

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

A Sociocognitive-Neuroeconomic Model of Social Information Communication: To Speak Directly or To Gossip

Permalink

<https://escholarship.org/uc/item/81f7v9b7>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 40(0)

Authors

Lee, Jeungmin

Kralik, Jerald D

Jeong, Jaeseung

Publication Date

2018

A Sociocognitive-Neuroeconomic Model of Social Information Communication: To Speak Directly or To Gossip

Jeungmin Lee^{1,2}, Jerald D. Kralik^{1,*}, Jaeseung Jeong^{1,2,*}

(*jerald.kralik@gmail.com, jsjeong@kaist.ac.kr)

¹Department of Bio and Brain Engineering, College of Engineering

²Program of Brain and Cognitive Engineering, College of Engineering

Korea Advanced Institute of Science and Technology (KAIST), Daejeon, 34141, South Korea

Abstract

Communication is a powerful means to disseminate social information, and gossip is an effective way of obtaining updated information about others. However, without a comprehensive theoretical framework of social communication, it is difficult to predict *a priori* when and why social information will be disseminated. There are general theories of human social interaction, however, they do not sufficiently capture the sociocognitive components underlying human decision-making in social settings. Therefore, we have developed a model of social communication, enabling the characterization of specific conditions under which social information will be spread: for example, when an agent should directly communicate with the target of the information, gossip it to others, or simply do nothing. We describe the model, the methods used to generate model predictions, and then list nine predictions derived from it as the current results. We next plan to test the predictions empirically and develop the model computationally.

Keywords: decision-making; theory of mind; social neuroscience; multi-agent system; artificial social intelligence

Introduction

People live in a complex, multi-agent world, and as such, sophisticated social intelligence is needed. Indeed, to make accurate predictions about others requires having a model of their minds — their beliefs, goals and intentions — and humans have evolved the ability to do so. Moreover, proper long-term social interaction also requires mental accounting of what you owe to others (from their help) and what others owe to you (from your help or their hindrance). At the same time, rich descriptions of all possible agents become computationally intractable, so our mental models are also necessarily limited. One strategy, for example, is to maintain more comprehensive models of those closest to us.

In any case, gathering information about others is crucial to maintain accurate models of them. Accordingly, communication is a critical means by which agents share and update information about each other. In fact, statistical assessments have found that sharing social information consumes a significant portion of daily conversation: ~65% (Dunbar, 2003). Furthermore, an important component of this social information exchange involves events that were not seen by others. When we describe events to others about an absent target, we call it *gossip* (see Foster, 2004).

Several studies on gossip suggest many reasons for it, such as social influence, information sharing, cultural learning, and social bonding (Baumeister, Zhang, & Vohs, 2004;

Beersma & Van Kleef, 2011; Dunbar, 2003; Ellwardt, Labianca, & Wittek, 2012; Fernandes, Kapoor, & Karandikar, 2017; Foster, 2004; McAndrew & Milenkovic, 2007; Wu, Balliet, & Van Lange, 2016). Recent evidence also shows that, at least under some conditions, gossip can be a more efficient and effective tool than direct punishment for promoting and maintaining cooperative behavior (Wu, Balliet, & Van Lange, 2016). However, without a comprehensive theoretical framework of social communication, it is difficult to predict *a priori* when and why social information will be disseminated, and whether gossip will occur.

There are general theoretical frameworks for human social interaction, most notably social exchange theory (see Foster, 2004). However, current theories do not sufficiently capture the sociocognitive components underlying human decision-making in social settings. Moreover, they are normally not specified well enough to make accurate *a priori* predictions about specific human social interactions, such as whether someone will actually gossip in a given situation. Therefore, here we present our model of social communication, developed to characterize the specific conditions under which different types of social communication will occur: for example, when an agent should directly communicate with the target of the information, gossip to others, or simply do nothing. The overall goal of our research program is to produce a comprehensive model of human social intelligence. To do so, we believe explicit modeling at multiple levels of analysis is necessary. In particular, a general, more qualitative theoretical framework should first be provided to layout the critical causal factors and their general interactions from a more top-down perspective. From this theoretical model, important predictions can be derived to generate hypotheses for further empirical research, which in turn test the model. After this, computational specification can occur in a more informed, meaningful, and understandable way, ultimately contacting neural circuitry.

