

Trusting in Others' Biases: Fostering Guarded Trust in Collaborative Filtering and Recommender Systems

Jo Ann Oravec

Collaborative filtering is being used within organizations and in community contexts for knowledge management and decision support as well as the facilitation of interactions among individuals. This article analyzes rhetorical and technical efforts to establish trust in the constructions of individual opinions, reputations, and tastes provided by these systems. These initiatives have some important parallels with early efforts to support quantitative opinion polling and construct the notion of "public opinion." The article explores specific ways to increase trust in these systems, albeit a "guarded trust" in which individuals actively seek information about system foibles and analyze the reputations of participants.

Introduction

Individuals have sought advice from sources such as matchmakers and oracles in many societal contexts. Matchmakers have arranged marriages for potential partners for thousands of years, playing trusted roles in their societies. The skills of able matchmakers include understanding the circumstances, characteristics, and expectations of each party, as well as the affordances and constraints of the community in which the match is made. The legacy of the oracle is also long and continuing, with individuals seeking direction and

Jo Ann Oravec is an associate professor in the College of Business and Economics at the University of Wisconsin at Whitewater. She received her MBA, MA, MS, and Ph.D. degrees from the University of Wisconsin at Madison. She taught computer information systems and public policy at Baruch College of the City University of New York and also taught in the School of Business and the Computer Sciences Department at UW-Madison. In the 1990s, she served as the chair of the Privacy Council of the State of Wisconsin, the nation's first state-level council dealing with information technology and privacy issues. She has written several books (including *Virtual Individuals*, *Virtual Groups: Human Dimensions of Groupware and Computer Networking*, Cambridge University Press and *Home as Information Space: Electronic Commerce and the Domestication of Computer Networking*, forthcoming). She has worked for public television and developed software along with her academic ventures. She may be reached at <oravecj@uww.edu>.

Knowledge, Technology, & Policy, Fall 2004-Winter 2005, Vol. 17, No. 3-4, pp. 106-123.

advice on weighty matters from mysterious and often partly-disguised sources. Collaborative filtering and recommender systems are continuing in these traditions, rendering specific advice along with forms of online interpersonal interaction. They facilitate information exchange based on analyses of the decision behaviors or expressed preferences of the parties involved. Providing advice and establishing linkages among individuals (in whatever context) are often activities that involve the collection and analysis of information by intermediaries, based on profiles of individuals and stereotypes of ideal matches. They are also activities in which forms of bias and differences in perspective are tightly integrated (as people express their opinions within the context of a complex system), making the very notion of “trust” problematic. Participants in effect are trusting that the biases and opinions of others will be usefully and appropriately reflected in the systems.

This article analyzes the development of collaborative filtering and recommender systems and relates how trust in them has been supported through rhetorical and technological strategies in ways comparable to quantitative opinion polling. Rhetorical approaches to the study of information technology have gained significance as researchers explore the influences of how systems are depicted and promoted on their eventual adoption and utilization (Gurak, 1997; Haynes, 2002; Horton, 1990; Lee, 1992; Tovey, 1996). The results of these studies can aid managers and developers in deciding how systems should be characterized in their various stages to particular audiences. This article attempts to expand this literature by exploring how a new kind of system is strategically framed in rhetorical terms, comparing it with the emergence of opinion polls. The article contrasts two models of these framing approaches, one in which an oracle-style “black box” approach is taken and another in which aspects of system mechanics and social network structures are selectively revealed to participants. Cases of “reputation building” within collaborative filtering systems are given special attention as examples of how rhetorical means are utilized to increase levels of participation and bolster trust. The article argues that developers should guard against participant cynicism, providing accessible means for the redress of assaults upon the systems’ integrity. Such critical incidents play roles in the development of computer system genres, providing users with topics of concern about the systems and presenting developers with specific problems to overcome. Developers can take steps to support a “guarded trust” in the systems by specifically addressing user concerns about such incidents. They can also proactively present projected scenarios of deception or manipulation so that system participants have advance warning. With guarded trust, individuals are empowered to utilize specific means to make informed and cautious judgments as to how and when to apply the system’s advice given various scenarios of use.

Trust has become a strategic resource in many organizational and societal contexts (Armstrong, 2001; Houston, 2001; Putnam, 1996). As a consequence, organizations have often used various rhetorical strategies and specific marketing campaigns to reinforce trust in their products and services. For example, in 2000 Bridgestone/Firestone attempted to reestablish trust in its tires through advertising imagery and slogans (Davis and Halliday, 2000), although

the difficulties in mounting effective campaigns of this type are tremendous (Garfield, 2001). After the Tylenol tampering scare in the United States in the late 1970s, Johnson & Johnson used a variety of strategies (including personal appeals from the CEO) for consumers to regain trust in its product (Pearson and Misra, 1997). The widespread proliferation of “tamper-proof” packaging was a by-product of this incident; consumers were warned about current and projected tampering hazards and advised not to purchase items for which the protective seals were broken. This article describes how comparable kinds of rhetorical and technological approaches can be coupled to support participant trust in computer genres, as developers and participants work together to identify particular hazards and provide specific means for prevention.

Automating “Word of Mouth”

The expanding varieties of collaborative filtering and recommender systems have some common themes: through collection and analysis of information about individuals and their preferences, sets of individuals with comparable or compatible characteristics are matched (if only in their virtual, profiled forms) and recommendations or opinions are shared. For example, in one mode of collaborative filtering (the kind used in MovieLens, described below), individuals are given recommendations based on the opinions or other expressions of the group members to which their profiles are the closest match. The assumption of this system is that individuals would probably prefer the movies, books, ideas, or other items that their virtual cohorts would like. According to Stenmark (2000): “such systems are able to anticipate what items a user is likely to be interested in, and can thus, in a hopefully intelligent way, recommend such items” (p. 12). Some collaborative filtering systems (such as the one Amazon.com uses for book recommendations) primarily rely on the passive collection of information about subjects’ buying habits, web surfing activity, and other behaviors, which is then used to produce profiles (Ansari, Essegaiier, and Kohli, 2000). In other forms of these systems (discussed in a later section), “reputations” are key constructs; for example, individuals with wider social networks or higher-rated contributions to the system can increase their influence on system results by gaining higher reputation scores (Kollock, 1999; Resnick, Zeckhauser, Friedman, and Kuwabara, 2000; Zacharia, Moukas, and Maes, 1999; Zacharia and Maes, 2000).

