

The Etiology of Social Change

Kathleen M. Carley, Michael K. Martin, Brian R. Hirshman

Institute for Software Research, Carnegie Mellon University

Received 12 September 2008; received in revised form 12 May 2009; accepted 13 May 2009

Abstract

A fundamental aspect of human beings is that they learn. The process of learning and what is learned are impacted by a number of factors, both cognitive and social; that is, humans are boundedly rational. Cognitive and social limitations interact, making it difficult to reason about how to provide information to impact what humans know, believe, and do. Herein, we use a multi-agent dynamic-network simulation system, Construct, to conduct such reasoning. In particular, we ask, What media should be used to provide information to most impact what people know, believe, and do, given diverse social structures? All simulated agents are boundedly rational both at the cognitive and social level, and so are subject to factors such as literacy, education, and the breadth of their social network. We find that there is no one most effective intervention; rather, to be effective, messages and the media used to spread the message need to be selected for the population being addressed. Typically, a multimedia campaign is critical.

Keywords: Dynamic network analysis; Multi-agent simulation; Information diffusion; Bounded rationality; Mass media

1. Primacy of agent and network

A fundamental aspect of human beings is that they learn. They do not always learn what is right, they do not always understand or believe what they learn, and they do not necessarily act on what they learn. Nevertheless, they do learn—and as such, what they know, believe, and do tends to change.

The process of learning and the content of what is learned are impacted by a number of factors. These factors are both social and physiological (Carley & Newell, 1994). This fundamental idea was described by Simon (1982) as bounded rationality. Bounded rationality, as is often forgotten, has both a cognitive and a social component. On the cognitive side,

Correspondence should be sent to Kathleen M. Carley, Professor, Institute for Software Research, Carnegie Mellon University, Pittsburgh, PA 15213. E-mail: kathleen.carley@cs.cmu.edu

humans cannot recall or process all information related to even moderately complex tasks. On the social side, humans because of their social position do not have access to all information related to even moderately complex tasks.

The implications of this simple idea are not obvious in most settings, particularly as the degree of variance in the cognitive and the social side is different and the two sides interact. Cognitive variations in recall, processing, and risk taking all exist and have been documented in laboratory settings (i.e., Anderson, 2005). However, variations in cognition often are not measured at the population level. Such variance, however, may play out indirectly by influencing such things as the level of literacy, which varies considerably across the population in the United States (Kirsch et al., 2002) or the highest level of education attained. Moreover, variations in attributes such as literacy, Internet access, and education are also a function of opportunities which are correlated with—if not caused by—position in the underlying social structure. For example, leaving aside human pathologies and cases of brain damage, there seems to be less variance in the mechanics of human cognition than in the way groups structure themselves and so the types and amount of access people have to information (i.e., Dunbar, 1993). Nevertheless, individuals who can process more information tend to become more highly educated and more highly educated, individuals tend to have a broader social network. In general, the friendship or networking process is a social process influenced by a person's sociodemographic position (Lazarsfeld & Merton, 1954), social comparison (Festinger, 1954), and dissonance (Festinger, 1957).

What humans do is not simply a function of what they know, but also what they believe. Humans communicate both ideas and beliefs when communicating. These ideas and beliefs can have enormous impacts on actions. As individuals learn information, as they learn what others believe, their beliefs alter through a process often referred to as social influence (Friedkin, 1998). Consequently, humans may still hold erroneous beliefs—even when they have only correct information at their disposal. However, even when individuals hold the same beliefs and have the same information, they may act differently. Actions may be influenced by considerations other than knowledge and beliefs, such as access to resources such as equipment or money. Two individuals might have the same knowledge and beliefs, but one may be unable to act due to a lack of resources. On the other hand, two individuals may take the same action but do so for entirely different reasons. Thus, although coupled, knowledge, beliefs, and actions are only loosely coupled (Moore, 1985; White, 1992; Wooldridge & Jennings, 1995); but changes at any level can lead to the cascading evolution of new organizational forms (White, 2008).

Another critical factor influencing human behavior is the way in which information is sought after and provided (Case, 2002; Kuhlthau, 1999; Moorthy, Ratchford, & Talukdar, 1997). For example, humans gain access to information from other humans, from newspapers, the Internet, and a wide variety of other sources (Case, 2002). The source of the message can be critical (Marchionni, 1995). Individuals may trust different sources to different extents and for different reasons (Kramer, 1999; McAllister, 1995). Further, even for the same type of source, such as a fellow human, some are trusted more than others; for instance, one might trust one's best friend more than a stranger. Furthermore, some sources have greater reach and so interact with and send information and beliefs to more others

(Kaufer & Carley, 1993, 1994). For example, an opinion leader who communicates via television will reach more of the population than one who simply goes around and talks to people. Thus, when trying to influence others, a variety of information scenarios can occur; positive messages sent by an opinion leader through the Web may have different effects than negative messages sent by a different opinion leader in a newspaper due both to the message content, the expressiveness of the medium, and the end user's trust in the medium. Understanding and assessing the impact of such media can be complicated because the relative impact of such interventions is affected by the underlying sociocognitive milieu.

Herein we ask: If we both take the social and cognitive constraints seriously, how can information be provided to groups of humans to have the most impact on what they know, believe, and do? What kinds of media, or scenarios for providing information and belief, are most effective? A secondary question we address is to what extent is the choice of media dependent on the community structure as evinced by the social network topology. What kinds of networks are more or less receptive to information provision via particular kinds of media, and what can that tell us about media selection and information diffusion more generally?

We address these issues using the Construct simulation model (Carley, 1990, 1991; Hirshman, Martin, & Carley, 2008a; Schreiber & Carley, 2004a, 2004b). Construct is a powerful and highly veridical computer model that captures both cognitive and social constraints on human behavior (see Section 2). In a series of virtual experiments, we use Construct to examine the effect of various information-providing media on populations with different social structures defined in terms different network topologies (see Section 3). To overcome the standard problem in response surface analysis of simulation results, the TeraGrid (Catlett et al., 2007) was used so that a sufficiently large population of data was generated to ensure robustness of results and to enable sufficient variation in underlying cognitive differences across agents.

We focus our discussion of the results around three key questions (Section 4): Which type of media is most effective in general? Which type of media is most effective by education level? Which type of media is most effective for different social network topologies? Other factors, such as age, race, gender, risk taking, forgetting, and so on are controlled. The virtual experiments generate outcomes on three key features: Who knows the information being communicated? Whose beliefs change as a function of this communication? Whose activity changes as a function of this knowledge and beliefs? The results (Section 5) indicate fundamental effects due to the manipulated factors, and nonadditive impacts of the media selected to provide information.

Note: We focus on societies that differ in network topology rather than societies that differ in sociodemographic makeup for several reasons. First, most simulation work on diffusion coming out of the social sciences assumes a random network, yet empirical work on topology shows that the impact of topology can be profound (Carley, Lee, & Krackhardt, 2001). Our analysis sought to determine whether these effects influenced the effectiveness of the media used. Second, in real-world applications, the sociodemographic makeup of a population is generally easier to determine than the topology of the network: The first is gathered by the census, while the second requires a questionnaire about relations (N vs. N^2

bits of information). For large communities, gathering accurate network data is particularly difficult. Hence, by varying the network structures, we provide a means for estimating the range of effectiveness that a media would be likely to have if the underlying topology is unknown, and we provide a way of identifying which media to use if one could determine the appropriate type. Finally, since all populations have some fraction of the population in each sociodemographic class, the impact of higher or lower proportions per class can be inferred by assessing each class separately. In fact we do this in the Results, where we explore variations in the impact of the media by educational level. The same cannot be done for network topology. In other words, in a community all sociodemographic classes are present but only one network. This allows us to explore the impact of both network topology and sociodemographic differences.

2. Construct model

Construct¹ 3.5—hereafter referred to as simply Construct—is a multi-agent dynamic-network simulation model for examining the co-evolution of agents and the sociocultural environment (Carley, 1990, 1991, 1999). Using Construct, we can examine the evolution of networks and the processes by which information moves around a social network (Carley, 1995; Hirshman, Carley, & Kowalchuck, 2007a). Construct captures dynamic behaviors in groups, organizations, and populations with different cultural and technological configurations (Schreiber & Carley, 2004a, 2004b). In Construct, groups and organizations are complex systems. The variability of human, technological, and organizational factors in such systems is captured through heterogeneity in information processing capabilities, knowledge, and resources. The user can characterize this heterogeneity by specifying a set of agents (Hirshman, Carley, & Kowalchuck, 2007a, 2007b) and the social and knowledge networks in which they are embedded (Hirshman et al., 2007b). Multiple nonlinearities in the system generate complex temporal behavior on the part of the agents.