In the current paper, we describe our top-level theoretical model and predictions derived from it. In the following sections we first describe the details of the model, then the methods used to generate model predictions; we then list nine predictions derived from it. In ongoing work, we have recently conducted a behavioral study with human participants to test the predictions, and are now specifying the model computationally.

The General Model of Social Communication

1. Agents

The central problem-solving agent (A_1) is the focal agent of the social communication process in our model: the one who takes the information input and decides what to do based on the set of possible actions and expected outcomes. To avoid the clunky “his/her” we will denote A_1 as female.

Receiver (A_2) is the agent (or agents) to whom A_1 may transmit information, i.e., communicate with. A_2 can be characterized by their relationship to A_1 , e.g., closeness (family, friends to strangers), relative status, and other group memberships (such as classmates, coworkers, colleagues), with such characteristics influencing A_1 ’s communication decisions. Once the information transmission from A_1 to A_2 is completed, A_2 then becomes the next central problem-solving agent in a subsequent state, and then must decide whether to communicate to receivers, and so on.

Target (A_3) is also an essential agent in the process. A_3 is the subject of the information that A_1 is contemplating. In other words, A_3 is the agent who took the action that caused the initial state change. What A_3 has done, i.e., the details of the event surrounding A_3 ’s action — which we call *event* or *scenario* — will influence A_1 ’s decision. Moreover, like A_2 , A_3 can be categorized based on their relationship to A_1 (such as closeness, status, and other group memberships), which should also influence A_1 ’s decisions.

Information source (A_0) is another agent in the system. The central agent, A_1 , receives the information about the target agent, A_3 , via either direct observation or via another source — A_0 is the other possible source. Thus, the role of A_0 is to pass the information about A_3 to A_1 . A_0 can be a person who has made a decision to spread the information directly to A_1 ; A_0 can also be a person who decided to spread the information to many unspecified individuals via various means such as social or mass media, books, or internet forums. A_1 will evaluate the information received from A_0 based on reliability and credibility of A_0 (i.e., trustworthiness); the outcome of this evaluation will determine whether A_1 continues to assess the information.

The role of each actual agent, then, is not fixed. An individual who is A_1 in one state, for example, can become A_0 , A_2 , or A_3 in subsequent states. All the agents in the system have their own goals to achieve, and to approach their goals, the agents need to have a good understanding of the others and self. That is, they have to consider what the other agents in the system are doing in the current state and are going to do in future states. Because each agent is attempting to move toward her own goal state, the system is highly complex and dynamic, which in turn makes it challenging for the agents to build an accurate mental representation of each other; therefore, constant updating of information regarding every agent in the system is necessary. We capture this in the current model by focusing on the central agent, A_1 , and her “mind”, which also includes her models of the other agents’ minds as well as her model of her own mind.

The mind of the agents is thus currently represented in the model by what we describe for the central agent A_1 , but indeed all agents would have the same model architecture for their mind. And for A_1 to take the proper action, she must consider the minds of the others, which we describe below. Our general mind architecture is described throughout sections 2-5, but in general it follows the basic ‘sensation/perception $\rightarrow$ cognition (including problem representation and decision-making) $\rightarrow$ action’ circuit (Gazzaniga, Ivry, & Mangun, 2013; Glimcher & Fehr, 2014; Kralik, 2017). Each component of this circuit is influenced by both (a) longer-term, more stable characteristics, such as personal background (e.g., family and educational), personal traits (like personality in general, such as extroversion vs. introversion), cultural (e.g., individualism vs. collectivism), political, and sexual identities, social traits (e.g., morality and prosociality orientation), and longer-term goals and interests; and (b) more short-term, ephemeral characteristics, such as current immediate goals and mood. Our focus here, however, is directed toward the basic sensation-to-action circuit, with particular emphasis on social cognition.

2. Goals of the Agents

From A_1 ’s perspective, the goals of all the agents in the system are valuable pieces of information to make the best decisions. However, the most important goal that directly drives A_1 to choose a specific action is that of A_1 herself.

