} = K_{[a]}. \tag{4.1.1}$$

That is, kin groups are assigned a dominance rank at random. A lower $Ra_{K_{[a]}}$ imply a higher rank. Thus, if $Ra_{K_{[a]}} < Ra_{K_{[b]}}$, then the dominance rank of all of a's kin is higher than that of all of b's kin. Finally, within each kin group, individuals receive ranks at random too. Overall, these steps represent an implementation of the arrow from K to Ra, resulting in a dominance hierarchy that is stratified by matrilines of various sizes.

Next, sampling effort varies across individuals and dyads:

$$f_{S_{[a]}}: \quad S_{[a]} \sim \text{Uniform}(1, 2.5) \tag{4.1.2}$$

$$f_{S_{|a,b|}}: \quad S_{|a,b|} = S_{[a]} + S_{[b]}. \tag{4.1.3}$$

Sampling effort, genetic relatedness, and dominance rank difference all affect the interaction rate through $\tau_{[a,b]}$:

$$f_\tau: \quad \tau_{[a,b]} \sim \begin{cases} \text{Normal}(\log(S_{|a,b|}) + 0.6^* \cdot Re_{|a,b|} & \\ +0.8 \cdot (Ra_{[b]} - Ra_{[a]}), 0.5) & \text{if } R_{[a]} < R_{[b]} \\ \text{Normal}(\log(S_{|a,b|}) + 0.6^* \cdot Re_{|a,b|} & \\ +0 \cdot (Ra_{[b]} - Ra_{[a]}), 0.5) & \text{if } R_{[a]} \geq R_{[b]} \end{cases} \tag{4.1.4}$$

The main novelty of f_τ, compared to SCM 3, regards the causal effect of rank; or rather, the effect of the difference in ranks $Ra_{[b]} - Ra_{[a]}$. Here, individuals increase their interaction rate as a function of how much higher in the hierarchy their partner is: a larger increase for a larger difference. The effect is asymmetrical, and individuals

[8]Note that we removed the indices for simplicity, but the path splits in two in the middle.

are not impacted by how much lower in the hierarchy their partner is. The remaining structural equations are identical to those in SCM 3:

$$f_\gamma : \quad \gamma_{[a]} \sim \text{Normal}(0, 0.5) \tag{4.1.5}$$
$$f_\rho : \quad \rho_{[a]} \sim \text{Normal}(0, 0.5) \tag{4.1.6}$$
$$f_m : \quad m_{[a,b]} = \exp(-1.2 + \gamma_{[a]} + \rho_{[b]} + \tau_{[a,b]}) \tag{4.1.7}$$
$$f_y : \quad y_{[a,b]} \sim \text{Poisson}(m_{[a,b]}). \tag{4.1.8}$$

The network of simulated observations $y_{[a,b]}$, as well as their distribution, are shown in Figure S14.

Coming back to our backdoor problem, let us now imagine two directed dyads: $[1, 2]$ and $[3, 4]$. We assume that the exogenous (random) causes of $m_{[1,2]}$ and $m_{[3,4]}$ are set to zero. Suppose that the first dyad is composed of kin ($Re_{|1,2|} = 1$) who occupy adjacent ranks in the dominance hierarchy ($\Delta Ra_{|1,2|} \simeq 0$). Thus,

$$m_{[1,2]} = \exp(-1.2 + 0.6^* \cdot 1 + 0.8 \cdot 0) \simeq 0.5.$$

Suppose that the second dyad is composed of non-kin ($Re_{|3,4|} = 0$), who are further apart in the dominance hierarchy: $\Delta Ra_{|3,4|} \simeq 0.7$. Therefore,

$$m_{[3,4]} = \exp(-1.2 + 0.6^* \cdot 0 + 0.8 \cdot 0.7) \simeq 0.5.$$

Despite the positive effect of relatedness on grooming rate, $m_{[1,2]} \simeq m_{[3,4]}$, because the two dyads were not comparable to one another with respect to rank. The same applies in the rest of the population, where kin dyads are closer in the dominance hierarchy on average compared to non-kin dyads, thereby making the effect of relatedness unidentifiable by directly comparing kin and non-kin. Instead, to identify the effect of interest, we need to block the path created by rank by conditioning on it (we apply the backdoor criterion; see Pearl et al., 2016). That is, to estimate our estimand, we need to build an estimator that can address the following question: *"once we know the difference in rank between a and b, what association remains between genetic relatedness and grooming rate?"*.

Backdoor criterion
Graphical criterion providing a sufficient adjustment set (*i.e.* set of variables to include in a statistical model) for causal identification. Given a treatment X and an outcome Y, a set of variable Z satisfies the backdoor criterion if it does not include descendants of X (*i.e.* nodes caused by X), and if it blocks all paths starting with an arrow pointing into X.

3.1.4. Statistical model

Statistical model 4 is identical to statistical model 3, except that it includes dominance rank.

$$y_{[a,b]} \sim \text{Poisson}(m_{[a,b]}) \tag{4.2.1}$$
$$m_{[a,b]} = \exp(D + G_{[a]} + R_{[b]} + \hat{T}_{[a,b]}) \tag{4.2.2}$$
$$\hat{T}_{[a,b]} = \log(S_{|a,b|}) + T_{[a,b]} + b^*_{Re} \cdot Re_{|a,b|} + \begin{cases} b_{Ra_1} \cdot (Ra_{[b]} - Ra_{[a]}) & \text{if } R_{[a]} < R_{[b]} \\ b_{Ra_2} \cdot (Ra_{[b]} - Ra_{[a]}) & \text{if } R_{[a]} \geq R_{[b]} \end{cases} \tag{4.2.3}$$

The stratification by $\Delta Ra_{|a,b|}$ allows for asymmetrical effects. b_{Ra_1} should recover the effect of 0.8, and b_{Ra_2} should recover 0 (equation 4.1.4). Conditioning on rank

closes the backdoor path mentioned earlier, and thus, b^*_{Re} should recover the unbiased effect of 0.6. Finally, $G_{[a]}$, $R_{[a]}$, and $T_{[a,b]}$ are varying-effects, as in SCM 3:

$$\begin{pmatrix} G_{[a]} \\ R_{[a]} \end{pmatrix} \sim \text{MVNormal}\left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} s_G^2 & c_{GR}s_Gs_R \\ c_{GR}s_Gs_R & s_R^2 \end{pmatrix}\right] \tag{4.2.4}$$

$$\begin{pmatrix} T_{[a,b]} \\ T_{[b,a]} \end{pmatrix} \sim \text{MVNormal}\left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} s_T^2 & c_{TT}s_T^2 \\ c_{TT}s_T^2 & s_T^2 \end{pmatrix}\right]. \tag{4.2.5}$$

The full specification of the model, as well as prior predictive simulations, can be found in section S1.7.2 and S1.7.3. We also fitted the synthetic data to an estimator that was *not* stratified by rank—*i.e.* without the second and third lines of equation 4.2.3—, so that we could visualise the confounding effect of the backdoor path.

3.1.5. Posterior model (simulated data)

Statistical model 4 successfully recovered our estimand by deconfounding the biasing effect of dominance rank (Figure 6b). Looking at the left panel, we see that the marginal posterior distributions of b^*_{Re} are concentrated on the target value of 0.6. This is not the case for the panel on the right, where the model that was not adjusted by rank gives a consistently biased estimate, close to 0. These results confirm that given our causal assumptions, a statistical model that correctly blocks the backdoor path between $Re_{|a,b|}$ and $m_{[a,b]}$ is *necessary* to accurately estimate our estimand. More generally, the results validate our estimator for the effect of genetic relatedness in female Assamese macaques; insofar, at least, as we believe that our causal assumptions are good approximations of the data-generating process (we will come back to this issue in the discussion). We can thus turn to updating this estimator with empirical data.

3.2. Empirical data

The empirical observations that we fit to statistical model 4 were collected as part of a long-term research project on wild Assamese macaques at Phu Khieo Wildlife Sanctuary, in Northeastern Thailand (Figure 7a). We focused on the adult females of three social groups. The data collection took place between July 2017 and July 2018. The animals were fully habituated to the presence of researchers. Observers recorded the monkeys' behaviour using continuous focal sampling protocols of 40 minutes, during which they recorded all instances of dyadic interactions between adult females, including grooming and submissive behaviours, from which dominance ranks were computed. Submissive behaviours were also recorded *ad libitum*.

We defined $S_{[a]}$ as the number of 12 hours periods that each individual was observed for (hereafter "days" of observations), summed over sampling protocols. The average sampling effort was equal to 6 ± 1.4 days per individual (Figure S18). As in SCM 4, $S_{|a,b|} = S_{[a]} + S_{[b]}$ (Figure S19).[9] Field workers recorded when a macaque would

[9] This is an oversimplification. In practice, we split the sampling effort for each day, and only counted days when the two individuals were alive (one individual died during the study period).

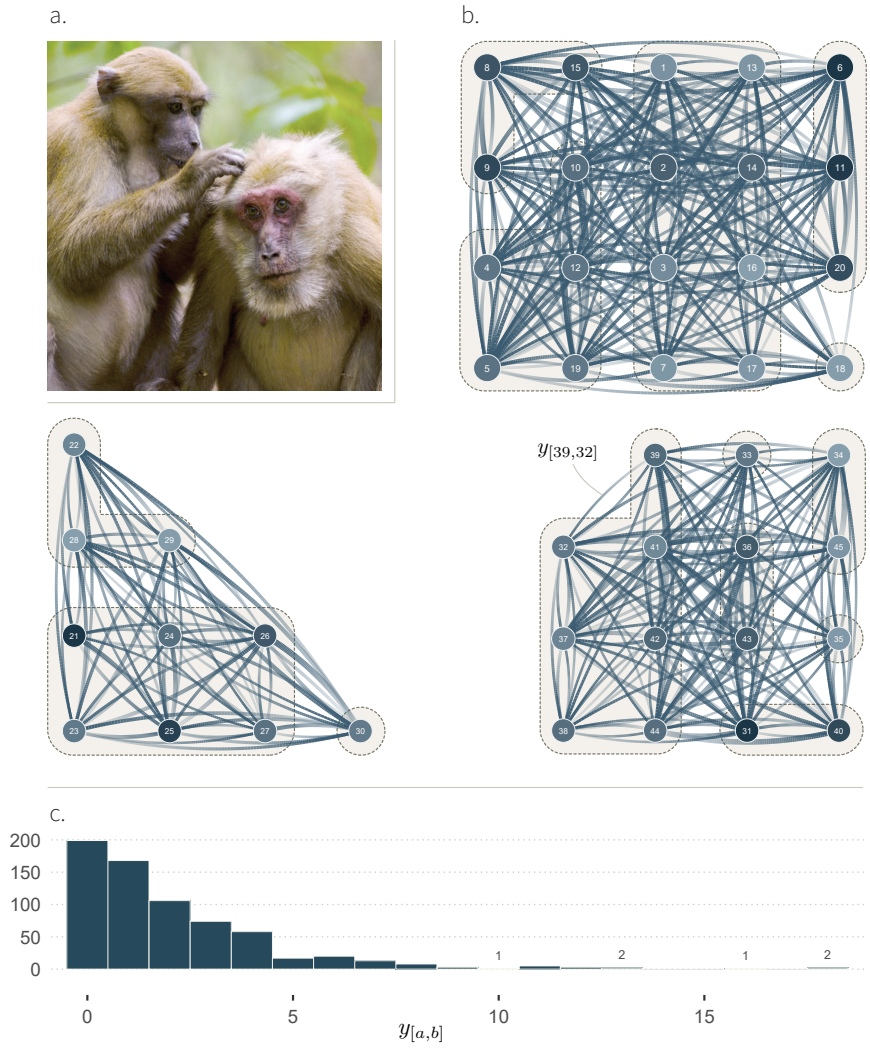

Figure 7 – *Empirical system used to illustrate our causal framework: female Assamese macaques.* **a.** Two Assamese macaques grooming. **b.** The graphs show 45 individuals across three groups of respectively 20, 10, and 15 females (we ignore the males). The edges represent the number of observed directed grooming bouts $y_{[a,b]}$ among them. Their width indicates the number of observed interactions: from 0, for no edge, to 18, for the thickest edges. The transparency gradient of the edges corresponds to the direction of the interaction ($y_{[a,b]}$ or $y_{[b,a]}$): the white end of an edge shows the giver, and its darker end shows the receiver. We highlight, as an example, the observed edge $y_{[39,32]}$ (1 interaction) from individual 39 to individual 32. The colour of the node depicts the individual dominance rank: lighter for low ranks (subordinates), and darker for higher ranks (dominants). Kin group are further highlighted by dashed outlines. This network corresponds to the third level of abstraction, in Figure 1. **c.** Distribution of observed directed grooming bouts $y_{[a,b]}$.

start, and when it would *stop*, to groom another individual. We then defined a directed grooming bout as a grooming "start" from a to b that took place at least 30 minutes following the previous grooming "stop" from a to b. We obtained a total of 481 directed bouts across 700 directed dyads. We counted bouts within each directed dyad to obtain $y_{[a,b]}$—which we show in Figure 7b-c (see also Figure S20). We established the pedigree from which kin groups $K_{[a]}$ were determined by combining observed birth events, with microsatellite data (De Moor et al., 2020; Sukmak et al., 2014). As in SCM 4, we dichotomised maternal kinship: $Re_{|a,b|} = 1$ if a and b belonged to the same matriline, and $Re_{|a,b|} = 0$ if a and b were born in different matrilines. We modeled each individual's rank $Ra_{[a]}$ from submissive behaviours, using Elo-score point estimates (details in Sadoughi et al., 2024; see also Albers and de Vries, 2001; Elo, 1978).

3.2.1. Posterior model (empirical data)

In this section, we describe the posterior model obtained by fitting the empirical data to statistical model 4. To start, we show the marginal posterior distributions for the estimator's fixed effects (Figure 8a). This figure is similar to figures we showed above (*e.g.*, Figure 6b), but differs in two ways. First, the red dots (*i.e.* the target values) are absent. The "true" structural parameters of the world are, of course, unknown; they are what we are trying to estimate. Second, there is only one posterior distribution per parameter. This is the case because we fitted the model to one empirical data set, instead of several iterations of a simulation.

We observe that 96% of b^*_{Re}'s posterior distribution is concentrated above 0, with a mean of 0.22 (Figure 8a). It means that the most plausible parameter values for b^*_{Re} are positive, given the data and the model. This marginal distribution cannot be intelligibly interpreted by itself in terms of the causal effect of interest nor, therefore, in terms of the species' biology. In fact, the computation of causal effects often requires more than just a slope. They are function of the whole statistical model—an issue we had so far ignored (McElreath, 2020).

Generalised linear models (GLM), like statistical model 4, are built-in interaction devices (Rohrer & Arslan, 2021; Spake et al., 2023). The effect of one parameter on the outcome always depends on (*i.e.* interacts with) the value of the other predictors. Imagine a simple GLM with a slope of value 2 quantifying the causal effect of a binary variable $X \in \{0, 1\}$ on a rate μ, for i observations:

$$\mu_i = \exp(\alpha + 2 \cdot X_i)$$

Now suppose that the intercept α can take two different values: 0 and 1. In the case where $\alpha = 0$, the causal effect of X is the following: $\exp(0 + 2 \cdot 1) \simeq 7.4$, minus $\exp(0 + 2 \cdot 0) = 1$, which is equal to 6.4. If, however, $\alpha = 1$, then the average effect is $\exp(1 + 2 \cdot 1) \simeq 20$, minus $\exp(1 + 2 \cdot 0) \simeq 2.7$, equal to $\simeq 17.4$. So, the same slope implies two very different effects: 6.4 and 17.4. With an exponential "inverse-link" function, causal effects on the outcome scale are multiplicative: they are larger when the baseline value (here, the value of α) increases.

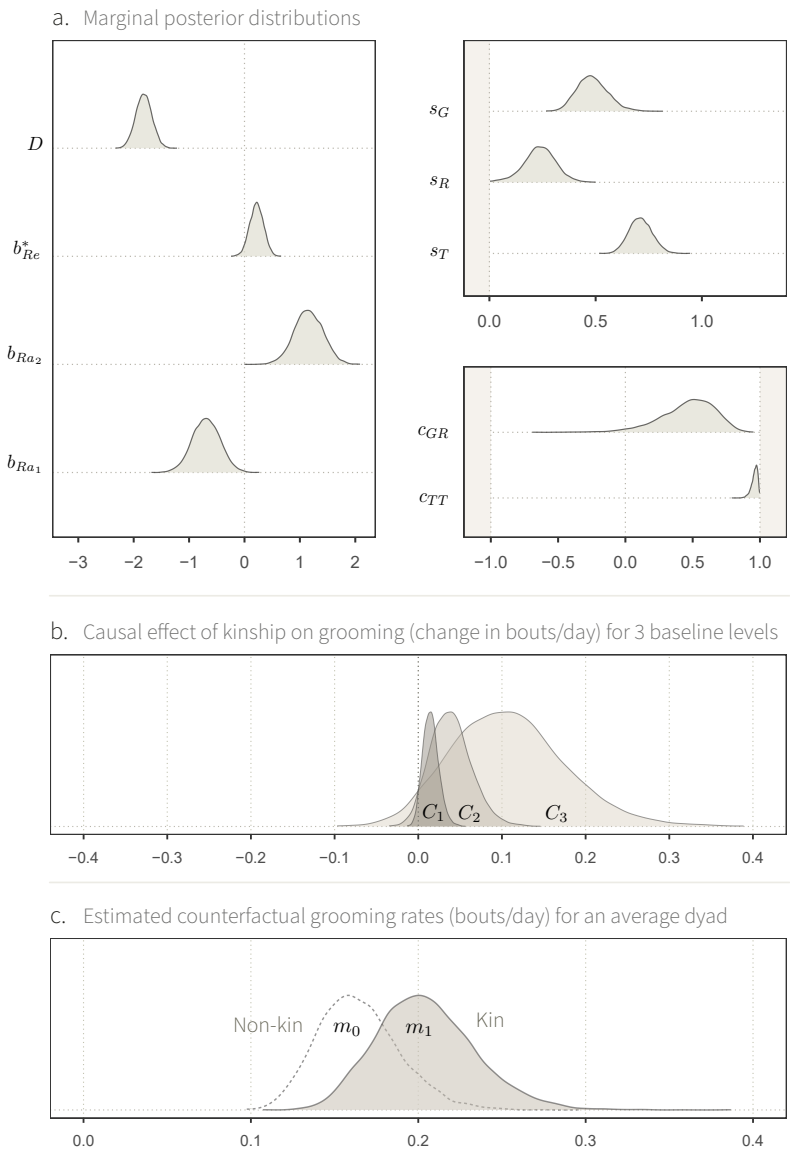

Figure 8 – *Posterior model obtained by updating statistical model 4 with empirical data from female macaques.* **a.** Marginal posterior distributions for the estimator's fixed effects. **b.** Conditional average relatedness effect C_j on inferred grooming rate $m_{[a,b]}$ (posterior mean contrast). We show this effect for three hypothetical dyads, respectively varying in their "baseline" level of grooming rate $\hat{\psi}_j$ (see main text)—from darkest to lightest: low-, medium-, and high-baseline. **c.** Counterfactual rates m_0 and m_1 for a hypothetical average directed dyad, where the two individuals occupy the same rank in the dominance hierarchy. The difference between the two distributions represents the causal effect of genetic relatedness for that dyad.

