

Network Inference, Error, and Informant (In)Accuracy: A Bayesian Approach

Carter Butts*

Department of Social and Decision Sciences
Center for the Computational Analysis of
Social and Organizational Systems
Carnegie Mellon University

January 10, 2000

v0.1b

This is a draft document. Please do not copy or cite.

Abstract

Much, if not most, social network data is derived from informant self-reports; past research, however, has indicated that such reports are in fact highly inaccurate representations of social interaction. In this paper, a family of hierarchical Bayesian models is developed which allows for the simultaneous inference of informant accuracy and social structure in the presence of measurement error. In addition to point estimation of posterior quantities of interest (such as individual error rates, graph and node level indecies, and the criterion graph), it is shown that the models here considered provide a quantification of the posterior uncertainty associated with said quantities. Implications of the Bayesian modeling framework for improved data collection strategies and the validity of the criterion graph are also discussed.

Keywords: informant accuracy, measurement error, hierarchical Bayesian models, network inference, data collection strategies

1 Introduction

1.1 Self-Report Data and Informant Accuracy

The ultimate foundation of any scientific endeavor is the data from which inferences are to be drawn. In the field of social network analysis, the quality

*This material is based upon work supported under a National Science Foundation Graduate Fellowship, and was supported in part by the Center for the Computational Analysis of Social and Organizational Systems and the Institute for Complex Engineered Systems at Carnegie Mellon University.

of this foundation has long been in question. A large literature on informant accuracy in network analysis - most notably that written by and in response to the work of Bernard, Killworth, and Sailer (BKS) - has debated the question of the extent to which the most common form of network data (informant self-reports) can be considered to represent anything beyond the cognitions of the informants in question¹ (see, for instance, Killworth and Bernard (1976; 1979/80), Bernard and Killworth (1977), Bernard et al. (1979/80; 1984), Freeman et al. (1987), Hildum (1986), Romney and Faust (1982), and Romney et al. (1986)). While the dust has not yet settled on this debate, the present state of knowledge within the field of network analysis strongly favors the position that informant self-reports contain considerable noise at the dyadic level (Bernard et al., 1984; Krackhardt and Kilduff, 1999). Whether the level of error² observed is severe enough to warrant the assertion that self-report data is unusable for inference of interaction patterns appears to depend on a number of factors (including the time scale of the interaction, the application to which the data is to be put, and (arguably) the degree of optimism on the part of the researcher (Freeman et al., 1987; Freeman, 1992; Krackhardt and Kilduff, 1999; Bernard et al., 1979/80)), but unquestionably there is a non-trivial level of uncertainty inherent in self-report network data.

The recognition that self-reports of interaction are not in fact exact proxies for observed behavior has led to two general approaches to the interpretation of self-report data. The first approach, which we shall here call the "classical" or "criterion/error" perspective, conceptualizes the data generation process in terms of a hypothesized "real" (or criterion) structure which is imperfectly reported by informants. This perspective is most clearly in line with the broader literature on informant accuracy (the notion of "accuracy" itself requiring the assumption of a criterion from which informant reports are assumed to deviate in some fashion), and is followed both by BKS and by later work seeking to identify the determinants of (in)accurate perception by socially embedded actors (see, for instance, Killworth and Bernard (1976), Krackhardt (1990), Calloway et al. (1993)). The second approach, here referred to as the "cognitivist" perspective, has focused on identifying the determinants of actors' perceptions of social structure *per se* (e.g., Krackhardt (1987), Carley and Krackhardt (1996), Freeman et al. (1987), Hammer (1985), Hildum (1986)). This perspective has tended to de-emphasize the role of the criterion structure, and indeed has in some cases questioned the relevance or even meaningfulness of such a construct (e.g., Krackhardt (1987)). From a cognitivist point of view, the language of "error" and "accuracy" reflects a clear bias towards third-party observability (which may or may not be relevant to dyadic perceptions, particularly for private or tacit interactions) and non-attributive relations (Hildum, 1986). From

¹Bernard, Killworth and Sailer (1979) go so far as to posit that "We are now convinced that cognitive data about communication can not be used as a proxy for behavioral data," a position which is rejected by researchers such as Romney et al. (1986) and Romney and Faust (1982).

²In the sense of deviation from externally observable interaction patterns; we shall consider this issue in more detail presently.

the classical perspective, by contrast, these are critical considerations: inference regarding the criterion structures which are the objects of traditionalist theorizing depends on the accuracy of informant reports, and the minimization of error is thus of central concern (Bernard et al., 1979/80).

It is beyond the scope of this paper to conduct a detailed consideration of the assumptions, applicability, and justification of the classical and cognitivist perspectives on self-report data. For our purposes, we shall adopt the language and assumptions of the classical approach, including the notions of the criterion graph, informant accuracy, and reporting error. Though the usefulness of that which follows for any given application is dependent upon the validity of these assumptions, we do not take them to be universally unproblematic; this issue is revisited in the discussion section below.

1.2 The Enduring Problem of Error in Social Network Analysis

While the informant accuracy question, then, has been considered from a variety of perspectives, at least four basic problems confront us with respect to dealing with error in social network analysis (from a classical point of view):

1. Determining the extent of error in existing data
2. Determining the mechanisms by which error is produced
3. Finding means of collecting higher quality data
4. Minimizing and accounting for the uncertainty associated with existing data in network analyses

All of these problems, of course, are of critical importance to social network analysis. Without knowing the degree of uncertainty with which we should regard present data, we have no way of evaluating the reliability (or even validity!) of present or past work. Without knowing how or why errors are produced, we are at a loss to predict which data will be most heavily compromised. With no means of collecting clean data, we continue to be vulnerable to error and uncertainty, and without techniques for minimizing and assessing that uncertainty we are unable to draw appropriate inferences. This paper, then, will touch at least briefly on all four issues, although the primary emphasis will be on the fourth problem.

1.3 A Bayesian Approach

Given that social network analysis is heavily dependent upon self-report data, and that informants are known to err in their reporting of network ties, how are we to proceed? As noted above, a number of approaches are possible; in this paper, our primary interest is in the development of inferential techniques which can address the informant accuracy problem. Such techniques, clearly,

must simultaneously accomplish two goals: first, they must infer the criterion graph from informant reports; and second, they must infer the accuracy of each informant. Though solving both such problems at once would seem a difficult task, it is inescapable; inferences regarding informant accuracy will necessarily affect inferences regarding the criterion graph, and vice versa. This implies, further, that whatever approach is employed must scale effectively to high-dimensional inference problems, as the number of informants (whose accuracy is in question) grows on order $|V(G)|$ and the number of ties in the criterion graph grows on order $|V(G)|^2$. Given the high-dimensionality of the inference problem, the data available is likely to be modest. In the best standard case - a CSS - we have access to $|V(G)|$ observations per arc and $|V(G)|^2$ observations per actor; while the total number of data points thus grows on order $|V(G)|^3$, the small size of most CSSs provides us with relatively few observations per parameter. Our technique, then, must not only scale well, but must also be data-efficient. Methods which are justified only in the large-N limit are unlikely to be of practical value for this particular problem.

All of the above factors suggest the efficacy of a hierarchical Bayesian modeling approach to the network inference/informant accuracy problem (Gelman et al., 1995). Hierarchical Bayesian models can readily represent complex interrelated stochastic processes, scale well, and are not dependent on limit arguments for their justification. Furthermore, use of the Bayesian paradigm permits us to draw direct inferences regarding posterior probabilities, and grounds our inferential framework on an axiomatic basis³. Finally, the hierarchical Bayesian modeling framework readily facilitates expansion and modification of existing models to account for new information or to take advantage of the features of particular situations. Given these advantages, we here employ the aforesaid approach in examining the network inference/informant accuracy problem; while classical alternatives could be pursued as well, they are not considered here.

One further manner in which the present approach differs from some others is that its emphasis is as much on accounting for uncertainty as on reducing it. Given that we analyze imperfect data in an imperfect world, we generally cannot claim to be completely certain of the quantities with which we work; a realistic approach to data analysis, then, cannot afford to simply brush problems of uncertainty under the proverbial rug. Arguably, however, traditional social network analysis has done exactly that: by failing to account for error in network data, network analysts have put the quality of their inferences at risk. Even where the models presented here do not provide substantial uncertainty reduction, then, they may be of use in their ability to provide a concrete treatment of uncertainty in network data.

³For an examination of some of the foundational strengths of the Bayesian paradigm, see Robert (1994).

2 Bayesian Models of the Network Inference/Informant Accuracy Problem

As we have argued, an integrated, formal approach is required to simultaneously address the problems of network inference and informant accuracy. In this section, then, we develop a family of hierarchical Bayesian models which allow for inference on these two problems, and demonstrate the use of these models in determining posterior quantities of interest.