In the model, all agents’ actions are presumed to be in pursuit of goals, with each agent’s action choice depending on what they think will provide the best outcome for themselves in pursuit of their goals. Goal pursuit in a multi-agent environment, however, often requires interaction with others, especially when help (i.e., cooperation) is needed or conflict arises; and the expected outcomes of these interactions must be factored in. In our model, we assume that benefits (or costs) to others ultimately translate to benefits (or costs) for the actor via what we call *social value* or *social equity*. For example, if agent A_1 helps A_2 achieve their goal, the benefit given to A_2 should translate into social equity to A_1 for future help when needed. In this way, reciprocal altruism is achieved (and fairness upheld). Thus, in general, there are two types of *value* — nonsocial and social — that must be taken into account when considering actions in pursuit of any given goal.

At the same time, the goal itself (being pursued by a given agent) can be either nonsocial or social. **Nonsocial goals** are perhaps more typically studied even in multi-agent systems, such as cooperating or competing in pursuit of positive rewards (e.g., money) in game-theoretic scenarios like Prisoner’s Dilemma. In contrast, **social goals** may also be pursued directly, such as attempting to build friendships (translating in the model to acquiring social equity). Both nonsocial and social goals have **primary** and **secondary** elements (Gazzaniga, Ivry, & Mangun, 2013; Glimcher & Fehr, 2014). Primary social goals are more innately specified, in which we are “wired” to desire them (both regarding wanting and liking). For social goals, in general, it is a natural

tendency to seek social interaction, with primary reward resulting from it (Barak & Feng, 2016). Secondary elements are based more on learning and strategy.

Focusing more specifically on social information communication, such instances are often instigated by actions taken by a target of interest, A_3 , that A_1 comes to know about and then must determine what to do: e.g., whether to tell A_2 , i.e., gossip. This information about A_3 normally precipitates a new problem for A_1 to solve. For instance, consider the scenario of A_3 cheating on a final exam. The new problem posed to A_1 may now be one of fairness: i.e., that the resulting state is now unfair, with the goal being to reestablish it. A_1 then must determine the proper course of action: i.e., whether to confront A_3 directly, tell others (i.e., A_2), or do nothing. In future development we elaborate these concepts of problem solving and goal pursuit to the solving and pursuit of multiple problems and goals simultaneously.

3. Input Information to A_1

There are two general types of information input to A_1 .

An event consists of three main components: target, content, and valence (see Figure 1). The *target* specifies A_3 , the agent who has become a subject of information by taking an action in a previous state. Again, A_3 can be characterized by various factors such as closeness to A_1 (e.g., a close friend, a casual acquaintance, a total stranger), and various identities, typically from A_1 's perspective (e.g., a classmate, a colleague, a family, a celebrity). *Content* is the actual body of information that describes the target's action, and often sets up a new problem to be solved that in turn requires its own problem representation (e.g., cheating on exam, with problem of fairness to rectify). The content also often divulges aspects of A_3 that enables A_1 to use in her model of A_3 's mind. If A_3 cheated, e.g., it likely evokes a concern of selfishness and lack of empathy. Finally, *valence* indicates the polarity (i.e., positive or negative) of the information.




Figure 1: Two types of information input to A_1 : Event & Context. The *event* includes information related to who did what, and why it occurred; The *context* includes additional details such as setting.

Once information is generated by A_3 taking an action, A_1 receives the information by one of two means: (1) direct observation or (2) via A_0 . The information source itself will influence A_1 's decision as well, discussed further below.

Context is comprised of two aspects. The first is the main multi-agent environment that A_1 finds herself in, most notably, that A_2 is present. The second includes other specific factors, such as location (e.g., school, workplace, restaurant, conference), or the type of information source (e.g., is it a newsmagazine or a tabloid magazine; is it from a reliable source or a random internet forum?).

4. Information Processing by A_1

Information about the event and context are thus received by A_1 as the input stimuli. A_1 then engages in a series of internal processes to determine what action to take based on the newly obtained information about A_3 , which we now describe and is depicted in Figure 2.




