

What Makes a Search Engine Good for Democracy? Public Opinion Polling and the Evaluation of Software^{***}

Jo Ann Sison and Warren Sack
University of California, Santa Cruz

Abstract: We propose one possible set of criteria for evaluating software – specifically search engines – according to their usefulness for deliberative democracy. We then describe a user study of the search capabilities of three, existing, online archives (Google Groups, Omgili, or Technorati) of threaded, conversational data. Our study measures the capabilities of these search engines according to the proposed criteria. We conclude by stating the software design implications of the study.

INTRODUCTION

The essential need ... is the improvement of the methods and conditions of debate, discussion and persuasion. That is the problem of the public. John Dewey, *The Public and its Problems* (1927)

For at least a quarter century (see Hiltz and Turoff 1978, p. 195) many have been excited about the possibilities of computer networks as a means of facilitating democratic participation. Reviewing the area in the mid-1990s, sociologist Manuel Castells noted that local democracies appeared to be flourishing around the world and that “*When electronic means are added to expand participation and consultation by citizens, new technologies contribute to enhanced participation in local government*” (Castells 1997: 350). Collections, such as Tsagarousianou, Tambini and Bryan’s *Cyberdemocracy: Technology, Cities and Civic Networks*, documented how these experiments in local, online democracy were progressing in Amsterdam, Athens, Berlin, Bologna, Manchester, Santa Monica, and elsewhere (Tsagarousianou, Tambini and Bryan 1998).

At the national level, there has been less interest in citizen-to-citizen communication and more emphasis on delivering government services and documents (e.g., tax documents, forms for zoning petitions, legal codes, etc.) to citizens via the Internet. Moreover, it appears to be the case that – as municipal websites become more and more common – even local governments seem most intent on supporting a one-way, “services” model of information delivery rather than many-to-many deliberative discussion.

In fact, some social scientists argue that the potential of information and communication technologies lies more in the realm of delivery of governmental services rather than as facilitators for new forms of interaction between citizens and governments. Harvard political scientist Pippa Norris and her colleagues

performed a content analysis of national government departmental websites from 191 nations for a chapter of the *United Nations World Public Sector Report* entitled “Deepening Democracy via E-Governance.” In conclusion, Norris writes “*On balance, therefore, the new communication and information technologies have greater potential for deepening ... democracy, by strengthening government transparency, and by improving public satisfaction with the delivery of routine public services, more than by stimulating new forms of civic activism [emphasis added].*” (Norris, 2003, p.19). Others have come to similar conclusions (cf., Chadwick and May, 2003; Fountain, 2001). Norris summarizes these findings like this: “*...many commentators suggest that e-governance has succeeded mostly in its managerial technocratic functions of improved service delivery for routine matters such as registering for transportation permits, access to land registries, or tax payments, delivering efficiency gains by streamlining labor-intensive bureaucratic transactions, but that it has commonly largely failed in its participatory or consultative functions[emphasis added].*” (Norris, 2003, p. 3).

These social science findings can be understood in at least two ways by computer scientists: (1) computer technologies will never be capable of supporting the demands of deliberative democracy; or, (2) the current state of the art has not measured up to the needs of deliberative democracy, but that a future technology might address those needs. We propose that the second of these two interpretations is wiser. Because social scientists -- especially political scientists and policy experts -- usually treat computer technologies, even software, as “black boxes” (i.e., as components that are seen as immutable and impenetrable to internal examination), these social science studies of democracy and digital technologies are composed from the perspective of a user, not a designer, of technology. Consequently, we accept these findings, not as a predicament, but rather as a challenge: *What new computer technologies can be designed to facilitate democratic deliberation between citizens and between citizens and governments?*

The first step towards the development of better, democratic technology is the articulation of criteria that can be used to evaluate existing information and communication technologies (ICT) for democratic potential. The legal scholar, Larry Lessig, argues persuasively for this approach when he advocates that we build our constitutional values of free speech into the foundational architectures of software. This idea of “value-based design” of ICTs is gaining momentum (e.g., Introna and Nissenbaum, 2000). But, Lessig also touches on the main difficulty of this approach. Consider, for instance, the criteria used in computer science to determine if a piece of software is good or bad. If a piece of software is fast and