2.1 Stochastic Network Formalism

In order to proceed with a formal model of the data generation process, we must first develop a formalism for the criterion graph⁴ itself. While a number of possibilities exist in this regard, we here treat the criterion graph as a random loopless digraph of fixed order such that the existence of each arc, (i, j) is an independent event which occurs with some probability $p((i, j) \in E(G))$. Such a structure is referred to as a *Bernoulli graph* in the literature (Wasserman and Faust, 1994), and may be conveniently treated as a random variable via its adjacency matrix:

$$A_{ij} = \begin{cases} 1 & \text{if } (i, j) \in E(G) \\ 0 & \text{if } (i, j) \notin E(G) \end{cases} \quad (1)$$

Obviously, the cells of the adjacency matrix A of digraph G serve as indicator variables for the arcs of G , and are thus independently Bernoulli distributed with $p(A_{ij} = 1) = p((i, j) \in E(G))$. In general, we shall represent the distribution of A (that is, the matrix of arc probabilities) by the Bernoulli parameter matrix Θ , and observations pertaining to A by the data matrix \mathbf{Y} . Note that, while we assume that the *realizations* of arcs within the model are independent, *inferences* regarding the distribution of A will usually assume only conditional independence. Thus, it will generally be the case that information regarding a particular set of arcs will affect inference regarding other arcs⁵; we require, however, that this enter only through the parameter matrix Θ . This assumption, while not especially stringent, is subject to question: this is left as a topic for future research.

2.1.1 General Informant Accuracy Model

As has been pointed out, one cannot address the problem of informant accuracy without having some prior conception of the manner in which the hypothesized criterion structure is related to the social structure which is, in fact, provided by informant accounts. This is not an unproblematic question: the relationship in question has been argued to be mediated by cognitive mechanisms, instrument

⁴Throughout this paper, we shall treat the criterion graph as a loopless digraph; the term may be used generically to refer to any sort of graph, however (e.g., simple, signed, valued, etc.) which acts as the criterion for informant reports.

⁵E.g., indirectly via the informant accuracy parameters.

design, informant experience, and even the informant's position in the social structure itself (Freeman, 1992; Freeman et al., 1987; Carley and Krackhardt, 1996; Krackhardt, 1990; Hildum, 1986; Krackhardt and Kilduff, 1999). Plainly, one can propose models with varying degrees of sophistication, which take into account a greater or lesser number of influences on accuracy, and which model those influences in more or less sophisticated ways. In this paper, our primary purpose is to introduce a reasonably simple family of models which are easily utilized in network research; our secondary purpose is to formulate these models in such a way as to make them as compatible as possible with the prior empirical work on informant accuracy within the context of social network analysis. For this reason, we shall utilize a general informant accuracy framework which follows that implicitly utilized (we shall argue) by Bernard, Killworth, Sailer, and others in their informant accuracy studies. Although this framework is reasonably flexible, it is not asserted that this approach will be optimal in all cases; as will be discussed, some applications may be better suited by alternative formulations.

If we are to treat informant accuracy in such a way as to maintain compatibility with prior work, it behooves us to consider the models which have been (explicitly or implicitly) invoked in the past. Fortunately, the most common approach has been described fairly clearly (albeit verbally) by BKS. With respect to the basic structure of the problem, the version articulated here is quite typical: "Social structure is assumed to be built up out of the interaction of the members of the group. Then, the plausible leap is made whereby the answers to [a] sociometric question reflect the pattern of interactions" (Bernard et al., 1979/80). Thus, it is assumed that there exists a criterion network which is formed from "actual" interactions between actors, and accuracy is measured with respect to the degree that informant responses on sociometric instruments match this network. More explicitly, BKS state, "At its simplest level, network data are 'accurate' if, when *i* says he talked to *j* by some amount, then he did" (Bernard et al., 1979/80). Though this statement may seem to be self-evident at first blush, its implications are non-trivial. For instance, the BKS model requires that the "actual" network be in some sense verifiable – and hence defined – beyond the individual perceptions or testimonials of those involved in the proposed relation⁶. This sets the scope of the accuracy question (from this perspective) in such a way as to exclude purely ascriptive relations (e.g., individual perceptions of influence, love, or power, reputation and prestige⁷, and

⁶BKS, indeed, repeatedly question the meaningfulness of any structure which does not have this feature, even going so far as to argue that networks are "real" insofar as they "have some external correlate like performance in problem solving" (Killworth and Bernard, 1979/80). It is not our intention to make arguments for or against this position in this paper, but the issue is obviously one of great relevance to the continued theoretical and methodological development of network analysis.

⁷Note that unlike dominance and power, which can be fairly readily defined beyond their ascriptions (though *perceptions* of such are always ascriptive), reputation and prestige are by nature ascribed relations.

some notions of friendship⁸), at least on a first-order basis⁹. As the BKS statement makes clear, the focus here is not on *accounting for* informant responses to sociometric instruments, but rather on determining the degree to which these responses are consistent with a pre-defined observable (but unobserved) relation. Although the former has been considered to be an interesting question in its own right at least since Krackhardt's introduction of the cognitive social structure concept (Krackhardt, 1987), this is not our primary interest in this case¹⁰.

The base assumption which is made, then, is that each observer is exposed to the underlying "real" network and reports observed ties which are, with some probability, erroneous¹¹. As there are by assumption two states which may be taken by each arc (present and absent), there are clearly two ways in which observers can report incorrectly: an observer may report present ties as being absent (false negatives) or may report absent ties as being present (false positives). We then model the data generation process via the following Bernoulli mixture,

$$p(Y_{ij}|A_{ij}, e^+, e^-) = \begin{cases} Y_{ij}e^+ + (1 - Y_{ij})(1 - e^+) & \text{if } A_{ij} = 0 \\ Y_{ij}(1 - e^-) + (1 - Y_{ij})e^- & \text{if } A_{ij} = 1 \end{cases} \quad (2)$$

where Y_{ij} is the informant report of the value of the arc indicator function, A_{ij} is the value of that indicator function for the criterion structure (i.e., the ij th cell of the adjacency matrix), and e^+ and e^- represent the probabilities of false positives and false negatives (respectively) for that arc.

The above provides the core likelihood for a single observer/arc/criterion network combination. In general, of course, we are interested in reports concerning multiple arcs, by multiple observers, on (and in) criterion networks about which we are uncertain.

2.1.2 Assignment of Network Priors

As indicated, our network formalism is the Bernoulli graph, with distribution given by the parameter matrix Θ . In order to draw inferences regarding the criterion graph, then, we must specify initial values for Θ which reflect our prior information regarding the network of interest. This must, of necessity, depend on the particular problem at hand, the existing literature regarding said problem, etc. Nevertheless, it is useful to provide some basic strategies for assigning

⁸See, for instance, Carley and Krackhardt (1996).

⁹E.g., one can still ask whether a given actor's account of the ascriptions of others matches their actual *ascriptions*, but this is a second-order rather than a first-order question.

¹⁰Though, as we shall see, the models presented here can potentially shed some light on this question as well.

¹¹Note that we are unconcerned at present with whether the errors in question are the result of inaccurate *observation* by the informant, errors in the cognitive *coding* of the relation, or difficulties with *retrieval* of this information upon exposure to the appropriate instrument. These error sources are folded together in the present model, and we are concerned only with the aggregate result.

network priors, which may be utilized in a range of contexts. Two heuristics, then, are here presented which are expected to have wide applicability; each researcher, however, should be careful to select network priors which are accurate depictions of his or her prior information, and should avoid blind reliance on pre-packaged choices.

The first, and perhaps most obvious, heuristic for assignment of network priors is that of an uninformative distribution on the arc set of G . Such a prior distribution is given simply by

$$\Theta_{ij} = 0.5 \forall i, j \in V(G) \quad (3)$$

and corresponds to a uniform distribution on the set of all loopless digraphs of order $|V(G)|$. This prior is obviously somewhat attractive in that it has a clear interpretation, and that it does not depend on the knowledge of the researcher. In general, however, an uninformative prior on the set of digraphs is a poor reflection of one's prior information, and better options are available; researchers should question whether or not they are in fact completely ignorant as to the structure of the criterion graph before selecting such a prior.

One specific example of the weakness of an uninformative network prior is the matter of network density. While the uniform distribution over digraphs favors graphs of density approximately equal to 0.5¹², most social structures (particularly large ones) are considerably less dense. If a researcher has reason to believe that the relation he or she is studying is of a particular density, it behooves him or her to take this into account when selecting a prior for the criterion graph¹³. One possibility, then, for somewhat more informative network prior, is to choose a distribution of the form

$$\Theta_{ij} = d \forall i, j \in V(G) \quad (4)$$

where d is the median density of a set of other, similar networks examined in past research. While this is obviously only one of many possible density-based priors, it is clearly more informative than the assumption of a uniform distribution, and is nevertheless diffuse enough to avoid extreme sensitivity to initial assumptions. In situations for which more information is available, however, researchers should not hesitate to use it: priors including homophily or distance effects, for instance, might be reasonable in many situations, as might the posterior distribution of the same network at an earlier point in time. The topic of network prior selection is complex enough to warrant a detailed treatment on its own, and it is hoped that future work in this area will develop more sophisticated strategies for exploiting prior knowledge than can be discussed within the bounds of this paper.

¹²For reasons of combinatorics: there are far more structures of moderate density than of extreme density.

¹³For a discussion of the effects of size and density on network inference, see Anderson et al. (1999).

2.2 A Simple Model for Fixed, Known Error Probabilities

Before we begin a consideration of the more complex cases of network inference in the presence of unknown informant inaccuracies, it behooves us to begin with a simpler model. As an introduction to our basic modeling framework, then, we shall first develop a network inference model in which we assume that our error probabilities are both fixed and known¹⁴. Clearly, this does not reflect the context of most network research: we do not, in general, know the probability of error within our data to within an arbitrary degree of certainty. Nevertheless, there may be special cases - such as, for instance, automated data collection procedures whose error rates may be determined exogenously - for which the simple model may be applicable. Our reason for introducing it, however, is as much to elucidate the concepts involved as to provide a practical tool for social research.