Figure 2: Internal information processing of A_1 from sensory input to action selection based on sociocognitive considerations.

4-1. Sensation and Perception: A combination of the two types of information (i.e., events and context) is detected as sensory input and then perceptually processed.

4-2. Evaluation of veracity: After sensory and perceptual processing, higher-level cognition commences (Gazzaniga, Ivry, & Mangun, 2013). First, A_1 must assess whether the information warrants further processing: most notably whether the information can be trusted. To accurately estimate credibility and reliability of the input information, A_1 must access long-term memory (i.e., prior knowledge) about the source, target, context, and any other details that may increase the accuracy of the evaluation. If the information seems reliable, processing continues.

4-3. Representation of current state: Reliable information that has passed the evaluation process is now ready to be used for updating A_1 's internal model of A_3 's mind; A_1 must also assess and update the representation of the current state consisting of all agents (including self). As described under "The mind of the agents" section, all agents have their own problem representations, with each including all agents and models of their minds; and although these are not depicted for A_2 and A_3 , they nonetheless are represented by A_1 in her problem representation, which also illustrates the type of recursion that takes place in human social cognition (see Figure 2) (Barak & Feng, 2016; Gazzaniga, Ivry, & Mangun, 2013; Glimcher & Fehr, 2014). And because of the uniqueness of each individual, and the number of events that take place (including many out-of-sight of subsets of agents), the effort to maintain accurate models of others is daunting (likely providing evolutionary selection pressures leading to human high-level social cognition) (Dunbar, 2003).

4-4. Decision-making: As discussed, A_1 , like all agents, must select actions to reach a given goal, with an action policy that attempts to maximize expected value. For any given action, then, A_1 needs to estimate the consequences of all action options and to choose the action expected to return the best outcome. More specifically, the decision-making process is divided into two key sequential steps: *valuation* and *action selection* (Figure 2) (Glimcher & Fehr, 2014).

Valuation. In goal-directed decision-making, each action option is valued based on its expected outcome, with A_1 attempting to maximize value by reaching the current goal state efficiently. However, to determine this, A_1 must also consider what the other agents are expected to do, prior to, simultaneously with, and after A_1 's action. Regarding possible prior or simultaneous actions, game-theoretic strategies can be undertaken to first determine the others' possible actions, and to then value (and select) one's own based on these expectations (Glimcher & Fehr, 2014). For the current paper, A_1 does not expect the others to take further action prior to A_1 , and so she must only consider what they are expected to do subsequent to her action.

Currently, there are three possible options A_1 can take involving the social information about A_3 (i.e., the target): (1) A_1 talks directly to A_3 about the information; (2) A_1 tells others (A_2) about A_3 ; or (3) A_1 does nothing with the information. We consider each action option in turn.

Direct communication. With no intermediary agents between A_1 and A_3 , direct communication with A_3 has both pros and cons. Potential advantages include (a) an opportunity for A_1 to confirm the veracity of the information with the target directly (assuming honesty); (b) the ability for A_1 to obtain more circumstantial event details, including A_3 's action intent or stance on the event; and (c) a higher likelihood of influencing A_3 (vs. relying on others via indirect communication), giving A_3 an opportunity to correct matters without the possible repercussions via spread (and potential amplifications and distortions) across the social network.

Disadvantages of direct communication include the risk of A_1 being wrong about the information and hurting A_3 's feelings (i.e., decreasing social equity with A_3 and perhaps others based on A_3 's future actions). Even if the information is true, and particularly with negative information, direct criticism is often considered harsh, rude, or offensive, potentially leading to repercussions for A_1 by both A_3 and others if the information spreads (via A_3 or other witnesses). To minimize such risks, humans have evolved the ability for indirect communication (Dunbar, 2003; Foster, 2004).