Construct is the embodiment of constructivism, a mega-theory which states that the sociocultural environment is continually being constructed and reconstructed through individual cycles of action, adaptation, and motivation. Many social science theories and findings are part of the constructivist theoretical approach, including structuration theory (Giddens, 1984), social information processing theory (Salancik & Pfeffer, 1978), symbolic interactionism (Manis & Meltzer, 1978; Stryker, 1980), social influence theory (Friedkin, 1998), cognitive dissonance (Festinger, 1957), and social comparison (Festinger, 1954). In addition, a number of cognitive processes are embedded in the simulation, such as transactive memory (Wegner, 1986).

There are three key features of Construct 3.5 that make it ideal for our purposes. First, the experiment designer has complete control over which subagent models are used for interaction over the course of a run. Second, Construct contains a suite of agent models that enable diverse sociotechnical conditions to be modeled. Third, general agent characteristics can be easily configured a priori using empirical data or they can be based on hypothetical data. While additional information about the Construct interaction model can be found elsewhere

(e.g., Carley, 1991, Hirshman et al. 2007a), the core Construct agent dynamics are as follows.

Each time period of the simulation, agents take a variety of actions, including initiating an interaction, responding to requests for interactions, sending messages, engaging in tasks, and updating beliefs. For each agent, these actions tend to occur cyclically except for responding to media and making decisions, which may occur off cycle—see Fig. 1. Exactly which actions an agent can take, and how many can be done simultaneously, depends on the agent's sociocognitive nature. Each action takes a certain amount of time, typically a time period. Human agents use their preference for homophily or expertise (described in Section 2.1), their transactive memory of other agents' knowledge, their beliefs, their sociodemographic characteristics, their availability, and their recommendations from others in order to rank the importance of interacting other agents in their social network. Based on this ranking, the human agent may choose to initiate communication with one or more other agents. The type of agent—human, Web-page, etc—will determine whether the agent can initiate interaction and what the agent's information processing characteristics are.

If two agents select each other, then both agents will prepare a message and send it to the other. A message is a set of memes (Wilkins, 1998), and so consists of one or more instances of any or all of the following: knowledge, beliefs, and transactive memory about the knowledge or beliefs of third parties. Once prepared, the message is communicated to the interaction partner, where it may be modified, misinterpreted, or ignored based on the sociocognitive properties of the receiver. After receiving a message, processing it, and possibly learning from it, both parties may modify their beliefs or make any relevant decisions. This process then repeats for each agent during each time period.

All agents operate in the same "time frame" meaning that interventions and/or interruptions can occur at a particular time and all agents can respond to them. For example, a news advertisement can come out at time period three. Behavior in the model is not lock

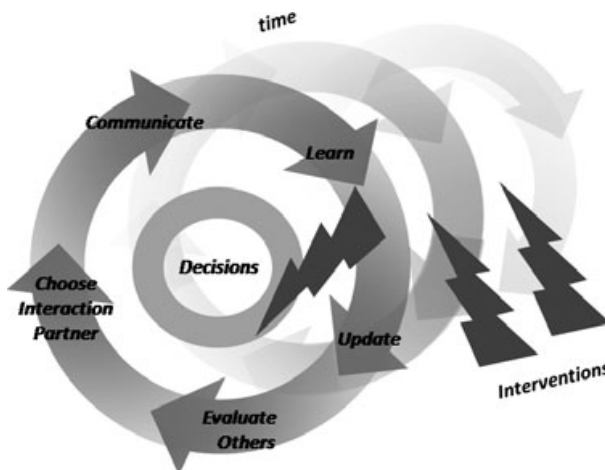


Fig. 1. Cycle of agent activity.

stepped, but day one is day one for all agents. Statistics, outputs, and decision information are gathered relative to these the time periods, as well as at the end of the simulation.

Although Construct was originally developed as a pure lock-stepped model with each agent interacting with each time period and then updating their memory, such behavior is no longer the case. As of version 2.5, Construct includes event-driven mechanisms, variable duration interaction, and fixed as well as mutable agent characteristics. In addition to these changes to update Construct with newer simulation technology, additional work has been performed on validating the core mechanisms independent of the exact technological mechanism in a variety of settings. Finally, the current version (3.5) is multithreaded.

The fundamental mechanisms in Construct have been scientifically validated (Carley, 1990; Carley & Hill, 2001; Carley & Krackhardt, 1996; Schreiber & Carley, 2004a, 2004b). It has been used to explain group mobilization (Carley, 1990), the impact of leadership (Schreiber & Carley, 2008), and the impact of the printing press (Kaufer & Carley, 1993). Directly germane to the current study, Construct has been used to compare and contrast different educational media by sociodemographic feature (Hirshman et al., 2008a) and the impact of media and opinion leaders in real cities (Hirshman, Martin, Bigrigg, & Carley, 2008b).

2.1. Agents

Agents are decision makers with varying information processing, sociodemographic, and access constraints. As such they may or may not be human (Carley, 2002). Within Construct, agents go about their business interacting, communicating, and learning each time period, as described in Fig. 1. As agents learn or acquire information, they may change their preferred interaction partners and modify what they are likely to communicate. These factors, in turn, influence what types of decisions are made by each agent. A variety of influences affect whom agents select as interaction partners, what they communicate with that partner, how much and how they communicate, whether they learn anything from that partner, and whether the message is learned accurately. Such influences include the agent's socio-demographic characteristics, information-processing characteristics, proximity, and current position in the social and knowledge networks. The agent model has been described in depth in other venues (e.g., Carley, 1990, 1991; Hirshman et al., 2007a, 2007b); thus, we concentrate here on both a high-level description and details of those components used for this simulation work.

Within Construct, agents both influence and are influenced by others. Agents who have influence over others can use that influence to escalate or de-escalate activity at a societal level by communicating information and/or beliefs. Social influence derives from shared attributes such as sociodemographic factors, shared knowledge, beliefs, and proximity (Carley, 1991). Social influence co-evolves with the spread of knowledge and beliefs. Consequently, in more heterogeneous populations where the lines of differentiation line up the chance of self-reinforcing beliefs at the group level is greater (Blau, 1977; Lau & Murnighan, 1998). Factors that are not influenced by the diffusion of information and beliefs include the agent's sociodemographic role (e.g., age, race, gender, level of education),

its basic cognitive limitations and information-processing capabilities (e.g., likelihood of forgetting, risk taking, amount of information and beliefs that can be communicated or processed, and whether the agent has transactive memory), the size of its sphere of influence (at least in the short term), and factors that have resulted from its sociocognitive interactions (e.g., literacy, access to newspapers, radio, and the Internet).

Within Construct, agents develop likelihoods of interacting with others based on relative similarity (RS, Equation 1) and relative expertise (RE, Equation 2) (Carley et al., 2001; Hirshman et al., 2007a). Relative similarity is a homophily-based mechanism (Carley, 1991; McPherson & Smith-Lovin, 1987) and derives from the idea that individuals are more likely to interact if they have more in common. Homophily-based interaction is a multicausal phenomenon due to ease of communication, shared understandings, and comfort. The RS of agent i and agent j , from i 's perspective, is characterized as

$$RS_{ij} = \frac{\sum_{k < K} (AK_{ik} \times AK_{jk})}{\sum_{j < I} \sum_{k < K} (AK_{ik} \times AK_{jk})} \quad (1)$$

where individual i 's relative similarity to j is determined in terms of sociodemographics, knowledge, and belief items K in the agent-to-knowledge matrix AK .

Of important note: An individual is most relatively similar to itself, and each period will have a reasonably high probability of choosing to "interact with itself" and to avoid communicating with others. Just because an agent has the highest RS with itself, however, does not mean that an agent will always interact with itself; indeed, due to the large number of other agents in the simulation, such avoidance of communication is relatively rare.

Relative expertise is a search-based mechanism and derives from the idea that individuals are more likely to interact if one has information that the other wants. The RE of agent j as judged by agent i is characterized as

$$\text{if } AK_{ik} = 0, \text{ then } X_{jk} = AK_{jk} \text{ else } X_{jk} = 0$$

$$RE_{ij} = \frac{\sum_{k < K} X_{jk}}{\sum_{j < I} \sum_{k < K} X_{jk}} \quad (2)$$

where individual i 's relative expertise to j is determined in terms of sociodemographics, knowledge, and beliefs (the set of K) items (Schreiber, 2006).

An agent is more likely to initiate interaction with another if the initiator thinks that the other has information it needs and/or it is similar to them. However, there is a curvilinear relation between this familiarity and expertise; to wit, as agents initially increase in similarity (homophily) they are more likely to realize the other has expertise they need but as they increase still further in similarity they realize that the other is so similar there is no specialized expertise.