The same principle applies to statistical model 4, where the effect of the slope b^*_{Re} on the rate $m_{[a,b]}$ differs across dyads—*i.e.* a larger effect for dyads with a higher "baseline" level of grooming interactions. We define these differences in baseline using a parameter $\psi_{[a,b]}$ ("psi"):

$$\psi_{[a,b]} = D + G_{[a]} + R_{[b]} + T_{[a,b]}$$

Where the distribution of $\psi_{[a,b]}$ captures the variation across dyads that is caused by the unobserved (or exogenous) network structuring features, like age, social group differences, "personality", or friendship—we show the distribution of $\psi_{[a,b]}$ in Figure S26. We roughly summarise it with three values: $\hat{\psi}_j \in \{-2.8, -1.9, -0.9\}$, which respectively represent dyads with a low-, medium-, and high-baseline levels of grooming rate 8b). Equivalently, we can think of $\hat{\psi}_1$, $\hat{\psi}_2$, and $\hat{\psi}_3$ as standing for negative, null, and positive deviations from the global intercept D. For instance, a directed dyad with baseline $\hat{\psi}_1 = -2.8$ implies that unobserved factors are causing a to groom b at a low rate, compared to the other dyads in the population. We then compute the causal effect of genetic relatedness C, for each of these three levels $\hat{\psi}_j$ (Figure 8b):

$$C_j^{(n)} = \exp(\hat{\psi}_j + 1 \cdot b_{Re}^{(n)}) - \exp(\hat{\psi}_j + 0 \cdot b_{Re}^{(n)}),$$

That is, we obtained a posterior distribution for the contrasts of interest C_j from manipulating the Markov Chain Monte Carlo (MCMC) samples of the posterior model: each contrast C_j is composed of 8000 posterior samples n: $C_j^{(1)}, \ldots, C_j^{(8000)}$. This causal effect is sometimes called *Conditional Average Treatment effect* (CATE), because its value depends on—is *conditional* on—the value of the other predictors. Our model shows an expected increase of, respectively, 0.01, 0.04 and 0.11 grooming bouts per day due to genetic relatedness, for low-, medium-, and high-baseline interaction levels, respectively (that is, about an order of magnitude change from low- to high-baseline). These values correspond to the posterior CATE means. Note, however, that the estimates' distributions are rather broad. Albeit unlikely, small negative effects are still plausible from the estimator's perspective.

Markov Chain Monte Carlo Algorithm to draw samples from—and thus, approximate—the joint posterior probability distribution of a Bayesian model.

On the last panel (Figure 8c), we show our causal estimate in a different way. We compare $m_{Re=1}$ (m_1) and $m_{Re=0}$ (m_0), where:

$$m_1^{(n)} = \exp(D^{(n)} + 1 \cdot b_{Re}^{(n)})$$
$$m_0^{(n)} = \exp(D^{(n)} + 0 \cdot b_{Re}^{(n)}).$$

Again, the n index indicates 8000 draws from the posterior distribution. These two estimates represent two counterfactual outcomes for a hypothetical average dyad. That is, they correspond to estimated grooming rates of a dyad, if the two individuals were unrelated (left) or if they were kin (right), given that everything else remains identical: $G_{[a]}$,$R_{[b]}$, and $T_{[a,b]}$, and the difference in rank between them are set to zero. Comparing the two distributions allows us to visualise the relative effect of genetic relatedness. The posterior estimates have a mean of 0.16 and 0.20 grooming bouts per day, respectively. 0.20 is equal to 1.24 times 0.16; so relatedness is

estimated to cause a relative increase of about 24% in grooming rate. Note that the difference between the two distributions roughly corresponds to the middle contrast in Figure 8b.

3.2.2. Summary & discussion

We do not wish to further discuss these results in terms of the female Assamese macaques's biology—it would go beyond the scope of the manuscript. We also do not, in any way, argue that our estimator is optimal and that it captures all necessary factors at play. Instead, we built this example to illustrate two general points. Our first point regards the logical implications of causal assumptions for the development of an estimator. By examining our DAG, we showed that it was *necessary* to block a backdoor path (by conditioning on dominance rank) to identify our estimand. Not doing so resulted in a consistently biased results. Furthermore, encoding the DAG as a SCM has made visible that we needed to account for the dyadic *difference* in rank between individuals—not only for the individuals' ranks—, and to allow for this effect to possibly change depending on whether the difference was positive or negative. These considerations arose from relatively consensual biological knowledge about the social system (*e.g.*, see Cheney & Seyfarth, 2007; Seyfarth, 1977). Yet, they were not obvious from intuition alone.

Translating domain expertise into formal assumptions allowed us to *logically* connect causal assumptions and their implications for the estimator structure, thereby providing a licence for this estimator. Failing to establish this connection places an analysis on loose grounds, for it becomes impossible to know under which conditions it can work, even in theory. For instance, most empirical studies of the effect of maternal kinship on affiliation in Cercopithecinae have not considered the biasing path of rank (*e.g.*, Schino, 2001; Seyfarth, 1980; Wu et al., 2018). It is, however, hard to know if the confounder was simply ignored, or if instead, the results hold under other assumptions—in which case, we do not know which ones.

Secondly, we discussed why computing causal estimates on the outcome scale usually involves not only a slope, but the rest of the posterior model too. We notably showed that in the population under study, the effect substantially changed across dyads depending on their baseline interaction level. These considerations are particularly important because they determine how to interpret evidence for causal effects in light of the species' biology.

We refer the interested readers to supplementary sections where we explore other aspects of this study. We show MCMC diagnostics in section S1.8.2 and posterior predictive checks in section S1.8.3. We also compute the effect of rank—which, like that of relatedness, cannot be interpreted from the marginal posterior distributions alone—in section S1.8.5. Finally, we introduce simulation study 4', where group sizes also affects interaction rates through a novel parameter δ: a group-level counterpart of γ, ρ, and τ (sections S1.9.1-S1.10).

4. General discussion

Behavioural ecologists study the phenotypic and ecological factors that affect the structure of animal societies. This body of research is grounded in theoretical models, whether verbal or formal, that provide potential *causal explanations* for patterns of behavioural variation. These models have highlighted the roles of processes at the individual level (*e.g.*, age, reproductive state; Seyfarth, 1977; Siracusa et al., 2022), at the dyad level (*e.g.*, dominance, kinship, homophily, friendship; Nicolás Carlock et al., 2023; Seyfarth, 1977; J. B. Silk, 2002), and at the supra-dyadic level that can shape the social relationships between individuals (first level of abstraction, Figure 1), and thus, the social network structure emerging from it. Apace with this, there has been growing interest in empirically studying these processes in wild and captive populations. However, the connection between empirical results and theoretical models is often unclear. Many empirical studies are causally ambiguous, and theoretical assumptions specifying how variables cause one another in the system are rarely spelled out (Arif & MacNeil, 2022; Byrnes & Dee, 2024; Grosz et al., 2020). In addition, social network data analyses rarely integrate the measurement process, *e.g.*, by assuming that the observed social interactions perfectly capture the underlying interaction network (Figure 1). Statistical models that are not built in light of these biological and sampling processes can produce noisy and biased estimates, yielding incorrect results—as we outlined in problems I-IV.

Here, we have advanced a general framework for studying the causes of animal social network structure that allows empiricists to translate their theoretical domain expertise into an analytical strategy. *Within* each simulation study, we showcased how researchers may, first, define causal assumptions at various levels (*e.g.*, individual, dyad, group), and second, how they can validate an estimator for the effect of interest given these assumptions. We highlighted how these models naturally deal with (I) the gap between the interaction network and the observed network (discussed in simulation 1), (II) variation in sampling effort across individuals and dyads (sim. 3), (III) counfounders (sim. 4) and (IV) unobserved network causes (sim. 1-3). *Across* the four studies, we built models layer by layer, and illustrated an iterative workflow for causal inference in social networks, which we summarise in Box 1. In doing so, we provided empiricists with reproducible analytical tools that they can build upon to specify causal and statistical models for their own study system (see our GitHub repository).

The steps in Box 1 represent the core components of a full workflow for causal inference in animal social networks. Yet, they omit certain complexities. First of all, there often exist several plausible causal models that are compatible with a researcher's domain expertise. In this case, it can be useful to run several analyses in parallel, and to compare their outputs—that is, the list branches out. Fitting empirical data to estimators of intermediate complexity can also be insightful to check whether we observe the conditional dependencies that are expected from the DAG (Pearl et al., 2016). Furthermore, predictive tools (*e.g.*, Information Criteria, Cross Validation) can be useful to compare causally consistent estimators that differ in their parametric specifications (see Scholz & Bürkner, 2022). Finally, we did not highlight common

Box 1

A simple workflow for causal inference in animal social networks

1. *Define the estimand*, *i.e.* the theoretical quantity that the analysis is designed to estimate (Lundberg et al., 2021).
2. *Draw a DAG*, specifying qualitative assumptions about the biological and sampling processes that generate the data. An adjustment set can be derived from this DAG, for instance using the backdoor criterion (see Pearl et al., 2016). Note that certain confounding paths might be automatically blocked by the varying-effects of the Social Relations Model (*e.g.*, section S1.9.1-S1.9.3).
3. *Translate the DAG as a SCM*, encoding plausible quantitative assumptions about the data generating process (McElreath, 2020; Pearl et al., 2016). The SCM should ideally encode a target value for the estimand, which can be recovered by an estimator. We recommend exploring the parameter space to better understand how the simulation behaves.
4. *Build an estimator* based on the adjustment set derived from the DAG, and following the functional assumptions encoded in the SCM. The basic architecture of the Social Relations Model is a reasonable starting point for many animal social network models (Redhead, McElreath, et al., 2023; Ross et al., 2023; Snijders & Kenny, 1999).
5. *Validate that the estimator can recover the estimand*, by fitting the synthetic data generated by the SCM to the estimator. To better understand how the estimator works, we recommend inspecting not only the marginal posterior distribution of the parameter(s) of interest, but also, plotting other parameters of the posterior distribution (*e.g.*, s_G, c_{TT}, $G_{[a]}$, $m_{[a,b]}$), as well as computing causal effects on the outcome scale.
6. *Loop back to* 1., adding one layer of complexity. The new layers may involve any of the step above, like refining the estimand, adding a variable in the DAG, or encoding a more realistic functional relationship. After repeating this cycle, if you believe that you have translated your domain expertise into assumptions that reasonably approximate the data generating process, and that your estimand is recoverable from the estimator, then go to the next step.
7. *Fit the empirical data to the estimator*, and compute causal effects from the joint posterior distribution.

aspects of a Bayesian workflow in the main text, like prior and posterior predictive checks, or MCMC diagnostics (see Gelman et al., 2020; McElreath, 2020). These procedures are, however, important. We showed how they can be conducted for social network models in the Supplementary Materials (*e.g.*, section S1.8.2-S1.8.3).

4.1. Future research avenues

4.1.1. Measurement model

Conducting causal inference in social networks forces us to carefully examine the mapping between latent constructs of interest and data. In this manuscript, we focused on two kinds of unobserved quantities: true interaction rates, and causal effects. However, empirical research often involves additional latent variables, whose true values are uncertain or unknown. Let us return to the macaque example. We notice that, in addition to the unobserved causal effect on the unobserved rate $m_{[a,b]}$, the very variable whose effect we studied—genetic relatedness—was not directly measured either. In the study, we *approximated* it as a binary variable using pedigree and microsatellite data, thereby discarding the continuous variation and uncertainty in the true genetic relatedness among individuals. Furthermore, we were interested in studying its effect on *affiliation*, a theoretical construct that we approximated with grooming rate (Figure 1). In doing so, we ignored the rich set of behavioural interactions that compose affiliation in this species. Finally, by using Elo-rating point estimates for dominance rank, we treated this variable as if it was directly measured. It is important to realise that whenever a variable is unobserved, its relationship with observed variables depends on theoretical assumptions, whether explicit or implicit. In this case, we implicitly assumed that the proxies captured the latent variables without any uncertainty or error, thereby making a *leap* between levels of abstractions. This is far from optimal, for it can lead to inefficient and inaccurate estimators.

Moving forward, we argue that such latent variables should be modelled explicitly when possible, as part of a *measurement model*: *i.e.* a joint set of assumptions defining the connection between observed and unobserved variables. Doing this is natural with Bayesian models, where latent variables are assigned a posterior distribution like any other parameter. In the empirical study, genetic relatedness could for instance be *estimated* alongside the other components of the model (Koller & Friedman, 2009). Similarly, dominance relationships and affiliation may be modelled as latent variables in multiplex network models (De Bacco et al., 2023; Neumann & Fischer, 2023a, 2023b; Redhead et al., 2024; Redhead, McElreath, et al., 2023; Young et al., 2020). In this regard, statistical model 4 can be considered an intermediate step towards more realistic estimators. These considerations further apply to traits like age, whose exact value is sometimes unknown, but which can be modelled given a plausible interval; *e.g.*, if an individual is known to be born between two population surveys (McElreath, 2020). Furthermore, we believe that explicit measurement models could clarify current debates in fields like animal personality research, where the link between the latent objects of study (personality, behavioural syndrome) and the observed variables (behavioural measures) is conceptually and analytically challenging (Carter et al., 2013; Goold, 2020; Martin et al., 2019).

4.1.2. Dynamical drivers of social network structure

Sometimes, the latent causes of network structure cannot be modelled directly, but instead, can be inferred from the multilevel structure of the statistical model. This is the case for dynamic processes like reciprocity, a plausibly important driver of social

network structure across social species. For instance, the Social Relations Model's c_{TT} parameter may be interpreted as quantifying the extent of *dyadic reciprocity* after (i) positing a specific mechanism for reciprocity (*e.g.*, Leimar & Bshary, 2024; McElreath, 2020) in a SCM or a more fine-grained agent based model, (ii) under the assumption that all confounding paths between $T_{[a,b]}$ and $T_{[b,a]}$ have been blocked, and (iii) after verifying that given (i) and (ii), the pattern of interest could be detected with c_{TT}. The same goes for c_{GR} as a tool to quantify generalised-reciprocity. These considerations not only apply to animal social network analysis, but also to several disciplines in the Social and Behavioural Sciences—*e.g.*, Psychology or Anthropology—where the Social Relations Model's parameters are interpreted in terms of latent processes (Back et al., 2023; Back & Kenny, 2010; Koster & Leckie, 2014; Redhead, Dalla Ragione, et al., 2023; Redhead, McElreath, et al., 2023). In any case, we wish to insist that such inferences require extreme care. As we saw earlier, most causal paths flowing through γ, ρ, and τ end up in c_{TT} and/or c_{GR}. Therefore, these paths would need to be all blocked for c_{TT} and c_{GR} to be interpret as meaningful signals of reciprocity. As always, the causal evidence will only be as strong as the causal assumptions are plausible. But remember, imperfect causal assumptions are still better than no explicit assumptions at all (Cartwright, 1995).

An area where the causal tools presented here may be particularly useful is to study the drivers and consequences of social network structure in the context of longitudinal data analysis (Imai & Kim, 2019; Rohrer & Murayama, 2023; Runge et al., 2023). In our studies, we always assumed stationary systems, where time did not matter. However, many network structuring processes are intrinsically dynamic. For instance, an individual's gregariousness might affect its health, reproductive success, and survival (Franz et al., 2013; Ostner & Schülke, 2018; Snyder-Mackler et al., 2020). These outcomes, in turn, can shape social relationships and the pool of individuals in the social network, thereby forming potential loops of reciprocal causation between social network structure and other phenotypic traits (Redhead & Power, 2022; Steglich et al., 2010). Another way for social network structure to affect itself over time regards triadic closure (Block, 2015; Nicolás Carlock et al., 2023). Consider three individuals, a, b, and c. If a and b are both connected to c at time t, they might become more likely to form a connection $a - b$ at time $t + 1$, *because* of their shared relationship with c (*"friends of friends become friends"*).

Social network structuring processes, along with the measurement procedure to capture them, are intrinsically causal. They all pose the inferential challenges that we have outlined throughout the paper: when not integrated into an analysis, they can lead to inferences that are simply wrong. In this context, we believe that establishing a logical connection between theory and data, using transparently justified estimators, is crucial. We hope that our proposed tools and workflow will inspire future empirical research in this effort.

Data and Code

All relevant code and data can be found on this GitHub repository.

Figures

The figures in this manuscript were generated using *R* (version 4.2.1), LaTeX (TikZ package), *Adobe Illustrator CC 2015*, or a combination of them. In R, we used the following packages: *ggplot2* (Wickham, 2011), *tidybayes* (Kay, 2020), patchwork (Pedersen, 2019), *ggraph* (Pedersen et al., 2017), and *igraph* (Csardi, 2010).

Acknowledgments

We thank Alice Hill, Ana Lucia Arbaiza Bayona, and Shivani for the productive Rethinking sessions, and for their comments on earlier versions of the manuscript.

Funding

This project was funded by the *"DFG Research Training Group 2070 Understanding Social Relationships*, Project Nr. 254142454". Daniel Redhead was supported by the 'Societal Transitions and Behavioural Change' sector plan from the ministry of Science and Culture of The Netherlands.

Conflicts of interest

None.

S1. Supplementary Materials

S1.1. Simulation study 1: random structuring features

S1.1.1. Observed interactions (y) simulated with SCM 1

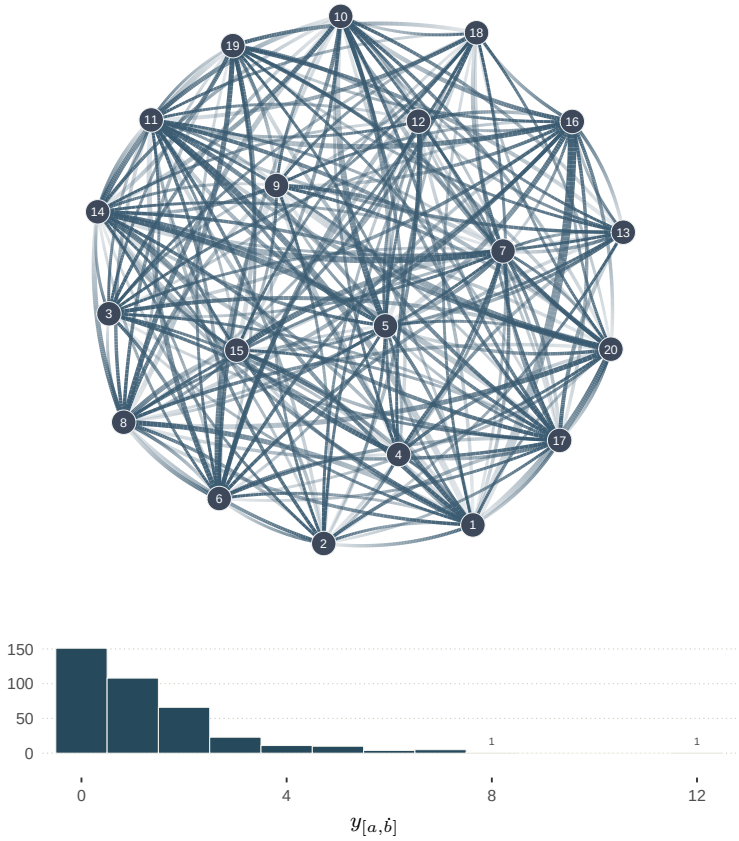

Figure S1 – *Network of observed interactions (y) generated with SCM 1.* The graph shows 20 individuals (nodes), and the number of observed interactions $y_{[a,b]}$ among them, where a and b correspond to individual indices: $a, b \in \{1, \ldots, 20\}$ and $a \neq b$. The width of the edges indicates the number of observed interactions: 1 interaction for the thinnest, and 12 interactions for the thickest (see distribution under the graph). No edges imply no observed interactions. The transparency gradient of the edges corresponds to the direction of the interaction ($y_{[a,b]}$ or $y_{[b,a]}$): the white end of an edge shows the giver, and its darker end shows the receiver. This network corresponds to the third level of abstraction, in Figure 1.