*** Appears in the Proceedings of *Tools for Participation: Collaboration, Deliberation, and Decision Support; Directions and Implications of Advanced Computing; Conference on Online Deliberation (DIAC-2008/OD2008)* Berkeley, CA., June 26-29, 2008: pp. 13-22.

efficient, then it is good. If it is slow or inefficient, then it is bad. Computer scientists all learn these criteria as undergraduates, right from the beginning, in the first course in the analysis of algorithms. Now, compare these criteria with some of the most important evaluative criteria employed by the designers of the U.S. Constitution: checks and balances between the legislative, executive, and judicial branches insure that a certain amount of inefficiency is built into our structures of governance and so, therefore, things do not run away from the will of the people. The bigger problem is this: the fundamental criteria of computer science and information theory are at cross-purposes with the democratic criteria we need. We need to employ democratic criteria in the design and evaluation of the ICTs that will underpin civil society of the next century.

In this paper we propose one possible set of criteria for evaluating ICTs – specifically search engines – according to their usefulness for deliberative democracy. We then describe a user study of the search capabilities of three, existing, online archives (Google Groups, Omgili, or Technorati) of threaded, conversational data. Our study measures the capabilities of these search engines according to the proposed criteria. We conclude by stating the ICT design implications of the study.

Defining Critical Criteria of Democracy

There are many ways of defining democracy. Deliberative democracy is just one of them. Following Norris (2003), *pluralist democracy* concerns elite competition between parties and large stakeholders; *representative democracy* is focused on electoral accountability of representatives; and, *deliberative democracy* takes citizen consultation and participation as central. Obviously, advocates of all these theories of democracy would see a public good in lower cost, more efficient, secure, and transparent electronic voting and government services. But, supporting deliberative democracy implies a very specific set of challenges: How can large discussions between citizens be facilitated? How can the government consult with its citizens? How can citizens find consensus and listen to differing opinions? How can citizens articulate a collective voice or a divergence of opinion? Central to a deliberative democracy is the means to facilitate high-quality debate and discussion in which citizens are able to argue for a position, listen to and substantively address the differences of other positions.

Political scientist James Fishkin states three conditions that make democratic deliberation possible: (1) political messages of substance can be exchanged at length; (2) there is opportunity to reflection on and debate the messages; and, (3) opinions stated in the messages can be interactively tested against rival arguments. (Fishkin, 1995). Fishkin's "rival arguments" are, in other words, principled differences of opinion (rather than just, for instance, biased prejudices). Principled difference is not simply good for

deliberation it is essential.

Recent empirical work demonstrating the necessity of difference of opinion in democratic exchange was done by the Joseph N. Cappella, Vincent Price, and Lilach Nir at the University of Pennsylvania and reported in their paper “Argument Repertoire as a Reliable and Valid Measure of Opinion Quality: Electronic Dialogue During Campaign 2000.” Cappella et al.’s study examined data taken from a year-long, multiwave survey of 1,684 Americans conducted during the 2000 presidential election campaign. Data was gathered in association with synchronous, real-time, moderated group discussions that were designed specifically to produce useful citizen deliberation. To measure what Cappella et al. term “argument repertoires” participants were asked to their opinion on several political issues.

Specifically, “argument repertoires” are defined as the range of arguments people hold both in support of and against their favored position on a particular political issue or toward some political object. *“For any given stated opinion on an issue, argument repertoire includes [1] the number of relevant reasons for the stated opinion and [2] the number of relevant reasons for the opposite opinion. For example, if people indicate that they are favorable toward the Republican party, the number of relevant reasons in support of their opinion gives the first part of their repertoire score. They are then asked the reasons why someone might be unfavorable toward the Republican party. The number of such relevant reasons provides the second part of their argument repertoire score.”* (Cappella et al., 2002).