When using construct, the researcher needs to specify the strength of each of these factors for agent-agent interaction. Herein, we set all human agents to use both logics and to at any time create a combined probability of interaction that is based on 60% similarity and 40%

expertise. In both cases, individuals are giving and receiving information and the overall tendency to give versus receive is about 60/40 as identified by Valente, Poppe, and Merritt (1996).

When setting up a virtual experiment in Construct the researcher needs to specify multiple parameters for each agent. This is often facilitated by the use of agent classes to parameterize multiple agents simultaneously. We next discuss the number of agents and the agent classes of this experiment (Section 2.2), the distribution of sociodemographic parameters for the agents of that class (2.3), the distribution of cognitive factors for each class (2.4), the sphere of influence for that class (2.5), and the access constraints for that class (2.6). While the full Construct model has a number of additional features that can be varied, such aspects were held constant for this simulation (Hirshman et al., 2007a).

2.2. Agent classes

In this study, we find it helpful to think in terms of two meta-classes of agents—human agents and media agents (which may or may not be human). Each time period, human agents may interact with other members of the general human population or with a media agent.

There are two classes of human agents: the general public, who will make decisions, and the opinion leader, who can help sway decisions. In the experiments performed, the opinion leader attempts to get the general public to act in one way while the media are designed to thwart it.

In this experiment, we consider five classes of media agents: newspaper advertisements (“ad”), publically accessible websites (“Web”), centers that have people in them that provide assistance when someone comes in physically or calls in via phone (“call”), radio advertisements (“radio”), and letters sent via postal mail (“mail”). Media agents differ from each other in terms of the time periods they are active and the length of the messages they send. All agents can communicate facts or beliefs, but the particular set transmitted depends on the agent’s knowledge or belief at the time. All media agents are passive—they cannot initiate communication with a human agent. Instead, they provide information only when the human agent selects to go to, listen to, or reads the information available through the media.

These particular media agents were chosen because they represent distinct forms of access to information. You might ask why we did not use television when it is so prevalent. The reason is that, within the characteristics we were varying, television and radio ads are identical. Thus, the reader can think of radio as radio/television ads.

The number of each type of agent, their activity level, and length of messages that they send in the simulation summarized in Table 1.

The initial knowledge and beliefs held by each of the media agents and the general public are described in Table 2. Over the course of the simulation, the human agents learn, however, the knowledge of the opinion leader and the media remain constant. The number of facts in each category, specified in the left-hand column of Table 2 is proportioned based

Table 1
Number of agents in agent class

Class	Number	Periodicity	Active Time Periods	Message Length
Humans	3,000	Continuous	104	1 fact, belief, transactive fact or belief, or social information
Opinion leader	1	Periodic every other time period	52	1 fact, belief, transactive fact or belief, or social information
Ad	1	New news ads are periodic	26	1–2 facts or beliefs
Web	1	Periodic access every other time	52	4 facts or beliefs
Call	1	Periodic access every fourth time	26	3 facts or beliefs
Radio	1	New ads are periodic	1 (per ad)	1–2 facts or beliefs
Mail	26	Periodic new mail	6	3 facts or beliefs

Table 2
Distribution of information and beliefs across media agents

Information and Beliefs	General Human		Media Agents (%)				
	Population (%)	Opinion Leader (%)	Ad	Web	Call	Radio	Mail
Activity exists (1 fact)	0	100	100	100	100	100	100
Activity know-how (6 facts)	0	100	10	33	10	10	10
Believe right (3 facts)	1	0	33	100	100	33	33
Believe not right (4 facts)	5	0	33	100	100	33	33
Believe worth doing (3 facts)	1	100	0	0	0	0	0
Believe not worth doing (3 facts)	5	0	33	100	100	33	33
General knowledge (500 facts)	20	20	10	2	5	10	10

on subject-matter expert's views of the relative amount of time it takes for the overall meta-concept—such as know-how knowledge for a task—to diffuse.

In order for a human agent to make a decision, an agent must recognize that the activity exists, must have sufficient know-how knowledge, and hold a positive view of at least one of the two beliefs. In order to have sufficient know-how information, an agent must learn at least three of the six know-how facts; considering that agents do not start with any of this information, they must learn it from, ultimately, the opinion leader or media. Additionally, we have modeled two beliefs here—one where the true belief is that one should not engage in the activity (believe not right), and one that is neutral as to whether there is some benefit to engaging in the activity (believe worth doing). In order to make the decision, agents must hold at least as many positive belief facts or negative belief facts, or they must be subject to social influence from their peers which convinces them that the decision is a good one.

The more facts per category—and hence the more complex the message—the longer it takes that category as a meta-concept to diffuse. However, it is important to note that all information related to the activity (20 total facts) is small relative to the amount of

simulated general knowledge (500 facts) so that most of the time the general population will not be communicating facts about the activity. Furthermore, the ratio of positive and negative facts associated with the belief influences whether the ultimate “correct” belief is positive, negative, or neutral. We have two beliefs here—one where the true belief is that one should not engage in the activity (the “believe right & believe not right” belief), and one that is neutral as to whether there is some benefit to engaging in the activity (the “believe worth doing & believe not worth doing” belief). Lastly, the amount of information associated with activity know-how and with any one belief is comparable so that both spread in a comparable amount of time. The more complex the know-how and the more complex the belief, the longer that information will take to spread and the lower the fraction of the population that will have the expertise or belief at any one time.

The advertisement intervention is meant to provide a small amount of knowledge and belief while also containing a large amount of general-knowledge information to encourage agents to examine it. Such behavior is typical of articles or advertisements in newspapers. Advertisements only exist for a few time periods during the simulation, reflecting relative infrequent publication. Advertisements can be expected to have a small impact on a variety of agents due to the infrequent interaction and small message conveyed; however, advertisements will be among the most common media that human agents access. Since the advertisement is in printed media, it can be subject to two different constraints: a cognitive constraint, literacy, and an access constraint, subscription access. Both of these constraints are described in Section 2.6.

In contrast to the advertisement, the website is designed to provide a large amount of belief information by proving a large number of reasons why the activity is inappropriate. In doing so, however, the website could potentially be scraped for knowledge information and thus serve a purpose that is contrary to what the designers intended. For this reason, resources such as the website can be two-edged swords, mostly decreasing but potentially increasing the behavior. Because the website is frequently available, it will be easily accessed; however, users accessing it may have literacy or Internet access issues.

The information call center is designed to answer questions, based on requests for information from those members of the general population who contact the center. Because the information center represents the actions of humans who work at the center it has associated with it more social knowledge than the website. Unlike the website, though it may be difficult to get to the center as it requires physical movement and thus may not be as favorable an interaction partner to some agents.

The radio advertisement is very similar to the print advertisement. It is designed to provide a small amount of information or beliefs but can reach a large number of agents in the general population. Unlike the advertisement, however, the radio advertisement is not affected by the literacy or access constraints as modeled in this experiment.

The postal mailing is designed to represent a piece of mail containing information meant to deter at-risk agents from engaging in the activity in question. It, too, has the same information content as the advertisement, but the way the general population interacts with it is unique. Only some human agents receive mail. However, whether the human agent reads the mail is up to the individual agent. For the next six time periods, the mail message resides

in the agent's mailbox. The general population agent then has a certain probability of checking its mail and learning the information in the letter. Agents who read the letter absorb some of the information contained in the letter.

The media of Tables 1 and 2 were chosen because they represent important forms of access to information. The list of media examined in this experiment is by no means exhaustive: for instance, media such as television advertisements or informational pamphlets are common in today's society. However, the media omitted have much in common with those that were modeled. For instance, radio advertisements share much in common with television advertisements: Both provide a small amount of information to a mass audience and are not impacted by the cognitive and access constraints of Section 2.6. Pamphlets are similar to the mail messages used in this experiment, since both convey a substantial amount of information but also require literacy. At the level of modeling performed in this experiment, it is not possible to distinguish how information is provided by these different types of media; future work should consider other sociocognitive lines along which these media differ.

2.3. Agent sociodemographics

In Construct, agents can have a set of nonevolving attributes that influence behavior. Herein, we consider those attributes to be sociodemographic characteristics. These attributes can be set based on census data, or based on other considerations. The researcher can in fact define any characteristics as agent attributes and then use these to effect interaction. The critical difference between attributes and knowledge/beliefs is that for an agent the attributes are fixed for the duration of the simulation; in contrast, the agent's knowledge and beliefs may change. Consequently, attribute-based effects on interaction tend to be stable, while variations in interaction are due to changes in knowledge and beliefs.