“Recommender system” is generally considered a more inclusive term than collaborative filtering, referring to a broad class of systems that involve the exchange of recommendations and opinions of which collaborative filtering is just one variety (Ansari et al., 2000). The Firefly system designed by a group of MIT researchers (and later acquired by Microsoft) was a pioneering effort, providing early demonstrations of collaborative filtering mechanisms on the web (Oravec, 1996). Current streams of research in this area are incorporating insights about social network construction and maintenance (Kautz, Selman, and Shah, 1997; Raghavan, 2001; Resnick, 2002). Some researchers and theorists in artificial intelligence and information retrieval are construing

collaborative filtering in terms of its prospects for reducing Internet-related information overload (Borchers, Herlocker, Konstan, and Riedl, 1998; Wyner, 1998), thus incorporating it into many kinds of information management and e-commerce applications. However, without requisite levels of trust in the systems, their use may eventually be consigned to relatively trivial entertainment or commercial purposes.

MovieLens played an important role in efforts to increase popular acceptance of collaborative filtering and recommender systems, with a highly popular theme. With MovieLens (<http://www.movielens.umn.edu/>), users supply their own ratings of specific films, which are subsequently compared with profiles of others' ratings in order to obtain advice on movie selections. Frick (2000) reports from his personal movie-searching experience that it "asks you to rate movies from one star to five, and then compares your tastes with those of more than 50,000 people to come up with a list of films you'll be glad you came home with. When we gave thumbs up or down to ten films (loved *The Grifters*, hated *Meet Joe Black*), MovieLens picked one of our all-time favorites, *The Bank Dick*" (p. 34). Collaborative filtering mechanisms of the kind that are employed in MovieLens and Amazon.com are being extended and subsequently applied to areas including scientific peer review, political information dissemination, and distance education (Cambridge, 1999; Eysenbach, 1998; Holt, 2002). These mechanisms are often attempts to model "word of mouth" interactions, defined as the communications that individuals undertake when they glean information in informal contexts on various problem areas (Shardanand and Maes, 1995; Szymanski and Henard, 2001). News of an exciting new play or musical work travels quickly, spread in water cooler conversations and even the increased foot traffic around a performance hall. Collaborative filtering system developers often work on the assumption that word-of-mouth interactions are generally conducted with those of similar tastes and lifestyles. From the perspective of consumer research, Holt (2002) declares that "consumers want to author their lives, but they increasingly are looking for ghostwriters to help them out" (p. 86).

Much as opinion polling and related research served to make public opinion and the "general public" a more comprehensible entity (as discussed in the next section), collaborative filtering and recommender systems are helping to make a range of online social network notions and applications more explicit and intelligible. Investigators have developed research constructs for "word-of-mouth" and shown how it plays critical roles in information exchange on important topics. Word-of-mouth is also considered a highly effective mode of advertising, and thus related marketing applications involving collaborative filtering have become widespread (Clemons and Wang, 2000; Wyner, 2000). Organizations have explored collaborative filtering systems as ways of mapping expertise and connecting organizational participants who might not be aware of others' interests, such as in travel planning (Hwang and Fesenmaier, 2001). It can supplement the word-of-mouth exchanges that spontaneously occur in face-to-face settings (Maybury et al., 2000; Stenmark, 2000). Less serious applications also are proliferating; the system Jester attempts to match participants' tastes in jokes with those of others in order to deliver hu-

mor online that has a high chance of being regarded as funny (“Punch Line,” 1999).

Along with word-of-mouth, reputation-related notions are implemented in a variety of ways in collaborative filtering and recommender systems. The very notion of “reputation” in the context of collaborative filtering systems can be seen as a rhetorical attempt to bolster trust in the systems. The label of “popularity index” or a comparable phrase could also be aptly utilized in some cases to refer to the calculations that are made; however, the concept of popularity in itself does not have the rhetorical force of “reputation” in ethical realms. The notion of reputation has long roots in Western society, especially in business and exchange contexts; people “build” reputations through time, often through procedures established by professional groups or through tradition. Reputation has played increasingly important roles in markets for information products, since they are inherently characterized by imperfect information concerning the quality of what is being purchased (Stiglitz, 2000; Whitmeyer, 2000).

Systems that incorporate reputational aspects often require complex streams of input and attempt to shape the behavior of system participants to maintain these streams. For example, they try to increase levels of participation by creating incentives for individuals to input accurate data and monitor their reputation indices. Individuals invest time and effort in establishing reputations, which creates opportunity costs for those who choose to leave particular systems. For example, Plastic.com editors assign users with high reputation levels positions as “moderators,” which allows them to have substantial clout in decision making on how participant contributions are evaluated; however, their clout does not travel with them to other systems. The rhetoric of individual responsibility—for example, responsibility for reputation maintenance—thus gives some incentive to maintain high levels of participation. However, reputational notions can run counter to overall system interests when individual reputation is pursued at all costs, thus undermining some pertinent social considerations.

Comparative Perspectives on Trust in Opinion-Related Measures

Trust-related issues are of growing importance for organizations. For example, confidence in equity markets is failing in part because of recent business scandals, and workplace trust is decreasing as loyalty to organizations diminishes with downsizing warnings (Baba, 1999; “Backlash against Business,” 2002; Bigley and Pearce, 1998). Concerns about trust that are related to Internet usage as a whole and the specific applications of particular system genres are expanding as incidents of online fraud, deception, and scams increase, as documented by the US Federal Trade Commission (FTC) and other agencies (Katyal, 2001; McMaster, 2001). Collaborative filtering efforts may eventually serve to raise levels of trust and participation in community life as a significant portion of everyday social, artistic, and intellectual activity moves online (Resnick, 2002). Certain forms of collaborative filtering have the potential to reverse negative trends associated with the anonymity of various

kinds of Internet interactions, as participants are given some incentive to reveal personal information and establish reputations in the context of community activity. Individuals can also examine how their reputations are portrayed in the systems over time, possibly increasing trust by adding a kind of check to the system itself (feedback that is generally not available with opinion polling and other thicker-grained barometers of public sentiment).

There is considerable precedent for strategic initiatives involving the fostering of trust in specific opinion and advice systems. In past decades, many individuals and organizations placed significant levels of confidence in opinion polling and generalized notions of the “public,” even though the notions are hazy at best and the statistical tools for capturing a snapshot of public opinion are imprecise (Herbst, 1993). Such widespread confidence has afforded opportunities for kinds of statistically-infused discourse about expressions of the “public mind” that would not otherwise be feasible. For example, there are lengthy discussions about what “the people” want on nearly every television or radio public affairs forum, bolstered with polling data. However, this confidence also has drawbacks as powerful individuals or organizations often attempt without much vocal opposition to manipulate polling results and reflexively influence the expressed opinions of large numbers of individuals. The very notion of the public as a monolithic entity has often been questioned, and the ways that the data collected in polling are often given unwarranted significance is heavily criticized (Sanders, 1999). Many pollsters have consequently attempted to support the institution of polling through public information campaigns, pointing out how its various weaknesses can be mitigated.

