

## Why We Need an Epistemic Model of Social Networks

Using examples and unstructured intuitions that highlight the importance of knowledge and of beliefs, both individual and mutual, to the outcomes of social situations and interpersonal relations, we argue for the usefulness of explicit epistemic models of human interactions and networks. We introduce the notions of an epistemic state—that is, a link between individuals and propositions they may know or believe—and of an epistemic tie—that is, a connection between individuals' epistemic states: if Alpha knows Beta knows Gamma knows that the park is closed after dark, then there is a set of epistemic ties connecting Alpha, Beta, Gamma, and the proposition that the park is closed after dark, which is part of the epistemic structure of the situation. We show how the structure of epistemic networks—epinets—formed by such links among individuals and their beliefs is relevant to the dynamics of human interactions, and how the dynamics of these networks are critical elements of complex interpersonal narratives.

---

What must human agents know about what other humans—with whom they are connected—know in order for the resulting patchwork of ties among them to *function* as a social network? Suppose that an anonymous survey of the friendship network of a seven-person executive team reveals that Beth, Harry, and Martha form a clique, with each describing the others as “friends.” We designate the triad as a clique, rather than as a patchwork of ties, because we expect these three to exhibit some special forms of cohesion that may be evidenced by, for example, an above-average ability to coordinate, collaborate, communicate, and collude. In other words, we expect the triad to *function* as a clique: we expect each member to know—and know that the other two know—sensitive information about an event of mutual importance, or we expect that such sensitive information will quickly propagate within the triad.

What each knows of and about the others and their knowledge is the “epistemic glue” of the clique; it is what allows Beth to react to an unforeseen disaster in ways she knows Harry and Martha will find justifiable, and it is

what allows her to make sense of their intentions based on observing their reactions and knowing what they think about what she knows. The grammar is somewhat complicated, but its complexity closely tracks that of the phenomena we expect a clique to exhibit. This epistemic superstructure is what makes the clique a clique—an identifiable network substructure with very specific expected properties—rather than a patchwork of ties and connections that can offer no further insight or predictive value.

At a more fundamental level, what must human agents know or believe about what others know or believe for their interactions to have joint or shared sense and meaning and to lead to stable patterns of interpersonal behavior? Game theory has contributed a basic canon of coordination, cooperation, and collaboration “games” that require coherent mutual beliefs (players’ beliefs about other players’ beliefs, about other players’ beliefs about their own beliefs, and so on) whose “epistemic structures” can be analyzed to arrive at the preconditions for coordination, cooperation, collaboration, and even coherent communication. However, these neat analytical structures come at the cost of oversimplifying what humans believe and how they believe it as well as what they know and how they know it. States of knowing, like “oblivion” (not knowing and not knowing you do not know) or “forgetting” (knowing but not recalling what you know), are ruled out by assumptions such as those of “common priors” and “common knowledge,” even though these states are all important to the unfolding of real human interactions. Moreover, because the event spaces of game theory do not admit interpretations and shadings, the resulting analyses lack the subtlety required to understand that humans interpret “Can I pray while I smoke?” very differently from the way they interpret “Can I smoke while I pray?” The conjunctive “while” functions very differently in first-order logic from the way it functions in plain English.

The contemporary importance of epistemic moves and games to an understanding of social interactions is clear from the direction of technological progress and innovation. *Homo sapiens* is *Homo communicans* and makes use of the full range of methods for passing information-bearing signals to shape, control, and predict the social milieu of being in the world. Consider the “cc” (carbon copy) and “bcc” (blind carbon copy) functions of everyday e-mail, which act as levers for shaping the informational structure of an interaction: “cc” creates pools of mutual knowledge about the contents of a message and serves as an aggregation tool; “bcc” oppositely brackets cliques that are “in

the know” from individuals outside a circle of trust or power. But these are just the rudiments: new technologies allow senders to control a message after they have sent it—and possibly delete it—making it possible for them to deny ever having sent such a message even though they know the recipients know its contents; to “hack” into each other’s e-mail servers to access a critical message without the message’s author or recipient knowing that the hacker knows its contents; and to encrypt a message so that only intended recipients can decode it on the basis of access to the public or private key with which it has been encoded. The complexity of “interactive epistemology” has multiplied over the past few decades and continues to do so. A new language and new models are needed to understand the epistemic glue of human social interactions.