In this paper, we find it useful to consider six different sociodemographic attributes: age, gender, race, income, parent, and education. These attributes and their distributions are described in Table 3. The sociodemographic attributes are used to set the baseline interaction ties that exists independently of agent knowledge. The greater the overlap in agent sociodemographic attributes, the more likely the agents will interact, as part of the homophily effect (Smith-Lovin, McPherson, & Cook, 2001). These attributes were selected as they are ubiquitous general aspects of human behavior. Moreover, most empirical social science work looking at individual differences, tends to control for three or more of these attributes. Using attributes is one of the ways in which cognitive and external social factors can be accounted for in the simulation. In this case, education is the only "cognitive" related sociodemographic attribute. That is, for human beings, their cognitive capabilities may influence the extent to which they finish various levels of education; however, it has a subtle and indirect effect.

Two classes of agents, general human agents and opinion leaders, have these sociodemographic attributes. Media agents could be targeted so that they were aimed to match and so interact with humans with different attributes. Specifically, media were designed to target the agents who had either the lowest or second-lowest level of income and education. Thus,

Table 3
Sociodemographic attributes, values, and distribution

Attribute	Number of Values	Values (% of Human Agent Population)
Age	5	0–29 (20%); 30–39 (20%); 40–49 (20%), 50–64 (20%), 65+ (20%)
Gender	2	Male (50%); Female (50%)
Race	5	White (60%); African American (15%); Hispanic (10%); Asian (10%); Other (5%)
Income	6	0–15 k (40%); 15 k–30 k (35%); 30 k–50 k (15%); 50 k–80 k (6%); 80 k–120 k (3%); 120 k+ (1%)
Parent	2	Yes (50%); No (50%)
Education	4	Less than high school (40%); High school diploma (35%); College degree (30%); Graduate schooling (1%)

the opinion leaders and media would match any agent who had one of those two attribute values, but they would not match any other agent. These attributes were oversampled in the human agent population relative to the general population of the United States in order to better understand the effects of the cognitive limitations and information-processing capabilities described in Section 2.6.

The correlation between these attributes is also an important consideration. For a city's population these correlations could have been generated in one of three ways: (a) proportional to census data for a canonical U.S. city, (b) randomly, and (c) evenly. For the results we present here, we use proportions relevant to a canonical city with a high fraction of low-income individuals. Results can vary dramatically with the sociodemographic distribution as well as the correlation among sociodemographic variables; thus, the results we show should be taken as indicative of the inter-relationship, not as evidence that a particular media will always have a particular result (Hirshman, Martin, Birukou, Bigrigg, & Carley, 2008c).

2.4. Agent cognitive limitations and information-processing capabilities

Construct agents are complex. Two core features of Construct agents are information-processing capabilities and transactive memory. Agents are information-processing decision makers and so have one or more of these capabilities: initiate interaction, send messages, receive messages, and learn from messages. Agents can have both general and transactive memory (Wegner, 1986). An agent's general memory can contain both what information the agent knows—its facts—and what beliefs it holds. The transactive memory, on the other hand, contains the agent's representation of potential interaction partners—who knows what and who believes what. This transactive memory can be incorrect: the third party might not know the knowledge or hold the belief. Who knows what, as well as who knows who knows what, can be tracked by time period.

Table 5
Information-processing capabilities by agent class

Factors	General		Media Agents				
	Human Population	Opinion Leader	Ad	Web	Call	Radio	Mail
Initiate	Yes	Yes	No	No	No	No	No
Send	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Receive	Yes	Yes	No	No	No	No	No
Decide to take action	Yes	No	No	No	No	No	No
Learn	Yes	No	No	No	No	No	No
Change beliefs	Yes	No	No	No	No	No	No
Information atrophy	No	No	No	No	No	No	Yes
Message complexity	Very low	Very low	Low	High	Med	Low	Med
Supports multiple searches	Yes	Yes	Yes	Yes	No	No	Few

this the social network can be further broken down into a proximity-based network, a socio-demographic based network, and a knowledge/belief based network.

Construct allows the experiment designer to set networks as fixed or dynamic during the simulation. Moreover, the initial topology of such networks can be specified. This enables the impact of topology to be studied at the same time as the impact of information-based media. In this study, parts of the social network are fixed based on demographics and parts are dynamic based on changing knowledge and belief. The result is that the overall probability of interaction between each dyad is dynamic. In this paper, both knowledge networks and belief networks were dynamic and would change over the course of the simulation, whereas, the task assignment network was statically specified. However, we do not analyze this information in this paper.

The social network is of particular interest to this study. Empirical studies of social networks often form a network by asking individuals for the names of their interaction partners. The result is a snapshot of a network at a point in time as it is perceived. Based on this perspective it is tempting to think of networks as simple binary relations, two individuals either are or are not connected. Simulation, however, makes it obvious that the idea of a network is more amorphous. In Construct there are a number of ways to characterize the network of possible agent-agent interactions. All agents exist in a social network, and in this network the links among agents are probabilistic. These probabilities evolve over time, changing as agents increase in similarity and expertise. At any point in time, who is interacting with whom can be extracted in multiple ways: as a moving average, as probabilities, as a number of interactions in one particular time period, and as an interaction during the time periods of interest. This extraction represents the agent's sphere of influence that is, for whom the probability of interaction is likely to be nonzero.

In designing the simulation, the sphere of influence—the alters with which an ego's probability of interaction is nonzero—is the set of others who, in a short simulation, the agent is likely to interact with. In the full Construct model this sphere can grow and shrink; however, in this study we leave it fixed. Agents with greater reach—such as the opinion leader—have a larger sphere of influence, while most human agents have a relatively small sphere of

influence. Constraints on information access, as will be described in Section 2.6, can impact the effective size of an agent’s sphere of influence by making certain types of agents inaccessible. The size of the sphere of influence per agent class is described in Table 6; both the theoretical maximum is provided and the “in-practice” value determined from the experiments run. Note that the opinion leader is in every general human agent’s sphere of influence. Additionally, for each human agent whether a media agent is in the human agent’s sphere of influence depends on whether the human agent has an access constraint that prevents interaction.

In Table 6, it is important to note a distinction between the theoretical maximum size of the sphere as well as the in practice value. The network underlying the sphere of influence is designed exogenously by the experimenter prior to the start of the simulation. However, the actual sphere of influence in practice is the set of partners with whom the individual agent interacts due to homophily, expertise, or sociodemographic similarity. Since agents often do not interact with all of their potential partners, the effective size of the interaction sphere in practice is often much smaller than the theoretical maximum.

For each pair of agents, the probability that they interact is a function of proximity, sociodemographics, knowledge, and beliefs, as described previously. Since the sociodemographics remained constant and beliefs and knowledge changed over time in this study, the overall probability of interaction contains both a fixed and a nonfixed component. Since these overall probabilities can change, we say that the social network is evolving as who actually interacts with whom will vary over the simulation run: The interaction partners of the early simulation periods could differ substantially from those of the later periods. This evolution can be observed in the changing likelihoods that the agent has for interacting with those in its sphere of influence; however, the size of the sphere of influence and the topology of the fixed portion of these probabilities do not change within a single simulation. Thus, as the probability of interaction increased for any pair of agents, that increase must have come relative to that of other agents in the interaction sphere and must mean that both agents are evolving to become relatively less similar to all other possible interaction partners.

In this work, we examine three different topologies for the fixed portions of these networks. The network can be random (Erdos & Renyi, 1959), cellular (Frantz & Carley, 2005), or small world (Newman, 2000; Watts & Strogatz, 1998). The accuracy of the simulated topology is extremely dependent on number of agents and the overall density.

Table 6
Sphere of influence size by agent class (maximum size, 3,000 agents)

Factors	General Human Population	Opinion Leader	Media Agents				
			Ad	Web	Call	Radio	Mail
Sphere of influence theoretical maximum	40 ± 10	3,000	3,000	3,000	3,000	3,000	3,000
Sphere of influence in practice	25 ± 10	250 ± 50	300 ± 75	100 ± 50	66 ± 20	250 ± 100	150 ± 50

For example, with 10 agents and a density of 0.5 it is not possible to get a cellular network—agents are too interconnected to exhibit cellular structures; similarly, a network with 100 agents and a density of 1.0 is not random, cellular, or small world because everyone is uniformly connected to everyone. Since our populations have more than 3,000 agents including intervention agents and opinion leaders, we selected densities and sphere of interaction sizes to ensure that the topologies examined were good representations of that topology, that is, networks that were truly random, cellular, or small world.

The key differences in the random, small world, and cellular network topologies are differences clustering and the ratio of internal to external ties. In a random network the links are distributed independently and identically. In the small world network, each agent has a few links and a few agents have many links. In contrast, in a cellular network the agents are clustered into a few cells and mostly communicate with other cell members while only one or two members per cell interact with anyone in other cells. The placement of these ties affects the diffusion of information throughout the society and has the potential to lead to different rates of diffusion among and between different agents, as we explore in the work.