S1.1.2. Full description of statistical model 1

Below, we write the full statistical model of simulation study 1, as we have implemented it in *Stan*. It differs from the less detailed equations $1.2.1 - 1.2.4$ (see main text) in three aspects. First, each dyad has two directions (from a to b, and from b to a), which simultaneously run in the same Markov Chain. Second, the Multivariate Normal adaptive priors are parameterised as Cholesky factors, for computational reasons. This parameterisation is mathematically equivalent to—but, in this context, more efficient than—the centred parameterisation shown in equations $1.2.3 - 1.2.4$ (see McElreath, 2020; Nicenboim et al., 2021). Third, we show the (here, weakly regularising) priors and hyper-priors. These differences also apply to the other statistical models described in the manuscript.

$$
\begin{aligned}
y_{[a,b]} &\sim \text{Poisson}(m_{[a,b]}) \\
y_{[b,a]} &\sim \text{Poisson}(m_{[b,a]})
\end{aligned}
$$

$$
\begin{aligned}
m_{[a,b]} &= \exp(D + G_{[a]} + R_{[b]} + T_{[a,b]}) \\
m_{[b,a]} &= \exp(D + G_{[b]} + R_{[a]} + T_{[b,a]})
\end{aligned}
$$

$$
\begin{aligned}
\begin{pmatrix} G_{[a]} \\ R_{[a]} \end{pmatrix} &= \begin{pmatrix} s_G & 0 \\ 0 & s_R \end{pmatrix} \times L_{\text{ind}} \times \begin{pmatrix} z_{G_{[a]}} \\ z_{R_{[a]}} \end{pmatrix} \\
\begin{pmatrix} T_{[a,b]} \\ T_{[b,a]} \end{pmatrix} &= \begin{pmatrix} s_T & 0 \\ 0 & s_T \end{pmatrix} \times L_{\text{dyad}} \times \begin{pmatrix} z_{T_{[a,b]}} \\ z_{T_{[b,a]}} \end{pmatrix}
\end{aligned}
$$

$$
\begin{aligned}
D, z_{G_{[a]}}, z_{R_{[a]}}, z_{T_{[a,b]}}, z_{T_{[b,a]}} &\sim \text{Normal}(0,1) \\
s_G, s_R, s_T &\sim \text{Exponential}(1) \\
L_{\text{ind}}, L_{\text{dyad}} &\sim \text{LKJ Cholesky}(2).
\end{aligned}
$$

S1.1.3. Prior distributions (simulation study 1)

Using R, we drew n samples for each population-level parameter, or fixed effect (Figure S2):

$$
\begin{aligned}
s_G^{(n)}, s_R^{(n)}, s_T^{(n)} &\sim \text{Exponential}(1) \qquad\qquad n \in \{1, \ldots, 10000\} \\
c_{GR}^{(n)}, c_{TT}^{(n)} &\sim \text{LKJ}(\eta = 2, 2) \\
D^{(n)} &\sim \text{Normal}(0, 1)
\end{aligned}
$$

Where the *sub*scripts distinguish the parameters from one another (they are not indices), and the *super*script (n) is an index for the prior draws. For instance, $s_G^{(1)}$ is the first draw for the parameter s_G, $s_G^{(2)}$ is the second, and so on. We executed more in depth prior predictive checks for statistical model 4 (see section S1.7.3).

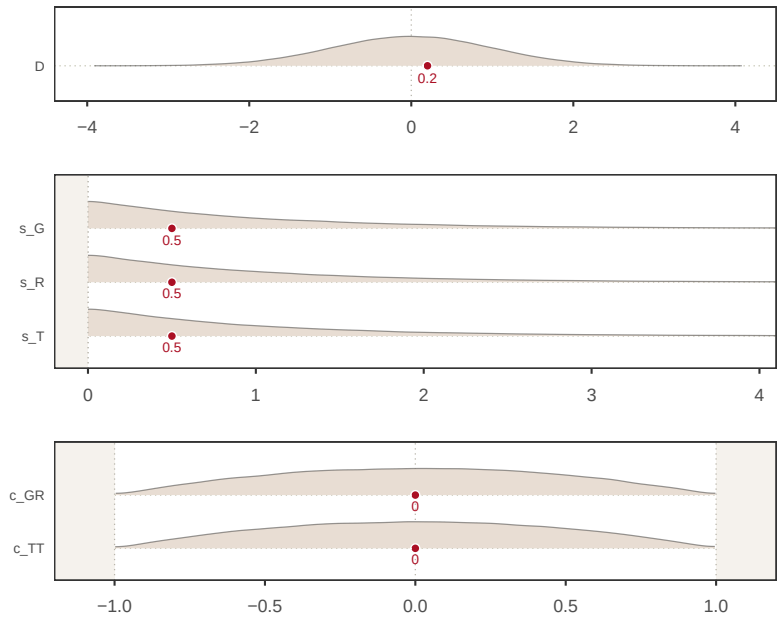

Figure S2 – *Prior distributions for the fixed effects of statistical model 1.* The prior probability density of a range of parameter values (*x-axis*) is described for several parameters (*y-axis*) of statistical model 1. As for Figure 3, the red dots represent the target values, derived from SCM 1.

S1.1.4. Posterior distributions of m (simulation study 1)

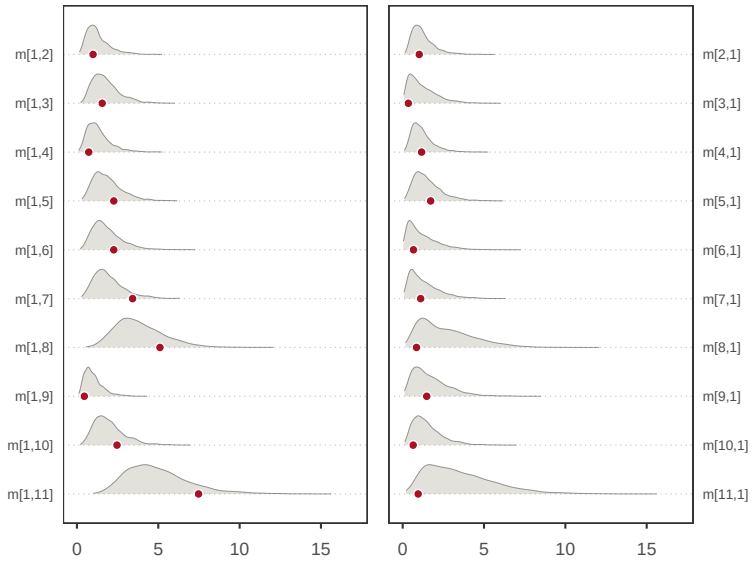

Figure S3 – *Marginal posterior distributions for a subset of true rates m (posterior model 1).* The posterior probability density of a range of parameter values (x-axis) is described for 20 parameters $m_{[a,b]}$ (y-axis). These parameters correspond to the third level of abstraction of Figure 1. The panel on the left shows the parameters in one direction (a-to-b), and the panel on the right shows them in the other direction (b-to-a). The posterior densities were obtained by updating statistical model 1, with the data generated in SCM 1. The red dots represent the target values, as they are encoded in the SCM.

S1.1.5. Joint distribution for individual-level parameters (simulation study 1)

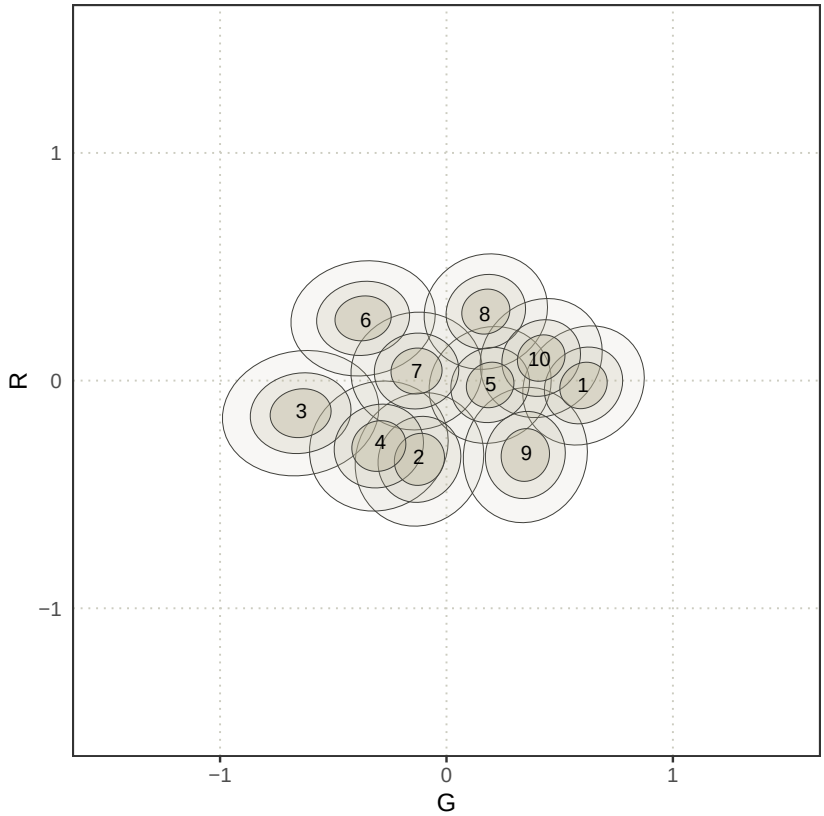

Figure S4 – *Joint posterior distributions of a subset of individual-level parameters G and R (posterior model 1).* The ellipses approximate 50%, 35%, and 10% percentile intervals, for the posterior probability of joint individual-level parameters (10 individuals shown). A darker colour and a narrower ellipse indicate a higher probability density. Individual indices are shown near the mode of the ellipses. The variances in $G_{[a]}$ and $R_{[a]}$ are respectively estimated by the s_G and s_R parameters (Figure 2c). c_{GR} further quantifies the (lack of) correlation between $G_{[a]}$ and $R_{[a]}$. Figure 3a shows the posterior distributions of $G_{[a]}$ and $R_{[a]}$, by *marginalising* (*i.e.* averaging) over the other parameters.

S1.2. Simulation study 2: individual-level features

S1.2.1. Observed interactions (y) simulated with SCM 2

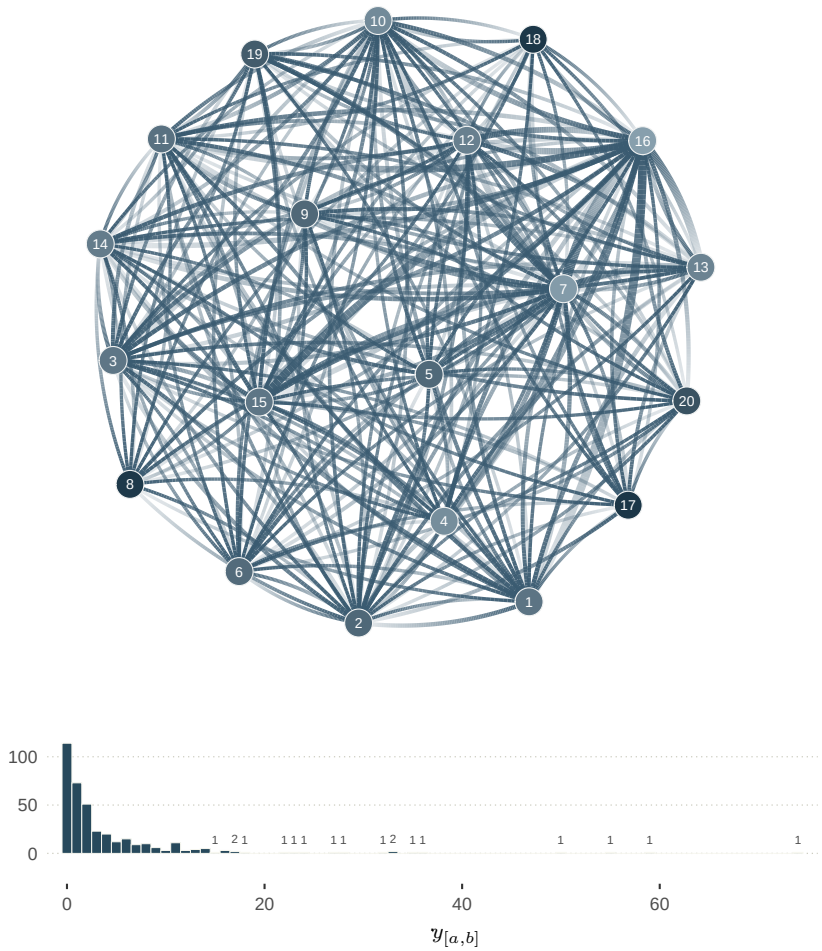

Figure S5 – *Network of observed interactions (y) generated with SCM 2.* The graph shows 20 individuals (nodes), and the number of observed interactions $y_{[a,b]}$ among them, where a and b correspond to individual indices: $a, b \in \{1, \ldots, 20\}$ and $a \neq b$. The node colour shows the value of $X_{[a]}$: light blue for lower values (*e.g.*, younger individuals), and darker for higher values (*e.g.*, older individuals). The width of the edges indicates the number of observed interactions (see distribution under the graph), where no edges imply no observed interactions. The transparency gradient of the edges corresponds to the direction of the interaction ($y_{[a,b]}$ or $y_{[b,a]}$): the white end of an edge shows the giver, and its darker end shows the receiver. This network corresponds to the third level of abstraction, in Figure 1.

S1.2.2. Full description of statistical model 2

$$y_{[a,b]} \sim \text{Poisson}(m_{[a,b]})$$
$$y_{[b,a]} \sim \text{Poisson}(m_{[b,a]})$$

$$m_{[a,b]} = \exp(D + \hat{G}_{[a]} + \hat{R}_{[b]} + T_{[a,b]})$$
$$m_{[b,a]} = \exp(D + \hat{G}_{[b]} + \hat{R}_{[a]} + T_{[b,a]})$$

$$\hat{G}_{[a]} = G_{[a]} + b_G \cdot X_{[a]}$$
$$\hat{G}_{[b]} = G_{[b]} + b_G \cdot X_{[b]}$$

$$\hat{R}_{[a]} = R_{[a]} + b_R \cdot X_{[a]}$$
$$\hat{R}_{[b]} = R_{[b]} + b_R \cdot X_{[b]}$$

$$\begin{pmatrix} G_{[a]} \\ R_{[a]} \end{pmatrix} = \begin{pmatrix} s_G & 0 \\ 0 & s_R \end{pmatrix} \times L_{\text{ind}} \times \begin{pmatrix} z_{G_{[a]}} \\ z_{R_{[a]}} \end{pmatrix}$$
$$\begin{pmatrix} T_{[a,b]} \\ T_{[b,a]} \end{pmatrix} = \begin{pmatrix} s_T & 0 \\ 0 & s_T \end{pmatrix} \times L_{\text{dyad}} \times \begin{pmatrix} z_{T_{[a,b]}} \\ z_{T_{[b,a]}} \end{pmatrix}$$

$$D, b_G, b_R, z_{G_{[a]}}, z_{R_{[a]}}, z_{T_{[a,b]}}, z_{T_{[b,a]}} \sim \text{Normal}(0, 1)$$
$$s_G, s_R, s_T \sim \text{Exponential}(1)$$
$$L_{\text{ind}}, L_{\text{dyad}} \sim \text{LKJ Cholesky}(2)$$

S1.2.3. Joint distribution for individual-level parameters (simulation study 2)

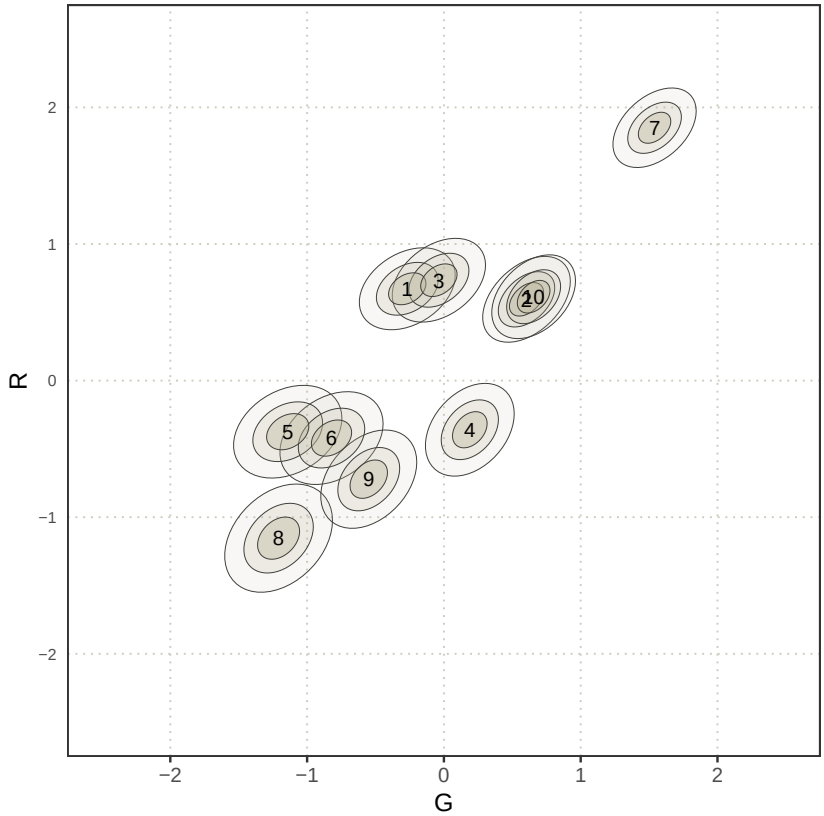

Figure S6 – *Joint posterior distributions of a subset of individual-level parameters G and R (posterior model 2).* The ellipses approximate 50%, 35%, and 10% percentile intervals, for the posterior probability of individual-level parameters (10 individuals shown). A darker colour and a narrower ellipse indicate a higher probability density. Individual indices are shown near the mode of the ellipses. Note that this plot shows the distribution described by s_G, s_R and c_{GR} (see Figure 4c).

S1.3. Simulation Study 2': categorical individual-level features

This simulation study is a variation of simulation study 2. Its graphical causal model is identical (see Figure 4a-b). However, its quantitative implementation differs. Below, we encode $X_{[a]}$ as categorical variable with two levels, *e.g.* sex, in a SCM (section S1.3.1-S1.3.2), and show that a well-specified estimator (section S1.3.3) can recover its structural parameters (section S1.3.4).

S1.3.1. Structural Causal Model 2'

$$
\begin{aligned}
f_X: &\quad X_{[a]} \in_R \{1, 2\} \\
f_\gamma: &\quad \gamma_{[a]} \sim \begin{cases} \text{Normal}(-0.5, 0.5) & \text{if } X_{[a]} = 1 \\ \text{Normal}(+0.5, 0.5) & \text{if } X_{[a]} = 2 \end{cases} \\
f_\rho: &\quad \rho_{[a]} \sim \begin{cases} \text{Normal}(-0.5, 0.5) & \text{if } X_{[a]} = 1 \\ \text{Normal}(+0.5, 0.5) & \text{if } X_{[a]} = 2 \end{cases} \\
f_\tau: &\quad \tau_{[a,b]} \sim \text{Normal}(0, 0.5) \\
f_m: &\quad m_{[a,b]} = \exp(0.2 + \gamma_{[a]} + \rho_{[b]} + \tau_{[a,b]}) \\
f_y: &\quad y_{[a,b]} \sim \text{Poisson}(m_{[a,b]}).
\end{aligned}
$$

Here, f_X states that the value of X for each individual is random.