Although we are interested in building intuitive, yet precise, models of this epistemic glue, we are assuredly not pioneers of the epistemic dimension of social interactions. Nuanced treatments of epistemic structures and effects have appeared in the fields of artificial intelligence (Fagin et al. 1995), epistemic game theory (Aumann 1989), and analytic philosophy (Kripke 2011). Nor are we the first to point out that social networks (and social structures more generally) require descriptions sensitive to differences between what social agents think, and what we think social agents think, about such structures (Krackhardt 1987). What we are after is a tool kit for modeling, measuring, and manipulating the *epistemic glue* of human interactions and networks in ways that are as accessible to social network analysts as they are engaging to logicians, epistemic game theorists, and artificial intelligence researchers. We are building an application—an “app”—as much as we are theorizing, modeling, or philosophizing.

### **The Epistemic Structure of a “Friendship Tie”**

Because we are building an app as much as a theory, we need to become intimate with “user requirements”—that is, the kinds of uses to which our modeling tool kit may be put. To that end, consider the friendship tie between Beth and Harry in our earlier example. Beth “knows” Harry: she sees him daily, is familiar with his latest setbacks and successes, works with him on a joint project, and sees him socially about every other week. That is her longhand unpacking of the shorthand answer “Harry is my friend,” which she gave on the anonymous survey. Now what we want to know is this: when Beth needs

Harry to convey to her, quickly and covertly, a sensitive piece of information she believes he has received from one of his acquaintances with whom Beth has no connection, what must Beth know about Harry for her to have good reason to believe that he will come through with it?

The minimal set of beliefs about Harry that Beth needs to rationalize her expectations may include the following: she believes that Harry knows the information is useful to her, that it is important to her that their office mate, Martha, does not know Beth has come to know it, and perhaps that Harry knows Beth will not divulge her source after he has passed along the information. Complications can arise: if Beth knows Harry's boss knows of Harry's ties to Beth and is monitoring Beth's actions to detect any sign that Harry has leaked to Beth the information he was entrusted to hold in confidence, then Beth may have to believe Harry knows of this threat and trusts in her integrity and competence not to "blow his cover." Alternatively, she may have to believe Harry does *not* know about this threat (in which case she may choose to inform him of it as befits the level of trust she assumes they share).

In each case, there is a structure to the knowledge that these social agents "share" that is both intelligible and intuitive, although it grows quickly in logical complexity with the addition of new information and people. The structure of this "epistemic glue" is rendered intuitive and intelligible by the recursive and interactive nature of what this "social knowledge," as it should properly be called, relates to, which is often knowledge *about* knowledge: Beth's knowledge about Harry's knowledge, which includes Harry's knowledge about Beth's knowledge about Harry's knowledge, and so on. The structure of the epistemic glue is "interactive": it links not only an agent's mind to a proposition but also one agent's mind to a proposition via another agent's mind: Beth knows Harry knows that his boss is monitoring Beth's actions for any sign of information leaked by Harry.

Of course, it is not only knowledge but beliefs, conjectures, and even barely articulated hunches that we want to capture and address with our language. Beth may not know—by any acceptable use of the term "knowing"—that Harry knows that the piece of information he possesses should be transmitted to her in a way that guards against eavesdropping—but she may simply believe it for reasons having to do with a complex of other prior beliefs. Harry may merely "sense" that Beth needs him to transmit the information covertly, without really having fully articulated that hunch as a proposition.