Indirect communication: Gossip. Alternatively, A_1 may choose to tell other agents (i.e., A_2) about the information in A_3 's absence. That is, A_1 can instigate gossip with others. We note three major advantages of gossiping. First, it can lower the risk of confrontation, threat, and retaliation from A_3 , while at the same time affecting A_3 's behavior via social influence, and in particular, via reputation (Wu, Balliet, & Van Lange, 2016). Second, it informs others (A_2) so that they can update their world models (and thus increase their

accuracy), which can enable A_2 to either (a) avoid A_3 , (b) help rectify the problem produced by A_3 's actions, such as righting injustices (e.g., if A_3 were found cheating on exam), or even (c) rectify A_2 's own behavior via social comparison and self-evaluation (e.g., to help see why certain actions are wrong or others to emulate). And again, these benefits to A_2 are expected to return value to A_1 via social equity (for the future). Third, informing A_2 may in fact enable A_1 to obtain more information about A_3 , to help decipher the information about A_3 (e.g., with respect to cultural norms, mores, rules, and laws), and determine whether something should be done about it — that is, A_1 can seek others for advice and consultation. Indeed, this interaction with A_2 can also help A_1 improve the accuracy of her own world model.

Despite the many advantages to gossiping, there are indeed risks. For example, if the information is negative in valence (like catching them cheating or shoplifting), there could be a threat of retaliation against A_1 for damaging A_3 's reputation. Moreover, gossip generally has a bad reputation in and of itself. Thus, A_1 could in fact become a notorious gossip, known to “talk behind peoples' backs”; as a result, social equity of A_1 could significantly drop. Therefore, sophisticated estimation and prediction of all possible outcomes must be attempted, but of course cannot be fully achieved. For example, how many steps into the future (and number of behaving agents) that can be simulated is necessarily limited. Such factors show how challenging it is to make good decisions in our multi-agent world (and how modeling is necessary as scientists to better understand the multiple factors and their interactions).

Do nothing. The last option is for A_1 to do nothing, and this choice can be strategic as well. Since the other options (i.e., direct and indirect communication) may both entail significantly high risks of yielding a net negative outcome, keeping the information private may be the safest option. It is also possible that the information is not significant enough in the current state to instigate communication, but yet may be so in the future, given that the system is constantly updating and modifying. At some point when the stored information becomes useful, A_1 can retrieve it (from long-term memory) and repeat the decision-making process again to decide whether or not to use it.

Action selection. After the valuation process, A_1 will have calculated the expected outcomes of all action options: to confront A_3 directly, do nothing, or tell A_2 — i.e., to gossip. At this stage, then, A_1 determines which action to select based on their relative valuations.

5. Take Action, Face Outcome, and Learn

Next, A_1 actually executes the action chosen. If A_1 chooses to tell others about the information, the receiver A_2 comes into play, adding more complexity to the system (see Figure 3). That is, there will be consequences after taking the action that now depend on A_2 . A_1 thus needs to monitor the actual outcome, assess it, and compare it to the expected outcome. The outcome may be close enough to expectation, moving A_1 closer to her goal; yet it is also possible that the outcome does




Figure 4: The complete model of social information communication. Brain region abbreviations: fusiform gyrus (FG), posterior superior temporal sulcus (pSTS), medial prefrontal cortex (MPFC), temporoparietal junction (TPJ), amygdala (AMG), insula, right temporal lobe (rTL), interparietal sulcus (IPS), anterior paracingulate cortex (aPCC), medial precuneus (med. precuneus), anterior insula (ant. insula), anterior cingulate cortex (ACC), orbitofrontal cortex (OFC), ventral premotor cortex (vPMC), dorsolateral prefrontal cortex (dlPFC), posterior parietal cortex (PPC), posterior superior temporal cortex (pSTC), caudate nucleus (CD), ventral striatum (VS), mesolimbic dopamine (DA) system, ventral tegmental area (VTA), premotor cortex (PMC), basal ganglia (BG).

standards to justify their position and/or attempting to raise one's relative position by lowering theirs (at least within the gossipers' community), and our model also predicts the same:

Prediction 5: There should be greater spreading of negativity about celebrities.

Although again limited, there is evidence consistent with Predictions 4 & 5 (Foster, 2004): e.g., one study found that people spread positive information about allies and negative information about potential enemies, including strangers and those with high status (McAndrew & Milenkovic, 2002).

At the same time, the model highlights the potential ramifications of spreading negative scenarios about those in position to retaliate, and thus the following is also predicted:

Prediction 6: Negativity should be reduced for ingroup targets.