2.2.1 Assignment of Priors

Within the simple network inference model, it is assumed that only the criterion graph is uncertain; therefore, priors are assigned only to the network variable, A . As indicated, the prior distribution of the criterion graph is performed via the parameter matrix Θ . The specific form, then, is as follows:

$$\Theta_{ij} \sim B(\theta_{ij}) \quad (5)$$

$$p(\Theta) = \prod_{i=1}^N \prod_{j=1}^N p(\Theta_{ij}) \quad (6)$$

where $N = |V(G)|$. Choice of the values of θ_{ij} is discussed in the above section on network priors; as mentioned, these values should reflect the researcher's prior knowledge regarding the structure of the criterion graph. Note that the joint prior on the criterion graph is obviously the product the priors of the individual arcs. This follows from the assumption of conditional independence of ties within the Bernoulli graph.

2.2.2 Assumed Likelihood

As we are assuming in this case that our error probabilities are fixed and known, the base likelihood for an observed arc follows straightforwardly from Equation 2:

$$p(Y_{ij} | \Theta_{ij}, e^+, e^-) = \frac{(1 - \Theta_{ij})(Y_{ij}e^+ + (1 - Y_{ij})(1 - e^+)) + \Theta_{ij}(Y_{ij}(1 - e^-) + (1 - Y_{ij})e^-)}{(1 - \Theta_{ij})(1 - e^+) + \Theta_{ij}(1 - e^-)} \quad (7)$$

¹⁴This is arguably something of a redundancy: from a Bayesian point of view, the stochasticity of the error probabilities is inherently a measure of uncertainty. The designation of "fixed and known", then, is employed only to make salient the assumption of determinism for non-Bayesian readers.

Observe that this is just the mixture from Equation 2, factoring in our prior uncertainty regarding the criterion graph. To find the joint likelihood of the data, then, we invoke the previously stated assumption of conditional independence of observations; this implies that the joint likelihood is the product of the individual arc likelihoods. Formally, the joint likelihood of the simple model is then given by

$$p(\mathbf{Y} | \Theta, e^+, e^-) = \prod_{i=1}^N \prod_{j=1}^N p(Y_{ij} | \Theta_{ij}, e^+, e^-) \quad (8)$$

(where \mathbf{Y} is the observed data matrix and $p(Y_{ij} | \Theta_{ij}, e^+, e^-)$ is the arc likelihood of Equation 7).

2.2.3 Computation of the Posterior

Generally, we would identify the posterior to within a normalizing constant by using the standard result that $p(\theta|y) \propto p(\theta)p(y|\theta)$ for random variables y and θ . In this case, however, we can straightforwardly apply Bayes' Rule to each arc in order to find the exact posterior¹⁵. The posterior probability for the existence of a given arc, then, is

$$p(\Theta_{ij} | Y_{ij}, e^+, e^-) = Y_{ij} \frac{\Theta_{ij} (1 - e^-)}{\Theta_{ij} (1 - e^-) + (1 - \Theta_{ij}) e^+} + (1 - Y_{ij}) \frac{\Theta_{ij} e^-}{\Theta_{ij} e^- + (1 - \Theta_{ij}) (1 - e^+)} \quad (9)$$

and the joint posterior of the criterion graph is hence

$$p(\Theta | \mathbf{Y}, e^+, e^-) = \prod_{i=1}^N \prod_{j=1}^N p(\Theta_{ij} | Y_{ij}, e^+, e^-) \quad (10)$$

The above, of course, provides the criterion posterior for a single group of observations. Additional observations can be added straightforwardly in the usual fashion (i.e., by repeating the above computation with the initial posterior as prior and with the new observations as \mathbf{Y}). It is, in fact, not even necessary given the assumptions made here for the error probabilities e^+ and e^- to be identical for all observations, or for one to consider an entire observation matrix, \mathbf{Y} , at once; though convenient, these assumptions can be relaxed so long as e_{ijk}^+ and e_{ijk}^- are known for each arc (i, j) and observation k , and Θ_{ij} is updated only once per observation¹⁶. In this case, then, the arc posterior is given by

$$p(\Theta_{ij} | \mathbf{Y}_{ij}, \mathbf{e}_{ij}^+, \mathbf{e}_{ij}^-) = \quad (11)$$

¹⁵This is, in essence, possible because we do not allow inference on any given arc to effect our inference regarding any other arc; no information is shared across observations, and hence we can subdivide the problem into a number of simple Bernoulli mixtures.

¹⁶This assumes that the pattern of data collection is ignorable with respect to Θ .

$$\prod_{k=1}^{n_{ij}} \left(Y_{ijk} \frac{\Theta_{ij} (1 - e_{ijk}^-)}{\Theta_{ij} (1 - e_{ijk}^-) + (1 - \Theta_{ij}) e_{ijk}^+} + (1 - Y_{ijk}) \frac{\Theta_{ij} e_{ijk}^-}{\Theta_{ij} e_{ijk}^- + (1 - \Theta_{ij}) (1 - e_{ijk}^+)} \right)$$

and the joint posterior of the criterion graph is thus simply the product of Equation 11 over all arcs.

It should be noted that, for most practical applications, the joint posterior of the criterion graph is not needed (though its computation is obviously quite straightforward where required). To draw from the posterior of the criterion graph in this case, one need only draw each (i, j) independently with probabilities as given by Equation 9 above. This makes estimation of posterior graph properties (see Section 2.5 below) particularly easy for the simple error model; this ease, however, comes at the expense of any inference regarding error probabilities (which are presumed known on an *a priori* basis) and of any shared information across arcs. The simple error model, then, is generally useful only when the accuracy of observers can be established *ex ante*, a condition which is rarely encountered in a typical social network analysis setting. It nevertheless serves as a reasonable prelude to the more sophisticated models which follow.

2.3 Pooled Error Probabilities (Single Observer Model)

Having considered a simple model of network inference in the special case for which error probabilities are fixed and known, we can now proceed to a more sophisticated - and useful - model. In particular, one reasonable elaboration of the fixed error probability model is to consider the case in which error probabilities are uncertain, but assumed to be constant across observations¹⁷. Such a pooled error probability model readily represents a situation in which a single observer or informant (e.g., an ethnographer) provides reports on all ties within a given network (e.g., the observers in the data sets of (Bernard et al., 1979/80)); we may be willing to presume that his or her accuracy is more or less constant across cases, but we are uncertain regarding the degree of accuracy itself.

2.3.1 Assignment of Priors

Unlike the simple error model, the single observer model treats error probabilities as uncertain; hence, we must assign priors to them. While a wide range of distributions are possible here, we have elected to represent the error probabilities e^+ and e^- as being drawn independently from two Beta distributions. Specifically:

$$e^+ \sim \text{Beta}(\alpha^+, \beta^+) \quad (12)$$

$$e^- \sim \text{Beta}(\alpha^-, \beta^-) \quad (13)$$

¹⁷Recall that this last was initially assumed for the fixed error probability model, but that we were able to relax this requirement due to our ability to decompose the problem into a series of arc inferences.

The choice of the Beta for the form of the error prior requires some justification; indeed, it may not be appropriate in all cases. In general, the Beta distribution is the conjugate prior for the Binomial likelihood, with the uninformative prior distribution on $[0,1]$ being a special case¹⁸. As a result, the outcome of any set of success/failure Bernoulli experiments which begins with an uninformative prior distribution will be Beta distributed, and the Beta is thus a logical choice when one's prior information can be summarized in this form. In general, prior knowledge regarding informant accuracy is likely to come from error rates derived from earlier studies (e.g., BKS), and hence it is not unreasonable to suppose that, for many researchers, the Beta *will* be an appropriate representation by this argument. Those drawing prior knowledge from other sources may need to consider alternative forms; these are not pursued within this paper.

Given the choice of Beta distributions for e^+ and e^- , there remains the question of selecting the four hyperparameters α^+ , β^+ , α^- , and β^- . While these could, in principle, be themselves drawn from a hyperprior distribution, we here treat them as known from prior data. As noted, the choice $\alpha = \beta = 1$ provides an uninformative prior for the Beta distribution; however, it is *strongly* recommended that researchers avoid this choice of prior. The reason for this is simple: the assumption of a uniform distribution of error parameters is highly unreasonable for most applications (given previous research in this area) and leads to highly counterintuitive (and improbable) inferences. Note, for instance, that (by Bayes Rule) the condition $e^+ + e^- > 1$ leads to a condition of *perverse inferences*, in which informant testimony causes one to update one's belief in the *opposite* direction of the report. Clearly, this is an unlikely event: even at their worst, it is hard to imagine that most informants' reports would be *negatively* informative¹⁹! Under a uniform prior, however, the *a priori* probability of such an occurrence is 0.5 – unacceptably high for most applications. In general, then α and β parameters should be chosen so as to cause the distribution of error probabilities to remain sensible, and to prevent the perverse inference condition from being highly probable. In examining a number of stylized facts from the BKS studies (i.e., apparent tendencies in error rates across relations), a prior of *Beta*(3, 5) for both false negative and false positive error parameters has thus far seemed reasonable for communication-like relations; individual researchers, however, should base their choices of hyperparameters on the particular data available.