Today most politicians take at least some notice of public opinion polls, and the polls are also used in the construction of the “general public” by academics, journalists and news commentators, although other tools such as focus groups are increasingly used (Bennett and Edelman, 1988; Wilkinson and Kitlinger, 2000). Although trust in these polls varies widely among segments of the population and in different contexts, the notion that there is a way to ascertain the public’s will has been inextricably linked to polling. Polling has been supported in its significance by the existence of mass markets and general elections, both of which are rooted in the preferences of large populations of individuals. For example, politicians are concerned about what “the majority” thinks in order to determine what it will do on Election Day. The mechanisms that produce public opinion polls (and their creation of a general public) have been refined and reinforced through the past decades, often in ways that support power and privilege (Gamson, 1989; Sobel, 1998). For example, organizations can purchase the services of polling firms to collect data on certain questions; by framing these questions in certain ways the organizations involved can often affect their results, although the entire process can have the imprimatur of trustworthy scientific precision and detachment. Trust in opinion polls has been bolstered by rhetoric, by associating polls with the “impartial, precise, specific, and scientific” (Larson, 1999), although aspects that are very non-scientific are often incorporated into their development.

Discourse about the “public” as an entity has been tightly coupled with polling for most of the past century, as a good share of political discussion has been relegated to the “machinery of polling” (Peters, 1995: 20). Trust in polling is often reinforced through the selective use of explanatory mechanisms. Along with polling numbers, some information about levels of confidence in the polls and perhaps the number of individuals surveyed is often conveyed in popular discourse. However, complete information about polling mechanics is infeasible if not impossible to provide, and the users of polling data must have some confidence in pollsters. Some level of institutional trust has been established over time since there are a number of polling organizations, each of which provides checks and balances to the others; for example, polls can be taken by both parties in a competitive situation (such as an election). However, trust in how polls are reported and polling data are used is often misplaced. For instance, Sobel (1998) relates how “the media tended to ‘frame’ stories in a way that simplified polling data” in order to favor a particular policy choice (p. 20). Polling mechanisms can be given some kinds of tests at least in electoral contexts (albeit incomplete ones), as pre-election and exit polls are compared with voting results—so that opinion polling mechanisms can be analyzed and possibly improved for their predictive capacity.

Accounts of how polling results have been manipulated have often diminished the reputation of polls, and subsequently provided the stimulus for changes in data collection and reporting protocols. Obvious failures of polling as a predictive device can also serve to discredit the opinion poll genre itself, sometimes fostering “deep skepticism” (Wring, 1998: 49). Many individuals have developed forms of “guarded trust” in polling, compensating for its limitations and keeping vigilant to new hazards and schemes as they emerge. Collaborative filtering mechanisms are of far more recent origins than those of opinion polling, but their bases in trust and institutional reputation are comparable. Participants need basic information about how the systems operate—however, it is impossible to provide complete information about the extensive mechanics of how personal data are collected and calculations performed. With polling systems, individuals are coaxed to trust in a “general public” and large-scale survey numbers; however, trust in collaborative filtering systems depends on more complex aspects, often based on recently-developed profiling techniques and small numbers of participants. Opinion polling is rooted as an institution on the collection of information about populations of individuals through sampling and other techniques. In contrast, collaborative filtering has the social network as the unit of interest, for which there is less of a research and applications track record, although it is indeed growing (Granovetter, 1985; Schneider and Teske, 1997).

Conceptualizing “Trust” in Collaborative Filtering Contexts

Trust is generally construed as a valuable commodity, one that can be placed both in individuals and in systems as a whole (Koehn, 1996; Lewicki and McAllister, 1998; Houston, 2001). It has aspects that are rooted in emotion and image as well as more rational dimensions linked to the analysis of be-

havior over time; it can thus be influenced by rhetoric as well as bolstered by technical support. In the previously-discussed case of opinion polling, without some level of trust in the systems involved, individuals would be far more reticent to participate. Lack of such trust has already become a factor in many polling contexts, as individuals either lie to pollsters or refuse to participate (Barone, 1996). Many efforts have been mounted by opinion polling researchers to take these factors into account when analyzing polling results (Herbst, 1993).

Trust is a rich construct; in the academic literature it has been given “a confusing potpourri of definitions applied to a host of units and levels of analysis” (Shapiro, 1987: 624). Sitkin and Roth (1993) outline four approaches toward the literature on trust, treating it (1) as an attribute of individuals, (2) as a behavior, (3) as a situational feature, and (4) as a form of institutional arrangement. In the context of collaborative filtering, all four approaches can be useful. However, an epistemological approach toward trust is also especially relevant, considering trust as intimately related to information collection and processing: the systems circumscribe the views that individuals receive of each other’s contributions, so the quality of information (especially about the profiles of participants) is of special importance. A number of profiles can be involved in the calculation of recommendations of entertainment products or books. Participants must consider the issues of whether these characterizations are appropriate and reasonably accurate, and whether the system that uses them provides the requisite information for individuals to make decisions.

Some systems are explicit as to specific dimensions of their social network composition, but thus open themselves up to concerns about fairness and potential manipulation. In a system such as Plastic.com (a current issues discussion system that incorporates some recommendation and filtering mechanisms), individuals are given information about the reputations of those who submit contributions and comments. Participants build up “plastic karma” for themselves by having their suggested stories published and through posting comments that are highly rated by moderators. Participants must have trust that their own reputations as well as others’ are properly maintained and disseminated; there has been considerable discussion in Plastic.com of possible abuses of the moderation system, demonstrating the difficulties developers face with trust considerations.