2.6. Constraints on information access

Cognitive and social factors combine to determine the level of information access that individuals may have. We examine three different information access mechanisms: literacy, Internet access, and newspaper readership (Hirshman et al., 2008a). Within Construct, these access mechanisms affect whether agents can interact with a specific media agent and get information from a specific forum. These mechanisms have been implemented as switches that the researcher can enable or not, depending on the research question.

In Construct, agents can be literate or not, as set by an experimenter-controlled switch. The literacy mechanism affects all media that require reading printed material. This means that printed advertisement, websites, and information sent in letters via the postal system are information on potentially affected. When literacy as an information access parameter is enabled, illiterate agents can still access these media; however, they do not learn all the information and beliefs conveyed in the message and they may even mislearn information. A small level of mislearning is implemented as the literature on literacy shows that literacy is in part a matter of degree, which often leads the illiterate individual to misinterpret what is being read. Literate agents are unaffected by enabling the literacy mechanism and receive the full information from these media. When the mechanism is disabled, all agents receive the full information.

In Construct, agents can surf the Web or not—and those that do have access to Internet-based media. When the Internet access constraint is enabled, agents lacking Web access cannot read information posted on websites and therefore change their knowledge or beliefs? Agents with Internet access can read such information, and use this information to affect subsequent interactions with other non-Web agents. When the mechanism is disabled, all agents can read information from websites.

In Construct, agents also have the ability to read newspapers and access the information contained in them. The newspaper access mechanism affects all media that require physical newsprint such as advertisements in newspapers and specialized articles by opinion leaders. When newspaper access is enabled, agents lacking newspaper subscriptions cannot read articles published in the paper. Agents who are newspaper readers, though, can still read such information. When the mechanism is disabled, all agents can read information printed in newspapers.

It is important to note that these mechanisms interact. For example, if an agent is illiterate but has a newspaper subscription, that agent may read news articles but do so with error. On the other hand, if an agent is literate but does not have access to the Internet, it still cannot read Web pages so the literacy parameter has no effect. These interactions are critical to effecting realistic sociocognitive behavior.

For each agent class, the researcher must exogenously specify whether an access constraint applies and the probability that an agent in that class is constrained. In this study access constraints only apply to general public human agents. That is, neither the opinion leader nor the media agents are constrained. In this study, the probability that an agent is illiterate, cannot access the Web, or does not read a newspaper was derived from sociodemographic attributes and national averages. A series of formulas, one for each constraint, were derived from national data in order to determine the probability that an agent is constrained based on its modeled sociodemographic characteristics (see Hirshman et al., 2008a for details).

3. Virtual experiment: Network structure and media

We conducted a virtual experiment to examine the impact of the way information was communicated on the effectiveness in reaching populations. Our underlying scenario is that there is some behavior that we want to discourage, for example, drinking under the influence, sharing needles, or engaging in tax fraud. In order to engage in this activity individuals must know how to do it and they must believe they should, in the manner described in Section 2.2. We assume that there is some local opinion leader trying to convince members of the population to engage in the counter activity, for example, that drinking and driving is acceptable, that the disease tests are dangerous, or that the fraudulent operation is legal. We, as the experiment designers, want to see how information can be provided to counter this opinion leader's message with a well-designed and effective contrary message. We also assume that the opinion leader is basically working through a grassroots approach—that is, dialogs with individuals one-on-one or conducting small group meetings. We also assume, consistently with the literature on beliefs, that once individuals decide to engage or not engage in an activity they quit actively seeking expertise on how to engage and their belief stabilizes.

The virtual experiment we conducted is described in Table 7. We focused on two factors that impact the way information is communicated: the media used for the media, and the topology of the network. There are assuredly many other factors, and indeed we have

Table 7
Virtual experiment

Variable	Number of Values	Values (Number of Agents)
Media scenarios	12	
Baseline	1	None (3,001)
Single	5	ad (3,002), Web (3,002), radio (3,002), call (3,002), mail (3,027)
Bundles	6	ad + Web (3,003), radio + Web (3,003), mail + Web (3,028), mail + radio + Web (3,029), all but ad [call + mail + radio + Web] (3,030), all [news-ad + call + mail + radio + Web] (3,031)
Network topology	3	random, cellular, small world

Note: Overall design $12 \times 3 = 36$ cells, 104 time periods per run, and 50 or 100 replications per cell depending on the number needed to get a well-formed average with minimal variance.

explored many of these in other settings. Herein, however, we limit the primary scope to these two factors.

We ran a baseline case with none of the media present. This was used to measure the amount of change in the population just based on the effect of an opinion holder. Next, in addition to the opinion leader, we explored the impact of each of the media in isolation. Finally, six different multimedia combinations were examined. Three of the cases were simple combinations with the website: the advertisement and the website, the radio advertisement and the website, and the mailing and the website. Three other super-media were constructed containing larger sets of media—the mailing, radio ad, and website; the information call center, mailing, radio ad, and website; and the print ad, center, mailing, radio ad, and website (all of the media together).

While the network topologies are varied across the different media scenarios, it is important to note that (a) the density for all three topologies was kept the same, and (b) the number of agents was kept the same for all three topologies. Thus, any experimental results are a function of topology, not the number of relations or edges in the graphs. It is important to note that at certain densities and number of agents these three network topologies can become indistinguishable. To avoid this problem we used a density around .0133, which is within the realm seen in human networks (e.g., see the densities of networks available on the Web) with 3,000 agents.

In this virtual experiment, we chose to look at the most realistic agents possible: The literacy and access mechanisms described in Section 2.6 were all enabled. These variations are based on national averages as described. To determine which agents were literate, had Internet access, or obtained a newspaper on a regular basis, data were gathered from a number of places, including the National Center for Statistics, the Pew Research center (Pew Internet, 2007), and Newspaper Association of America (The Newspaper Association of America, 2008), and other related studies (Livingston, 2007). Results were gathered for aggregate population groups, including breakdowns by gender, age, race, education, and income, and then used to model information access in Construct (Hirshman et al., 2008a). These national averages were then used to set the characteristics of the agents in the virtual

experiments. A secondary study was conducted to understand the individual effects of each of these mechanisms. These are described in Hirshman et al. (2008a). A key result of these access constraints is to slow down the rate of diffusion, and in cases where multiple constraints coexist in a subpopulation to put that subpopulation at risk of not receiving critical information.

While media and network topology are our key independent variables, there are additional factors that we control for such as age and education by keeping the distribution constant across the experiment conditions. These are described in Table 3. What this means is that we can analyze the results in terms of the impact on populations with different sociodemographic categories. Herein, we focus on the impact on individuals at different educational levels, as education has both a cognitive and a social component.

Similarly, there are a large number of potential dependent variables, including the overall shape of the social and knowledge networks and the way in which they evolve. Herein, we focus on three key outcomes: knowledge, belief, and activity. In particular, we examine the impact of media and topology on the fraction of the population that has the know-how to engage in activity, the belief that they should engage in an activity, and that actually engage in that activity. Results are presented as a percentage change from the baseline case where there is no media. Since the baseline condition for each network topology is fundamentally different, the impact of media under different network topologies is contrasted with the baseline for no media for that network topology.

4. Strategies for media

We find that on average, any media is better than none regardless of the network topology in terms of social outcomes. Since our series of interventions seek to suppress an activity, we are looking for media combinations that are most effective in suppressing the fraction of the population that takes part in an activity and believes that it should take part in the activity. However, our desired effects on know-how or expertise are more nuanced. In general, less expertise will reduce the involvement in the activity, which is what we desire. While full knowledge may increase the belief that this activity should be done, individuals with complete know-how may still refrain from engaging in the activity since they may be more aware of any negative repercussions. Thus, for expertise, the socially desirable outcome, is not clear. For this analysis, however, we treat suppression in the spread of know-how knowledge as a positive social outcome. Thus, on average, any media tends to suppress the spread of expertise, the spread of beliefs, and the engagement in a nondesirable activity. We now explore these results in more detail.

A second overarching result is that the ability of message passing to effect change, regardless of media choice, is strongest for information, then beliefs, and finally activity. In other words, a simple but short media campaign may educate the population, but it is unlikely to affect behavior in the short run. In some simulations virtually everyone will learn the information but less than 5% may change their behavior.

4.1. Media effects by education

We find that media, on average, have the most impact on more highly educated individuals in terms of expertise and behavior, and comparable impacts on beliefs. We also find that the impact of media is strongest on behavior, slightly less strong on beliefs, and weakest on expertise. We also find that while any media is relatively effective in spreading information and so impact the fraction of the population with know-how, to be effective at altering beliefs and activities a multimedia strategy is needed. While we present and analyze the results here from the random network, comparable results were observed for networks with similar topologies.