Given the choice of priors for the two error parameters, the form of the network prior is as for the simple error model, specifically:

¹⁸In fact, there are no less than three uninformative priors for this particular problem: the uniform distribution ($\alpha = \beta = 1$), the Jeffrey's prior ($\alpha = \beta = 1/2$), and the improper prior which is uniform in the natural parameter of the exponential family ($\alpha = \beta = 0$) (Gelman et al., 1995). We assume the uniform distribution as an uninformative prior unless otherwise indicated.

¹⁹This has other implications as well: for instance, the posterior construction of two basic scenarios such that in one scenario informant reports are nearly all *reversed* in implication.

$$\Theta_{ij} \sim B(\theta_{ij}) \quad (14)$$

$$p(\Theta) = \prod_{i=1}^N \prod_{j=1}^N p(\Theta_{ij}) \quad (15)$$

Here again, the reader is directed to the previous discussion regarding choice of network priors.

2.3.2 Assumed Likelihood

While our model has been changed by the introduction of uncertainty regarding our error parameters, the likelihood of the data is unaltered from our previous formulation. Thus, our arc likelihood is given by

$$p(Y_{ij} | \Theta_{ij}, e^+, e^-) = \Theta_{ij} (Y_{ij} (1 - e^-) + (1 - Y_{ij}) e^-) + (1 - \Theta_{ij}) (Y_{ij} e^+ + (1 - Y_{ij}) (1 - e^+)) \quad (16)$$

and the joint likelihood for the entire data set is, as before, the product of the arc likelihoods:

$$p(\mathbf{Y} | \Theta, e^+, e^-) = \prod_{i=1}^N \prod_{j=1}^N p(Y_{ij} | \Theta_{ij}, e^+, e^-) \quad (17)$$

Note that while the assumption of uncertainty in error parameters does not alter our likelihood, it will nevertheless affect posterior inference (as we shall see presently). Intuitively, the reason for the former is that the likelihood of the observed data already conditions on the error parameters, and hence treats them as “fixed”. The role of stochasticity in errors, then, is seen in the computation of the posterior.

2.3.3 Computation of the Posterior

In the simple error model, we were able to express the posterior of the criterion graph in very simple terms; in the single observer model, this is complicated by the stochasticity of the error parameters, e^+ and e^- . To derive the joint posterior, then, we invoke the more common (proportionality) form of Bayes law, as follows:

$$p(\Theta, e^+, e^- | \mathbf{Y}) \propto p(\Theta) p(e^+) p(e^-) p(\mathbf{Y} | \Theta, e^+, e^-) \quad (18)$$

$$= \left(\prod_{i=1}^N \prod_{j=1}^N p(\Theta_{ij}) \right) p(e^+) p(e^-) \left(\prod_{i=1}^N \prod_{j=1}^N p(Y_{ij} | \Theta_{ij}, e^+, e^-) \right) \quad (19)$$

$$= \left(\prod_{i=1}^N \prod_{j=1}^N B(\theta_{ij}) \right) (Beta(e^+ | \alpha^+, \beta^+)) (Beta(e^- | \alpha^-, \beta^-)) \left(\prod_{i=1}^N \prod_{j=1}^N p(\mathbf{Y} | \Theta, e^+, e^-) \right) \quad (20)$$

Now, the impact of uncertainty in the error parameters becomes clear: we cannot consider the probability of the criterion graph without considering the probability of the error parameters which give rise to that inference, and neither can we consider the probability of a particular pair of error parameters without taking into account the probability of the structural inference such a choice would induce. As asserted, the two problems are inseparable; this intuition is laid bare in the form of the joint posterior.

Given, then, that Equation 18 does not lend itself to direct computation, how may we employ it in practice? One approach, which is facilitated by the form of the posterior, is to simulate posterior draws from the joint posterior using a Gibbs sampler, and to in turn use these draws to estimate posterior quantities of interest. The Gibbs sampler - a Markov Chain Monte Carlo (MCMC) method - works by taking a series of draws from the full conditionals of the posterior, constructing a Markov chain whose equilibrium distribution converges to that of the joint posterior. Details concerning the use of MCMC techniques can be found in Gelman et al. (1995); our discussion here will focus exclusively on the conditional distributions which are necessary for its implementation.

The first of the conditional distributions we require is that of the criterion graph, conditional on the realizations of the two error parameters and of the data matrix. This conditional probability reduces to the simple product of the arc posteriors, which we derived in Equation 9 for the simple error model. The joint conditional of the criterion graph, then, is given by

$$p(\Theta | e^+, e^-, \mathbf{Y}) \sim \prod_{i=1}^N \prod_{j=1}^N \frac{\theta_{ij} (Y_{ij} (1 - e^-) + (1 - Y_{ij}) e^-)}{\theta_{ij} (Y_{ij} (1 - e^-) + (1 - Y_{ij}) e^-) + (1 - \theta_{ij}) (Y_{ij} e^+ + (1 - Y_{ij}) (1 - e^+))} \quad (21)$$

which lends itself quite readily to simulation by drawing each arc as an independent Bernoulli random variable with the probabilities given by the factors of Equation 21.

For the conditional probability of e^+ , we can follow a similar strategy; in particular, we can use the conjugate property of the Beta distribution to find the form of the conditional posterior directly. Note that the parameters of the conditional posterior are simple counts of successful and unsuccessful identifications of the nonexistence of ties, combined with the prior parameters. Specifically:

$$p(e^+ | \Theta, e^-, \mathbf{Y}) \sim \text{Beta} \left(\alpha^+ + \sum_{i=1}^N \sum_{j=1}^N (1 - \Theta_{ij}) Y_{ij}, \beta^+ + \sum_{i=1}^N \sum_{j=1}^N (1 - \Theta_{ij}) (1 - Y_{ij}) \right) \quad (22)$$

The computation of the above is quite straightforward. e^- , similarly, has a direct interpretation in terms of inference on the existence of ties; the conditional posterior is given by

$$p(e^- | \Theta, e^+, \mathbf{Y}) \sim \text{Beta} \left(\alpha^- + \sum_{i=1}^N \sum_{j=1}^N \Theta_{ij} (1 - Y_{ij}), \beta^- + \sum_{i=1}^N \sum_{j=1}^N \Theta_{ij} Y_{ij} \right) \quad (23)$$

In both cases, one would simulate draws from the posterior by taking independent Beta draws with the parameters given by Equations 22 and 23; this is readily accomplished using most statistical computing packages, either directly or by drawing from $\frac{\chi_{2\alpha}^2}{\chi_{2\alpha}^2 + \chi_{2\beta}^2}$ (where the two 2α -degree χ^2 variables reflect the same draw and the α and β draws are independent). By iteratively drawing from the conditional posteriors of the criterion graph, e^+ , and e^- , one can then simulate draws from the joint posterior distribution. (For a detailed description of the algorithm involved, see the reference above.)

2.4 Multiple Error Probabilities (Multiple Observer Model)

Having seen models for both fixed and pooled error probabilities, we are now ready to proceed to the more general case in which our network data is generated by a process involving multiple uncertain error probabilities. The canonical example of such a situation is that of a cognitive social structure, in which each actor within the structure reports the full set of relations among actors (presumably with uncertain error probabilities which vary by informant). Such a data set provides us with multiple observations across both arcs and actors; thus, we have a fair amount of leverage in drawing inferences regarding both informant accuracy and the criterion graph. Within the text that follows, then, we shall assume that our data takes the form of a CSS except as indicated otherwise. While the model considered can be applied to other data structures, this provides the example which is most likely to be of use to researchers in the field.

2.4.1 Assignment of Priors

The priors required for the multiple observer model are, as one would expect, a simple generalization of those employed in the single observer model. Whereas before we had only two sets of hyperparameters, here we have twice as many sets as the number of actors; the e^+ priors, for instance, under the assumptions of conjugacy and conditional independence, are given by

$$e_i^+ \sim \text{Beta}(\alpha_i^+, \beta_i^+), i \in \{1, 2, \dots, N\} \quad (24)$$

$$p(e^+) = \prod_{i=1}^N p(e_i^+) \quad (25)$$

Note that the hyperparameters are not required to be identical for all actors. Indeed, if individuating information (e.g., from past observation of a particular informant) is available, it should clearly be employed. Again with the same assumptions, we take the form of the e^- priors to be identical:

$$e_i^- \sim \text{Beta}(\alpha_i^-, \beta_i^-), i \in \{1, 2, \dots, N\} \quad (26)$$

$$p(\mathbf{e}^-) = \prod_{i=1}^N p(e_i^-) \quad (27)$$

It is worth pointing out, with respect to the above, that implicit in our assumptions is the requirement that inferences individual error parameters affect each other only indirectly via the estimation of the criterion graph. A reasonable extension of this model might relax this assumption by taking the hyperparameters for the individual actors as being drawn from a hyperprior distribution; one candidate for such a distribution would be a Gamma, which has a number of desirable properties in such a role. An extension of this type, of course, would add yet another layer to the hierarchical model, and will not be pursued here. Nevertheless, it is a promising direction for future research.