As the purposes to which collaborative filtering systems are placed become more important, concerns about whether trust is warranted are growing. Internet-based systems in which intermediaries and intelligent agents handle information transactions are often targets of abuse, both by individuals who stand to gain economically from the abusive activity and those who are attacking the systems for pleasure (Katyal, 2001). For example, eBay and other auction sites are associated with the greatest number of Internet-related crime reports among e-commerce websites (Hughes and Costello, 2001; Nadel, 2001). Pressure on system developers to build forms of bias and fraud detection into their systems to protect participants is increasing as concern by the public increases (and is subsequently noticed by authorities). For example, eBay

managers have “been scrutinized for their failure to effectively police self-serving ‘customer reviews’ posted by the sellers and their friends” (Urban et al., 2000). System designers cannot change human nature and eliminate entirely the potentials for misconduct. However, they can take steps to develop systems that give their participants the information they need to interpret the actions of others and to act appropriately themselves. A kind of guarded trust can be nurtured by designers by providing participants with appropriate notice of the potential fallibilities of the systems and supplying them with the kinds of information that would allow them to make informed choices about the likely behaviors and motivations of participants. Such conditional trust can increase users’ effectiveness in employing system results and perhaps increase their overall levels of participation.

Fostering the development of guarded trust involves mitigating complex factors such as the following, generally by making participants aware of them and providing special help features if and when they emerge:

1. little historical experience: there is still relatively scant experience associated with the systems, so their reputation for trustworthiness as well as usefulness has yet to be supported. The terms “collaborative filtering” and “recommender systems” are still not in common parlance;
2. system complexity and opacity: complex calculations are generally involved in producing results, the statistical reasoning behind which extends beyond most of their participants’ training and background;
3. potential for developer bias: possibilities for bias in the generation or presentation of results abound, as system developers can possibly characterize results in misleading ways without detection (for example, to increase the chances that recommendations will favor a particular product). Developers can also skew rating system results by creating fictitious raters. Such potentials would be especially problematic in commercially-sponsored systems that recommend certain products or services;
4. potential for participant bias: in some cases, participants in the systems can also bias the results that others receive by submitting input in strategic ways. For example, Amazon.com’s system for ranking reviewers has been reportedly beset by “gangs of review bashers out there casting votes without reading the reviews, determined to ‘win’ a higher ranking at another’s expense,” (Mark McAllister, quoted in Hamilton, 2001); and
5. system noise and participant error: confusion or lack of knowledge about the systems by participants submitting information can also result in systemic errors that may have the appearance of bias (the “butterfly ballot” syndrome). Research by McDonald and Ackerman (1998) shows that many collaborative filtering and recommender systems do not distinguish between different levels of knowledge (from novice to expert); thus providing extra explanatory help for novice users can be of special value.

Pertaining to the issue of providing explanatory help, we can contrast two approaches. In the first, the system is construed as a kind of “black box” (dispensing advice like an oracle, but not providing glimpses into system mechanics). In the second, such dimensions as reputation and social networks are more explicitly outlined (a “matchmaker” approach). Black box strategies

can be useful in situations where participants may want a quick suggestion as to what to purchase or read, and are unlikely to feel the need to explore the reasons that the suggestion was provided. For example, Amazon.com provides various recommendations to users but does not provide extensive details as to how profiles were composed and analyzed to produce them. In contrast, matchmaker strategies are more explicit than the black box approaches in how personal data are manipulated to produce advice and how system mechanics proceed (such as Herlocker, Konstan, and Riedl, 2000; Tatemura, 1999). These approaches may be preferable in situations where recommendations pertain to matters of significance and where individual reputation is involved.

System developers have growing levels of experience in developing explanation capabilities, although approaches that help explain social networks (as in matchmaker strategies) are just emerging. Expert systems (ES) designers in the realm of artificial intelligence have often incorporated explanation features, and research results have shown that “explanation facilities can make ES-generated advice more acceptable to users” (Ye and Johnson, 1995). In the context of collaborative filtering, developers can provide means for participants to receive more easily digestible information about how rankings and reputations are calculated and updated (Herlocker, Konstan, and Riedl, 2000); for example, they can present this information through a visual interface that supports querying (Tatemura, 1999). Some collaborative filtering systems are embedded in larger systems (as part of the “personalization” of the web), and are thus designed to be somewhat invisible to participants. Making sure that some of these less visible collaborative filtering systems offer forms of explanation capability (or at least provide overviews of their functions) is a step toward making collaborative filtering more acceptable overall as a genre. Such efforts may have positive side effects for collaborative filtering systems overall as individuals begin to learn about and get a better sense of the genre. However, substantial research in what specific kinds of explanations are most useful to participants in various contexts is still required.

Critical incidents involving manipulation and deception in the context of collaborative filtering are just emerging, and developers can assist users in projecting the wide range of potential hazards based on their own preliminary experience with the systems. Such incidents play important roles in discourse on computer systems; public response to various incidents related to privacy, viruses, and various attacks upon system integrity have served to shape legislation in these areas (Hussong, 2000). In the context of collaborative filtering, Dellarocas (2000a, b) describes how system participants can skew reputation information by providing unfairly high or low ratings, a problem of special importance in the beginning of a particular system’s evolution and in systems with a small number of participants. Developers can alert users when such skewing could be a factor. Such warnings may in some cases be akin to “crying wolf” but will serve to alert users to possible problems. They can also help convey the notion that developers are not hyping their systems but are fostering reasonable and conscientious utilization.

Proactive Initiatives to Increase Trust and Decrease Cynicism

Establishing trust in computer systems is a societal undertaking that incorporates the efforts of many educational, business, community, and professional groups, including librarians (Lynch, 2000), educators (Recker, Walker, and Wiley, 2000), legal analysts (Katy, 2001; Nadel, 2001), and marketing researchers (Holt, 2002). Collaborative filtering and recommender systems are among a variety of modes through which individuals can obtain expert advice; in this era of “reflexive modernity,” individuals can choose among an assortment of kinds of expert assistance (Giddens, 1990). Acceptance of the individualized “results” of these systems is essential to their success; the “psychological validity” of the recommendations the systems produce—the apparent appropriateness of the movie reviews, book recommendations, and other advice that are relayed by the systems—will be tested over time by participants. Participants will share information about the systems through various means, and systems that produce acceptable advice will probably be favored over time. In contrast, opinion polling and related tools that attempt to provide a snapshot of widespread trends have less high of a hurdle to jump toward acceptance, producing a broader-scale and less precisely tailored recommendation (as well as having a longer track record in terms of methodology).