Although studies have found evidence pertaining to Predictions 4-6, our model shows how specific factors will need to be isolated to clarify the true nature of the current findings: for example, whether higher positive spreading for ingroup is due to heightened positivity or lowered negativity (and *vice versa* for celebrities). Studies have yet to tease apart these factors sufficiently.

For scenario *content*, it is clear that more impactful events should be expected to generate more gossip, and thus:

Prediction 7: Dimensions related to morality are predicted to generate the most gossip, especially those involving more egregious threats, like harm and cheating.

Multiple studies have indeed found evidence for moral underpinnings of gossip (Fernandes, Kapoor, & Karandikar, 2017; Foster, 2004), although they have thus far focused narrowly on one or a few moral dimensions (e.g., fairness). Our predictions provide hypotheses for more comprehensive empirical studies, which we are currently undertaking. Indeed, although the moral dimensions are generally expected to generate more gossip than others, we also predict:

Prediction 8: Differences should be found among the morality domains themselves.

Finally, when assessing impact (for value maximizing), a more complete sociocognitive perspective shows that even more nondescript events could in principle provide high

value: e.g., if needed to maintain accurate models of those whose behavior requires high predictability. Such considerations lead to our final hypothesis:

Prediction 9: Basic social affairs should in some cases generate higher spreading rates for ingroup targets, in order to maintain accurate detailed knowledge about them.

In other words, we predict that *context matters*, also attesting to the importance of developing a comprehensive model of social interaction and communication to identify and specify potentially important context effects.

Along with our current empirical studies to test the model predictions, we are also developing the model computationally to specify the causal factors and interactions underlying the communication of social information at the next level of analysis, with the ultimate goal of providing a comprehensive understanding of human social intelligence.

References

- Barak, B. & Feng, G. (2016). Neurobiology of social behavior abnormalities in autism and Williams syndrome. *Nature Neuroscience*, 19, 647-655.
- Baumeister, R. F., Zhang, L., & Vohs, K. D. (2004). Gossip as cultural learning. *Review of General Psychology*, 8(2), 111-121.
- Beersma, B., & Van Kleef, G. A. (2012). Why People Gossip: An Empirical Analysis of Social Motives, Antecedents, and Consequences. *J. of Applied Social Psych.*, 42(11), 2640-70.
- Dunbar, R. I. M. (2003). The Social Brain: Mind, Language, and Society in Evolutionary Perspective. *Annual Review of Anthropology*, 32(1), 163-181.
- Ellwardt, L., Labianca, G., & Wittek, R. (2012). Who are the objects of positive and negative gossip at work? A social network perspective on workplace gossip. *Social networks*, 34, 193-205.
- Fernandes, S., Kapoor, H., & Karandikar, S. (2017). Do we gossip for moral reasons? The intersection of moral foundations and gossip. *Basic and applied social psychology*, 39(4), 218-230.
- Foster, E. K. (2004). Research on gossip: Taxonomy, methods, and future directions. *Review of General Psychology*, 8(2), 78-99.
- Gazzaniga, M. S., Ivry, R. B., & Mangun, G. R. (2013). *Cognitive neuroscience: the biology of the mind*. WW Norton & Company.
- Glimcher, P. W., & Fehr, E. (2014). *Neuroeconomics: Decision making and the brain*. Oxford: Academic Press.
- Haidt, J. (2007). The new synthesis in moral psychology. *Science*, 316, 998-1002.
- Kralik, J. D. (2017). Architectural design of mind & brain from an evolutionary perspective. *Proc. AAAI 2017 Fall Symposium: A Standard Model of the Mind*.
- Lee, D. (2008). Game theory and neural basis of social decision making. *Nature Neuroscience*, 11(4), 404-409.
- McAndrew, F. T. & Milenkovic, M. A. (2002). Of tabloids and family secrets: the evolutionary psychology of gossip. *J. of applied social psych.*, 32(5), 1064-1082.
- Wu, J., Balliet, D., & Van Lange, P. A. (2016). Gossip Versus Punishment: The Efficiency of Reputation to Promote and Maintain Cooperation. *Sci Rep*, 6, 23919.