S1.3.2. Observed interactions (y) simulated with from SCM 2'

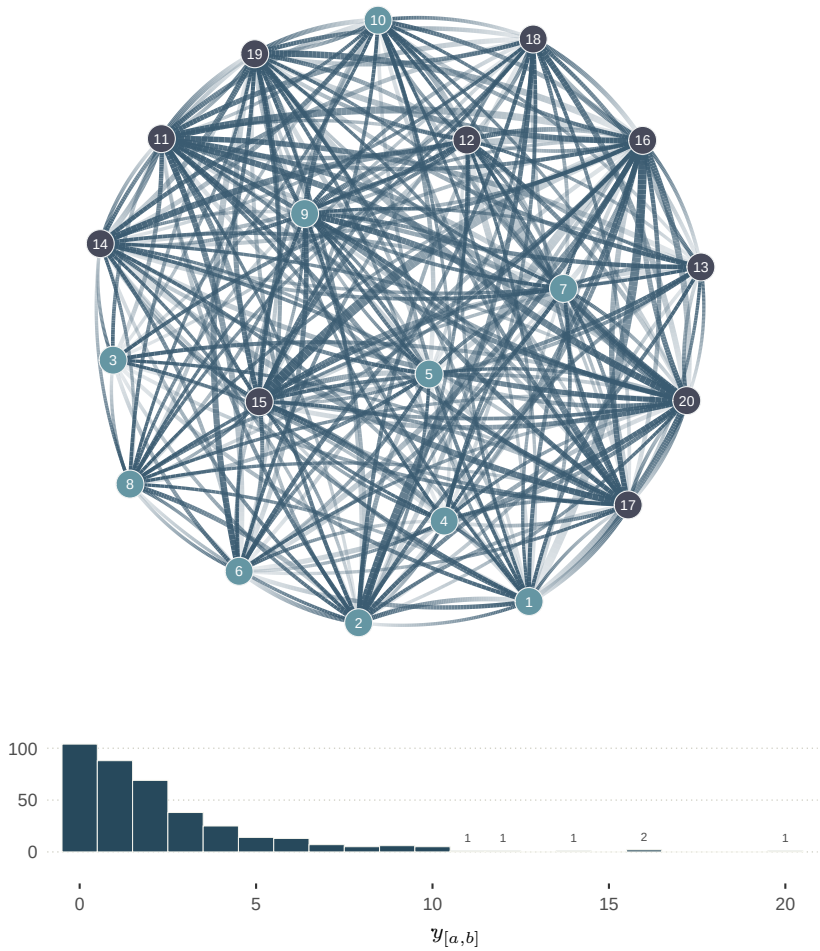

Figure S7 – *Network of observed interactions (y) generated with SCM 2'.* The graph shows 20 individuals (nodes), and the number of observed interactions $y_{[a,b]}$ among them, where a and b correspond to individual indices: $a, b \in \{1, \ldots, 20\}$ and $a \neq b$. The node colour shows the value of $X_{[a]}$: light blue for $X_{[a]} = 1$, and dark blue for $X_{[a]} = 2$. The width of the edges indicates the number of observed interactions (see distribution under the graph) where no edges imply no observed interactions. The transparency gradient of the edges corresponds to the direction of the interaction ($y_{[a,b]}$ or $y_{[b,a]}$): the white end of an edge shows the giver, and its darker end shows the receiver. This network corresponds to the third level of abstraction, in Figure 1.

S1.3.3. Statistical Model 2'

$$y_{[a,b]} \sim \text{Poisson}(m_{[a,b]})$$
$$y_{[b,a]} \sim \text{Poisson}(m_{[b,a]})$$

$$m_{[a,b]} = \exp(D + \hat{G}_{[a]} + \hat{R}_{[b]} + T_{[a,b]})$$
$$m_{[b,a]} = \exp(D + \hat{G}_{[b]} + \hat{R}_{[a]} + T_{[b,a]})$$

$$\hat{G}_{[a]} = L_{[X_{[a]}]} + G_{[a]}$$
$$\hat{G}_{[b]} = L_{[X_{[b]}]} + G_{[b]}$$

$$\hat{R}_{[a]} = M_{[X_{[a]}]} + R_{[a]}$$
$$\hat{R}_{[b]} = M_{[X_{[b]}]} + R_{[b]}$$

$$\begin{pmatrix} G_{[a]} \\ R_{[a]} \end{pmatrix} = \begin{pmatrix} s_G & 0 \\ 0 & s_R \end{pmatrix} \times \mathcal{L}_{\text{ind}} \times \begin{pmatrix} z_{G_{[a]}} \\ z_{R_{[a]}} \end{pmatrix}$$
$$\begin{pmatrix} T_{[a,b]} \\ T_{[b,a]} \end{pmatrix} = \begin{pmatrix} s_T & 0 \\ 0 & s_T \end{pmatrix} \times \mathcal{L}_{\text{dyad}} \times \begin{pmatrix} z_{T_{[a,b]}} \\ z_{T_{[b,a]}} \end{pmatrix}$$

$$D, L_{[1]}, L_{[2]}, M_{[1]}, M_{[2]}, z_{G_{[a]}}, z_{R_{[a]}}, z_{T_{[a,b]}}, z_{T_{[b,a]}} \sim \text{Normal}(0, 1)$$
$$s_G, s_R, s_T \sim \text{Exponential}(1)$$
$$\mathcal{L}_{\text{ind}}, \mathcal{L}_{\text{dyad}} \sim \text{LKJ Cholesky}(2)$$

The nested-brackets notation of $L_{[X_{[a]}]}$ and $M_{[X_{[a]}]}$ indicates that M and L have as many levels as the number of levels in X (*i.e.* here, two levels each), which are indexed based on the value of X for individual a. For instance, if $X_{[a]} = 1$, then $L_{[X_{[a]}]}$ captures the average deviation from D for individuals with $X = 1$; and inversely, if $X_{[b]} = 2$, then $L_{[X_{[b]}]}$ estimates the average deviation from D for individuals with $X = 2$.

S1.3.4. Posterior models (simulation study 2')

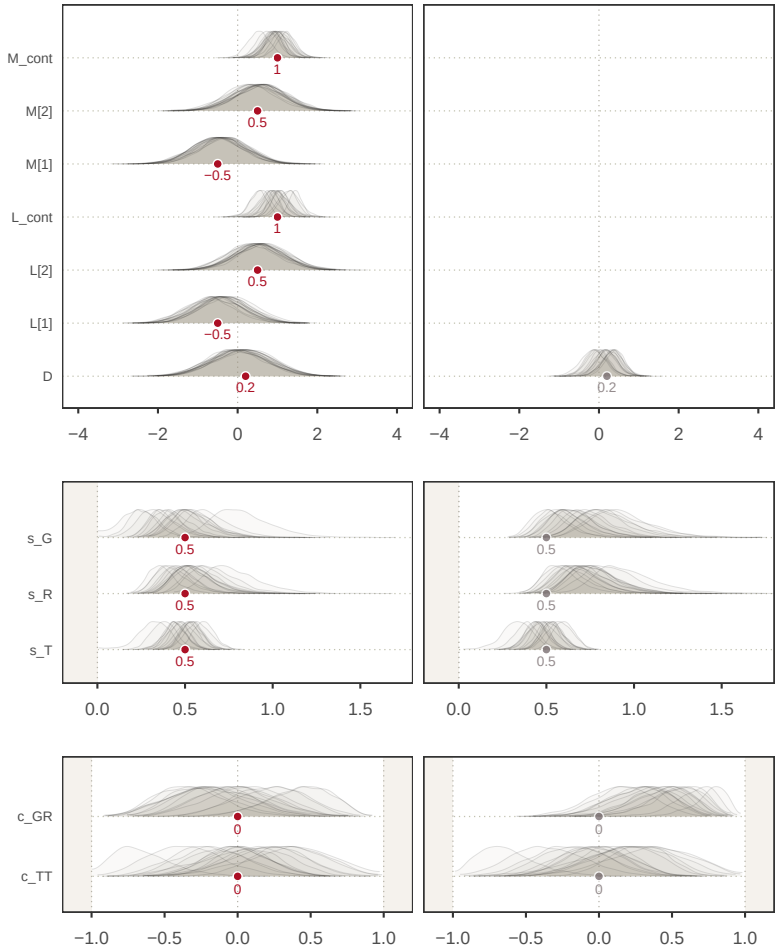

Figure S8 – *Marginal posterior distributions of the fixed effects of two different statistical models fitted to the data generated with SCM 2'.* Left: fixed effects of the social relations model adjusted by $X_{[a]}$ (see section S1.3.3). Right: fixed effects of the non-adjusted social relations model. "M_cont" and "L_cont" corresponds to the contrasts between $M_{[2]} - M_{[1]}$, and $L_{[2]} - L_{[1]}$, respectively.

S1.4. Simulation study 3: dyad-level features

S1.4.1. Observed interactions (y) simulated with SCM 3

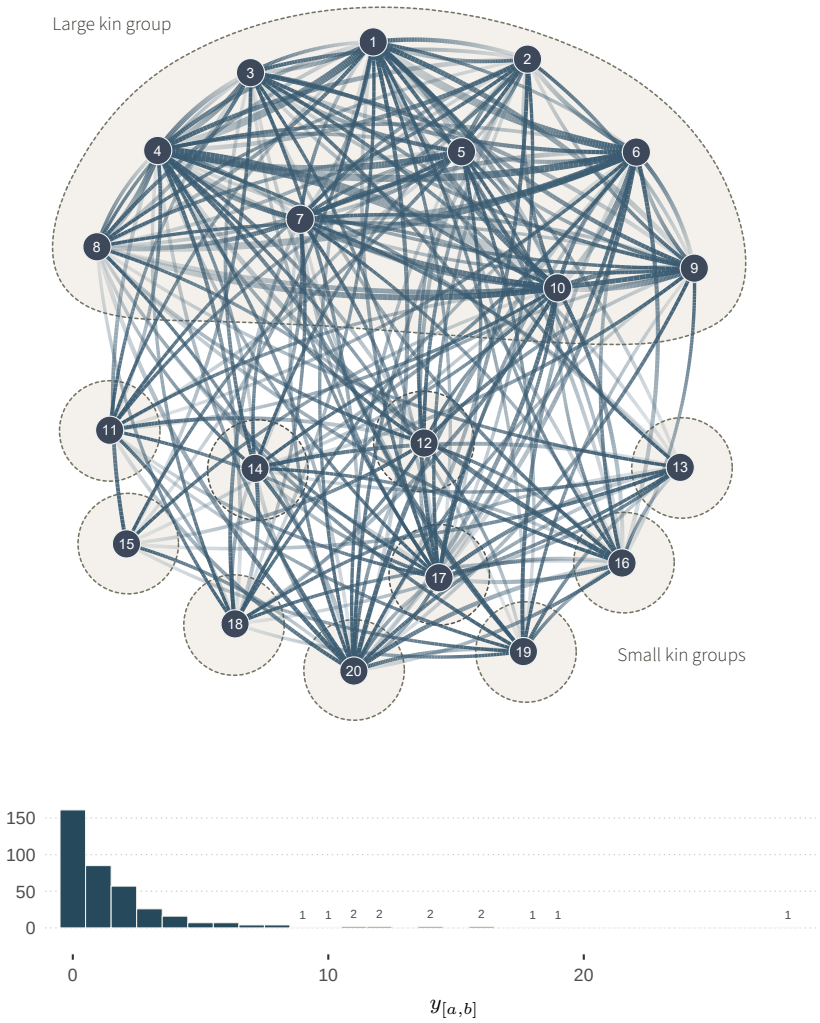

Figure S9 – *Network of observed interactions (y) generated with SCM 3.* The graph shows 20 individuals (nodes), and the number of observed interactions $y_{[a,b]}$ among them, where a and b correspond to individual indices: $a, b \in \{1, \ldots, 20\}$ and $a \neq b$. The width of the edges indicates the number of observed interactions (see distribution under the graph) where no edges imply no observed interactions. The transparency gradient of the edges corresponds to the direction of the interaction ($y_{[a,b]}$ or $y_{[b,a]}$): the white end of an edge shows the giver, and its darker end shows the receiver. Kin groups are highlighted by dashed outlines. This network corresponds to the third level of abstraction, in Figure 1.

S1.4.2. Full description of statistical model 3

$$y_{[a,b]} \sim \text{Poisson}(m_{[a,b]})$$
$$y_{[b,a]} \sim \text{Poisson}(m_{[b,a]})$$

$$m_{[a,b]} = \exp(D + G_{[a]} + R_{[b]} + \hat{T}_{[a,b]})$$
$$m_{[b,a]} = \exp(D + G_{[b]} + R_{[a]} + \hat{T}_{[b,a]})$$

$$\hat{T}_{[a,b]} = T_{[a,b]} + b_T \cdot X_{|a,b|} + \log(S_{|a,b|})$$
$$\hat{T}_{[b,a]} = T_{[b,a]} + b_T \cdot X_{|b,a|} + \log(S_{|a,b|})$$

$$\begin{pmatrix} G_{[a]} \\ R_{[a]} \end{pmatrix} = \begin{pmatrix} s_G & 0 \\ 0 & s_R \end{pmatrix} \times L_{\text{ind}} \times \begin{pmatrix} z_{G_{[a]}} \\ z_{R_{[a]}} \end{pmatrix}$$
$$\begin{pmatrix} T_{[a,b]} \\ T_{[b,a]} \end{pmatrix} = \begin{pmatrix} s_T & 0 \\ 0 & s_T \end{pmatrix} \times L_{\text{dyad}} \times \begin{pmatrix} z_{T_{[a,b]}} \\ z_{T_{[b,a]}} \end{pmatrix}$$

$$D \sim \text{Normal}(-1, 1)$$
$$b_T, z_{G_{[a]}}, z_{R_{[a]}}, z_{T_{[a,b]}}, z_{T_{[b,a]}} \sim \text{Normal}(0, 1)$$
$$s_G, s_R, s_T \sim \text{Exponential}(1)$$
$$L_{\text{ind}}, L_{\text{dyad}} \sim \text{LKJ Cholesky}(2)$$

S1.5. Simulation study 3: alternative parameterisation

In this section, we show an alternative parameterisation for the DAG, SCM, and statistical model of simulation study 3, where we do not explicitly define γ, ρ, and τ. Instead, we encode the causes of $m_{[a,b]}$ and $y_{[a,b]}$ with arrows pointing directly into them. Here, $Re_{|a,b|}$ affects $m_{[a,b]}$, and $S_{|a,b|}$ affects $y_{[a,b]}$. As a result, $m_{[a,b]}$ corresponds to the true rate of interactions from a to b per *one* time unit. This contrasts with the main text (equations 3.1.6-3.1.7), where $m_{[a,b]}$ defines the rate of interactions per $S_{|a,b|}$ time units. The synthetic data generated from the SCM, and the fixed-effect estimates obtained from the statistical model are identical to those in the main text. This new specification, where the biological variable affects the true rate m, and the sampling effort affects the observed interactions y, might be easier to connect to the three levels of abstraction of Figure 1. It would be a good choice if we were solely interested in deriving an adjustment set from the DAG and estimating a causal effect. In the main text, we use the parameterisation with γ, ρ, and τ, for it made causal assumptions very explicit, and it allowed us to show how individual- and dyad-level structuring features could affect $G_{[a]}$, $R_{[b]}$, and $T_{[a,b]}$ (Figure 5b).

S1.5.1. Directed Acyclic Graph

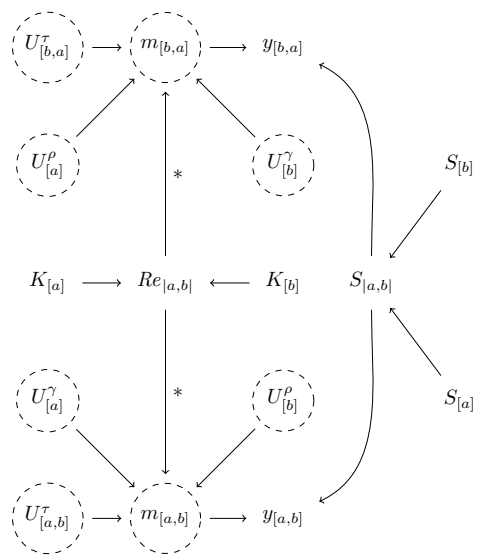

Figure S10 – *Alternative DAG for simulation study 3.*

Alternative parameterisation of SCM 3

$K_{[a]}$ is assigned to individuals following the description in the main text—*i.e.*, there are 10 individuals in one large kin group, and the other individuals are alone in their kin group.

$$
\begin{aligned}
f_{Re}: &\quad Re_{|a,b|} = \begin{cases} 1 & \text{if } K_{[a]} = K_{[b]} \\ 0 & \text{if } K_{[a]} \neq K_{[b]} \end{cases} \\
f_{S_{[a]}}: &\quad S_{[a]} \sim \text{Uniform}(1, 2.5) \\
f_{S_{|a,b|}}: &\quad S_{|a,b|} = S_{[a]} + S_{[b]}
\end{aligned}
$$

$$
\begin{aligned}
f_{U^\gamma}: &\quad U^\gamma_{[a]} \sim \text{Normal}(0, 0.5) \\
f_{U^\rho}: &\quad U^\rho_{[a]} \sim \text{Normal}(0, 0.5) \\
f_{U^\tau}: &\quad U^\tau_{[a,b]} \sim \text{Normal}(0, 0.5)
\end{aligned}
$$

$$
\begin{aligned}
f_m: &\quad m_{[a,b]} \sim \exp(-1.2 + U^\gamma_{[a]} + U^\rho_{[b]} + U^\tau_{[a,b]} + 0.8^* \cdot \tilde{Re}_{|a,b|}) \\
f_y: &\quad y_{[a,b]} \sim \text{Poisson}(m_{[a,b]} \cdot S_{|a,b|})
\end{aligned}
$$

Alternative parameterisation of statistical model 3

$$
\begin{aligned}
y_{[a,b]} &\sim \text{Poisson}(m_{[a,b]} \cdot S_{|a,b|}) \\
y_{[b,a]} &\sim \text{Poisson}(m_{[b,a]} \cdot S_{|a,b|})
\end{aligned}
$$

$$
\begin{aligned}
m_{[a,b]} &= \exp(D + G_{[a]} + R_{[b]} + T_{[a,b]} + b^*_T \cdot \tilde{Re}_{|a,b|}) \\
m_{[b,a]} &= \exp(D + G_{[b]} + R_{[a]} + T_{[b,a]} + b^*_T \cdot \tilde{Re}_{|a,b|})
\end{aligned}
$$

$$
\begin{aligned}
\begin{pmatrix} G_{[a]} \\ R_{[a]} \end{pmatrix} &= \begin{pmatrix} s_G & 0 \\ 0 & s_R \end{pmatrix} \times L_{\text{ind}} \times \begin{pmatrix} z_{G_{[a]}} \\ z_{R_{[a]}} \end{pmatrix} \\
\begin{pmatrix} T_{[a,b]} \\ T_{[b,a]} \end{pmatrix} &= \begin{pmatrix} s_T & 0 \\ 0 & s_T \end{pmatrix} \times L_{\text{dyad}} \times \begin{pmatrix} z_{T_{[a,b]}} \\ z_{T_{[b,a]}} \end{pmatrix}
\end{aligned}
$$

$$
\begin{gathered}
D \sim \text{Normal}(-1, 1) \\
b^*_T, z_{G_{[a]}}, z_{R_{[a]}}, z_{T_{[a,b]}}, z_{T_{[b,a]}} \sim \text{Normal}(0, 1) \\
s_G, s_R, s_T \sim \text{Exponential}(1) \\
L_{\text{ind}}, L_{\text{dyad}} \sim \text{LKJ Cholesky}(2)
\end{gathered}
$$

S1.6. Simulation Study 3': categorical dyad-level features

This simulation study is a variation of simulation study 3. Its causal structure (Figure S11) is identical to Figure 5a-b. However, its quantitative implementation differs. Here, the dyadic variable $X_{[a,b]}$ (equivalent of $Re_{|a,b|}$ in simulation 3) is a categorical variable, *e.g.* the combination of sexes for a given directed dyad. The effect of $X_{[a,b]}$ is parameterised as a two-by-two matrix, encoding an effect for each combination of sexes (section S1.6.2-S1.6.3). We show that a well-specified estimator (section S1.6.4) can recover the SCM's structural parameters (section S1.6.5).

S1.6.1. Directed Acyclic Graph 3'

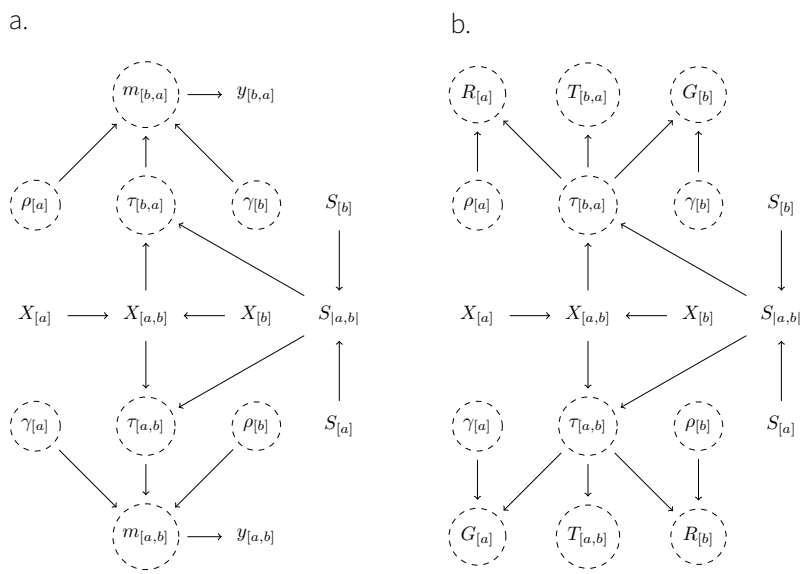

Figure S11 – *Directed Acyclic Graph for simulation study 3'.* $X_{[a]}$ represents the sex of an individual, a, and $X_{[a,b]}$, the combination of sexes for a given directed dyad. Males are encoded as 1, and females as 2. For example, if a is a male and b is female, then $X_{[a,b]} = X_{[1,2]}$. S represents the sampling effort for each individual ($S_{[a]}$) and dyad ($S_{|a,b|}$).