Developers often use rhetorical strategies to reinforce trust; rhetorical dimensions of collaborative filtering systems have often been considered by developers, even in choice of name. Collaborative filtering was once more commonly known as “social filtering”; the latter term was reportedly deemed to have unfortunate associations with “ethnic cleansing” (Lashkari, 1998). One of the early instances of the label “collaborative filtering” is in Goldberg, Nichols, Oki, and Terry (1992). Developers and protagonists of collaborative filtering have used such themes as intelligence, community, and reputation to buttress the levels of trust in their systems in rhetorical terms (see Table 1). Through linking collaborative filtering with notions that have positive societal connotations, developers have attempted to build powerful rhetorical associations that can reinforce trust. For example, participating in a “web of trust” has more positive implications than simply being a member of a less descriptive “social network” of some sort. “Web of trust” is a “conceptual blend” metaphor that merges the concepts of web and trust into a new, powerful (yet related) entity (Gozzi, 1999-2000). “Community” and “social capital” have been explored in many computer system genres (Werry, 1999; Houston, 2001) but have special meaning in systems that so tightly rely on social networks. Rhetoric alone cannot guarantee sustained participation in particular systems, but it may increase the likelihood that individuals will frame their experiences in ways that support the systems’ overall goals and talk about the systems using a particular set of terms and images.

Establishing trust in a new communications genre in the public sphere often involves examination of the negative aspects of the genre, as in the example of opinion polling outlined in a previous section. Critical incidents often arise that expose for public view how media genres can be exploited; they are

TABLE 1

Some Emergent Themes in Rhetorical Approaches to Collaborative Filtering	
Community and Social Capital	Resnick, 2002
Web of Trust	Khare and Rifkin, 1997
Intelligence/Intelligent Agency	Wyner, 1998 Carlson, 2000 (citing Riedl)
Reputation/Individual Responsibility	Zacharia, Moukas and Maes, 1999 Zacharia and Maes, 2000 Resnick, Zeckhauser, Friedman, and Kuwabara, 2000

disseminated by such means as rumor and broadcast media. For example, the genre of the TV game show suffered tremendous harm in its early stages because of a well-publicized “scandal,” albeit a mild one in proportion to many of today’s media ethics cases (Goldman, 1998). There have already been well-publicized incidents of manipulation of a popularity ratings feature of *Yahoo! News*, as users attempted to raise the rankings of particular photographs (“Media Manipulation,” 2001); Dellarocas (2000a, b) has provided descriptions of other related hazards. Satirical accounts of collaborative filtering’s benefits and drawbacks have already been generated (“Collaborative Filtering Goes Bad,” 2001). “Cynicism” in organizational and institutional contexts has many dimensions, one of which is the sense that even if problems with certain system features are discovered nothing substantive will be done. As various cases of abuse of collaborative filtering are recognized by participants, designers should examine their systems and determine what precautions to take so that participation levels are not adversely affected.

As the applications of collaborative filtering systems expand in kind and availability, individuals can utilize an assortment of these systems and move from system to system as their needs and interests change. In such a dynamic environment, trust factors can play important roles as they provide incentives to remain in a certain system or move from system to system. Increased cynicism and a “crisis of confidence” in a particular system (for example, widespread charges of bias in recommendations toward certain products or services) may trigger mass exodus. However, learning how to maneuver in a particular system and building a reputation takes some time for participants, and so there are opportunity costs in switching to other systems.

Trust issues also include how personal information about participants is handled outside of the immediate context of the system; developers should thus be straightforward with participants as to their policies in this regard. With some collaborative filtering systems, even casual users who want simple advice about particular life situations are often asked for identifying information, or such information is collected from them surreptitiously. Such infor-

mation collection has apparently stymied the application of these systems in certain realms: Lynch (2000) contends that the implementation of collaborative filtering and recommender systems in academic library contexts has been stalled because of privacy concerns. This is especially unfortunate because library applications—including potentially sensitive ones in health care— may be among the most useful involving collaborative filtering.

The varieties of collaborative filtering systems are expanding, adding to the pressure on developers to request increasing levels of personal information in efforts to build more complex (and possibly more effective) reputation models and more personalized results. Combining collaborative filtering systems with customer management systems (for example, in the ordering mechanisms of Amazon.com and a number of other retailers) provides special problems for privacy abuse as sensitive financial data are combined with rich opinion-related information. This increased reliance on personal information could backfire as some individuals would be reluctant to release such information, and thus choose to “invent” personal information or eschew the systems entirely. Consumers are becoming more wary of Internet privacy issues as well-publicized legal cases abound (Hussong, 2000; Tynan, 2000). “Lurking” and “free-riding” (not contributing information, but using some form of system results) are possible in some collaborative filtering systems, giving individuals the temptation to decrease levels of participation. Giving answers without requisite amounts of forethought is certainly possible as well, which can erode the overall quality of system results. Putting extra burdens on users (such as those involved in providing personal information) can apparently stymie user participation (Starr, Ackerman, and Pazzani, 1996).

Establishing trust among various parties involves some projection and analysis of the qualities of the individuals involved, as well as assumptions about their future behaviors. However, for many important purposes it also involves the ability to take the role of the other (Baier, 1985). Even the numerically-oriented *Wall Street Journal* warns us that trust involves more than formulas, and that in many forms of online interaction “a person’s voice or body language, for instance—are absent” (Weber, 2000: B1). The notion that profiles of individuals could serve critical functions in constructing and maintaining various social goods can be problematic. Extensive use of these profiles could add to the potential for “objectification” of individuals—particularly in organizational contexts—as decision makers decrease contact with real-life individuals (Cheney and Carroll, 1997). Many politicians have relied heavily on opinion polls, often eschewing other forms of political information, and have missed important trends as a consequence. In many sensitive contexts, use of information systems should thus be coupled with the fostering of personal connections through face-to-face social networking.

Conclusions and Implications

In a complex society trust often resides along with forms of distrust and even cynicism. Individuals can have guarded trust in complex systems that allow for a wide range of human behaviors to be expressed, from the altruistic

to the disreputable, using these systems in appropriate ways to further their interests. Developers of collaborative filtering and recommender systems have the difficult challenge of fostering guarded trust in the context of their systems: simply proclaiming the systems' best features and promoting their use is not sufficient to ensure their adoption. They must indeed encourage participation, yet also provide the kinds of information that would allow individuals to make well-reasoned judgments about the other participants in the system and their synthesized opinions and reputations. Also, systems that provide some protections against cynicism, with mechanisms to deal effectively with problems as they occur, can inspire user confidence over the long run. Developers can proactively explore with users various critical incidents in which the systems can be manipulated, guiding them toward ways to mitigate particular abuses and to be vigilant for new ones. It is shortsighted for developers to promote systems unreservedly (using rhetoric to distract from system foibles) or attempt to cover up potential modes of deception.