Cappella et al. define “relevant reasons” as *“reasons that are acknowledged in public discourse as plausible reasons”*; and, Cappella et al. do not attempt to distinguish accurate from inaccurate reasons in analyzing their data. Thus, even though they show how measures of argument repertoires positively correlate with measures of political knowledge, what is most crucial here is an understanding of opposing points of view on the political issues – opposing opinions that are not necessarily right or wrong, but, most importantly *“acknowledged in public discourse as plausible.”* Cappella et al. show that argument repertoire is a reliable and valid measure of opinion quality and thus a measure of the quality of deliberative discussion.

One can understand these criteria and results as a demonstration that deliberative democracy hinges on the ability of participants to be able to know and articulate – not only their own opinions about a subject – but also the opinions of those who might oppose them. Rephrased as a software design problem, these results imply a very specific user requirement (or, more specifically, a citizen requirement for the software). *User need: What are the differing opinions about subject X that are expressed in a large*

volume of conversational data (e.g., a database of archived newsgroup threads or blogs)?

Defining Consensus Bias

We propose a novel (to computer science) measure of evaluation that borrows from well-known findings in the literatures of public opinion polling, political psychology, and social cognition (see Kunda, 1999 for an overview of social cognition). These findings relate to the fact that people have a hard time estimating, *a priori*, the distribution of others' opinions about some given topic. For example, consider the following question and set of responses posed during a public opinion poll.

“Now I am going to ask you something different. I am going to read a list of issues and I want you to tell me whether, overall, you think the Democrats or the Republicans do a better job with this issue. If you do not know, just tell me and we will move on to the next item....The economy... Do you think the Democrats or the Republicans do a better job with that? (If Democrats/Republicans, ask:) Would that be much better or somewhat better?”

<i>Democrats much better</i>	23%
<i>Democrats somewhat better</i>	18%
<i>Republicans somewhat better</i>	19%
<i>Republicans much better</i>	19%
<i>Both (vol.)</i>	3%
<i>Neither (vol.)</i>	7%
<i>Don't know/Refused</i>	10%

[Survey by Democracy Corps, Conducted by Greenberg Quinlan Rosner Research, January 14-January 19, 2003, Data provided by The Roper Center for Public Opinion Research, University of Connecticut. [USGREEN.03DCJAN.R030] <http://www.ropercenter.uconn.edu>]

Looking at the column of percentages, one can see a distribution of opinion: some prefer the Democrats, others prefer the Republicans. Researchers in the area of social cognition have found the following: we tend to see more support for our opinions than do people holding the opposing opinions. This finding is called the *consensus bias effect* (Ross, Greene, House, 1977).

To show such a result with the question outlined above, one would have to ask another set of questions to have respondents estimate the distribution that we see in the table above: i.e., “What percentage of the population will respond ‘Democrats do a better job with the economy?’” and “What percentage of the population will respond ‘Republicans do a better job with the economy?’” Consensus bias is a *relative measure*. One compares, for instance, the responses to these two additional questions by members of the “Democrats better” group with the estimates of the “Republicans better” group. Consensus bias predicts that the “Democrats better” will estimate that a larger portion of the population will support their opinion than the percentage estimated by the “Republicans better” for the pro-Democrats position; and, of course,

the converse should apply as well: the pro-Republicans will guess that a larger percentage of the general population is pro-Republican than the percentage guessed by the pro-Democrats to be pro-Republican.

This result is considered a *consensus bias* because of the consistent relative underestimation of the support for opposing opinions (as compared with the opposition's estimate). Because consensus bias is a *relative measure* it is possible for one of the parties to be correct in its estimate of support if, in comparison with its estimate, its estimate is larger than the estimate of the opposition. Consensus bias can be understood as a predisposition that biases our estimates of the prevalence of our own opinions. In Kunda's words "*False consensus exists when people's own choices, attitudes, or beliefs bias their estimates of those of other people, leading them to view their own reactions as relatively common while viewing alternative reactions as relatively uncommon.*" (Kunda, 1999, p. 37).