S1.6.2. Structural Causal Model 3'

$$
\begin{aligned}
f_{S_{[a]}}: &\quad S_{[a]} \sim \text{Uniform}(1, 2.5) \\
f_{S_{|a,b|}}: &\quad S_{|a,b|} = S_{[a]} + S_{[b]} \\
f_{X_{[a]}}: &\quad X_{[a]} \in_R \{1, 2\} \\
f_{X_{[a,b]}}: &\quad X_{[a,b]} = [X_{[a]}, X_{[b]}] \\
f_{\psi}: &\quad \psi = \begin{pmatrix} -1 & 0.5 \\ 0 & 0.5 \end{pmatrix}
\end{aligned}
$$

$$
\begin{aligned}
f_{\gamma}: &\quad \gamma_{[a]} \sim \text{Normal}(0, 0.5) \\
f_{\rho}: &\quad \rho_{[a]} \sim \text{Normal}(0, 0.5) \\
f_{\tau}: &\quad \tau_{[a,b]} \sim \text{Normal}(\log(S_{|a,b|}) + \psi_{S_{[a,b]}}, 0.5)
\end{aligned}
$$

$$
\begin{aligned}
f_{m}: &\quad m_{[a,b]} = \exp(-1.2 + \gamma_{[a]} + \rho_{[b]} + \tau_{[a,b]}) \\
f_{y}: &\quad y_{[a,b]} \sim \text{Poisson}(m_{[a,b]})
\end{aligned}
$$

S1.6.3. Observed interactions (y) simulated with from SCM 3'

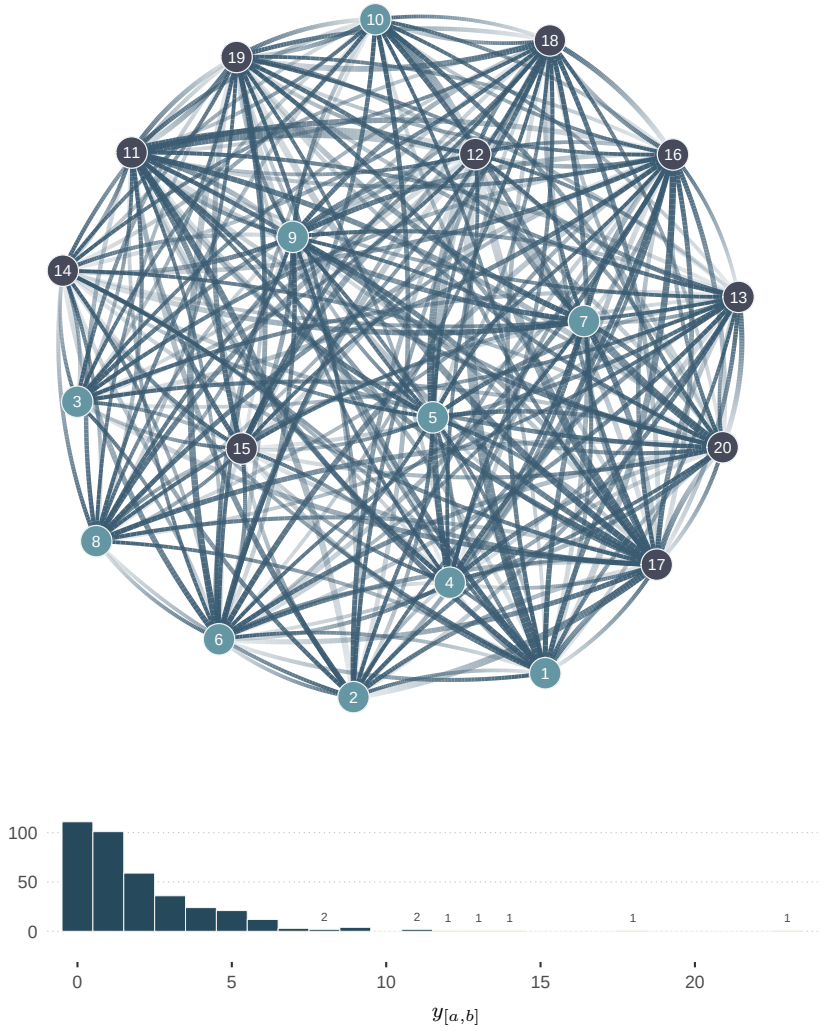

Figure S12 – *Network of observed interactions (y) generated with SCM 3'.* The graph shows 20 individuals (nodes), and the number of observed interactions $y_{[a,b]}$ among them, where a and b correspond to individual indices: $a, b \in \{1, \ldots, 20\}$ and $a \neq b$. The node colour shows the value of $X_{[a]}$: light blue for males, and dark blue for females. The width of the edges indicates the number of observed interactions (see distribution under the graph) where no edges imply no observed interactions. The transparency gradient of the edges corresponds to the direction of the interaction ($y_{[a,b]}$ or $y_{[b,a]}$): the white end of an edge shows the giver (start of the arrow), and its darker end shows the receiver (head of the arrow). This network corresponds to the third level of abstraction, in Figure 1.

S1.6.4. Full description of statistical model 3'

$$y_{[a,b]} \sim \text{Poisson}(m_{[a,b]})$$
$$y_{[b,a]} \sim \text{Poisson}(m_{[b,a]})$$

$$m_{[a,b]} = \exp(D + G_{[a]} + R_{[b]} + \hat{T}_{[a,b]})$$
$$m_{[b,a]} = \exp(D + G_{[b]} + R_{[a]} + \hat{T}_{[b,a]})$$

$$B = \begin{pmatrix} b_{[1,1]} & b_{[1,2]} \\ b_{[2,1]} & b_{[2,2]} \end{pmatrix}$$

$$\hat{T}_{[a,b]} = \log(S_{|a,b|}) + T_{[a,b]} + B_{X_{[a,b]}}$$
$$\hat{T}_{[b,a]} = \log(S_{|a,b|}) + T_{[b,a]} + B_{X_{[b,a]}}$$

$$\begin{pmatrix} G_{[a]} \\ R_{[a]} \end{pmatrix} = \begin{pmatrix} s_G & 0 \\ 0 & s_R \end{pmatrix} \times L_{\text{ind}} \times \begin{pmatrix} z_{G_{[a]}} \\ z_{R_{[a]}} \end{pmatrix}$$
$$\begin{pmatrix} T_{[a,b]} \\ T_{[b,a]} \end{pmatrix} = \begin{pmatrix} s_T & 0 \\ 0 & s_T \end{pmatrix} \times L_{\text{dyad}} \times \begin{pmatrix} z_{T_{[a,b]}} \\ z_{T_{[b,a]}} \end{pmatrix}$$

$$D \sim \text{Normal}(-1, 1)$$
$$b_j, z_{G_{[a]}}, z_{R_{[a]}}, z_{T_{[a,b]}}, z_{T_{[b,a]}} \sim \text{Normal}(0, 1)$$
$$s_G, s_R, s_T \sim \text{Exponential}(1)$$
$$L_{\text{ind}}, L_{\text{dyad}} \sim \text{LKJ Cholesky}(2)$$

S1.6.5. Posterior models (simulation study 3')

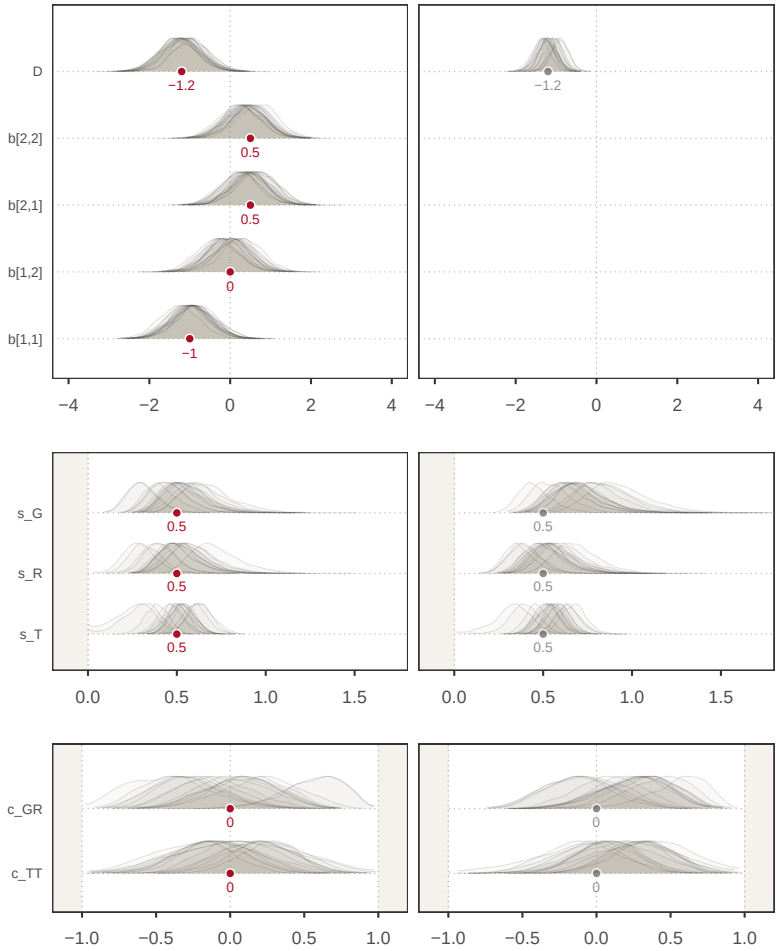

Figure S13 – *Marginal posterior distributions of the fixed effects of two different statistical models fitted to the data generated with SCM 3'.* Left: fixed effects of the social relations model adjusted by $X_{[a,b]}$ and $S_{|a,b|}$ (see section S1.6.4). Right: fixed effects of the estimator adjusted by $S_{|a,b|}$ only. This second statistical model is identical to the one in section S1.6.4, but without the $B_{X_{[a,b]}}$ parameters.

S1.7. Simulation Study 4: kinship in female macaques

S1.7.1. Observed interactions (y) simulated with SCM 4

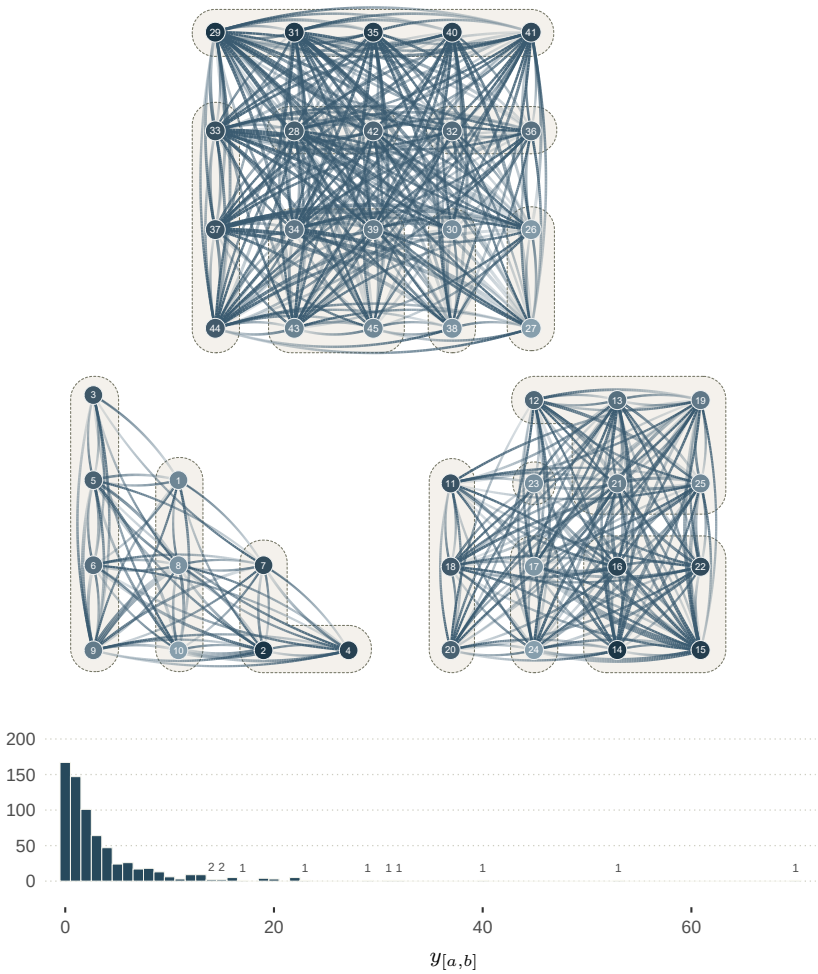

Figure S14 – *Network of observed interactions (y) generated with SCM 4.* The graph shows 45 individuals (nodes) across three groups of 10, 15, and 20 individuals, respectively. The edges show the interactions $y_{[a,b]}$ among them. Their width indicates the number of observed interactions (see distribution under the graphs) where no edges imply no observed interactions. The transparency gradient of the edges corresponds to the direction of the interaction ($y_{[a,b]}$ or $y_{[b,a]}$): the white end of an edge shows the giver (start of the arrow), and its darker end shows the receiver (head of the arrow). The colour of the node corresponds to the individual dominance rank: lighter for low ranks, and darker for higher ranks. Kin group are highlighted by dashed outlines. This network corresponds to the third level of abstraction, in Figure 1.

S1.7.2. Full description of statistical model 4

$$y_{[a,b]} \sim \text{Poisson}(m_{[a,b]})$$
$$y_{[b,a]} \sim \text{Poisson}(m_{[b,a]})$$

$$m_{[a,b]} = \exp(D + G_{[a]} + R_{[b]} + \hat{T}_{[a,b]})$$
$$m_{[b,a]} = \exp(D + G_{[b]} + R_{[a]} + \hat{T}_{[b,a]})$$

$$\hat{T}_{[a,b]} = T_{[a,b]} + b_{Re} \cdot Re_{|a,b|} + \log(S_{|a,b|}) + \begin{cases} b_{Ra_1} \cdot (Ra_{[b]} - Ra_{[a]}) & \text{if } Ra_{[a]} < Ra_{[b]} \\ b_{Ra_2} \cdot (Ra_{[b]} - Ra_{[a]}) & \text{if } Ra_{[a]} \geq Ra_{[b]} \end{cases}$$
$$\hat{T}_{[b,a]} = T_{[b,a]} + b_{Re} \cdot Re_{|a,b|} + \log(S_{|a,b|}) + \begin{cases} b_{Ra_1} \cdot (Ra_{[a]} - Ra_{[b]}) & \text{if } Ra_{[b]} < Ra_{[a]} \\ b_{Ra_2} \cdot (Ra_{[a]} - Ra_{[b]}) & \text{if } Ra_{[b]} \geq Ra_{[a]} \end{cases}$$

$$\begin{pmatrix} G_{[a]} \\ R_{[a]} \end{pmatrix} = \begin{pmatrix} s_G & 0 \\ 0 & s_R \end{pmatrix} \times L_{\text{ind}} \times \begin{pmatrix} z_{G_{[a]}} \\ z_{R_{[a]}} \end{pmatrix}$$
$$\begin{pmatrix} T_{[a,b]} \\ T_{[b,a]} \end{pmatrix} = \begin{pmatrix} s_T & 0 \\ 0 & s_T \end{pmatrix} \times L_{\text{dyad}} \times \begin{pmatrix} z_{T_{[a,b]}} \\ z_{T_{[b,a]}} \end{pmatrix}$$

$$D \sim \text{Normal}(-1.5, 1)$$
$$z_{G_{[a]}}, z_{R_{[a]}}, z_{T_{[a,b]}}, z_{T_{[b,a]}} \sim \text{Normal}(0, 1)$$
$$b_{Re}, b_{Ra_1}, b_{Ra_2} \sim \text{Normal}(0, 0.6)$$
$$s_G, s_R, s_T \sim \text{Exponential}(1)$$
$$L_{\text{ind}}, L_{\text{dyad}} \sim \text{LKJ Cholesky}(3)$$

S1.7.3. Prior predictive simulations

In this section, we describe the prior predictive checks that we conducted for statistical model 4. Using R, we started by drawing n samples for each of the nine population-level parameters (fixed effects):

$$
\begin{aligned}
s_G^{(n)}, s_R^{(n)}, s_T^{(n)} &\sim \text{Exponential}(1) \qquad\qquad n \in \{1, \ldots, 1000\} \\
c_{GR}^{(n)}, c_{TT}^{(n)} &\sim \text{LKJ}(\eta = 3, 2) \\
b_{Re}^{(n)}, b_{Ra_1}^{(n)}, b_{Ra_2}^{(n)} &\sim \text{Normal}(0, 0.6) \\
D^{(n)} &\sim \text{Normal}(-1.5, 1)
\end{aligned}
$$

Where the *sub*scripts distinguish the parameters from one another, and the *super*script (n) is an index for the prior draws. Using these parameters, we defined two prior variance-covariance matrices Σ: one for the individual-level and one for the dyad-level varying effects.

$$
\begin{aligned}
\Sigma_{\text{ind}}^{(n)} &= \begin{pmatrix} s_G^{(n)} s_G^{(n)} & c_{GR}^{(n)} s_G^{(n)} s_R^{(n)} \\ c_{GR}^{(n)} s_G^{(n)} s_R^{(n)} & s_R^{(n)} s_R^{(n)} \end{pmatrix} \\
\Sigma_{\text{dyad}}^{(n)} &= \begin{pmatrix} s_T^{(n)} s_T^{(n)} & c_{TT}^{(n)} s_T^{(n)} s_T^{(n)} \\ c_{TT}^{(n)} s_T^{(n)} s_T^{(n)} & s_T^{(n)} s_T^{(n)} \end{pmatrix}
\end{aligned}
$$

Next, we constructed the three groups containing, respectively, 10, 15, and 20 individuals (45 individuals in total), resulting in 340 dyads and 680 directed dyads across the three groups. For each prior sample n, we drew 20 individual-level parameters $G_{[a]}$, 20 individual-level parameters $R_{[a]}$, and 680 dyad-level parameters $T_{[a,b]}$ and $T_{[b,a]}$ (340 in each direction).

$$
\begin{aligned}
\begin{pmatrix} G_{[a]}^{(n)} \\ R_{[a]}^{(n)} \end{pmatrix} &\sim \text{MVNormal}\left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \Sigma_{\text{ind}}^{(n)}\right] \\
\begin{pmatrix} T_{[a,b]}^{(n)} \\ T_{[b,a]}^{(n)} \end{pmatrix} &\sim \text{MVNormal}\left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \Sigma_{\text{dyad}}^{(n)}\right]
\end{aligned}
$$

For each sample n, we also assigned a rank to each individual, and a value of genetic relatedness for each dyad:

$$
\begin{aligned}
Ra_{[a]}^{(n)} &\sim \text{Uniform}(0, 1) \\
Re_{|a,b|}^{(n)} &\in_R \{0, 1\}
\end{aligned}
$$

Finally, we deterministically obtained an expected rate $m_{[a,b]}^{(n)}$, per directed dyad and