In addition to the choice of prior distributions for the error parameters, of course, we are still left with the usual problem of defining priors for the criterion graph. The form for the criterion prior is unchanged from its previous incarnations:

$$\Theta_{ij} \sim B(\theta_{ij}) \quad (28)$$

$$p(\Theta) = \prod_{i=1}^N \prod_{j=1}^N p(\Theta_{ij}) \quad (29)$$

It is perhaps worth emphasizing that, though our data array contains N separate graphs, we are still interested in inferring a single criterion structure; hence the form of the criterion prior, above. While one could also attempt to utilize a very different notion of the criterion, drawing a different graph for each observer (and thereby reinterpreting the “error” parameters as reflecting propensity to misreport one’s cognitive beliefs, rather than to misreport an externally observable interaction pattern), this would simply reduce to the single observer model already considered. As stated in the introduction, we here follow the classical interpretation of self-report data, in which even cognitive social structures are seen as emerging from observation of a single criterion graph. Cognitivist uses of the present model are of potential interest, but are beyond the scope of this paper.

2.4.2 Assumed Likelihood

The assumed likelihood for the multiple observer model is a straightforward extension of the likelihood for the single observer model. Each observation is,

as usual, a Bernoulli mixture, with the distinction in this case being the use of separate error probabilities for each observer. Following Equation 7, then, the arc likelihood is given by

$$p(Y_{ijk} | \Theta_{ij}, e_k^+, e_k^-) = \Theta_{ij} (Y_{ijk} (1 - e_k^-) + (1 - Y_{ijk}) e_k^-) + (1 - \Theta_{ij}) (Y_{ijk} e_k^+ + (1 - Y_{ijk}) (1 - e_k^+)) \quad (30)$$

Assuming that our data set takes the form of a CSS (in which each actor acts as an observer), the joint likelihood of the data is simply the product of the individual arc likelihoods under our standard assumption of conditional independence. Formally, this is given by

$$p(\mathbf{Y} | \Theta, \mathbf{e}^+, \mathbf{e}^-) = \prod_{i=1}^N \prod_{j=1}^N \prod_{k=1}^N p(Y_{ijk} | \Theta_{ij}, e_k^+, e_k^-) \quad (31)$$

which can be seen quite readily to be a simple generalization of Equation 17 to the multiple observer case.

As noted, the above likelihood assumes that the data being analyzed takes the form of a CSS. In general, however, this is not required; any number of observers may be utilized, provided that the index k in the above equations counts over the number of observers. Thus, one could readily use the above model (with the given modification) to take into account reports of both participants and external observers in a given network. A caveat is in order here, however: the model presented above assumes that the data collection mechanism - including the selection of observers - is ignorable with respect to inferences on Θ , \mathbf{e}^+ , and \mathbf{e}^- . In the case of non-ignorable designs, the data generation model should be modified to reflect the consequences of the data collection procedure; for a more general discussion of the problems of ignorability, see Gelman et al. (1995).

2.4.3 Computation of the Posterior

Having determined the joint likelihood of the data, we are now in a position to write down the posterior. Using Equation 31 and Bayes' law:

$$p(\Theta, \mathbf{e}^+, \mathbf{e}^- | \mathbf{Y}) \propto p(\Theta) p(\mathbf{e}^+) p(\mathbf{e}^-) p(\mathbf{Y} | \Theta, \mathbf{e}^+, \mathbf{e}^-) \quad (32)$$

$$= \left(\prod_{i=1}^N \prod_{j=1}^N p(\Theta_{ij}) \right) \left(\prod_{i=1}^N p(e_i^+) \right) \left(\prod_{i=1}^N p(e_i^-) \right) \left(\prod_{i=1}^N \prod_{j=1}^N \prod_{k=1}^N p(Y_{ijk} | \Theta_{ij}, e_k^+, e_k^-) \right) \quad (33)$$

$$= \left(\prod_{i=1}^N \prod_{j=1}^N B(\theta_{ij}) \right) \left(\prod_{i=1}^N \text{Beta}(e_i^+ | \alpha_i^+, \beta_i^+) \right) \left(\prod_{i=1}^N \text{Beta}(e_i^- | \alpha_i^-, \beta_i^-) \right) \left(\prod_{i=1}^N \prod_{j=1}^N \prod_{k=1}^N p(\mathbf{Y} | \Theta, \mathbf{e}^+, \mathbf{e}^-) \right) \quad (34)$$

Analytically, the joint posterior given by Equation 32 is somewhat difficult to work with. However, as in the case of the single observer model, we can exploit

the form of the posterior to easily derive the full conditionals of the model, which in turn happen to be appropriate for the straightforward implementation of a Gibbs sampler. In the case of the criterion graph, for instance, the joint conditional probability of the posterior is given by

$$p(\Theta | e^+, e^-, Y) \sim \prod_{i=1}^N \prod_{j=1}^N \frac{\theta_{ij} \prod_{k=1}^N (Y_{ijk} (1 - e_k^-) + (1 - Y_{ijk}) e_k^-)}{\theta_{ij} \prod_{k=1}^N (Y_{ijk} (1 - e_k^-) + (1 - Y_{ijk}) e_k^-) + (1 - \theta_{ij}) \prod_{k=1}^N (Y_{ijk} e_k^+ + (1 - Y_{ijk}) (1 - e_k^+))} \quad (35)$$

Note that this can be seen simply as a straightforward application of Bayes' law to each arc, given the arc likelihood of Equation 30. Computationally, we exploit this structure by drawing each arc separately using the probabilities given by Equation 35. A similar exploitation is possible for the probability of false positives,

$$p(e^+ | \Theta, e^-, Y) \sim \prod_{k=1}^N \text{Beta} \left(\alpha_k^+ + \sum_{i=1}^N \sum_{j=1}^N (1 - \Theta_{ij}) Y_{ijk}, \beta_k^+ + \sum_{i=1}^N \sum_{j=1}^N (1 - \Theta_{ij}) (1 - Y_{ijk}) \right) \quad (36)$$

and of false negatives,

$$p(e^- | \Theta, e^+, Y) \sim \prod_{k=1}^N \text{Beta} \left(\alpha_k^- + \sum_{i=1}^N \sum_{j=1}^N \Theta_{ij} (1 - Y_{ijk}), \beta_k^- + \sum_{i=1}^N \sum_{j=1}^N \Theta_{ij} Y_{ijk} \right) \quad (37)$$

as each is conditionally distributed Beta with parameters given by counts of tie outcomes.

To implement the Gibbs sampler, then, we alternately take draws from the conditional posteriors of Θ , e^+ , and e^- given some set of initial conditions. This can be readily accomplished using standard statistical computing tools; see Gelman et al. (1995) for more details. Using the Gibbs sampler, we then simulate taking draws from the joint posterior. These draws can be used to estimate posterior quantities of interest in the usual fashion.

2.5 Estimating Network Variables from Posterior Distributions

We have now developed three different Bayesian models of the network inference/informant accuracy problem, and have shown in each case how we may simulate draws from the relevant posterior distribution. For certain applications (e.g., estimating the accuracy of particular actors) this may be sufficient; in general, however, our primary interest will be in various quantities which are derived from the criterion graph itself. Given that we are uncertain about

the criterion graph, these quantities will necessarily be random variables, with distributions which depend upon the posterior of the criterion. In the following section, then, we shall demonstrate the application of the network posterior to three standard inference problems: identification of the maximally probable criterion graph; estimation of graph and node level indices for the criterion graph; and metric inference between criterion graphs.

To illustrate these procedures, we will utilize posterior draws from a CSS of Krackhardt (1987) under the multiple observer model. The data set in question contains reported advice seeking relations among 21 management personnel within a high-tech firm, and is a "classic" CSS study. For purposes of illustration, the network prior for this analysis was chosen such that $\Theta_{ij} = 0.3$ for all arcs, and all individual error parameters were given $Beta(3, 5)$ priors. Posterior draws were taken using a Gibbs sampler, with three Markov chains being employed, each having a burn-in of 500 iterations. After burn-in, 500 draws were taken from each chain, and the resulting data points were randomly reshuffled to remove any dependence between adjacent observations²⁰. The 1500 posterior draws derived from this process were then used in the analyses below.