In general, media tend to suppress the expertise gained for all groups, except those with only a bachelor's degree (see Fig. 2). There are two countering factors at work: Individuals with higher education have, on average, more initial expertise even in association with this activity and they have larger social networks. As the level of education increases there is less for them to learn so the amount they learn is less likely to increase, and they have more others to learn from so the amount they learn is likely to increase; but they are just as likely if not more so to learn countervailing information and beliefs. For the most educated, there is a large decrease in the population that gains know-how in part because they quickly choose to not engage even without the know-how. We observed that the percentage change in the most highly educated agents fluctuates greatly due to the fact that there are so few of them included in the population in this virtual experiment. The chance of agents actively getting information through any media decreases with lower level of education. In part, this is due to the fact that factors like illiteracy and lack of access to the Internet or newspapers increase as education decreases. This means they are less likely to get information. Thus, there should be less change in the know-how of these least educated groups and indeed this is the case.

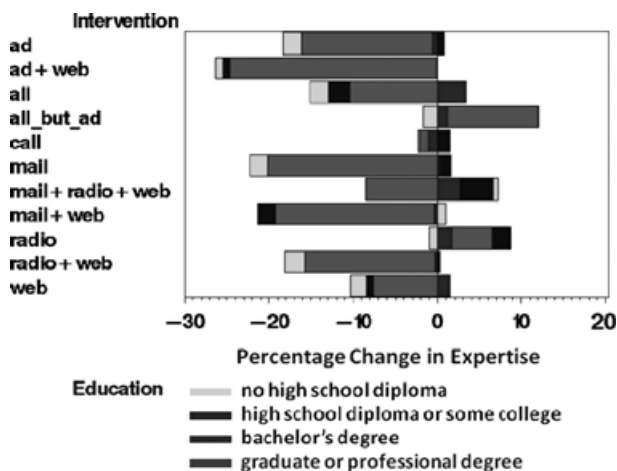


Fig. 2. The impact of media on the spread of expertise by educational level.

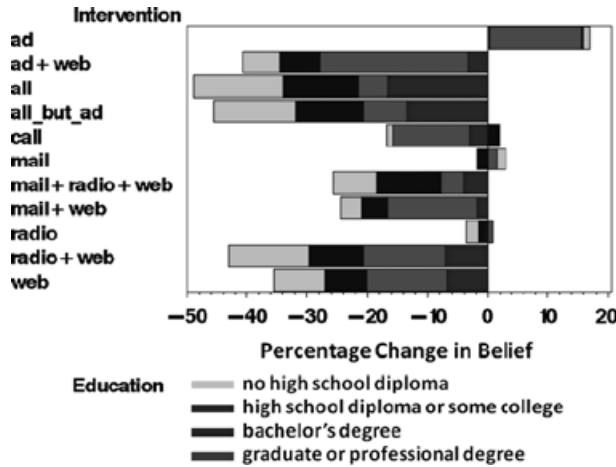


Fig. 3. The impact of media on the spread of beliefs by educational level.

In general, independent of educational level, all media tend to result in less of the population believing that they should engage in the activity (Fig. 3). The critical exception here is the news ad, which can actually lead to an increase in the number who hold one of the pro-activity beliefs. This is because the most educated are most likely to read newspapers and be unaffected by the literacy mechanism; the information that such agents learn in the newspaper is largely about the expertise needed to engage and information about the inadvisability of engaging in the activity. Additionally, by interacting with the newspaper, these highly educated people are not interacting with members of the population and so are less likely to come into contact with the beliefs of others. Consequently, they are more likely to set their own beliefs based on information rather than social influence. A similar effect can be seen with respect to mail and radio ads. The effects in these last two cases are less pronounced as less information is conveyed by these mediums. Information centers and the Web because they hold information on beliefs as well as other information have the opposite impact; that is, they actually serve to decrease the fraction of the population who believe they should engage in this activity. The critical point here is that the content of the message is at least as critical as the media by which it is communicated. Combinations of media are best at suppressing the diffusion of inappropriate beliefs.

In terms of activity, some media can actually lead to an increase in the fraction of the population engaged in the undesirable activity (Fig. 4). In other words, there are indeed unintended consequences for some media and media combinations. The tendency to engage in the activity, despite the media, is most pronounced for those with a bachelor's degree. In this case, this increase in behavior is due to the fact that the media is ineffective; for example, both radio and mail + Web media lead to minimal changes in this subpopulation's expertise or beliefs. All agents are learning not just from the media but from other humans with whom they interacted. Hence, for these college educated agents, social influence from other humans appears to be impacting their behavior. The result is an increase in




Fig. 4. The impact of media on the engagement in activity by educational level.

engagement in the activity. For the most highly educated, in the all-but-newspaper ad case they engage more in the activity despite a population level decrease in belief due to the even higher increase in expertise; this effect is magnified by the small numbers effect. Nevertheless, an important observation is that there can be many unanticipated consequences in terms of engagement in activities.

4.2. Network topology effects

We now explore the mitigating effects of network topology. In general, the high-level conclusions seen when looking at educational impacts still hold regardless of topology: The spread of expertise is more predictable than the spread of beliefs than the spread of activity. Further, regardless of network topology, media are typically more effective than single media, the impacts of the media are not cumulative, and there are unanticipated consequences particularly in terms of the impact of media on engagement in activities.

However, there are important differences by topology (Fig. 5). Small-world networks tend to vastly exacerbate the impact of media over that observed with random networks and to an extent over cellular networks. That is, if the real world were truly one with a small world, then the impact of media should be more pronounced than was seen when the random network topology was used in the last section. While the small-world model captures some aspects of real-world networks—such as the presence of communication hubs—it is not as clustered as real human networks. Cellular networks capture this clustering in a manner that is more similar to what is observed in the real world. Interestingly, the impact of clustering is much less regular than the impact of hubs vis-à-vis the effects of media. The reason for this is complex. First, cellular networks can spread information very rapidly between seemingly distant agents, so regardless of the media, information is moving more rapidly and so are beliefs. Secondly, when a member of the general population gets information from a

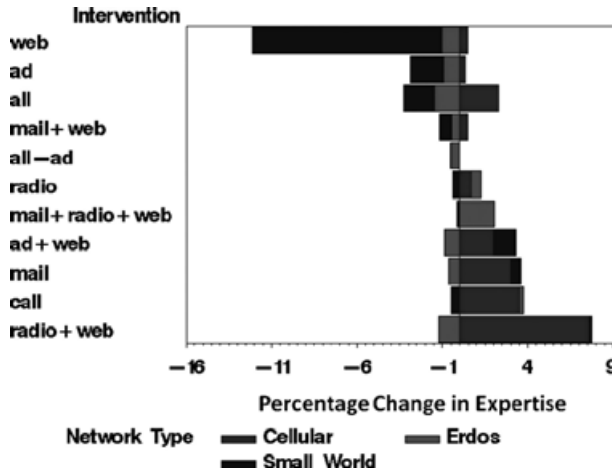


Fig. 5. The impact of media on the spread of expertise by network topology.

media agent that information diffuses to the cell very rapidly; however, the movement out of the cell to another group depends on the connectors happening to spread that specific information. If the connector happens not to send the information from the media agent, then diffusion can be slow between cells.

For the spread of expertise, Web-related media tend to greatly reduce the spread of know-how knowledge in small world networks but lead to a small increase in cellular networks. Basically, in a small world network, if a hub does not get information from the Web, then few agents will due to lack of familiarity. In contrast, in a cellular network, if any member of a cell gets information from the Web then many agents from the same cell will do so due to intercommunication between agents in the same cell. In cellular and small-world networks, the combination of radio + Web is particularly effective at spreading expertise. The familiarity that the general population has with the radio increases their likelihood of getting information from it that then increases their understanding of the additional expertise to be had from the Web. Then due to hubs in the small world network, and the clustering in the cellular network, the information spreads rapidly to the overall population.

In terms of beliefs, for all network topologies, complex media are more effective at changing the beliefs in the general population (Fig. 6). Most individual media will also lead to decreases in aggregate levels of belief, though individual media are less effective than bundles. Radio ads used in isolation seem to backfire when the network topology is small world. In that case, when the hubs are communicating a belief that one should engage in the activity and the ads are projecting the beliefs that one should not, the hubs will dominate. Thus, if the world were really a small-world network it would be critical to convince opinion leaders. For cellular networks, such an unintended consequence occurs to a lesser extent when mass mailings are used.

As for activities, most Web-related and complex media tend to result in fewer members of the general population engaging in the activity (Fig. 7). These decreases are




Fig. 6. The impact of media on the spread of beliefs by network topology.