Developers of collaborative filtering and recommender systems have other specific hurdles to overcome in supporting their genres. They must recognize and often counter the considerable rhetorical forces in society that more generalized notions related to "public opinion" currently have. They must present the case that (in some situations) personalized social networks and the delivery of tailored results are indeed a more effective way of obtaining advice than following wide-scale opinion polls. Most individuals still prefer recommendations rooted in broadly-based public opinion rather than the results of collaborative filtering mechanisms, which are based on the newer and more complex notions of "social networks" and "webs of trust." For example, despite the existence of individually-tailored collaborative filtering results to aid in making entertainment choices, many individuals still consult lists of top movie box office winners for movie recommendations, and the *New York Times* Best Seller List is still highly coveted in the realm of books. *Yahoo! News*' "Most E-mailed" feature (a measure of overall popularity) is still apparently preferred to many collaborative filtering efforts (Herz, 2001).

Collaborative filtering and recommender systems are not simply impersonal market mechanisms in which relatively-remote factors are involved. Rather, they are systems in which personal judgments are integral. Thus, intangible aspects such as respect for others' opinions and choices are often important elements, especially in small-scale systems, as individuals are called to give and receive subjective judgments with at least some degree of fairness. Systems in which users can inspire each other to remain diligent and focused (giving serious input, maintaining their reputations, keeping their ratings of specific items current, etc.) have a higher likelihood of attracting and maintaining high levels of participation. Users need to acquire a sense of potential kinds of manipulation and deception so as to take steps to protect themselves as well as to avoid such malfeasance in their own behavior.

Collaborative filtering and recommender systems still have many rhetorical and technical hurdles to overcome before they are widely adopted. However, by using such systems (in a context of guarded trust) to support our discover-

ies and decision-making we can often be more effective in information-related activities.

References

- Altheide, D. (2000). "Identity and the Definition of the Situation in a Mass-Mediated Context," *Symbolic Interaction*, 23(1), 1-27.
- Ansari, A., Essegai, S. and Kohli, R. (2000). "Internet Recommendation Systems," *Journal of Marketing Research*, 37(3), 363-375.
- Appelbaum, K. (1995). "Marriage with the Proper Stranger: Arranged Marriage in Metropolitan Japan," *Ethnology*, 34(1), 37-52.
- Armstrong, C. (2001). "Review of *Bowling Alone*," *Knowledge, Technology and Policy*, 13(4), 124-130.
- Baba, M. (1999). "Dangerous Liaisons: Trust, Distrust, and Information Technology in American Work Organizations," *Human Organization*, 58(3), 331-347.
- "Backlash against Business" (2002). *Economist*, 364(8280), 27-28.
- Baier, A. (1985). *Postures of the Mind: Essays on Mind and Morals*. Minneapolis: University of Minnesota Press.
- Barone, M. (1996). "Why Opinion Polls are Worth Less," *U.S. News & World Report*, 121(23), 52.
- Bennett, W.L., and Edelman, M. (1985). "Toward a New Political Narrative," *Journal of Communication*, 35, 156-171.
- Bigley, G. and Pearce, J. (1998). "Straining for Shared Meaning in Organization Science: Problems of Trust and Distrust," *Academy of Management Review*, 23(3), 405-422.
- Borchers, A., Herlocker, J., Konstan, J., and Riedl, J. (1998). "Ganging Up on Information Overload." *IEEE Computer*, 31(4), 106-108.
- Cambridge, D. (1999). "Supporting the Development of a National Constellation of Communities of Practice in the Scholarship of Teaching and Learning through the Use of Intelligent Agents." In Hoadley, C. and Roschelle, J. (eds.) *Proceedings of the Computer Support for Collaborative Learning (CSCL) 1999 Conference*, 12-15 December, Stanford University, Palo Alto, CA. Mahwah, NJ: Lawrence Erlbaum Associates.
- Carlson, S. (2000). "Web Site Uses Sophisticated Programming to Offer Movie Advice," *Chronicle of Higher Education*, 46(38), A53.
- Cheney, G. and Carroll, C. (1997). "The Person as Object in Discourses in and around Organizations," *Communication Research*, 24(6), 593-631.
- Clemons, E. and Wang, Y. (2000). "Special Issue: Technology Strategy for Electronic Marketplaces," *Journal of Management Information Systems*, 17(2), 5-8.
- "Collaborative Filtering Goes Bad" (2001). <<http://www.plastic.com/article.pl?sid=01/02/19/1413236&mode=thread>, accessed 26 June 2001.
- Davis, B. and Halliday, J. (2000). "Firestone Prepares Effort to Rebuild Damaged Brand," *Advertising Age*, 71(51), 8.
- Dean, J., Brandes, P., and Dharwadkar, R. (1998). "Organizational Cynicism," *Academy of Management Review*, 23(2), 341-353.
- Dellarocas, C. (2000a). Immunizing online reputation reporting systems against unfair ratings and discriminatory behavior. *Proceedings of the Second ACM Conference on Electronic Commerce*, Minneapolis, MN, October 2000.
- Dellarocas, C. (2000b). Mechanisms for coping with unfair ratings and discriminatory behavior in online reputation reporting systems. *Proceedings of the 21st International Conference on Information Systems (ICIS)*, Brisbane, Australia, December 2000.
- Eysenbach, G., Diepgen, T. and Coiera, E. (1998). "Towards Quality Management of Medical Information on the Internet: Evaluation, Labelling, and Filtering of Information," *BMJ: British Medical Journal*, 7171, 1496-1501.
- Fallows, J. (2000). "Can the Net Govern Itself?" *American Prospect*, 11(10), 48.
- Figallo, C. and Rhine, N. (2001). "Tapping the Grape Vine," *EContent*, 24(3), 38-43.
- Frick, R. (2000). "Pick of the Flicks," *Kiplinger's Personal Finance Magazine*, 54(12), 34.
- Gamson, W. (1989). "Media Discourse and Public Opinion: A Constructionist Approach," *American Journal of Sociology*, 95, 1-37.