Given the relatively common occurrence of consensus bias, it is clear that a variety of factors need to be overcome to achieve a high measure of argument repertoire (as described by Cappella et al., 2002) in which participants in discussion both can articulate their own opinion, but are also well-enough acquainted with opposing opinions to understand the difference and diversity of their opponents. Ideally, use of a search engine for deliberative democracy would provide citizens with the means to discover opposing opinions and gain adequate insight into the diversity and difference of their fellow citizens.

To evaluate the proposed system we propose to measure consensus bias to determine if a search engine can lower the power of its effect on users/citizens. This measure of search engine efficacy greatly differs from the usual criteria of information retrieval which dictates that search engines should be measured according to the criteria of precision and recall.¹

METHOD

1. Select a well-defined group of citizens; e.g., students at the University of California; regular participants in the Usenet newsgroup alt.politics.elections; etc.
2. Select a number of public opinion poll questions that can currently be found on the website of The Roper Center for Public Opinion Research at the University of Connecticut

¹ "Consider an example information request I (of a test reference collection) and its set R of relevant documents. Let $|R|$ be the number of documents in this set. Assume that a given retrieval strategy (which is being evaluated) processes the information request I and generates a document answer set A . Let $|A|$ be the number of documents in this set. Further, let $|Ra|$ be the number of documents in the intersection of the sets R and A The recall and precision measures are defined as follows. Recall is the fraction of the relevant documents (the set R) which has been retrieved; i.e., $Recall = |Ra| / |R|$. Precision is the fraction of the retrieved documents (the set A) which is relevant; i.e., $Precision = |Ra| / |A|$." (Baeza-Yates and Ribeiro-Neto, 1999: 75).

(<http://www.ropercenter.uconn.edu/>). An example of such a question is the one used above to illustrate the discussion of consensus bias.

3. Divide the selected group into three groups: two control groups and one experimental group.

A. Control Group 1/With Poll Results/Gold Standard: Members of this sub-group are presented with the poll questions with knowledge of the poll results. They are then asked to answer the poll questions and estimate the distribution of responses within the entire group (knowing the distribution of results for a, presumably, larger more diverse population that was queried for the original poll).

B. Control Group 2/No Information/Lead Standard: Members of this sub-group are presented with the poll questions with no knowledge of the poll results and asked to respond to the question and estimate the distribution of opinion within the entire group. They are given no means to investigate the distribution of the range of possible opinions.

C. Experimental Group/With Search Engine: Members of this sub-group are presented with the poll questions and asked to respond to the questions and estimate the distribution of opinion within the entire group. To help them with their estimates, members of this subgroup are given access to one of three existing search engines for online discussions.

Hypothesis: Sufficient usage of search engines will decrease measurable consensus bias. Control Group 1 should have the smallest measure of consensus bias. Control Group 2 should have the largest. The Experimental Group – i.e., those with access to a search engine -- should perform better than Control Group 2 and slightly less well than Control Group 1.

Participants

Ninety-three undergraduates (28 women, 65 men, mean age = 20.1 years) at the University of California at Santa Cruz participated in the study for partial fulfillment of course requirements for an Introduction to Digital Media class. On the demographic survey, participants were asked to rate themselves on an ordered response scale with the following responses: 1 = very liberal, 2 = liberal, 3 = conservative, 4 = very conservative. Participants who chose not to answer this item were excluded from the analysis. The large majority of participants labeled themselves as liberal ($n = 66$), as compared to those labeling themselves as very liberal ($n = 15$) or conservative ($n = 8$). No participant chose the label of very conservative. Because there were not adequate numbers of participants in each political grouping, political affiliation was not included as a factor in the statistical analyses.

Materials

The first session of the study took place in a large-group lecture where participants completed the following two surveys: 1) a survey to collect demographic information and assess computer and Internet usage and 2) a multiple choice survey assessing political opinion on the topics of global warming, energy policy, and stem cell research. The political opinion survey items were selected from different national surveys from 2005 using the Roper Poll database.

The second session of the study was an individual lab session in which participants were presented with the same political opinion survey for a second time in an online format. After making their consensus estimates, participants completed an exit questionnaire with open-ended questions regarding the task and participant perceptions of online discussions.