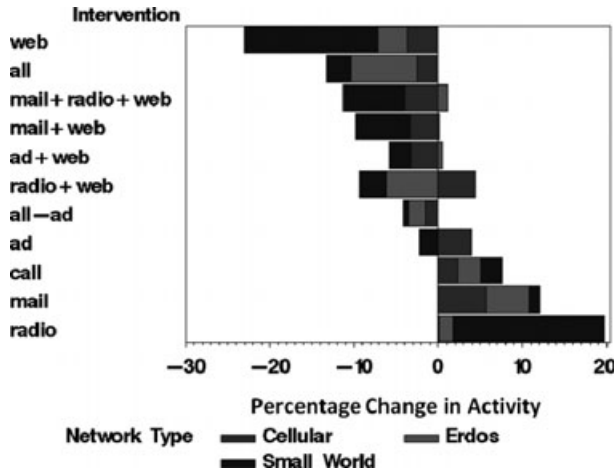


Fig. 7. The impact of media on the engagement in activity by network topology.

partially a result of the fact that substantially fewer people will have the requisite knowledge and beliefs, as seen in the previous graphs. Radio media will have severe unintended consequences in either random or small-world networks, partially due to the fact that knowledge and belief both increase when this medium is active. Finally, we find that cellular topologies tend to resist the impact of media on activities. That is, regardless of the medium used, there is less change in the general population's engagement in the activity than for other network topologies. This suggests that if the real world were cellular, then it will take not only more media, but media over a longer time period to effect change in behavior.

5. Sensitivity analysis

To what extent are the exact results reported herein robust? On the surface, this question sounds straightforward to answer. Unfortunately, the answer is complex. Construct is a simulation with many sources of stochastic variability. Thus, the primary threat to the reliability of Construct outcomes (similar to any simulation) is path dependence due to model components that stochastically influence information diffusion. Although we did not perform statistical analyses on variances, we ran 50–100 replications of each condition to guard against outcomes due to stochastic path dependence. Based on informal inspections of Construct data in these and many other virtual experiments, we know that the cumulative effect of stochastic model components on variability tends to be relatively small when compared to the variability created by model components that systematically influence information diffusion. Thus, error variance tends to be quite small relative to systematic (effect) variance, especially when given 50 or more replications. Therefore, the value of examining Construct data (and indeed when using any simulation as a tool for thought), lies more in the relative ordering of noticeable differences than in statistical notions of confidence.

We report means, but do not report error bars and we believe in this case their inclusion would be inappropriate. Error bars would be especially tricky to produce as we are actually reporting the mean as a relative change to baseline. And calculating error bars for deviation from the baseline and for all pairwise differences would make the visuals in the graphs we provide incomprehensible. While daunting issues and computationally prohibitive issues, these are not the core problem with error bars. Rather, error bars are misleading as they depend on the number of replications—for large-scale systems you get stability in differences of means long before the error bars calm down—so noticeable differences tend to be real differences. In other words, as one increases replications the variance about the mean decreases due to minimization in stochastic variations. Further, as one increases replications, differences in means tend toward zero, unless there are systematic variations. As a result, small differences are tending toward insignificance and high variability and large differences toward significance and low variability.

6. Conclusion

We examined the issue of how best to provide information to alter behavior on the part of a population. This was done using the Construct model that takes both social and cognitive factors into account. Our results demonstrate that the most effective media for spreading information are not necessarily the most effective at spreading beliefs or altering behavior. Moreover, the most effective media will differ by subpopulation due to access constraints such as illiteracy, Internet penetration, and newspaper access. In general, we also find that the impacts of media are not cumulative, and that as one moves from the knowledge to the belief to the activity space, the prevalence of unintended consequences increases.

Clearly, additional controlled experiments to assess variations in parameters are called for. Key additional experiments would consider alternative network structures, and

structures with both hubs and clusters. Another key experiment to consider is one in which an examination is made of the role of local opinion leaders with and without media access. Further experiments are planned to address such issues and to examine further interactions between individual cognitive capabilities in conjunction with social network factors.

The work has several limitations. One key limitation is that we are not looking at message transmission in concrete terms: Was the message auditory or visual? Was it from a stranger or from a friend? Was the message firsthand or retold through a chain of friends? For example, as noted, television and radio appear the same in this model; all messages are treated in an abstract form without regard to method of transmission. Future work should explore how to account for the differences. Future work will also address both the manner of transmission as well as the trust that agents have in the agent who initially conveys the message.

However, that being said, the conclusions we can draw here are intriguing. Such work suggests that, in general, a multimedia strategy is needed in order to produce the greatest effect. Second, the content of the message being conveyed in terms of information and beliefs needs to be considered regardless of the medium when creating an effective media strategy. Web-based media often appear more successful simply because they have more and different information. If the communication technology is poor for conveying beliefs, as one might argue for newspaper ads, then alternative media such as the Web should be used, perhaps in addition to the Web. Third, simple media may backfire and have unintended consequences. For example, radio advertisements actually had the opposite of the desired effect; that is, they increased the spread of beliefs that one should engage in the activity and more people did engage in the activity even though it was socially undesirable. The differences in the impact of the media depend on the amount and type of information available through the media, and the general public's familiarity with and understanding of potential message types provided by the media are critical.

When a sociocognitive model such as Construct is used, network topology plays a key role. We examined three stylized forms—random, small world, and cellular. The impact of media on knowledge, beliefs, and activity can vary dramatically based on the network topology. While results for small-world networks tend to look like those for random networks, just more extreme, the results for cellular are sometimes in the opposite direction. Differences in the role that hubs and clustering play in the spread of information account for this difference.

Construct is a simulation environment based on an abstraction of the real world, and as such can provide results and guidance that are as good as the assumptions on which it is based. The work that we have presented has suggested that the sociocognitive aspects of behavior are important considerations when designing and evaluating the effects of media. Future work must carefully consider both the cognitive and social aspects of behavior. As our cognitive and our social assumptions improve, simulations such as Construct will become increasingly powerful tools for examining the etiology of social change.

Note

1. The Construct system itself is freely downloadable from the CASOS website at <http://www.casos.cs.cmu.edu/projects/construct>. It is critical to note that we are using Construct 3.5, which is a completely refactored, updated, and expanded version of the 1990 Construct, to utilize new agent technology that combines rule-based and equation-based processing, high-speed sparse social-network processing techniques, event-based media, etc. and to be more tightly tied to empirical data as input.

Acknowledgments

This work was supported in part by the IRS project in Computational Modeling and the National Science Foundation IGERT in CASOS (DGE 997276). In addition Tera Grid access was provided by the National Science Foundation through the TeraGrid resources provided by Purdue University. Additional support for the fundamental development of Construct was provided in part by the Office of Naval Research (N00014-06-1-0104), the Internal Revenue Service, the National Science Foundation (SES-0452487), the Air Force Office of Sponsored Research (MURI: Cultural Modeling of the Adversary, 600322) for research in the area of dynamic network analysis, ARL Alion—STM1193391, ARI—W91-WAW07C0063, and the Air Force (FA865007C6769) for examining the unintended consequence of action using multi-agent simulation. Additional support was provided by CASOS—the center for Computational Analysis of Social and Organizational Systems at Carnegie Mellon University. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Office of Naval Research, the Internal Revenue Service, the Air Force Office of Sponsored Research, the U.S. Air Force, the National Science Foundation or the U.S. Government.

References

- Anderson, J. (2005). *Cognitive psychology and its implications* (6th ed.). New York: Worth Publishers.
- Blau, P. M. (1977). *Inequality and heterogeneity*. New York: The Free Press of Macmillan Co.
- Carley, K. (1990). Group stability: A socio-cognitive approach. In E. Lawler, B. Markovsky, C. Ridgeway, & H. Walker (Eds.), *Advances in group processes: Theory & research*, Vol. VII. (pp. 1–44). Greenwich, CN: JAI Press.
- Carley, K. M. (1991). A theory of group stability. *American Sociology Review*, 56(3), 331–354.
- Carley, K. M. (1995). Communication technologies and their effect on cultural homogeneity, consensus, and the diffusion of new ideas. *Sociological Perspectives*, 38(4), 547–571.
- Carley, K. M. (1999). On the evolution of social and organizational networks. In S. B. Andrews, & David Knotee (Eds.), *Research on the sociology of organizations*, vol. 16 (pp. 1–30). Greenwich, CT: JAI Press.
- Carley, K. M. (2002). Smart agents and organizations of the future. In L. Lievrouw & S. Livingstone (Eds.), *The handbook of new media* (pp. 206–220). Thousand Oaks, CA: Sage.