- Garfield, B. (2001). "Firestone Bid to Rebuild Trust Comes Too Late for the Wary," *Advertising Age*, 72(18), 35.
- Giddens, A. (1990). *The Consequences of Modernity*. Cambridge: Polity.
- Gladwell, M. (1999). "The Science of the Sleeper," *New Yorker*, 75(29), 48-55.
- Goldberg, D., Nichols, D., Oki, B., and Terry, D. (1992). "Using Collaborative Filtering to Weave an Information Tapestry," *Communications of the ACM*, 35(12), 61-70.
- Goldman, D. (1998). "The Quiz Show Scandal," *Biography*, 2(12), 20-22.
- Gozzi, R. (1999-2000). "The Power of Metaphor," *ETC: A Review of General Semantics*, 56(4), 380-405.
- Granovetter, M. (1985). "Economic Action and Social Structure: The Problem of Embeddedness," *American Journal of Sociology*, 91, 481-510.
- Gurak, L. (1997). *Persuasion and Privacy in Cyberspace*. New Haven, CT: Yale University Press.
- Hamilton, J. (2001). "A Wide Web of Advice," *Business Week*, 3716, EB18-EB20.
- Haynes, C. (2002). "Playing Rhetoric: Language, Computer Games, and Hacker Pedagogy," *Computer Games in School Symposium*. Grafisk Multimedia Pedagogik Institute-Stockholm University, Sweden, June 2002.
- Herbst, S. (1993). *Numbered Voices: How Opinion Polling Has Shaped American Politics*. Chicago: University of Chicago Press, 1993.
- Herlocker, J., Konstan, J., and Riedl, J. (2000). "Explaining Collaborative Filtering Recommendations," *CSCW '00 Proceedings*, 2-6 December 2000, Philadelphia, PA.
- Herz, J. (2001). "Evolution of the Online Species," *Brandweek*, 42(10), SR12-13.
- Holt, D. (2002). "Why Do Brands Cause Trouble? A Dialectical Theory of Consumer Culture and Branding," *Journal of Consumer Research*, 29(1), 70-91.
- Horton, W. (1990). "Visual Rhetoric for Online Documents," *IEEE Transactions on Professional Communication*, 33(3), 108-113.
- Houston, D. (2001). "Trust in the Networked Economy: Doing Business on Web Time," *Business Horizons*, 44(2), 38-45.
- Hughes, R. and Costello, D. (2001). "The Electronic Yard Sale," *Wall Street Journal*, 237(107), W1.
- Hussong, S. (2000). "Medical Records and Your Privacy: Developing Federal Legislation to Protect Patient Privacy Rights," *American Journal of Law & Medicine*, 26(4), 453-475.
- Hwang, Y. and Fesenmaier, D. (2001). "Collaborative Filtering: Strategies for Travel Destination Bundling," In Payline, J., Wöber, K. and Fesenmaier, D. (eds.) *Information and Communication Technologies in Tourism 2001: Proceedings of the International Conference of ENTER 2001*. Montreal, Canada, 167-175.
- Katyal, N. (2001). "Criminal Law in Cyberspace," *University of Pennsylvania Law Review*, 149(4), 1003-1115.
- Kautz, H., Selman, B. and Shah, M. (1997). "ReferralWeb: Combining Social Networks and Collaborative Filtering," *Communications of the ACM*, 40(3), 63-65.
- Khare, R. and Rifkin, A. (1997). "Weaving a Web of Trust," *World Wide Web Journal*, 2(3), 77-112.
- Koehn, D. (1996). "Should We Trust in Trust?" *American Business Law Journal*, 34(2), 183-204.
- Kollock, P. (1999). "The Production of Trust in Online Markets," In Lawler, E.J., Macy, M., Thyne, S., and Walker, H.A. (eds.) *Advances in Group Processes (Vol. 16)*. Greenwich, CT: JAI Press.
- Larson, S. (1999). "Public Pinion in Television Election News: Beyond Polls," *Political Communication*, 16(2), 133-146.
- Lashkari, Y. (1998). "Subject: RE: Social Filtering or Collaborative Filtering?" <<http://www.sims.berkeley.edu/resources/mailling-lists/collab/0110.html>> (accessed June 26, 2001).
- Latour, B. (1999). *Pandora's Hope: Essays on the Reality of Science Studies*. Cambridge, MA: Harvard University Press.
- Lee, C. (1992). "The Rhetoric of an On-Line Document Retrieval System," *Proceedings of the ACM Tenth International Conference on Document Retrieval Systems*, 25-45.
- Lewicki, R. and Mcallister, D. (1998). "Trust and Distrust: New Relationships and Realities," *Academy of Management Review*, 23(3), 438, 21.
- Lynch, C. (2000). "From Automation to Transformation," *Educause Review*, 35(1), 60-69.
- Maes, P. (1994). "Agents That Reduce Work and Information Overload," *Communications of the ACM*, 37(7), 31-40.
- Maybury, M., D'Amore, R., and House, D. (2000). "Automating the Finding of Experts," *Research Technology Management*, 43(6), 12-16.