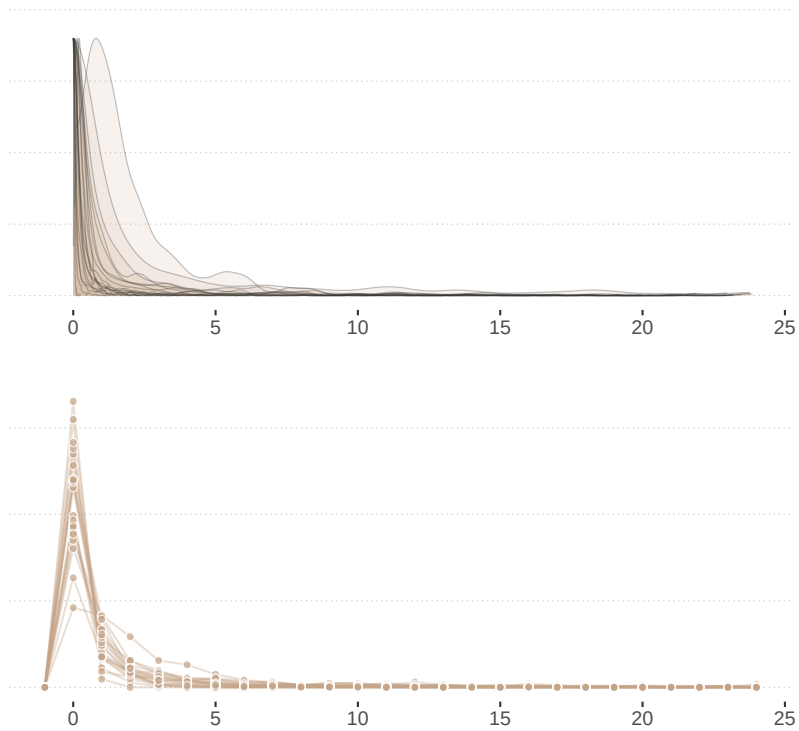

Figure S15 – ***Top:*** *Distribution of prior average interaction rates* $m^{(n)}_{[a,b]}$. Each slab represents the distribution of 680 directed rates, for one prior draw n. We show the distribution for 20 prior samples n, *i.e.* 20 slabs. ***Bottom:*** *Distribution of prior predictions* $y^{(n)}_{[a,b]}$. Each sequence of dots connected by one line represent the counts of observed interactions $y_{[a,b]}$ across 680 directed dyads, for one prior sample n. Again, we show 20 distributions, one for each prior sample n.

per prior sample.

$$\hat{T}^{(n)}_{[a,b]} = T^{(n)}_{[a,b]} + b^{(n)}_{Re} \cdot Re^{(n)}_{|a,b|} + \begin{cases} b^{(n)}_{Ra_1} \cdot (Ra^{(n)}_{[b]} - Ra^{(n)}_{[a]}) & \text{if } R^{(n)}_{[a]} < R^{(n)}_{[b]} \\ b^{(n)}_{Ra_1} \cdot (Ra^{(n)}_{[b]} - Ra^{(n)}_{[a]}) & \text{if } R^{(n)}_{[a]} \geq R^{(n)}_{[b]} \end{cases}$$

$$m^{(n)}_{[a,b]} = \exp(D^{(n)} + G^{(n)}_{[a]} + R^{(n)}_{[b]} + \hat{T}^{(n)}_{[a,b]})$$

We show the distribution of $m^{(n)}_{[a,b]}$ and $m^{(n)}_{[b,a]}$ for 20 (out of 1000) prior draws in Figure S15 (top panel). From this rate, we drew a prior prediction y per directed dyad and per prior sample (Figure S15, bottom panel):

$$y^{(n)}_{[a,b]} \sim \text{Poisson}(m^{(n)}_{[a,b]}).$$

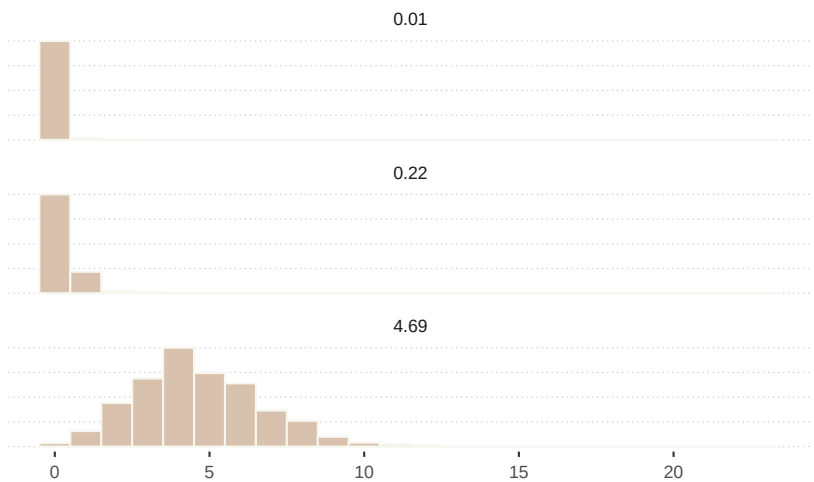

Figure S16 – Distribution of predicted observations $y^{(n)}_{[a,b]}$ for dyads with low-, medium-, and high-rates $\hat{m}_k$ of interactions.

To better understand what the prior model implies, we summarised the distribution of rates $m^{(n)}_{[a,b]}$ in terms of its 10%, 50%, and 90% percentiles. We call these summary values $\hat{m}_k$. They take the following values: $\hat{m}_1 = 0.01$, $\hat{m}_2 = 0.22$, $\hat{m}_3 = 4.69$. That is, the prior *average* interaction rates per directed dyad roughly ranges from 1 interaction every 100 days, to about 4 or 5 interaction per day. We then drew 1000 prior predictions y for each value $\hat{m}_k$ (Figure S16).

$$y^{(n)}_k \sim \text{Poisson}(\hat{m}_k).$$

In summary, our prior model is compatible with a broad range of dyadic grooming interaction rates: dyadic rates of almost zero, as well as rates of several interactions per day, are plausible before the estimator it is updated with data.

To visualise the prior distribution of slopes, we started by computing the variation across directed dyads due to unobserved network structuring features:

$$\psi^{(n)}_{[a,b]} = D^{(n)} + G^{(n)}_{[a]} + R^{(n)}_{[b]} + T^{(n)}_{[a,b]}$$

We selected 24 samples n (out of 1000), and computed an interaction rate $m^{(n,j)}_{[a,b]}$ for the whole range of a predictor variable $X^{(j)}$:

$$m^{(n,j)}_{[a,b]} = \exp(\psi^{(n)}_{[a,b]} + b^{(n)}_{Re} \cdot X^{(j)})$$

Where $X^{(j)}$ is a vector of size 101, ranging from 0 to 1. $X^{(1)} = 0$, $X^{(2)} = 0.01$, $X^{(3)} = 0.02$, and so on, until $X^{(101)} = 1$. X represents either $Re_{|a,b|}$ or $\Delta Ra_{|a,b|}$

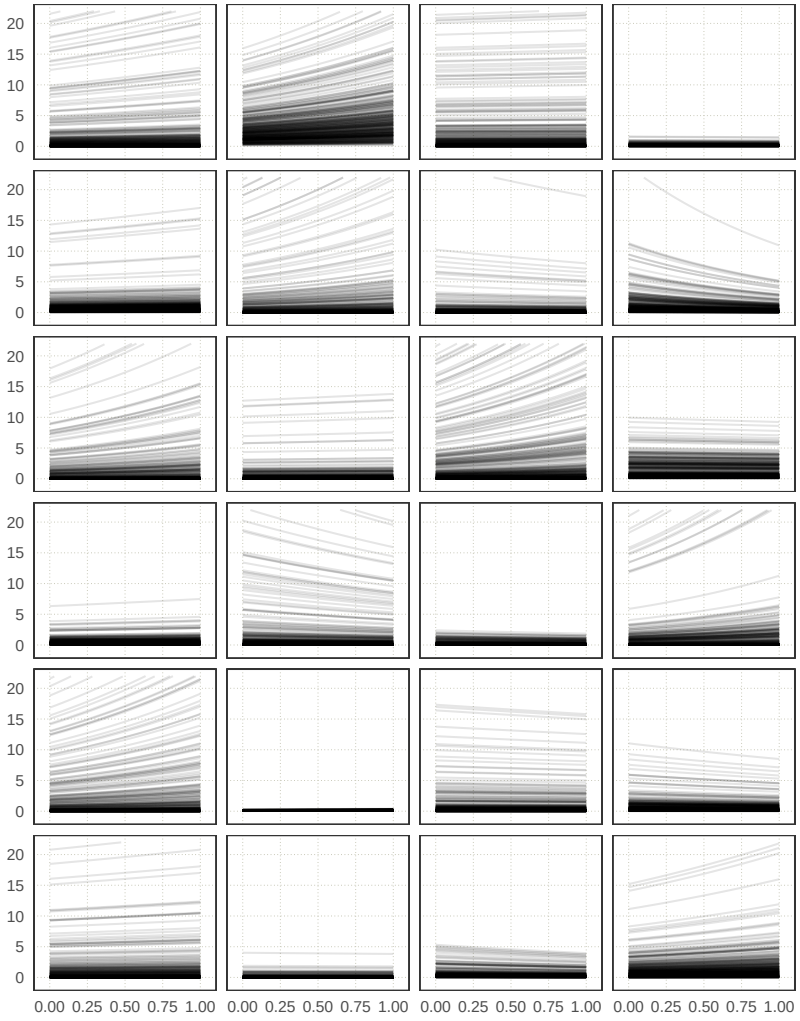

Figure S17 – *Distribution of prior slopes for* 24 *prior samples.* The x-axis represents the range of a dyad-level predictor variable (difference in rank or genetic relatedness), the y-axis represents $m_{[a,b]}$. Each panel corresponds to one prior sample n, and contains 680 slopes (one for each directed-dyad).

(in the case of $Re_{|a,b|}$, only the first and last values of the vector apply). Note that $b_{Re}^{(n)}$ is interchangeable with $b_{Ra_1}^{(n)}$ and $b_{Ra_2}^{(n)}$, for they have the same prior distribution. We plot the resulting prior slopes in Figure S17. Within each panel, we can visualise the difference across dyads that exists for a given prior sample n—it is, in this regard, complimentary to Figure S15. Across the 20 panels, we observe the variation implied by the prior probability distribution of the slope.

S1.8. Empirical Study

The plots shown in this section either represent the empirical data directly (section S1.8.1), or show the output of models that were updated with the empirical data (sections S1.8.2-S1.8.5).

S1.8.1. Distribution of sampling effort and observed interaction rates

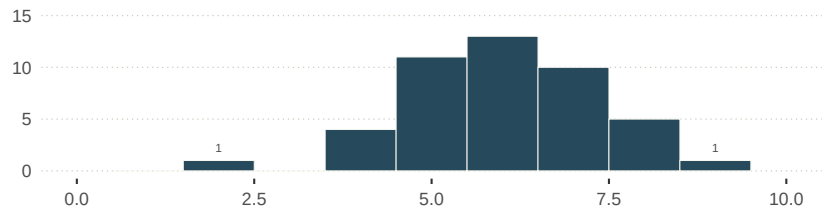

Figure S18 – *Distribution of sampling effort $S_{[a]}$, in full days, across 45 individuals.* The unit of $S_{[a]}$ is 12 hours. The lowest value of $S_{[a]}$ is 1.82 days (22 hours), and the highest is 9.02 days (108.3 hours).

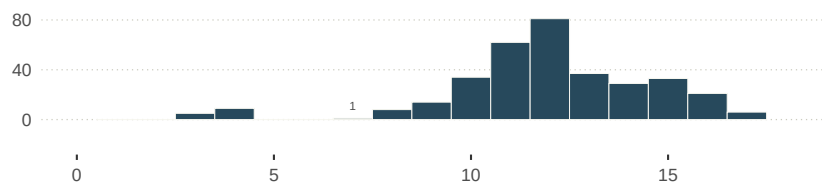

Figure S19 – *Distribution of sampling effort $S_{|a,b|}$, in full days, across 340 dyads.* The unit of $S_{|a,b|}$ is 12 hours. The lowest value of $S_{|a,b|}$ is 3.1 days (37.6 hours), and the highest is 17.3 days (207.9 hours).

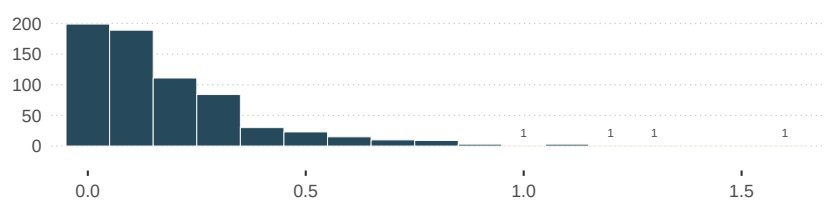

Figure S20 – *Distribution of observed interaction rates:* $(y_{[a,b]}/S_{|a,b|})$. The unit of $S_{|a,b|}$ is 12 hours.

S1.8.2. MCMC diagnostics

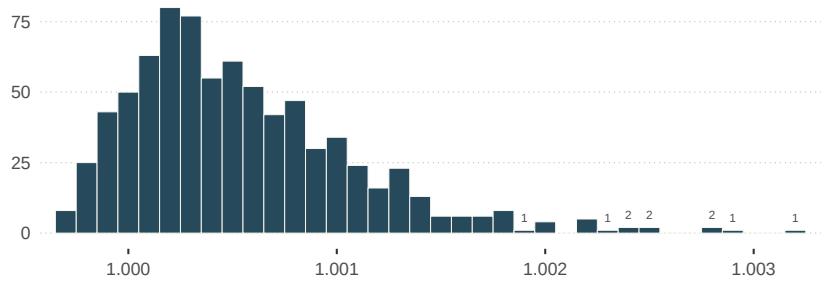

Figure S21 – *Distribution of Gelman-Rubin $\hat{R}$ across all parameters of statistical model 4 fitted to the empirical data.*

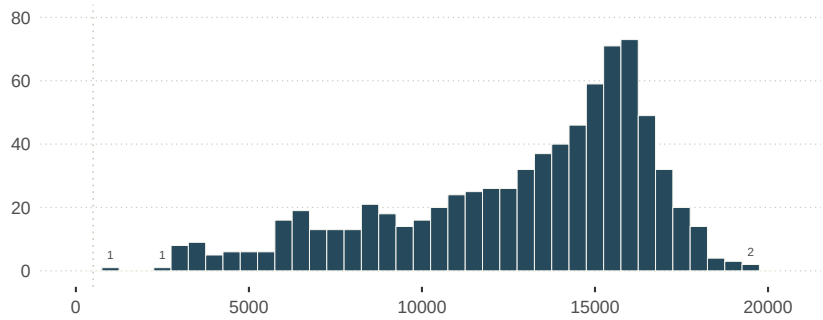

Figure S22 – *Distribution of the effective number of MCMC samples (bulk), across all parameters of statistical model 4 fitted to the empirical data.* The vertical line marks 500.

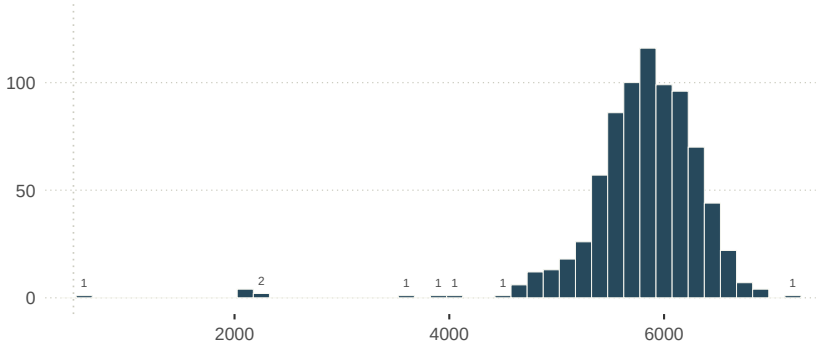

Figure S23 – *Distribution of the effective number of MCMC samples (tail), across all parameters of statistical model 4 fitted to the empirical data.* The vertical line marks 500.

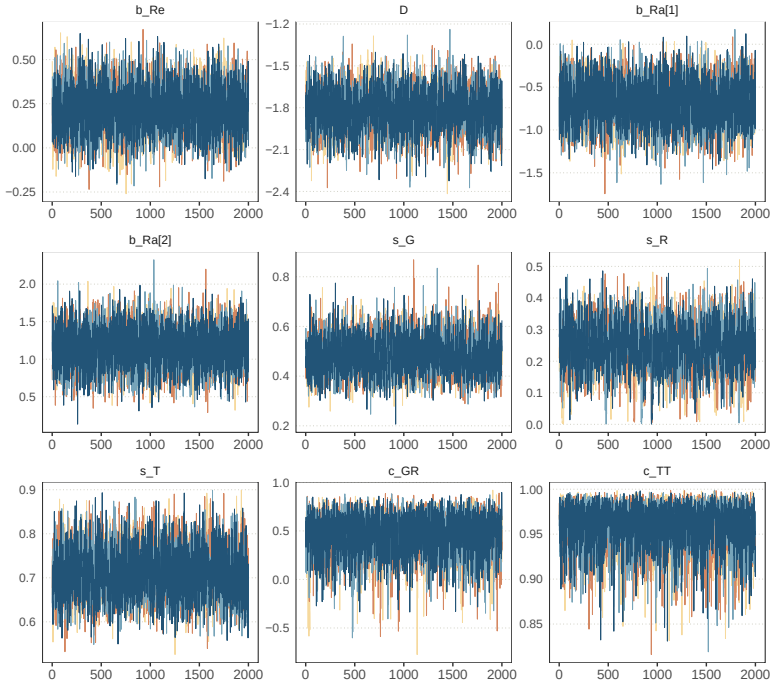

Figure S24 – *MCMC traceplots for the model's fixed effects.* Each of the four chains (four different colours) was run for 4000 iterations: 2000 warm-up draws, and 2000 post warm-up draws (shown here).

S1.8.3. Posterior predictive checks

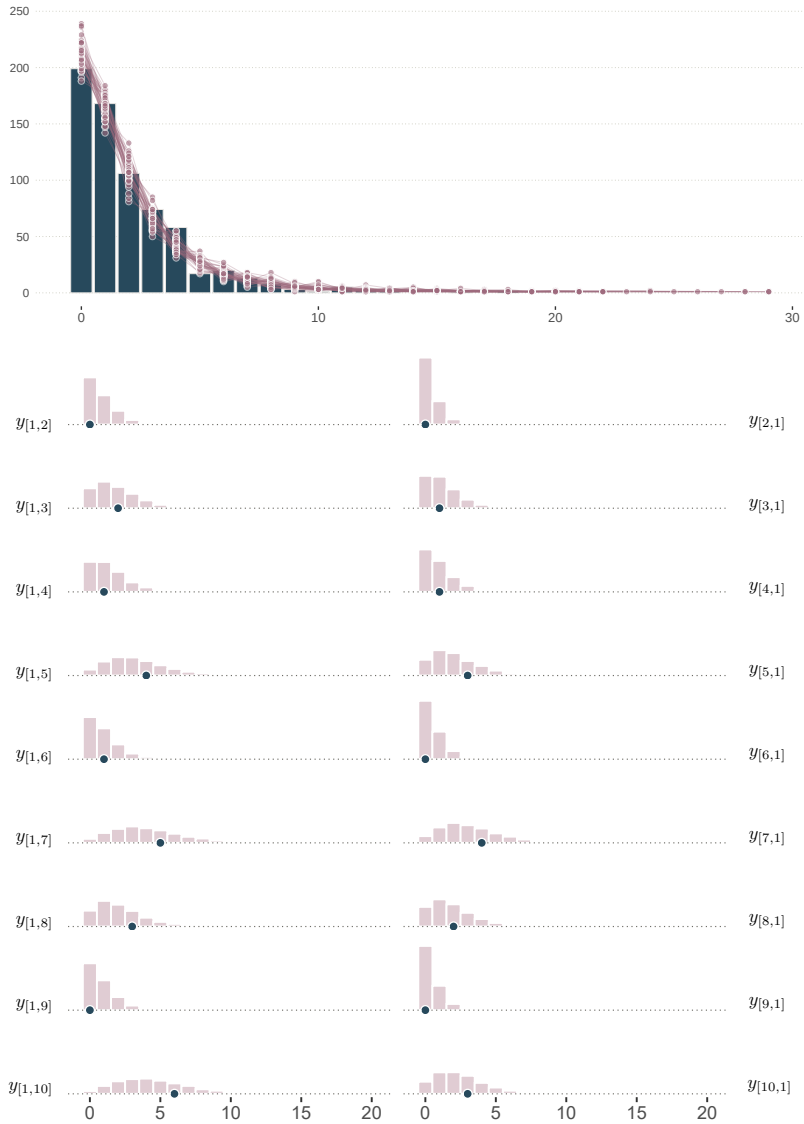

Figure S25 – Top: *Posterior predictive distribution of* $y_{[a,b]}$. The empirical distribution is shown in blue (same as in Figure 7c), and the posterior predictive distribution across all dyads is shown in pink. Predictions from the same posterior sample (40 samples in total) are connected by a line. Bottom: *Dyad-specific posterior predictive distribution of* $y_{[a,b]}$*, for 9/340 dyads.* The pink histograms represent posterior predictions for each dyad. The observed values of $y_{[a,b]}$ are shown in blue.