Design and Procedure

Participants were randomly selected to make their consensus estimates in three conditions: a control condition, a search engine-condition, and a poll results condition. Control condition participants received no information to assist in making their estimates. Search engine-participants were instructed to use a search engine to read online discussions to help them make their consensus estimates. Within the search engine condition, participants were randomly assigned to one of the following three search engines: Google Groups, Omgili, or Technorati. Participants in the poll result condition were shown the national poll results for each of the survey items when making their estimates.

Participants in all groups were asked to estimate the percentage of students in their Introduction to Digital Media class that chose each survey response. Because the estimates for each survey item response should total to 100 percent, to assist participants in making their estimates, participants received a prompt alerting them if their estimates did not sum to 100 percent. Participants were allotted one hour to complete the estimating task for the 18 survey items, and participants were alerted to try to finish their estimates if they were not done ten minutes before the session was scheduled to end. Because participants in the control and poll results conditions typically finished before the one hour session was up, they were also randomly assigned to use one of the three search engines to estimate consensus for four random survey items. After completing the estimating task, participants completed a written exit questionnaire and were then debriefed regarding the goals of the study.

RESULTS

The General Consensus Bias Effect

The consensus bias effect is defined here as occurring when participants who endorse a survey response provide higher estimates of consensus than participants not endorsing that response (Ross, et al. 1977; Krueger & Zeiger, 1993). The following analyses will compare estimates made for survey responses that are chosen, or endorsed, versus the estimates for responses which are not endorsed by a participant.

The following analysis looks at the relative levels of consensus estimate for endorsed responses versus non-endorsed responses in the three experimental conditions, collapsed across survey items. A 2 (endorsement: endorsed or non-endorsed) x 3 (condition: control, search engine, and poll results) ANOVA with consensus estimates as the dependent variable showed a significant main effect for endorsement, where the mean consensus estimate for endorsed survey responses ($M = 36.28$, $SE = .92$) was significantly higher than the mean consensus estimate for non-endorsed responses ($M = 17.89$, $SE = .26$), $F(1, 89) = 245.68$, $p < .001$, but there were no other significant main effects or interactions. The consensus bias effects were replicated across all experimental conditions.

Our study also resulted in several other results. We do not have the space, in this paper, to provide a detailed description of these other results, but we anticipate a longer, journal article will soon be forthcoming. Below we shortly describe some of these results.

Aside from relative levels of consensus estimates, we were also interested in the accuracy of estimates made when response options are endorsed or non-endorsed. We expected that participants without information to make their estimates should be the least accurate of the three groups and should have the largest difference scores; participants with some source of information to use in the estimation task (i.e., reading online discussions or poll results) should have lower difference scores than the control condition. We found that participants in the poll results condition made more accurate estimates as compared to the participants in either the control or search engine condition.

We were interested in seeing how closely the survey results from our UCSC student sample resembled the survey results from the national samples, because if national survey data is greatly divergent from the local results, this necessarily affects the utility of the national data in helping participants make estimates for the local population of UCSC students. For the most part, we found that the national data was helpful to participants in their efforts to estimate the distribution of their classmates' responses.

We also examined the effects of opinion popularity on estimation bias. Opinion popularity interacts strongly with endorsement ($F(1, 1072) = 60.39, p < .001$), and this interaction corroborates the result of Mullen and Hu (1988) where consensus bias differs for those who choose the majority option versus the minority option.

DISCUSSION

The main goal of this study was to see if consensus bias effects would be reduced when participants have different information sources, such as online discussions or national poll results to aid in making their estimates. As expected, consensus bias effects were observed across all conditions of the experiment; participants who endorse responses to survey items will give higher estimates of consensus than participants not endorsing those same responses. However, participants who had access to national poll results made significantly more accurate consensus estimates relative to the control and search engine groups. Searching and reading online conversations were not as effective as national poll results were in helping participants estimate consensus for political opinions. Our study was able to use the relative differences in the amount of consensus bias in the three conditions to demonstrate that not all information sources are equally useful in the task of understanding consensus. Aside from finding the consensus bias effect in all conditions of the experiment, this study also replicated the effect of response popularity on consensus estimates for different items. The general effect across items was that when participants endorsed an option that only a minority of students endorsed, those participants overestimated the proportion of students selecting that option, and vice versa, although the underestimation of majority-endorsed choices was not as extreme. Experimental condition did not seem to affect this general finding.