2.5.1 Maximum Probability Networks

Possibly the most obvious question to ask, given a set of draws from the posterior of the criterion graph, is that of the maximum probability network: that is, the particular criterion graph which is most probable given the posterior distribution. This is derived fairly trivially from the posterior, as follows:

$$A_{ij}^* = \begin{cases} 1 & \text{if } \Theta_{ij} \geq 0.5 \\ 0 & \text{if } \Theta_{ij} < 0.5 \end{cases} \quad (38)$$

Note that we are able to convert the problem of finding the maximum probability graph into the problem of finding the maximally probable state for each arc due to the assumption of independence between arcs. For our illustrative data set, then, our estimated posterior for the criterion graph is given by

²⁰S code (written for the R statistical computing system) to fit this model is available from the author.

	a1	a2	a3	a4	a5	a6	a7	a8	a9	a10	a11	a12	a13	a14	a15	a16	a17	a18	a19	a20	a21
a1	0.00	0.54	0.00	0.00	0.00	1.00	0.66	0.35	0.00	0.96	0.76	0.98	0.99	0.00	0.96	1.00	1.00	0.94	0.91	1.00	0.64
a2	0.34	0.00	0.00	0.17	0.00	1.00	0.00	1.00	0.00	0.28	0.52	0.99	1.00	1.00	0.01	1.00	1.00	0.00	0.97	1.00	0.14
a3	0.68	0.98	0.00	0.99	0.00	1.00	0.96	0.78	0.15	0.93	1.00	0.99	0.99	0.07	0.95	1.00	1.00	0.26	0.99	0.99	0.19
a4	0.10	0.71	0.99	0.00	0.94	1.00	0.99	0.80	0.99	0.94	0.97	0.99	1.00	0.99	0.99	1.00	1.00	0.53	0.99	1.00	0.92
a5	0.00	0.47	0.83	0.02	0.00	1.00	0.94	0.00	0.90	0.01	0.94	0.71	1.00	0.00	0.03	1.00	1.00	0.99	0.24	0.99	0.98
a6	0.00	0.99	0.90	0.92	0.99	0.00	0.99	0.99	0.96	0.99	1.00	0.99	1.00	0.99	0.06	1.00	1.00	1.00	0.99	1.00	0.99
a7	0.00	0.99	0.02	0.91	0.99	1.00	0.00	0.99	0.00	0.99	1.00	0.00	1.00	0.99	0.99	1.00	1.00	0.82	0.95	0.99	0.64
a8	0.08	0.98	0.50	0.11	0.60	1.00	0.96	0.00	0.99	0.98	1.00	0.82	1.00	0.05	1.00	1.00	1.00	0.90	0.99	1.00	0.03
a9	0.00	0.92	0.07	0.01	0.00	1.00	0.93	0.01	0.00	0.00	0.95	0.00	1.00	0.00	0.99	0.76	1.00	0.99	0.06	1.00	0.00
a10	0.00	0.98	0.00	0.06	0.81	1.00	0.99	0.86	0.00	0.00	1.00	0.44	1.00	0.99	1.00	1.00	1.00	0.84	0.99	1.00	0.96
a11	0.00	0.98	0.00	0.01	0.00	1.00	0.99	0.95	0.00	0.79	0.00	1.00	1.00	0.99	0.02	1.00	1.00	0.94	0.99	0.99	0.52
a12	0.02	0.99	0.93	0.99	0.98	1.00	0.99	1.00	0.99	0.99	1.00	0.00	1.00	0.99	1.00	1.00	1.00	0.99	0.99	1.00	0.99
a13	0.00	0.01	0.94	0.00	0.00	0.73	0.97	0.00	0.00	0.00	0.08	0.00	0.00	0.00	0.03	0.19	1.00	0.00	0.00	0.99	0.00
a14	0.00	0.99	0.00	0.00	0.99	1.00	0.00	0.00	0.10	0.00	1.00	1.00	1.00	0.00	0.00	1.00	1.00	0.00	0.92	0.00	0.00
a15	0.00	0.01	0.99	0.00	0.07	1.00	0.97	0.01	0.40	0.04	1.00	0.99	1.00	0.54	0.00	1.00	1.00	0.61	0.30	0.99	0.00
a16	0.39	0.97	0.98	0.30	0.87	1.00	0.50	0.63	0.99	0.96	1.00	0.99	1.00	0.96	0.99	0.00	1.00	0.98	0.99	1.00	0.00
a17	0.00	0.99	0.00	0.59	0.00	1.00	0.99	0.99	0.01	0.99	1.00	0.99	1.00	0.00	1.00	1.00	0.00	1.00	0.99	0.42	0.98
a18	0.00	0.90	0.00	0.00	1.00	0.99	0.00	0.98	0.71	0.00	0.99	0.99	1.00	0.94	0.95	1.00	1.00	0.00	0.43	0.54	0.00
a19	0.00	0.98	0.05	0.00	0.00	0.99	0.93	0.01	0.67	0.00	0.91	0.01	1.00	0.99	0.89	0.97	1.00	0.00	0.00	0.97	0.00
a20	0.00	0.56	0.00	0.00	0.00	1.00	0.99	0.01	0.00	0.93	0.94	0.00	1.00	0.00	0.99	0.96	1.00	0.00	0.93	0.00	0.00
a21	0.89	0.99	0.53	0.41	0.88	1.00	0.97	1.00	0.96	0.97	1.00	0.99	1.00	0.98	0.99	1.00	1.00	0.99	0.99	0.94	0.00

and thus the maximum probability graph is

	a1	a2	a3	a4	a5	a6	a7	a8	a9	a10	a11	a12	a13	a14	a15	a16	a17	a18	a19	a20	a21	
a1	0	1	0	0	0	1	1	0	0	1	1	1	1	0	1	1	1	1	1	1	1	
a2	0	0	0	0	0	1	0	1	0	0	1	1	1	1	0	1	1	1	0	1	1	0
a3	1	1	0	1	0	1	1	1	0	1	1	1	1	0	1	1	1	0	1	1	0	
a4	0	1	1	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
a5	0	0	1	0	0	1	1	0	1	0	1	1	1	0	0	1	1	1	0	1	1	
a6	0	1	1	1	1	0	1	1	1	1	1	1	1	1	0	1	1	1	1	1	1	
a7	0	1	0	1	1	1	0	1	0	1	1	0	1	1	1	1	1	1	1	1	1	
a8	0	1	0	0	1	1	1	0	1	1	1	1	0	1	1	1	1	1	1	1	0	
a9	0	1	0	0	0	1	1	0	0	0	1	0	1	0	1	1	1	1	0	1	0	
a10	0	1	0	0	1	1	1	1	0	0	1	0	1	1	1	1	1	1	1	1	1	
a11	0	1	0	0	0	1	1	1	0	1	0	1	1	0	1	1	1	1	1	1	1	
a12	0	1	1	1	1	1	1	1	1	1	1	0	1	1	1	1	1	1	1	1	1	
a13	0	0	1	0	0	1	1	0	0	0	0	0	0	0	0	0	1	0	0	1	0	
a14	0	1	0	0	1	1	0	0	0	0	1	1	1	0	0	1	1	0	1	0	0	
a15	0	0	1	0	0	1	1	0	0	0	1	1	1	1	0	1	1	1	0	1	0	
a16	0	1	1	0	1	1	1	1	1	1	1	1	1	1	1	0	1	1	1	1	0	
a17	0	1	0	1	0	1	1	1	0	1	1	1	1	0	1	1	0	1	1	0	1	
a18	0	1	0	0	1	1	0	1	1	0	1	1	1	1	1	1	1	0	0	1	0	
a19	0	1	0	0	0	1	1	0	1	0	1	0	1	1	1	1	1	0	0	1	0	
a20	0	1	0	0	0	1	1	0	0	1	1	0	1	0	1	1	1	0	1	0	0	
a21	1	1	1	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	0	

It should be noted that the maximum probability graph is a point estimate, and as such necessarily discards much of the full information of the posterior distribution. For many applications, then, it may be more prudent to use the posterior distribution, rather than the maximum probability graph. Nevertheless, when it is commonly desirable to have some particular estimate of the criterion structure, the maximum probability graph will often be the logical choice²¹.

²¹In fact, selection of point estimates should be performed via a formal decision procedure; this is, however, beyond the scope of this paper.

2.5.2 Graph and Node Level Indices

The vast majority of classical network analysis is founded on the substantively important secondary indices such as degree, betweenness, and closeness (at the nodal level) and centralization, hierarchy, and connectedness (at the graph level). Many of these indices are highly sensitive to minor changes in network structure; yet, traditional network analysis has generally examined these indices under the assumption of error-free data. One useful application of the network posterior, then, is to permit estimation of graph and node-level indices in the presence of measurement error. The quantification of uncertainty, in particular, is useful here: even where the posterior is diffuse, knowledge of this fact may illuminate subsequent analysis. This last is particularly true of *comparisons* of network indices, in which changes in the uncertainty associated with point estimates may lead to substantively distinct conclusions.

By way of illustration, then, we here provide summaries of posterior intervals for Freeman degree centralization and nodal degree for the Krackhardt advice network. Note how, even for a simple measure such as degree, quantification of the uncertainty associated with the criterion graph can clearly affect posterior inference. In addition to the simple summaries shown here, it is fairly trivial to compute quantities such as, for instance, the posterior probability that actor 4 has a higher Freeman degree than actor 5 (in this case, the probability in question is approximately 0.798). Such an approach is clearly more powerful than classical methods (which can only examine relative likelihoods), and is obviously more robust than traditional network techniques which treat the data under examination as error-free. The Bayesian modeling approach then, is useful not only because of the ease with which it allows us to construct theoretically motivated models, but also because of the inferential uses to which the output of those models can be put.

