Research Statement

Gaurav Sood

Stanford University

September 19, 2011

My research deals centrally with political knowledge. Who knows what about politics (and why), and the effect political knowledge (or misinformation) has on exercise of citizenship, are vital questions in the study of democratic politics. However, considerable conceptual and operational challenges in measurement of political knowledge, among other things, have made definitive progress of these questions hard. In my research, I clarify some of the operational and conceptual issues around measurement of political knowledge (and its complement, misinformation) and propose some solutions. I also analyze why some people learn (mechanism that *causes* knowledge) more about politics than others. As part of my exploration, I discover that some experimental counterfactuals of an 'informed public' may be misleading as many respondents are left well short of 'full information'. In related research, I investigate mechanisms behind acquisition and spread of misinformation. Outside this core research in political knowledge, I study how media and campaigns result in affective (rather than ideological) polarization, and how style of news delivery (e.g. satire) affects range of political behaviors via differential affective arousal.

The 'MAO' of Learning

Acquisition of knowledge has been thought to depend on motivation, ability, and opportunity (Luskin, 1990), or 'MAO' as a Google search revealed. In the real world, all these factors are largely endogenous, and hard to manipulate. For example, provision of good schools may enhance the ability to comprehend political affairs, but doing so randomly seems infeasible given politics.

Similarly providing people with a circle of friends who greatly value political information may lead them to be instrumentally motivated to acquire political information, though it is not particularly feasible to do so. However some attempts at providing experimental counterfactuals of an 'informed public' provide unique opportunities to study how, at least in the short term, exogenous variation in motivation, and opportunity impact learning.

'Deliberative Polling' brings together a random sample, provides it with 'carefully balanced' briefing materials, offers it the opportunity (which many accept) to participate in moderated discussion in diverse small groups, and question a 'balanced' panel of experts (and/or politicians), with the aim of eliciting 'considered opinions' (Luskin, Fishkin and Jowell, 2002). In short, Deliberative Polling (exogenously) varies motivation (groups vary in size, diversity, gender composition, etc.), and opportunity (groups vary in knowledge etc.), and certain contextual factors, as they relate to learning.

I exploit this exogenous variation in opportunity and motivation in Deliberative Polling to investigate how changes in these variables impact learning among individuals of varying ability, and other relevant background traits. For estimation, I pool data from sixteen Deliberative Polls (a considerable challenge) and model learning via a (Bayesian) hierarchical model. Analyses show, expectedly, that ability crucially moderates how much people learn. However motivation and opportunity can have a combined effect well in excess of that of ability. And that attitude congenial learning is the norm in self-directed learning, though discussion in diverse groups can stimulate learning of attitudinally uncongenial information.

This research enriches our understanding of acquisition of political knowledge in the social world in at least three important ways. One, it gives us well identified estimates of both social and individual level predictors of learning. Concerns about ecological validity of such inferences (and thus generalizability) are but natural. However, there are reasons to believe that the problems are not grave. A large literature in psychology amply corroborates that even when people are randomly assigned to small groups, they form strong group identities (Billig and Tajfel, 1973). Hence in some crucial ways, learning in small groups provides a close facsimile to the (counterfactual) 'real world' social ecosystem and dynamics within it.

Two, it answers how do we design groups that maximize learning among participants. Until now most of the literature on group design has focused solely on designing groups in a manner that increase the propensity that the group (under variety of voting schemes) makes a 'correct decision' (say juries). Very few, if any, have focused on how to design groups to maximize knowledge gain among participants. This study gives us some guidance about optimal group composition to achieve that purpose. In doing so, it also gives us a richer understanding of small group dynamics.

And three, the research provides a sobering account of some of the well known experimental counterfactuals of 'informed opinion'. Many people are left stranded well short of 'full information', even after being treated in elaborate experiments designed to impart information. In short, we have yet to answer how a 'fully informed' public may look like, and it isn't clear if we can do so with current treatments.