It seems that a combination of cognitive (e.g., selective exposure to information) and motivational effects (i.e., deviance concern) results in the consensus bias observed in this study. However, the cognitive informational effects seem to be more influential in this particular study, when one considers that only participants with poll results were able to make more significantly accurate consensus estimates as compared to controls. *Although searching for and reading online discussions might have been useful in understanding the opinions of others, and thus act to reduce consensus bias, our study found that this was not the case.*

In their current state, search engines for online conversation simply point users to where people are talking about a given topic, but do not help the user understand the aggregation of opinions contained within the different conversations. Despite some of the problematic characteristics of online conversation, we feel that it should be possible for a specialized search engine to provide the user with a

macro-view of opinions on a given political topic. Contemporary technologies of corpus-base computational linguistics, social network analysis, pattern recognition and machine learning, information retrieval, and information visualization could be applied to the computation and display of opinion “clusters” that would assist users to summarize and visualize the range of opinions. We have been working on the design of such a system that uses conversational data from newsgroups, blogs, and discussion lists (Sack et al., 2005). However, designing and implementing such a system is not simply a matter of assembling a set of well-known technologies.

Unfortunately, most previous computer technologies have been evaluated using criteria that are tangential to, if not at cross-purposes with, the criteria of deliberative democracy. This paper is an attempt to demonstrate how one might evaluate software using ideas and criteria from political psychology, public opinion polling and social cognition.

Acknowledgements

This work has been supported by Award #0416353 (“An Interface and Search Engine for Deliberation”) from the National Science Foundation, Directorate for Computer and Information Science and Engineering. Any opinions, findings and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect those of the National Science Foundation.

References

- Cappella, J.N., Price, V., and Nir, L. (2002). Argument Repertoire as a Reliable and Valid Measure of Opinion Quality: Electronic Dialogue During Campaign 2000. *Political Communication*, 19:73–93.
- Castells, M. (1997). *The Information Age: Economy, Society and Culture; Volume II: The Power of Identity*. Oxford: Blackwell.
- Chadwick, A. and May, C. (2003) Interactions between states and citizens in the age of the Internet: ‘e-government’ in the United States, Britain and the European Union. *Governance* 16(2): 271-300.
- Davis, R. (2005). *Politics online: Blogs, chatrooms, and discussion groups in American democracy*. Taylor & Francis Group: New York.
- Dewey, J. (1927) *The Public and its Problems*. New York: H. Holt & Co.
- Drenteal, P., and Moren-Cross, J.L. (2005). Social capital and social support on the web: the case of an internet mother site. *Sociology of Health & Illness*, 27, 920-943.
- Fishkin, J. (1995). *The voice of the people: public opinion and democracy*. New Haven, CT: Yale University Press.
- Fountain, J.E. (2001). *Building the Virtual State: Information Technology and Institutional Change*. Washington DC: Brookings Institution Press.
- Gilovich, T. (1990). Differential construal and the false consensus effect. *Journal of Personality and Social Psychology*, 59, 623-634.
- Griffin, D.W., and Ross, L. (1991). Subjective construal, human inference, and human misunderstanding. *Advances in Experimental Social Psychology*, 24, 319-359.
- Introna, L., and Nissenbaum, H. Shaping the web: Why the politics of search engines matters, *The Information Society*, 16(3):1-17, 2000.