S1.8.4. Posterior distribution of baseline levels $\psi_{[a,b]}$

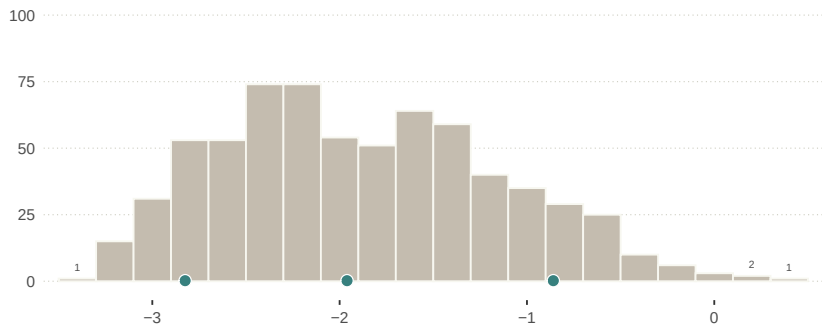

Figure S26 – *Distribution of $\psi_{[a,b]}$'s posterior means.* The blue dots indicate, from left to right: ψ_1, ψ_2, and ψ_3.

S1.8.5. Effect of dominance rank

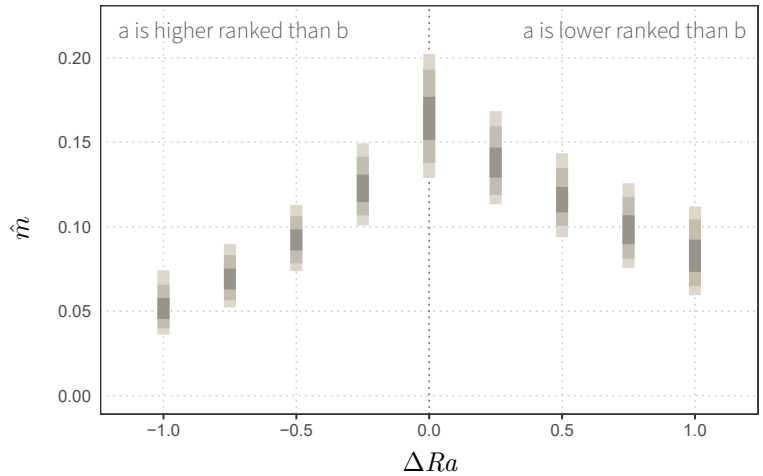

Figure S27 – *Effect of dominance rank on m.* Posterior draws of rates $\hat{m}_k$ (bouts/day), for average dyads only varying in ΔRa. The intervals of increasing density respectively contain 89%, 75%, and 40% of posterior samples.

To estimate the effect of dominance rank, we compute nine posterior distributions $\hat{m}_k^{(n)}$ (Figure S27):

$$\hat{m}_k^{(n)} = D^{(n)} + \begin{cases} b_{Ra_1}^{(n)} \cdot \Delta Ra_k & \text{if } \Delta Ra_k < 0 \\ b_{Ra_2}^{(n)} \cdot \Delta Ra_k & \text{if } \Delta Ra_k \geq 0, \end{cases}$$

Where n corresponds to 1000 posterior samples and ΔRa_k are differences in dominance rank ranging from -1 to 1, for $k \in \{1, \ldots, 9\}$. Note that we slightly abuse notation, here. We also named some parameters $\hat{m}_k$ in the section dedicated to prior predictive simulations (section S1.7.3) but $\hat{m}_k$ refer to a *different* set of parameters here.

How to interpret Figure S27 given the assumptions encoded in the statistical (section S1.7.2) and causal (Figure 6a) models:

The parameters $\hat{m}_k$ represent counterfactual rates for an average dyad ($G_{[a]} = R_{[b]} = T_{[a,b]} = 0$) composed of non-kin ($Re_{|a,b|} = 0$), in 9 different worlds changing only in their values of ΔRa. These rates allow us to isolate the causal effect of dominance rank, even if it is generically confounded by relatedness in nature. When an individual, a, faces a higher-ranked individual, b, its rate of grooming decreases: the bigger the distance in rank, the bigger the decrease (right panel). Likewise, when an individual, a, faces an individual, b, that is lower ranked, its grooming rate decreases as well, such that a larger distance implies a larger decrease (left panel). Together, this model suggests that rank causes individuals to groom one another more frequently when they are proximate in the dominance hierarchy.

S1.9. Simulation Study 4': group-level effects

Below, we present an slightly more realistic version of simulation study 4. The new elements of simulation study 4' regard group size $GS_{\text{gr}_{|a,b|}}$. This group-level variable is now assumed to affect interaction rates, such that in smaller groups, dyads tend to have a higher rate of social interactions. We conclude this section by fitting the resulting estimator (statistical model 4') to the empirical dataset.

S1.9.1. Causal Models 4'

$\delta_{\text{gr}_{|a,b|}}$ encodes the causal effect of group-level variables like group size on $m_{[a,b]}$ (Figure S28a). δ is thus analogous to γ, ρ, and τ, but for group-level effects.

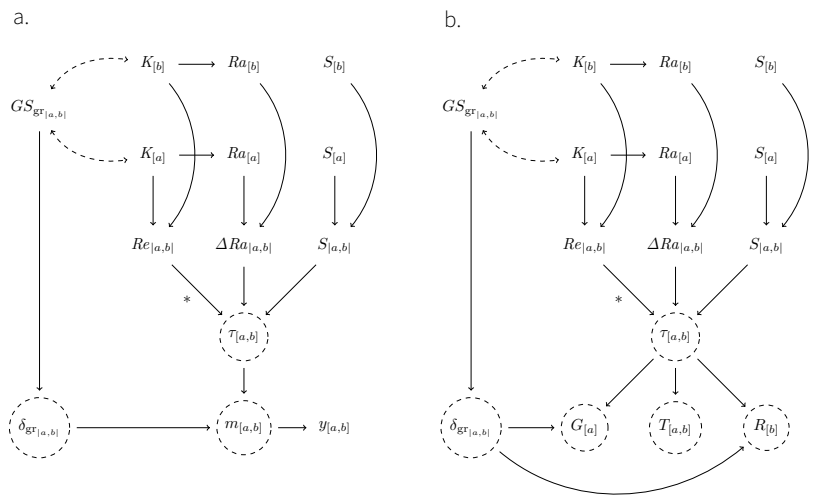

Figure S28 – *Directed Acyclic Graph for simulation study 4'.* The graph is similar to Figure 6a, except for the effect of group size $GS_{\text{gr}_{|a,b|}}$. Here, $GS_{\text{gr}_{|a,b|}}$ determines the value of $\delta_{\text{gr}_{|a,b|}}$, which then affects the rate $m_{[a,b]}$ for all dyads in the group of individuals a and b.

SCM 4' is very similar to SCM 4. The only differences regard the global intercept, and are encoded through f_δ and f_m:

$$f_\delta: \quad \delta_{\text{gr}_{|a,b|}} = \begin{cases} -0.8 \text{ if } GS_{\text{gr}_{|a,b|}} = 10 \\ -1.2 \text{ if } GS_{\text{gr}_{|a,b|}} = 15 \\ -1.6 \text{ if } GS_{\text{gr}_{|a,b|}} = 20 \end{cases}$$

$$f_m: \quad m_{[a,b]} = \exp(\delta_{\text{gr}_{|a,b|}} + \gamma_{[a]} + \rho_{[b]} + \tau_{[a,b]}).$$

S1.9.2. Observed interactions (y) simulated with SCM 4'

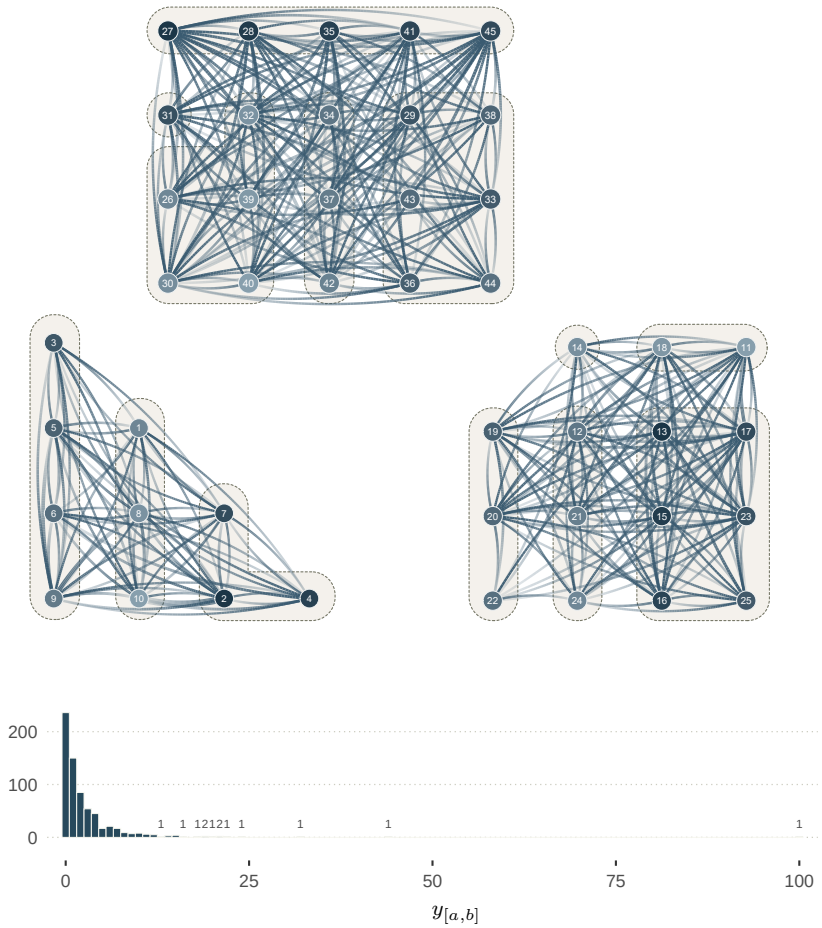

Figure S29 – *Network of observed interactions (y) generated with SCM 4'.* The graph shows 45 individuals (nodes) across three groups of 10, 15, and 20 individuals, respectively. The edges show the interactions $y_{[a,b]}$ among them. Their width indicates the number of observed interactions (see distribution under the graphs) where no edges imply no observed interactions. The transparency gradient of the edges corresponds to the direction of the interaction ($y_{[a,b]}$ or $y_{[b,a]}$): the white end of an edge shows the giver (start of the arrow), and its darker end shows the receiver (head of the arrow). The colour of the node corresponds to the individual dominance rank: lighter for low ranks, and darker for higher ranks. Kin group are highlighted by dashed outlines. This network corresponds to the third level of abstraction, in Figure 1.

S1.9.3. Full description of statistical model 4'

$$y_{[a,b]} \sim \text{Poisson}(m_{[a,b]})$$
$$y_{[b,a]} \sim \text{Poisson}(m_{[b,a]})$$

$$m_{[a,b]} = \exp(d_{[\text{gr}_{|a,b|}]} + G_{[a]} + R_{[b]} + \hat{T}_{[a,b]})$$
$$m_{[b,a]} = \exp(d_{[\text{gr}_{|a,b|}]} + G_{[b]} + R_{[a]} + \hat{T}_{[b,a]})$$

$$\hat{T}_{[a,b]} = T_{[a,b]} + b_{Re} \cdot Re_{|a,b|} + \log(S_{|a,b|}) + \begin{cases} b_{Ra_1} \cdot (Ra_{[b]} - Ra_{[a]}) & \text{if } Ra_{[a]} < Ra_{[b]} \\ b_{Ra_2} \cdot (Ra_{[b]} - Ra_{[a]}) & \text{if } Ra_{[a]} \geq Ra_{[b]} \end{cases}$$
$$\hat{T}_{[b,a]} = T_{[b,a]} + b_{Re} \cdot Re_{|a,b|} + \log(S_{|a,b|}) + \begin{cases} b_{Ra_1} \cdot (Ra_{[a]} - Ra_{[b]}) & \text{if } Ra_{[b]} < Ra_{[a]} \\ b_{Ra_2} \cdot (Ra_{[a]} - Ra_{[b]}) & \text{if } Ra_{[b]} \geq Ra_{[a]} \end{cases}$$

$$\begin{pmatrix} G_{[a]} \\ R_{[a]} \end{pmatrix} = \begin{pmatrix} s_G & 0 \\ 0 & s_R \end{pmatrix} \times L_{\text{ind}} \times \begin{pmatrix} z_{G_{[a]}} \\ z_{R_{[a]}} \end{pmatrix}$$
$$\begin{pmatrix} T_{[a,b]} \\ T_{[b,a]} \end{pmatrix} = \begin{pmatrix} s_T & 0 \\ 0 & s_T \end{pmatrix} \times L_{\text{dyad}} \times \begin{pmatrix} z_{T_{[a,b]}} \\ z_{T_{[b,a]}} \end{pmatrix}$$

$$d_{[1]}, d_{[2]}, d_{[3]} \sim \text{Normal}(-1.5, 1)$$
$$z_{G_{[a]}}, z_{R_{[a]}}, z_{T_{[a,b]}}, z_{T_{[b,a]}} \sim \text{Normal}(0, 1)$$
$$b_{Re}, b_{Ra_1}, b_{Ra_2} \sim \text{Normal}(0, 0.6)$$
$$s_G, s_R, s_T \sim \text{Exponential}(1)$$
$$L_{\text{ind}}, L_{\text{dyad}} \sim \text{LKJ Cholesky}(3)$$

S1.9.4. Posterior models: synthetic data (simulation study 4')

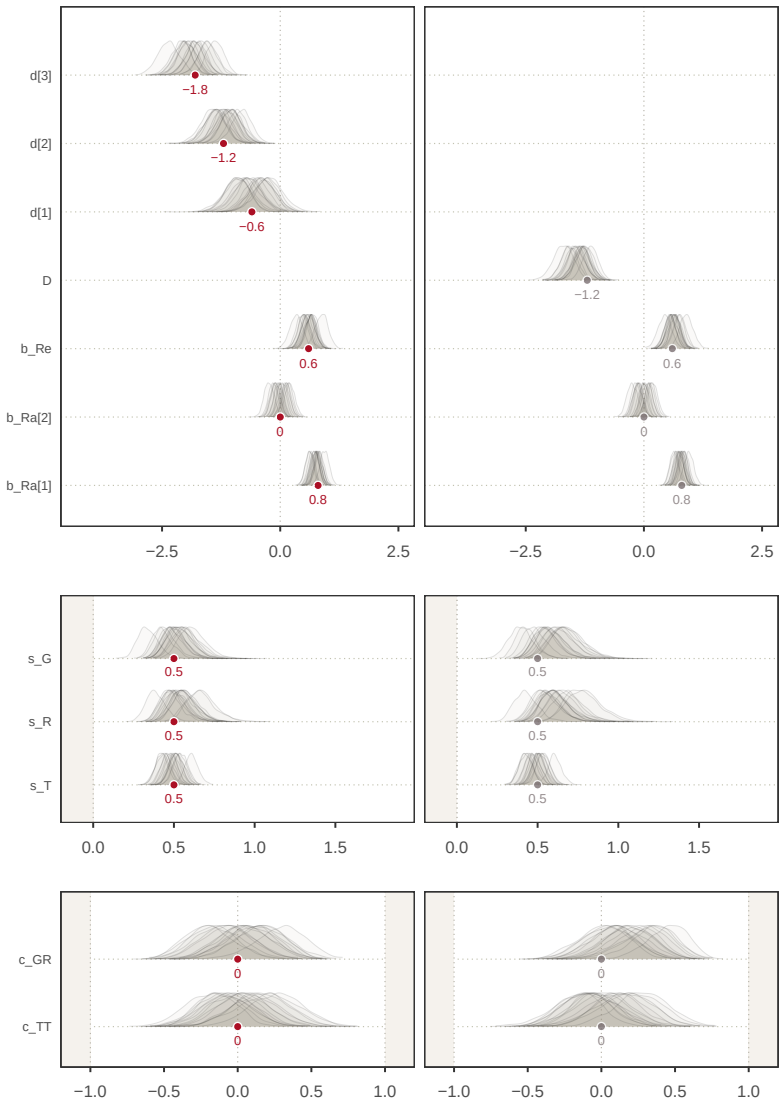

Figure S30 – *Simulation study 4'.* **b.** *Left*: fixed effects of the social relations model adjusted by $Re_{|a,b|}$, $S_{|a,b|}$, $\Delta Ra_{|a,b|}$, and group gr (statistical model 4'). *Right*: fixed effects of the social relations model adjusted by $Re_{|a,b|}$, $S_{|a,b|}$, and $\Delta Ra_{|a,b|}$ (statistical model 4). The target values of the fully-adjusted model are shown in grey.

The path that we wish to estimate ($Re_{|a,b|}$ to $m_{[a,b]}$) is confounded by group size $GS_{\mathrm{gr}_{|a,b|}}$. However, an examination of Figure S30 (right panel) tells us that statistical model 4 could recover the true causal effect of $Re_{|a,b|}$, even if it did not include group size. This is because in SCM 4', $GS_{\mathrm{gr}_{|a,b|}}$ causes individuals belonging to larger groups to give, and receive, fewer interactions. These *inter-individual* differences (unit level) are fully captured by G and R (Figures S28b and S30, right panel), thereby blocking the biasing effect of group size on the rates m (sub-unit level).

S1.10. Empirical Study'

Here, we show the posterior distributions of statistical model 4' (see sections S1.9.1-S1.9.4), fitted to the empirical data.

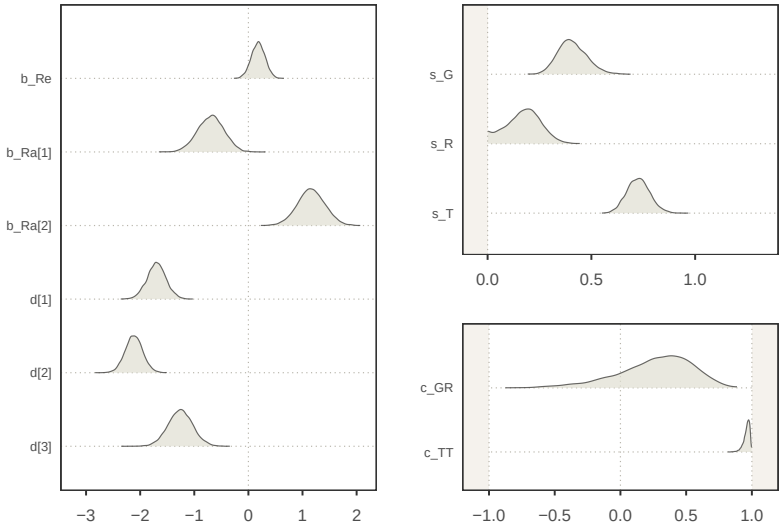

Figure S31 – *Marginal posterior distribution of the fixed effects of statistical model 4', updated with empirical data from female macaques.*

References

Albers, P. C., & de Vries, H. (2001). Elo-rating as a tool in the sequential estimation of dominance strengths. *Animal Behaviour*, *61*(2), 489–495.