- McCloskey, D. (1992). "The Art of Forecasting: From Ancient to Modern Times," *CATO Journal*, 12(1), 23-44.
- McDonald, D. and Ackerman, M. (1998). "Just Talk to Me: A Field Study of Expertise Location," *Proceedings of CSCW '98*. Seattle: ACM Press, 315-324.
- McMaster, M. (2001, July). "Too Close for Comfort," *Sales and Marketing Management*, 42-48.
- "Media Manipulation, Carnival, American Puritanism, and Yahoo! News" (2001). <<http://www.plastic.com/article.pl?sid=01/03/02/1622245&mode=thread>> (Accessed 26 June 2001.)
- Nadel, M. (2001). "The Consumer Product Selection Process in an Internet Age: Obstacles to Maximum Effectiveness and Policy Options," *Harvard Journal of Law and Technology*, 14(1), 182-262.
- Oravec, J. (1996). *Virtual Individuals, Virtual Groups: Human Dimensions of Groupware and Computer Networking*. New York: Cambridge University Press.
- Pearson, C. and Misra, S. (1997). "Managing the Unthinkable," *Organizational Dynamics*, 26(2), 51-65.
- Pennock, D., Horvitz, E., and Giles, C.L. (2000). "Social Choice Theory and Recommender Systems: Analysis of the Axiomatic Foundations of Collaborative Filtering," *Proceedings of Seventeenth National Congress of AAAI/IAAI*, 729-734.
- Peters, J. (1995). "Historical Tensions in the Concept of Public Opinion," In Glasser, T. and Salmon, C. (eds.), *Public Opinion and the Communication of Consent*. New York: Guilford Press, 3-32.
- Phelps, J., Nowak, G., and Ferrell, E. (2000). "Privacy Concerns and Consumer Willingness to Provide Personal Information," *Journal of Public Policy & Marketing*, 19(1), 27-42.
- "Punch Line" (1999). *IIE Solutions*, 31(10), 66.
- Putnam, R. (1996). *Bowling Alone: The Collapse and Revival of American Community*. New York: Simon & Schuster.
- Raghavan, P. (2001). "Social Networks on the Web and in the Enterprise." *Lecture Notes in Computer Science, LNAI 2198*. New York: Springer-Verlag, 58-60.
- Recker, M., Walker, A. and Wiley, D. (2000). "An Interface for Collaborative Filtering of Educational Resources." *Proceedings of the 2000 International Conference on Artificial Intelligence (IC-AI'2000)*, 26-29 June 2000, Las Vegas, Nevada, 317-323.
- Resnick, P. (1997). "Filtering Information on the Internet," *Scientific American*, 276(3), 62-65.
- Resnick, P. and Varian, H. (1997). "Recommender Systems: Introduction to the Special Section." *Communications of the ACM* 40(3), 56-58.
- Resnick, P., Zeckhauser, R., Friedman, E., and Kuwabara, K. (2000). "Reputation Systems," *Communications of the ACM*, 43(12), 45-48.
- Resnick, P. (2002). "Beyond Bowling Together: SocioTechnical Capital," In Carroll, J. (Ed.) *Human-Computer Interaction in the New Millennium*. Reading, MA: Addison-Wesley.
- Sanders, L. (1999). "Democratic Politics and Survey Research," *Philosophy of the Social Sciences*, 29(2), 248-281.
- Schneider, M. and Teske, P. (1997). "Networks to Nowhere: Segregation and Stratification in Networks of Information about Schools," *American Journal of Political Science*, 41(4), 1201-1224.
- Selznick, P. (1994). *The Moral Commonwealth: Social Theory and the Promise of Community*. Berkeley, CA: University of California Press.
- Shapiro, S. (1987). "The Social Control of Impersonal Trust," *American Journal of Sociology*, 93, 623-658.
- Shardanand, U. and Maes, P. (1995). "Social Information Filtering: Algorithms for Automating Word of Mouth," *ACM Conference on Human Factors in Computing Systems (CHI-95)*, 210-217.
- Sitkin, S. and Roth, N. (1993). "Explaining the Limited Effectiveness of Legalistic "Remedies" for Trust/Distrust," *Organization Science*, 4, 367-392.
- Sobel, R. (1998). "Portraying American Public Opinion toward the Bosnian Crisis," *Harvard International Journal of Press/Politics*, 3(2), 16-34.
- Starr, B., Ackerman, M., and Pazzani, M. (1996). "Do-I-Care: A Collaborative Web Agent," *CHI-96 Proceedings, Short Papers*. New York: ACM Press, 256-259.
- Stenmark, D. (2000/2001). "Leveraging Tacit Organization Knowledge," *Journal of Management Information Systems*, 17(3), 9-25.
- Stiglitz, J. (2000). "The Contributions of the Economics of Information to Twentieth Century Economics," *Quarterly Journal of Economics*, 115(4), 1441-1479.

- Szymanski, D and Henard, D. (2001). "Customer Satisfaction: A Meta-Analysis of the Empirical Evidence," *Journal of the Academy of Marketing Science*, 29(1), 16-36.
- Tatemura, J. (1999). "Visual Querying and Explanation of Recommendations from Collaborative Filtering Systems," *Intelligent User Interfaces*. New York: ACM, 189.
- Tovey, J. (1996). "Computer Interfaces and Visual Rhetoric: Looking at the Technology," *Technical Communication Quarterly*, 5(1), 61-76.
- Troman, G. (2000). "Teacher Stress in the Low-Trust Society," *British Journal of Sociology of Education*, 21(3), 331-354.
- Tynan, D. (2000). "Privacy 2000: In Web We Trust?" *PC World*, 18(6), 103-112.
- Urban, G., Sultan, F., and Qualls, W. (1999). "Design and Evaluation of a Trust-Based Advisor on the Internet. Cambridge, MA: MIT Sloan School of Management working paper, <<http://ebusiness.mit.edu/research/papers/forum/>>.
- Urban, G., Sultan, F., and Qualls, W. (2000). "Placing Trust at the Center of your Internet Strategy," *Sloan Management Review*, 42(1), 39-49.
- Uslaner, E. (2000). "The Internet and Social Capital," *Communications of the ACM*, 43(12), 60-64.
- Warren, K. (2001). *In Defense of Public Opinion Polling*. Boulder, CO: Westview Press.
- Weber, T. (2000). "To Build Virtual Trust, Web Sites Develop "Reputation Managers." *Wall Street Journal—Eastern Edition*, 236(10), B1.
- Werry, C. (1999). "Imagined Electronic Community: Representations of Virtual Community in Contemporary Business Discourse," <http://firstmonday.org/issues/issue4_9/werry/#note1> (accessed 27 January 2001).
- Whitmeyer, J. (2000). "Effects of Positive Reputation Systems," *Social Science Research*, 29(2), 188-208.
- Wilkinson, S. and Kitzinger, C. (2000). "Clinton Faces Nation: A Case Study in the Construction of Focus Group Data as Public Opinion," *Sociological Review*, 48(3), 408-425.
- Williams, J. (2002). "6 Hot Technologies with a Purpose," *Library Journal*, 127(2), 51.
- Wring, D., (1998). "The Media and Intra-Party Democracy: 'New' Labour and the Clause Four Debate in Britain," *Democratization* 5(2), 42-61.
- Wyner, G. (1998). "Collaborative Filtering: Research or IT?" *Marketing Research*, 10(3), 35-38.
- Wyner, G. (2000). "Life (on the Internet) Imitates Research: More Internet Processes Feel Like Research Processes," *Marketing Research*, 12(2) 38-39.
- Ye, L. and Johnson, P. (1995). "The Impact of Explanation Facilities on User Acceptance of Expert Systems Advice," *MIS Quarterly*, 19(2), 157-173.
- Zacharia, G. and Maes, P. (2000). "Trust Management through Reputation Mechanisms," *Applied Artificial Intelligence* 14(9), 881-907.
- Zacharia, G., Moukas, A., and Maes, P. (1999). "Collaborative Reputation Mechanisms in Electronic Marketplaces," *Proceedings of the 32nd Hawaii International Conference on System Sciences*, 323-327.
- Zimmerman, J. and Kurapati, K. (2002). "Exposing Profiles to Build Trust in a Recommender," *Proceedings of CHI 2002*, Minneapolis, MN, April 2002. AMC Press, 608-609.

Copyright of Knowledge, Technology & Policy is the property of Transaction Publishers and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.