- Kelleher, T., and Miller, B. M. (January 2006). Organizational blogs and the human voice: Relational strategies and relational outcomes. *Journal of Computer-Mediated Communication*, 11. Retrieved March 23, 2007, from <http://jcmc.indiana.edu/vol11/issue2/kelleher.html>
- Krueger, J. (2000). The projective perception of the social world: A building block of social comparison processes. In J. Suls & L. Wheeler (Eds.), *Handbook of social comparison: Theory and research* (pp. 323-352). New York: Kluwer Academic/Plenum Publishers.
- Krueger, J., & Zeiger, J.S. (1993). Social categorization and the truly false consensus effect. *Journal of Personality and Social Psychology*, 65, 670-680.
- Kunda, K. (1999). *Social Cognition*. Cambridge, MA: MIT Press.
- Luskin, R.C., Fishkin, J.S, and Jowell, R. (2002). Considered opinions: Deliberative polling in Britain. *British Journal of Political Science*, 32, 455-487.
- Marks, G., & Miller, N. (1985). The effect of uncertainty on consensus judgments. *Personality and Social Psychology Bulletin*, 11, 165-177.
- Marks, G., & Miller, N. (1987). Ten years of research on the false-consensus effect: An Empirical and theoretical review. *Psychological Bulletin*, 102, 72-90.
- Mullen B., Atkins, J.L., Champion, D.S., Edwards, C., Hardy, D., Story, J.E., et al. (1985). The false consensus effect: A meta-analysis of 115 hypothesis tests. *Journal of Experimental Social Psychology*, 21, 262-283.
- Mullen, B., & Hu, L. (1988). Social projection as a function of cognitive mechanisms: Two meta-analytic integrations. *British Journal of Social Psychology*, 27, 333-356.
- Norris, P. (2003). "Deepening Democracy via E-Governance," Draft chapter for the *UN World Public Sector Report* (online: <http://ksghome.harvard.edu/~pnorris.shorenstein.ksg/articles.htm>).
- Robinson, R.J., Keltner, D., Ward, A., & Ross, L. (1995). Actual versus assumed differences in construal: "Naïve realism" in intergroup perception and conflict. *Journal of Personality and Social Psychology*, 68, 404-417.
- Ross, L., Greene, D., & House, P. (1977). The false consensus effect: An egocentric bias in social perception and attribution processes. *Journal of Experimental Social Psychology*, 13, 279-301.
- Ross, L., & Ward, A. (1996). Naïve realism in everyday life: Implications for social conflict and misunderstanding. In E.S. Reed, E. Turiel, & T. Brown (Eds.), *Values and Knowledge* (pp.103-136). Mahwah, NJ: Laurence Erlbaum.
- Sack, W., Kelly, J., and Dale, M. (2005) "Searching the Net for Differences of Opinion," in *Proceedings of the Second Conference on Online Deliberation: Design, Research, and Practice*, Stanford University.
- Sanders, G.S., & Mullen, B. (1983). Accuracy in perceptions of consensus: Differential tendencies of people with majority and minority positions. *European Journal of Social Psychology*, 13, 57-70.
- Sherman, S.J., Presson, C.C., Chassin, L., Corty, E., & Olshavsky, R. (1983). The false consensus effect in estimates of smoking prevalence: Underlying mechanisms. *Personality and Social Psychology Bulletin*, 9, 197-207.
- Sherman, D.K., Nelson, L.D., & Ross, L.D. (2003). Naïve realism and affirmative action: Adversaries are more similar than they think. *Basic and Applied Social Psychology*, 25, 275-289.
- Sia, C., Tan, B. C. Y., & Wei, K. (2002). Group polarization and computer-mediated communication: Effects of communication cues, social presence, and anonymity. *Information Systems Research*, 13, 70-90.
- Siegel, J., Dubrovsky, V., Kiesler, S., & McGuire, T.W. (1986). Group processes in computer-mediated communication. *Organization Behavior and Human Decision Processes*, 37, 157-187.
- Tsagarousianou, R., Tambini, D., and Bryan, C. (1998). *Cyberdemocracy: Technology, Cities and Civic Network*. New York: Routledge.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185, 1124-1131.

Wetzel, C.G., & Walton, M.D. (1985). Developing biased social judgments: The false-consensus effect. *Journal of Personality and Social Psychology, 49*, 1352-1359.