Altmann, J. (1974). Observational study of behavior: Sampling methods. *Behaviour*, *49*(3-4), 227–266.

Arif, S., & MacNeil, M. A. (2022). Predictive models aren't for causal inference. *Ecology Letters*, *25*(8), 1741–1745.

Back, M. D., Branje, S., Eastwick, P. W., Human, L. J., Penke, L., Sadikaj, G., Slatcher, R. B., Thielmann, I., van Zalk, M. H., & Wrzus, C. (2023). Personality and social relationships: What do we know and where do we go. *Personality Science*, *4*(1), e7505.

Back, M. D., & Kenny, D. A. (2010). The social relations model: How to understand dyadic processes. *Social and Personality Psychology Compass*, *4*(10), 855–870.

Berrie, L., Arnold, K. F., Tomova, G. D., Gilthorpe, M. S., & Tennant, P. W. (2022). Depicting deterministic variables within directed acyclic graphs (DAGs): An aid for identifying and interpreting causal effects involving tautological associations, compositional data, and composite variables. *arXiv preprint*.

Blitzstein, J. K., & Hwang, J. (2019). *Introduction to probability*. Crc Press.

Block, P. (2015). Reciprocity, transitivity, and the mysterious three-cycle. *Social Networks*, *40*, 163–173.

Brask, J. B., Koher, A., Croft, D. P., & Lehmann, S. (2023). Linking social network structure and function to social preferences. *arXiv preprint arXiv:2303.08107*.

Bürkner, P.-C. (2017). brms: An R package for Bayesian multilevel models using Stan. *Journal of statistical software*, *80*, 1–28.

Byrnes, J. E., & Dee, L. E. (2024). Causal inference with observational data and unobserved confounding variables. *bioRxiv*, 2024–02.

Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M. A., Guo, J., Li, P., & Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of statistical software*, *76*.

Carrington, P. J., Scott, J., & Wasserman, S. (2005). *Models and methods in social network analysis* (Vol. 28). Cambridge university press.

Carter, A. J., Feeney, W. E., Marshall, H. H., Cowlishaw, G., & Heinsohn, R. (2013). Animal personality: What are behavioural ecologists measuring? *Biological reviews*, *88*(2), 465–475.

Carter, A. J., Lee, A. E., & Marshall, H. H. (2015). Research questions should drive edge definitions in social network studies. *Animal Behaviour*, *104*, E7–E11.

Cartwright, N. (1995). Précis of Nature's capacities and their measurement.

Češnovar, J., Bales, B., Morris, M., Popov, M., & Lawrence, M. (2021). Cmdstanr: a lightweight interface to Stan for R users. *R package version 0.3. 0*.

Chapais, B. (1983). Dominance, relatedness and the structure of female relationships in rhesus monkeys. *Primate social relationships: An integrated approach*, 208–219.

Cheney, D. L., & Seyfarth, R. M. (2007). *Baboon metaphysics: The evolution of a social mind*. University of Chicago Press.

Chepko-Sade, B. D., & Sade, D. S. (1979). Patterns of group splitting within matrilineal kinship groups: A study of social group structure in *Macaca mulatta* (cercopithecidae: Primates). *Behavioral Ecology and Sociobiology*, 67–86.

Croft, D. P., James, R., & Krause, J. (2008). *Exploring animal social networks*. Princeton University Press.

Csardi, G. (2010). Package 'igraph'. *CRAN1–187*.

Davies, N. B., Krebs, J. R., & West, S. A. (2012). *An introduction to behavioural ecology*. John Wiley & Sons.

De Bacco, C., Contisciani, M., Cardoso-Silva, J., Safdari, H., Lima Borges, G., Baptista, D., Sweet, T., Young, J.-G., Koster, J., Ross, C. T., Redhead, D., & Power, E. (2023). Latent network models to account for noisy, multiply reported social network data. *Journal of the Royal Statistical Society Series A: Statistics in Society*, *186*(3), 355–375.

De Moor, D., Brent, L. J., Silk, M., & Brask, J. (2024). Layers of latency in social networks and their implications for comparative analyses. *EcoEvoRxiv*.

De Moor, D., Roos, C., Ostner, J., & Schülke, O. (2020). Female assamese macaques bias their affiliation to paternal and maternal kin. *Behavioral Ecology*, *31*(2), 493–507.

Duboscq, J., Micheletta, J., Perwitasari-Farajallah, D., Engelhardt, A., & Neumann, C. (2023). Investigating the relationship between sociality and reproductive success in wild female crested macaques, *Macaca nigra*. *International Journal of Primatology*, 1–21.

Ellis, S., Snyder-Mackler, N., Ruiz-Lambides, A., Platt, M. L., & Brent, L. J. (2019). Deconstructing sociality: The types of social connections that predict longevity in a group-living primate. *Proceedings of the Royal Society B*, *286*(1917), 20191991.

Elo, A. (1978). *The rating of chess players, past and present*. Arco, New York.

Farine, D. R. (2017). A guide to null models for animal social network analysis. *Methods in Ecology and Evolution*, *8*(10), 1309–1320.

Farine, D. R., & Carter, G. G. (2022). Permutation tests for hypothesis testing with animal social network data: Problems and potential solutions. *Methods in Ecology and Evolution*, *13*(1), 144–156.

Farine, D. R., & Whitehead, H. (2015). Constructing, conducting and interpreting animal social network analysis. *Journal of animal ecology*, *84*(5), 1144–1163.

Feller, A., & Gelman, A. (2015). Hierarchical models for causal effects. *Emerging Trends in the Social and Behavioral Sciences*, 1–16.

Franz, M., Schülke, O., & Ostner, J. (2013). Rapid evolution of cooperation in group-living animals. *BMC evolutionary biology*, *13*, 1–11.

Fu, F., Nowak, M. A., Christakis, N. A., & Fowler, J. H. (2012). The evolution of homophily. *Scientific reports*, *2*(1), 845.

Gelman, A., Vehtari, A., Simpson, D., Margossian, C. C., Carpenter, B., Yao, Y., Kennedy, L., Gabry, J., Bürkner, P.-C., & Modrák, M. (2020). Bayesian workflow. *arXiv preprint arXiv:2011.01808*.

Gerber, L., Connor, R. C., Allen, S. J., Horlacher, K., King, S. L., Sherwin, W. B., Willems, E. P., Wittwer, S., & Krützen, M. (2022). Social integration influences fitness in allied male dolphins. *Current Biology*, *32*(7), 1664–1669.

Goold, C. (2020). Evaluating measurement models in animal personality: Operational, latent variable and network approaches. *OSF Preprints*. https://doi.org/https://doi.org/10.31219/osf.io/w4urf

Greenland, S. (2022). The causal foundations of applied probability and statistics. In *Probabilistic and Causal Inference: The Works of Judea Pearl* (pp. 605–624).

Grosz, M. P., Rohrer, J. M., & Thoemmes, F. (2020). The taboo against explicit causal inference in nonexperimental psychology. *Perspectives on Psychological Science*, *15*(5), 1243–1255.

Hart, J., Weiss, M. N., Brent, L. J., & Franks, D. W. (2022). Common permutation methods in animal social network analysis do not control for non-independence. *Behavioral Ecology and Sociobiology*, *76*(11), 151.

Hart, J., Weiss, M. N., Franks, D., & Brent, L. (2023). Bison: A Bayesian framework for inference of social networks. *Methods in Ecology and Evolution*, *14*(9), 2411–2420.

Hinde, R. A. (1976a). Interactions, relationships and social structure. *Man*, 1–17.

Hinde, R. A. (1976b). On describing relationships. *Journal of child psychology and psychiatry*, *17*(1), 1–19.

Hinde, R. A. (1983). *Primate social relationships: An integrated approach*. Blackwell Science Ltd.

Imai, K., & Kim, I. S. (2019). When should we use unit fixed effects regression models for causal inference with longitudinal data? *American Journal of Political Science*, *63*(2), 467–490.

Kalbitz, J., Schülke, O., & Ostner, J. (2017). Triadic male-infant-male interaction serves in bond maintenance in male assamese macaques. *PLoS One*, *12*(10), e0183981.

Kay, M. (2020). tidybayes: Tidy Data and Geoms for Bayesian Models (R package version 3.0.2).

Kenny, D. A., & La Voie, L. (1984). The social relations model. In *Advances in experimental social psychology* (pp. 141–182). Elsevier.

Koller, D., & Friedman, N. (2009). *Probabilistic graphical models: Principles and techniques*. MIT press.

Koster, J. M., & Leckie, G. (2014). Food sharing networks in lowland Nicaragua: An application of the social relations model to count data. *Social Networks*, *38*, 100–110.

Krause, J., James, R., Franks, D. W., & Croft, D. P. (2015). *Animal social networks*. Oxford University Press, USA.

Kummer, H. (1978). On the value of social relationships to nonhuman primates: A heuristic scheme. *Social Science Information*, *17*(4-5), 687–705.

Leimar, O., & Bshary, R. (2024). Social bond dynamics and the evolution of helping. *Proceedings of the National Academy of Sciences*, *121*(11), e2317736121.

Lundberg, I., Johnson, R., & Stewart, B. M. (2021). What is your estimand? Defining the target quantity connects statistical evidence to theory. *American Sociological Review*, *86*(3), 532–565.

Martin, J. S., Massen, J. J., Šlipogor, V., Bugnyar, T., Jaeggi, A. V., & Koski, S. E. (2019). The EGA+GNM framework: An integrative approach to modelling behavioural syndromes. *Methods in Ecology and Evolution*, *10*(2), 245–257.

McElreath, R. (2020). *Statistical rethinking: A Bayesian course with examples in R and Stan* (Second). Chapman; Hall/CRC.

Mielke, A., & Samuni, L. (2021). Accuracy and precision of social relationship indices. *bioRxiv*, 2021–04.

Neumann, C., & Fischer, J. (2023a). (De)composing sociality: Disentangling individual-specific from dyad-specific propensities to interact. *bioRxiv*. https://doi.org/10.1101/2023.08.15.549768

Neumann, C., & Fischer, J. (2023b). Extending Bayesian Elo-rating to quantify the steepness of dominance hierarchies. *Methods in Ecology and Evolution*, *14*(2), 669–682.

Newman, M. (2018). *Networks*. Oxford university press.

Nicenboim, B., Schad, D., & Vasishth, S. (2021). An introduction to Bayesian data analysis for cognitive science. *Under contract with Chapman and Hall/CRC statistics in the social and behavioral sciences series*.

Nicolás Carlock, J. R., Boyer, D., Smith Aguilar, S. E., & Ramos Fernández, G. (2023). Strength of minority ties: The role of homophily and group composition in a weighted social network. *Journal of Physics: Complexity*.

Nowak, M. A. (2006). Five rules for the evolution of cooperation. *Science*, *314*(5805), 1560–1563.

Ogburn, E. L., Sofrygin, O., Diaz, I., & Van der Laan, M. J. (2022). Causal inference for social network data. *Journal of the American Statistical Association*, 1–15.

Ostner, J., & Schülke, O. (2018). Linking sociality to fitness in primates: A call for mechanisms. In *Advances in the study of behavior* (pp. 127–175). Elsevier.

Pearl, J. (2001). Bayesianism and causality, or, why I am only a half-Bayesian. In *Foundations of Bayesianism* (pp. 19–36). Springer.

Pearl, J. (2009). *Causality*. Cambridge university press.

Pearl, J., & Bareinboim, E. (2014). External validity: From do-calculus to transportability across populations. *Statistical Science*, *29*(4), 579–595.

Pearl, J., Glymour, M., & Jewell, N. P. (2016). *Causal inference in statistics: A primer*. John Wiley & Sons.

Pearl, J., & Mackenzie, D. (2018). *The book of why: The new science of cause and effect*. Basic books.

Pedersen, T. L. (2019). Package 'patchwork'. *R package http://CRAN. R-project. org/package= patchwork. Cran*.

Pedersen, T. L., Pedersen, M., LazyData, T., Rcpp, I., & Rcpp, L. (2017). Package 'ggraph'.

Puga-Gonzalez, I., Ostner, J., Schülke, O., Sosa, S., Thierry, B., & Sueur, C. (2018). Mechanisms of reciprocity and diversity in social networks: A modeling and comparative approach. *Behavioral Ecology*, *29*(3), 745–760.

R Core Team. (2022). R: A language and environment for statistical computing. https://www.R-project.org/

Rawlings, C. M., Smith, J. A., Moody, J., & McFarland, D. A. (2023). *Network analysis: integrating social network theory, method, and application with R*. Cambridge University Press.

Redhead, D., Dalla Ragione, A., & Ross, C. T. (2023). Friendship and partner choice in rural Colombia. *Evolution and Human Behavior*, *44*(5), 430–441.

Redhead, D., Gervais, M., Kajokaite, K., Koster, J., Hurtado Manyoma, A., Hurtado Manyoma, D., McElreath, R., & Ross, C. T. (2024). Evidence of direct and indirect reciprocity in network-structured economic games. *Communications Psychology*, *2*(1), 44.

Redhead, D., McElreath, R., & Ross, C. T. (2023). Reliable network inference from unreliable data: A tutorial on latent network modeling using STRAND. *Psychological Methods*.

Redhead, D., & Power, E. A. (2022). Social hierarchies and social networks in humans. *Philosophical Transactions of the Royal Society B*, *377*(1845), 20200440.

Rohrer, J. M. (2024). Causal inference for psychologists who think that causal inference is not for them. *Social and Personality Psychology Compass*, e12948. https://doi.org/https://doi.org/10.1111/spc3.12948

Rohrer, J. M. (2018). Thinking clearly about correlations and causation: Graphical causal models for observational data. *Advances in methods and practices in psychological science*, *1*(1), 27–42.

Rohrer, J. M. (2023). *Less casual causal inference for experiments and longitudinal data*. Universität Wien Research Talks. https://youtu.be/OsAo2ffbUAQ

Rohrer, J. M., & Arslan, R. C. (2021). Precise answers to vague questions: Issues with interactions. *Advances in Methods and Practices in Psychological Science*, *4*(2), 25152459211007368.

Rohrer, J. M., & Murayama, K. (2023). These are not the effects you are looking for: Causality and the within-/between-persons distinction in longitudinal data analysis. *Advances in methods and practices in psychological science*, *6*(1), 25152459221140842.

Ross, C. T., McElreath, R., & Redhead, D. (2023). Modelling animal network data in R using STRAND. *Journal of Animal Ecology*.

Runge, J., Gerhardus, A., Varando, G., Eyring, V., & Camps-Valls, G. (2023). Causal inference for time series. *Nature Reviews Earth & Environment*, *4*(7), 487–505.

Sadoughi, B., Mundry, R., Schülke, O., & Ostner, J. (2024). Social network shrinking is explained by active and passive effects but not increasing selectivity with age in wild macaques. *Proceedings of the Royal Society B*, *291*, 20232736.

Schino, G. (2001). Grooming, competition and social rank among female primates: A meta-analysis. *Animal behaviour*, *62*(2), 265–271.

Scholz, M., & Bürkner, P.-C. (2022). Prediction can be safely used as a proxy for explanation in causally consistent bayesian generalized linear models. *arXiv preprint arXiv:2210.06927*.

Seyfarth, R. M. (1977). A model of social grooming among adult female monkeys. *Journal of theoretical Biology*, *65*(4), 671–698.

Seyfarth, R. M. (1980). The distribution of grooming and related behaviours among adult female vervet monkeys. *Animal behaviour*, *28*(3), 798–813.

Shalizi, C. R., & Thomas, A. C. (2011). Homophily and contagion are generically confounded in observational social network studies. *Sociological methods & research*, *40*(2), 211–239.

Shipley, B. (2016). *Cause and correlation in biology: a user's guide to path analysis, structural equations and causal inference with R*. Cambridge university press.

Silk, J. B. (2002). Using the 'F'-word in primatology. *Behaviour*, *139*(2/3), 421–446.

Silk, J. B., Alberts, S. C., & Altmann, J. (2006). Social relationships among adult female baboons (*Papio cynocephalus*) ii. variation in the quality and stability of social bonds. *Behavioral Ecology and Sociobiology*, *61*, 197–204.

Silk, J. B., Cheney, D., & Seyfarth, R. (2013). A practical guide to the study of social relationships. *Evolutionary Anthropology: Issues, News, and Reviews*, *22*(5), 213–225.

Silk, M. J. (2023). Conceptual representations of animal social networks: An overview. *Animal Behaviour*, *201*, 157–166.

Siracusa, E. R., Higham, J. P., Snyder-Mackler, N., & Brent, L. J. (2022). Social ageing: Exploring the drivers of late-life changes in social behaviour in mammals. *Biology letters*, *18*(3), 20210643.

Smaldino, P. (2023). *Modeling social behavior: Mathematical and agent-based models of social dynamics and cultural evolution*. Princeton University Press.

Snijders, T. A., & Kenny, D. A. (1999). The social relations model for family data: A multilevel approach. *Personal Relationships*, *6*(4), 471–486.

Snyder-Mackler, N., Burger, J. R., Gaydosh, L., Belsky, D. W., Noppert, G. A., Campos, F. A., Bartolomucci, A., Yang, Y. C., Aiello, A. E., O'Rand, A., et al. (2020). Social determinants of health and survival in humans and other animals. *Science*, *368*(6493), eaax9553.

Spake, R., Bowler, D. E., Callaghan, C. T., Blowes, S. A., Doncaster, C. P., Antao, L. H., Nakagawa, S., McElreath, R., & Chase, J. M. (2023). Understanding 'it depends' in ecology: A guide to hypothesising, visualising and interpreting statistical interactions. *Biological Reviews*.

Steglich, C., Snijders, T. A., & Pearson, M. (2010). Dynamic networks and behavior: Separating selection from influence. *Sociological methodology*, *40*(1), 329–393.

Sukmak, M., Wajjwalku, W., Ostner, J., & Schülke, O. (2014). Dominance rank, female reproductive synchrony, and male reproductive skew in wild assamese macaques. *Behavioral Ecology and Sociobiology*, *68*, 1097–1108.

Talts, S., Betancourt, M., Simpson, D., Vehtari, A., & Gelman, A. (2020). Validating Bayesian Inference Algorithms with Simulation-Based Calibration. *arXiv*.

Tinbergen, N. (1963). On aims and methods of ethology. *Zeitschrift für tierpsychologie*, *20*(4), 410–433.

Van Duijn, M. A., Snijders, T. A., & Zijlstra, B. J. (2004). p2: a random effects model with covariates for directed graphs. *Statistica Neerlandica*, *58*(2), 234–254.

Wasserman, S., & Faust, K. (1994). Social network analysis: Methods and applications.

Weiss, M. N., Franks, D. W., Brent, L. J., Ellis, S., Silk, M. J., & Croft, D. P. (2021). Common datastream permutations of animal social network data are not appropriate for hypothesis testing using regression models. *Methods in Ecology and Evolution*, *12*(2), 255–265.

Westreich, D., & Greenland, S. (2013). The table 2 fallacy: presenting and interpreting confounder and modifier coefficients. *American journal of epidemiology*, *177*(4), 292–298.

Whitehead, H. (2008). *Analyzing animal societies: Quantitative methods for vertebrate social analysis*. University of Chicago Press.

Wickham, H. (2011). ggplot2. *Wiley interdisciplinary reviews: computational statistics*, *3*(2), 180–185.

Widdig, A., Nürnberg, P., Bercovitch, F. B., Trefilov, A., Berard, J. B., Kessler, M. J., Schmidtke, J., Streich, W. J., & Krawczak, M. (2006). Consequences of group fission for the patterns of relatedness among rhesus macaques. *Molecular Ecology*, *15*(12), 3825–3832.

Wu, C.-F., Liao, Z.-J., Sueur, C., Sha, J. C. M., Zhang, J., & Zhang, P. (2018). The influence of kinship and dominance hierarchy on grooming partner choice in free-ranging macaca mulatta brevicaudus. *Primates*, *59*, 377–384.

Young, J.-G., Cantwell, G. T., & Newman, M. (2020). Bayesian inference of network structure from unreliable data. *Journal of Complex Networks*, *8*(6), cnaa046.