Index	Min	1st Q	Median	Mean	3rd Q	Max
Degree Cent	0.2171	0.2882	0.3053	0.3042	0.3211	0.3908
Degree, a1	9	14	15	15.23	16.25	21
Degree, a2	21	25	26	26.40	28.00	32
Degree, a3	17	22	23	22.59	24.00	27
Degree, a4	17	22	23	23.37	25.00	29
Degree, a5	15	20	21	21.20	22.00	25
Degree, a6	34	37	38	37.51	38.00	39
Degree, a7	26	30	31	31.04	32.00	34
Degree, a8	21	25	26	26.36	27.00	31
Degree, a9	14	18	19	18.47	19.00	22
Degree, a10	22	26	27	26.72	28.00	30
Degree, a11	27	30	31	31.24	32.00	35
Degree, a12	29	32	33	32.82	33.00	35
Degree, a13	22	24	25	24.92	25.00	28
Degree, a14	18	20	21	20.52	21.00	23
Degree, a15	21	24	25	24.80	26.00	30
Degree, a16	30	34	35	35.37	36.00	39
Degree, a17	32	33	34	33.96	34.00	36
Degree, a18	19	24	25	25.24	26.00	29
Degree, a19	19	25	26	26.05	27.00	30
Degree, a20	24	26	27	27.11	28.00	30
Degree, a21	21	25	27	26.50	28.00	31

2.5.3 Metric Distances Between Graphs

An emerging alternative to the use of summary indices to compare social structures is comparison using metric distances (Banks and Carley, 1994; Butts and Carley, 1998; Butts, 1998), generally either the Hamming distance or a generalization such as the structural distance (Butts and Carley, 1998; Butts, 1998). Such distances may be employed directly as a measure of difference (Banks and Carley, 1994) or indirectly to facilitate procedures such as cluster analysis of graphs (Butts, 1998). While past work in this area has often assumed that the social structures to be compared are perfectly known, it is possible to generalize the approach to the case in which the structures to be compared are Bernoulli graphs. In such a case, it is fairly trivial to derive posterior estimates for the moments of the Hamming distance between the structures in question. In particular, the expectation of the Hamming distance is given by

$$E(H(G_1, G_2)) = \sum_{i=1}^N \sum_{j=1}^N (\Theta_{1ij}(1 - \Theta_{2ij}) + (1 - \Theta_{1ij})\Theta_{2ij}) \quad (39)$$

and its variance is simply

$$\begin{aligned} \text{Var}(H(G_1, G_2)) = & \quad (40) \\ & \sum_{i=1}^N \sum_{j=1}^N (\Theta_{1ij} + \Theta_{2ij} - 4\Theta_{1ij}\Theta_{2ij} - \Theta_{1ij}^2 - \Theta_{2ij}^2 + 4\Theta_{1ij}^2\Theta_{2ij} + 4\Theta_{1ij}\Theta_{2ij}^2 - 4\Theta_{1ij}^2\Theta_{2ij}^2) \end{aligned}$$

Note that both of the above make use of the independence property of Bernoulli graphs in a very straightforward fashion. In the event that one wishes to examine the entire distribution of distances, one can simply estimate the distribution from the distribution of observed Hamming distances among posterior draws. The procedure is equivalent to that employed above for index distributions, and is fairly trivial to implement. For our illustrative advice network and the friendship network collected on the same population²², for instance, we find the following distance distribution:

Min	1st Q	Median	Mean	3rd Q	Max
185.0	215.0	233.0	229.4	242.0	273.0

As the maximum Hamming distance in this case is 441, we can see that only about half of the arcs are shared between these two structures; clearly they are quite distinct. Subsequent analysis might compare this distribution to distributions from null models, or might use the estimated Hamming distance as an input to a procedure such as cluster analysis. A number of possibilities exist, depending on the theoretical uses to which the data is to be put.

²²The same procedure was used to take posterior draws in this case as for the advice network

2.5.4 Estimated Informant Accuracy

While we may often be interested solely in the criterion graph (in which case we estimate individual accuracy parameters only because such are necessary for the previous problem), this is not always the case; in some cases, it is the accuracy of individuals which is of interest (e.g., Krackhardt, 1990). Given a series of draws from the posterior distribution, it is quite trivial to examine the posterior distribution of individual error parameters, and thereby to gain a sense of actors' ability to perceive their social surroundings. These distributions may in turn be used to create point estimates which may be employed in subsequent analyses, though as always it is preferable to use direct posterior draws for this purpose where possible so as to avoid losing the distributional information contained therein.

For our advice network, then, we find the following estimated posterior quantiles for each parameter:

Probability of False Positives (e^+)						
	Min	1stQ	Median	Mean	3rdQ	Max
a1	0.3556	0.4571	0.4844	0.4846	0.5118	0.6076
a2	0.6491	0.7421	0.7656	0.7642	0.7884	0.8600
a3	0.3800	0.4896	0.5184	0.5178	0.5454	0.6824
a4	0.1833	0.2794	0.3040	0.3042	0.3298	0.4567
a5	0.4419	0.5636	0.5939	0.5938	0.6233	0.7377
a6	0.1877	0.3063	0.3310	0.3311	0.3562	0.4533
a7	0.4659	0.5575	0.5852	0.5858	0.6137	0.7007
a8	0.2943	0.3806	0.4070	0.4077	0.4348	0.5526
a9	0.2358	0.3163	0.3421	0.3433	0.3688	0.4799
a10	0.2753	0.3643	0.3923	0.3913	0.4165	0.5378
a11	0.4120	0.5293	0.5570	0.5562	0.5841	0.6799
a12	0.1503	0.2283	0.2519	0.2535	0.2766	0.3863
a13	0.2554	0.3494	0.3772	0.3767	0.4015	0.5050
a14	0.5806	0.6980	0.7219	0.7222	0.7489	0.8375
a15	0.2384	0.3181	0.3444	0.3458	0.3725	0.4846
a16	0.2569	0.3592	0.3855	0.3859	0.4134	0.5108
a17	0.2075	0.2828	0.3069	0.3087	0.3320	0.4471
a18	0.6872	0.7824	0.8044	0.8035	0.8241	0.8919
a19	0.4124	0.5485	0.5790	0.5789	0.6089	0.7151
a20	0.2774	0.4058	0.4326	0.4330	0.4608	0.5685
a21	0.3633	0.4617	0.4884	0.4886	0.5150	0.6636

Probability of False Negatives (e^-)						
	Min	1stQ	Median	Mean	3rdQ	Max
a1	0.7489	0.8265	0.8440	0.8431	0.8605	0.9134
a2	0.3107	0.3879	0.4082	0.4088	0.4286	0.5268
a3	0.8759	0.9251	0.9360	0.9348	0.9465	0.9741
a4	0.7445	0.8165	0.8334	0.8322	0.8491	0.9028
a5	0.9120	0.9511	0.9593	0.9583	0.9668	0.9866
a6	0.6909	0.7579	0.7755	0.7745	0.7931	0.8564
a7	0.4038	0.4943	0.5145	0.5143	0.5361	0.6065
a8	0.7547	0.8242	0.8403	0.8395	0.8564	0.9181
a9	0.9281	0.9634	0.9709	0.9698	0.9771	0.9958
a10	0.8007	0.8708	0.8856	0.8846	0.8988	0.9453
a11	0.6780	0.7517	0.7714	0.7695	0.7877	0.8485
a12	0.7826	0.8494	0.8639	0.8635	0.8778	0.9251
a13	0.9400	0.9775	0.9831	0.9820	0.9876	0.9974
a14	0.4349	0.5095	0.5301	0.5299	0.5511	0.6332
a15	0.9130	0.9504	0.9591	0.9578	0.9666	0.9878
a16	0.8577	0.9112	0.9236	0.9227	0.9355	0.9721
a17	0.7733	0.8316	0.8477	0.8467	0.8626	0.9238
a18	0.4131	0.5001	0.5202	0.5203	0.5413	0.6185
a19	0.9083	0.9553	0.9647	0.9633	0.9728	0.9959
a20	0.8207	0.8808	0.8947	0.8938	0.9078	0.9470
a21	0.4860	0.5786	0.5964	0.5965	0.6159	0.6977

3 Discussion

We have, in previous sections of this paper, outlined a Bayesian modeling approach to (some aspects of) the problem of informant accuracy in social network analysis. In the course of this development, a number of deeper issues have arisen which, while important to the progress of our research, are somewhat tangential to the specific focus of the present work. Despite this, it behooves us to give at least some consideration to two particularly important issues: the concept validity of the criterion graph (on which the present enterprise obviously depends); and the problem of improving data collection for informant self-reports in social network research. This section, then, will initiate a discussion of these two matters which, if insufficient to resolve them, will hopefully serve to highlight them for subsequent research.

3.1 Concept Validity of the Criterion Graph

“If a group of 10 persons were all asleep and each person were dreaming of talking to at least one other person in the group, then is there a group structure to be uncovered?” (Killworth and Bernard, 1979/80)

As we have noted, the modeling approach utilized here depends critically upon the assumption that there exists *some* structure which accounts for the commonalities in informant responses. Rather than enmeshing ourselves in a larger debate about what it means to speak of “real” social structures (and whether some might be more “real” than others), we have taken the more restrained position of asserting the notion of the criterion graph purely as a useful construct which is hypothesized to account for shared variance. Even this, however, does not entirely extricate us from the dilemma raised by the Killworth

and Bernard comment above. If we find that we are dealing with relations which are purely ascriptive - which are defined purely in terms of actor reports - can we say that we are dealing with social networks at all? From the cognitivist perspective, the answer to this question is clearly yes: if the perceptions (or, from a slightly more behavioristic stance, reports) of individuals are of potential interest to us, then the question of whether an independently verifiable criterion structure can be said to exist is irrelevant. This raises some difficult epistemological (and hence methodological) questions, however. When *is* it sensible to speak of the existence (in at least a hypothetical sense) of the criterion graph? The question is not an idle one. If we, for instance, apply models which presume informant reports to be related to a central structure in situations for which such an assumption is invalid, then the inferences drawn from such an application will be highly misleading at best. On the other hand, ignoring the possibility of a criterion structure where one may be reasonably asserted may substantially limit our ability to draw predictive inferences regarding the social world. As BKS (1979/80) note, many processes of interest (e.g., diffusion of information) depend on behavioral, not cognitive, networks. Despite the gloomy prognosis of BKS, work by Romney et al. (1986), Romney and Faust (1982), and Freeman (1992) (among others) suggests that it should be possible to extract at least some useful information about such networks from informant reports. The present work is in this tradition, but recommends caution: we do not consider the existence of the criterion graph to be a trivial assumption, and recognize that the applicability of our approach depends upon the validity of the criterion concept. Further theoretical and empirical development of the foundations of network analysis per se - and, most importantly, of the *conceptual* foundations of our proposed subject matter - will be necessary if fruitful methodological work in this area is to continue.