- Carley, K. M., & Hill, V. (2001). Structural change and learning within organizations. In A. Lomi & E. R. Larsen (Eds.), *Dynamics of organizations: Computational modeling and organization theories* (pp. 63–92). Menlo Park, CA: MIT Press/AAAI.
- Carley, K. M., & Krackhardt, D. (1996). Cognitive inconsistencies and non-symmetric friendship. *Social Networks*, 18, 1–27.
- Carley, K. M., Lee, J., & Krackhardt, D. (2001). Destabilizing networks. *Connections*, 24(3), 31–34.
- Carley, K. M., & Newell, A. (1994). The nature of the social agent. *Journal of Mathematical Sociology*, 19(4), 221–262.
- Case, D. 2002. *Looking for information: A survey of research on information seeking, needs, and behavior* (2nd ed.). London: Academic Press.
- Catlett, C., Allcock, W. E., Andrews, P., Aydt, R., Bair, R., Balac, N., et al. (2007). TeraGrid: analysis of organization, system architecture, and middleware enabling new types of applications. HPC and Grids in Action, Ed. Lucio Grandinetti, IOS Press “Advances in Parallel Computing” series, Amsterdam.
- Dunbar, R. I. M. (1993). Co-evolution of neocortex size, group size, and language in humans. *Behavioral and Brain Sciences*, 16, 681–735.
- Erdos, P., & Renyi, A. (1959). On random graphs i. *Publicationes Mathematicae Debrecen*, 6, 290–297.
- Festinger, L. (1954). A theory of social comparison processes. *Human Relations*, 7, 117–140.
- Festinger, L. (1957). *A theory of cognitive dissonance*. Evanston, IL: Row, Peterson.
- Frantz, T., & Carley, K. M. (2005). *A formal characterization of cellular networks*. Carnegie Mellon University, School of Computer Science, Institute for Software Research International, Technical Report CMU-ISRI-05-109.
- Friedkin, N. (1998). *A structural theory of social influence*. New York: Cambridge University Press.
- Giddens, A. (1984). *The constitution of society: Outline of the theory of structuration*. Berkeley, CA: University of California Press.
- Hill, V., & Carley, K. M. (2008). Win friends and influence people: Relationships as conduits of organizational culture in temporary placement agencies. *Journal of Management Inquiry*, 17(4), 369–379.
- Hirshman, B. R., Carley, K. M., & Kowalchuck, M. J. (2007a). *Specifying agents in construct*. Carnegie Mellon University, School of Computer Science, Institute for Software Research, Technical Report, CMU-ISRI-07-107.
- Hirshman, B. R., Carley, K. M., & Kowalchuck, M. J. (2007b). *Loading networks in construct*. Carnegie Mellon University, School of Computer Science, Institute for Software Research, Technical Report, CMU-ISRI-07-116.
- Hirshman, B. R., Martin, M. K., Bigrigg, M., & Carley, K. M. (2008b). *The impact of educational media by socio-demographic attribute*. Carnegie Mellon University, School of Computer Science, Institute for Software Research, Technical Report, CMU-ISR-08-118.
- Hirshman, B. R., Martin, M. K., Birukou, A., Bigrigg, M., & Carley, K. M. (2008c). *The impact of educational media on real and stylized cities*. Carnegie Mellon University, School of Computer Science, Institute for Software Research, Technical Report, CMU-ISR-08-114.
- Hirshman, B. R., Martin, M. K., & Carley, K. M. (2008a). *Modeling information access in construct*. Carnegie Mellon University, School of Computer Science, Institute for Software Research, Technical Report, CMU-ISR-08-115.
- Kaufer, D. S., & Carley, K. M. (1993). *Communication at a distance: The effect of print on socio-cultural organization and change*. Hillsdale, NJ: Erlbaum.
- Kaufer, D. S., & Carley, K. M. (1994). Some concepts and axioms about communication: Proximate and at a distance. *Written Communication*, 11(1), 8–42.
- Kirsch, I., Jungebut, A., Jenkins, L., & Kolsted, A. (2002). *Adult literacy in America: A first look at the findings of the National Adult Literacy Survey* (3rd ed.). Washington, DC: National Center for Education Statistics.
- Kramer, R. (1999). Trust and distrust in organizations: emerging perspectives, enduring questions. *Annual Review of Psychology*, 50, 569–598.

- Kuhlthau, C. (1999). Inside the search process: Information seeking from the user's perspective. *Journal of the American Society for Information Science*, 42, 361–371.
- Lau, D., & Murnighan, J. K. (1998). Demographic diversity and faultlines: The compositional dynamics of organizational groups. *Academy of Management Review*, 23, 325–340.
- Lazarsfeld, P., & Merton, R. (1954). Friendship as a social process: A substantive and methodological analysis. In M. Berger, T. Abel, & C. Page (Eds.), *Freedom and control in modern society* (pp. 18–66). New York: Octagon.
- Livingston, G. (2007). *Latinos online* (pp. 1–23). Washington, DC: Pew Internet American Life Project.
- Manis, J. G., & Meltzer, B. N. (1978). *Symbolic interaction: A reader in social psychology*. Boston: Allyn & Bacon.
- Marchionni, G. (1995). *Information seeking in electronic environments*. New York: Cambridge University Press.
- McAllister, D. (1995). Affect- and cognition-based trust as foundations for interpersonal cooperation in organizations. *Academy of Management Journal*, 38, 24–59.
- McPherson, J. M., & Smith-Lovin, L. (1987). Homophily in voluntary organizations: Status distance and the composition of face-to-face groups. *American Sociological Review*, 52, 370–379.
- Moore, R. (1985). A formal theory of knowledge and action. In J. Hobbs & R. Moore (Eds.), *Formal theories of the commonsense world*, vol. 1 (p. 400). Norwood, NJ: Ablex Publishing.
- Moorthy, S., Ratchford, B., & Talukdar, D. (1997). Information search revisited: Theory and empirical analysis. *The Journal of Consumer Research*, 23, 263–277.
- Newman, M. (2000). Models of the small world. *Journal of Statistical Physics*, 101(3/4), 819–841.
- Pew Internet (2007). Pew Internet and American life project. Available at: www.pewinternet.org. Accessed on December 12, 2007.
- Salancik, G., & Pfeffer, J. (1978). A social information processing approach to job attitudes and task design. *Administrative Science Quarterly*, 23, 224–253.
- Schreiber, C. (2006). *Human and organizational risk modeling, critical personnel and leadership in network organizations*, Ph.D Thesis. Pittsburgh, PA: Carnegie Mellon University, School of Computer Science, Institute for Software Research, Computation, Organizations and Society.
- Schreiber, C., & Carley, K. M. (2004a). *Construct – a multi-agent network model for the co-evolution of agents and socio-cultural environments*. Carnegie Mellon University, School of Computer Science, Institute for Software Research International, Technical Report CMU-ISRI-04-109.
- Schreiber, C., & Carley, K. M. (2004b). Going beyond the data: Empirical validation leading to grounded theory. *Computational and Mathematical Organization Theory*, 10, 155–164.
- Schreiber, C., & Carley, K. M. (2008). Dynamic network leadership: Leading for learning adaptability. In M. Uhl-Bien, & R. Marion (Eds.), *Complexity leadership part I: Conceptual foundations* (pp. 291–332). Charlotte, NC: Information Age Publishing, Inc.
- Simon, H. (1982). *Models of bounded rationality*. Cambridge, MA: MIT Press.
- Smith-Lovin, L., McPherson, M., & Cook, J. (2001). Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27, 415–444.
- Stryker, S. (1980). *Symbolic interactionism*. Menlo Park, CA: Benjamin Cummings.
- The Newspaper Association of America (2008). *The daily and Sunday newspaper audience report 2007* (pp. 1–31). Arlington, VA: Interactive Market Systems, Inc.
- Valente, T. H., Poppe, P. R., & Merritt, A. P. (1996). Mass-media-generated interpersonal communication as sources of information about family planning. *Journal of Health Communication: International Perspectives*, 1(3), 247–266.
- Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of “small-world” networks. *Nature*, 393(6684), 409–410.
- Wegner, D. M. (1986). Transactive memory: A contemporary analysis of the group mind. In B. Mullen, & G. R. Goethals (Eds.), *Theories of group behavior* (pp. 185–205). New York: Springer-Verlag.

- White, H. C. (1992). *Identity and control: A structural theory of social action*. Princeton, NJ: Princeton University Press.
- White, H. C. (2008). *Identity and control: How social formations emerge*. Princeton, NJ: Princeton University Press.
- Wilkins, J. S. (1998). What's in a meme? *Journal of Memetics-Evolutionary Models of Information Transmission*, 2(1), 2–33.
- Wooldridge, M., & Jennings, N. (1995). Intelligent agents: Theory and practice. *Knowledge Engineering Review*, 10, 115–152.