3.2 Suggestions for Improved Data Collection

In the introduction to this paper, four general problems relating to the informant accuracy in network research were mentioned. As indicated, we have focused primarily on the last, namely the development of inferential techniques for quantifying (and hopefully reducing) the uncertainty inherent in this form of data. This endeavor, however, is strongly related to another: the development and deployment of data collection strategies which facilitate the reduction of uncertainty regarding quantities of interest. While a variety of issues are involved in this pursuit, we shall here constrain ourselves to a single matter, namely sampling strategies employed in eliciting network data from informants.

The standard procedure for eliciting informant reports regarding social structure is generally to provide each member of the social network with a survey instrument which elicits his or her ties to others. In some cases, this is extended by inquiring into the existence of incoming ties, or (as in ego net research) asking for ties among adjacent alters; nevertheless, these approaches are less uniformly deployed than the first. The sub-optimality of this standard procedure from the point of view of network inference can easily be appreciated by counting the

number of *repeated observations* accorded each arc: plainly, the standard procedure counts each arc but once, and even extending this to incoming ties provides only two observations. With so few observations, it is hardly surprising that our ability to infer social structure is so problematic! Given even minor deviations from perfect reporting, procedures which supply only one to two replications per arc are unlikely to provide sufficient data for reasonable inference on the criterion graph.

One alternative to this procedure is the elicitation of cognitive social structures. Though developed explicitly as a cognitivist tool, the CSS instrument is highly desirable from a classical perspective due to the fact that it provides a large number of repeated observations - $|V(G)|$, to be precise - on each arc. Further, the fact that the CSS elicits observations from all network members means that it can be considered to be an ignorable design so long as the complete data set is defined in terms of all participant observations. CSS data, then, is of much greater potential value to the network analyst than traditional data, particularly when employed in conjunction with inferential tools which allow for inference across arcs and across actors.

For all its benefits, there is a clear drawback to the CSS design: due to the fact that each informant is asked to report on all arcs, the number of items on a CSS instrument increases on the order of $|V(G)|^2$. This polynomial growth stands in sharp contrast to the linear growth of instrument complexity for traditional approaches, and severely limits the size of networks which can be examined in this fashion. A 50 node network, for instance, requires each subject to consider 2500 items for every relation examined; plainly this stretches the limits of informant endurance. Unfortunately, we are often interested in networks which are of even larger sizes, which all but eliminates the CSS from use in a wide range of settings.

Given the above problems, it may be sensible to consider an alternative to both the traditional and CSS data collection strategies, particularly for large networks. Such an alternative should be ignorable, should provide multiple observations on each arc, and should provide multiple observations on each informant, while maintaining linear complexity in network size. One data collection strategy which fulfills these requirements is what we shall here call an *M-replication balanced arc sampling design*; while we will not consider all of its properties in detail, we shall outline the basic procedure by which it may be employed.

The core intuition of the *M-replication balanced arc sampling design* is that if one desires to have M observations on each arc in a directed graph, one need only ask each informant to supply $M(|V(G)|)$ observations. (This follows from the fact that $|V(G)|$ subjects reporting $M(|V(G)|)$ arcs results in $M(|V(G)|^2)$ observations, enough to allow M per arc.) The challenge, then, is to allocate the arcs sampled in such a way as to maintain ignorability. One simple means of doing so is to randomly allocate arcs to instruments such that A) each informant is given exactly $M(|V(G)|)$ arcs on which to report, and B) each arc is reported on exactly M times. Such a design is then ignorable, as observations not included are missing at random (see Gelman et al. (1995)), and balanced

(as each informant and arc contribute equally to the joint likelihood). With the M -replication balanced arc sampling design, then, one can gain many of the advantages of a CSS design (albeit on a more limited basis) without incurring the same complexity penalty. Inferential methods such as those discussed here can then be employed to estimate the criterion graph, which can then be used for classical network analysis purposes.

4 Conclusion

The dual problems of network inference and informant accuracy pose central methodological challenges for network analysis. While the particular approach employed in dealing with these problems depends on underlying epistemological assumptions, models which assume that informant reports stem from a single criterion graph may be useful in a wide range of circumstances. Given the dimensionality and data efficiency challenges posed by simultaneous estimation of the criterion graph and individual accuracy parameters, a hierarchical Bayesian approach has much to recommend it. Here, we have developed a family of such models, and have shown how they may be applied to the analysis of network data, both for the purposes of direct estimation and for the quantification of uncertainty in derived quantities such as network indecies. We have discussed the assumptions implicit in the use of these or other models, and have suggested data collection strategies which will facilitate estimation of network variables. Clearly, the network inference/informant accuracy problem is a serious one, which will require a combination of theoretical, methodological, and empirical research to address. It is hoped that the present work will contribute to this development by providing a suite of theoretically motivated methods with clear applicability to basic problems of empirical research in the field of network analysis.

5 References

- Anderson, B.S., Butts, C., and Carley, K. (1999). "The Interaction of Size and Density with Graph Level Indecies." *Social Networks*, 21(3), 239-267.
- Banks, D., and Carley, K. (1994). "Metric Inference for Social Networks." *Journal of Classification*, 11, 121-149.
- Bernard, H.R., and Killworth, P. (1977). "Informant Accuracy in Social Network Data II." *Human Communication Research*, 4(1), 3-18.
- Bernard, H.R., Killworth, P., Kronenfeld, D., and Sailer, L. (1984). "The Problem of Informant Accuracy: The Validity of Retrospective Data." *Annual Review of Anthropology*, 13, 495-517.

- Bernard, H.R., Killworth, P., and Sailer, L. (1979/80). "Informant Accuracy in Social Network Data IV: A Comparison of Clique-Level Structure in Behavioral and Cognitive Network Data." *Social Networks*, 2, 191-218.
- Butts, C. (1998). "Cluster Analysis of Unlabeled Structures." ICES Research Report 88-04-98, Carnegie Mellon University.
- Butts, C., and Carley, K. (1998). "Canonical Labeling to Facilitate Graph Comparison." ICES Research Report 88-06-98, Carnegie Mellon University.
- Calloway, M., Morrissey, J.P., and Paulson, R.I. (1993). "Accuracy and Reliability of Self-Reported Data in Interorganizational Networks." *Social Networks*, 15, 377-398.
- Carley, K., and Krackhardt, D. (1996). "Cognitive Inconsistencies and Non-Symmetric Friendship." *Social Networks*, 18, 1-27.
- Freeman, L.C. (1992). "Filling in the Blanks: A Theory of Cognitive Categories and the Structure of Social Affiliation." *Social Psychology Quarterly*, 55(2), 118-127.
- Freeman, L.C., Romney, A.K., and Freeman, S.C. (1987). "Cognitive Structure and Informant Accuracy." *American Anthropologist*, 89, 310-325.
- Gelman, A., Carlin, J.B., Stern, H.S., and Rubin, D.B. (1995). *Bayesian Data Analysis*. London: Chapman and Hall.
- Hammer, M. (1985). "Implications of Behavioral and Cognitive Reciprocity in Social Network Data." *Social Networks*, 7, 189-201.
- Hildum, D.C. (1986). "'Competence' and 'Performance' in Network Structure." *Social Networks*, 8, 79-95.
- Killworth, P., and Bernard, H.R. (1976). "Informant Accuracy in Social Network Data." *Human Organization*, 35(8), 269-286.
- Killworth, P., and Bernard, H.R. (1979/80). "Informant Accuracy in Social Network Data III: A Comparison of Triadic Structure in Behavioral and Cognitive Data." *Social Networks*, 2, 10-46.
- Krackhardt, D. (1987). "Cognitive Social Structures." *Social Networks*, 9, 109-134.
- Krackhardt, D. (1990). "Assessing the Political Landscape: Structure, Cognition, and Power in Organizations." *Administrative Science Quarterly*, 35,

242-260.

Krackhardt, D., and Kilduff, M. (1999). "Whether Close or Far: Perceptions of Balance in Friendship Networks in Organizations." *Journal of Personality and Social Psychology*, 76(5), 770-82.

Robert, C.P. (1994). *The Bayesian Choice: A Decision-Theoretic Motivation*. New York: Springer-Verlag.

Romney, A.K., and Faust, K. (1982). "Predicting the Structure of a Communications Network from Recalled Data." *Social Networks*, 4, 285-304.

Romney, A.K., Weller, S.C., and Batchelder, W.H. (1986). "Culture as Consensus: A Theory of Culture and Informant Accuracy." *American Anthropologist*, 88(2), 313-338.

Wasserman, S., and Faust, K. (1994). *Social Network Analysis: Methods and Applications*. Cambridge: Cambridge University Press.