Measurement of Learning

One of the most basic questions one may want to answer is if someone learned a particular item over a course of a treatment. To answer this, we need to know whether the person knew the item before being treated, and whether she knew the item after the experiment. Ask once, ask twice. Rinse and repeat. If treatment conveys information not available outside, you need only ask once (post-treatment). However, even in this simple setup plenty can go wrong. For example, a person may 'guess luckily' in the pre-treatment wave, and learn the item in the treatment, and hence show no observable gain. Or she may observably 'lose' information as a result of lucky guessing in one wave, and unlucky guessing in post-treatment wave (or as a result of becoming misinformed by the treatment).

I propose a method to address these issues using Confirmatory Latent Class Analysis. Assuming respondents don't lose knowledge over the course of the experiment, and that they don't become misinformed, all '10' (correct answer in the pre-test followed by incorrect answer in the post-test) response patterns indicate lucky guessing on the pre test. Using this restriction, we can estimate proportion of respondents who guessed luckily on the post-test (a '01' pattern), and hence estimate number of true transitions from ignorance to knowledge.

Sometimes, one may want to answer a different question - how much do people learn (in

'total') over an informative process? We generally estimate learning via observed knowledge gain - the increase in the proportion of knowledge questions answered correctly. However taking the individual-level observed knowledge gain as a left- or righthand-side variable in analyses seeking to explain learning as a function of other variables or other variables as a function of learning often produces disappointing results. The observed post-event knowledge by itself often does better. We[1] show that this seeming anomaly can be explained by item sampling bias in knowledge batteries, the fact that rich get richer and ceiling effects. And that post-event knowledge can be more heavily correlated with true knowledge gain than observed knowledge gain.

Measurement of Political Knowledge, and Misinformation

The complement of knowledge is not simply ignorance (the absence of relevant cognition). There is also misinformation (the presence of incorrect cognition). But separating misinformation from ignorance (and knowledge) can be both conceptually and operationally challenging. I turn to closed-ended questions from a number of ANES studies, and some GSS studies, to estimate misinformation. Since guessing here is rampant, the trick is to distinguish misinformation from unlucky guessing. To do so, I estimate a trinomial logistic regression model for correct vs. incorrect vs. DK responses as a function of partisan disposition and general political knowledge, then compare the predicted response (the one with the highest predicted probability) with the actual response. I also fit a constrained latent class model to estimate the true proportions of various underlying distributions, and to confirm the results obtained via the predictive model. I find that Republicans are at once more knowledgeable and more misinformed than Democrats.

In ongoing work on this project, that ties measurement of political knowledge and misinformation, I rely upon novel probes and survey experiments to sift misinformation from ignorance. For instance, when people are encouraged to offer a 'Don't Know' response, they are less liable to give the misinformed answer than when prodded to give us their best guess. Similarly items asking people to provide probability assessments of whether each of the option is the correct one, diminishes proportion of respondents choosing the misinformed response.

[1] Jointly with Robert Luskin and Ariel Helfer

'Hidden Knowledge' Versus 'Veiled Ignorance'

Related to the task of clarifying the extent to which people are misinformed, I tackle another nascent controversy about to what degree are people informed. For many years, at least, there seemed to be a broad scholarly consensus that most people know remarkably little about politics. Even "extenuationists" arguing that most people manage to approximate their "full-information" attitudes and votes anyway, took the existence of widespread ignorance as given. Recently, however, a number of studies have suggesting that the public may be less ignorant than has been thought. Some suggest that respondents may be insufficiently motivated to recall political trivia (Prior and Lupia, 2008). Others point to the failure to discourage "don't know" (DK) responses (i.e., to encourage substantive answers) to knowledge items (Mondak 2001). Yet others suggest that verbal questions miss visual knowledge (Lang, 1995; Prior, 2008). I assessed these claims using a large number of knowledge questions in a series of original survey experiments, varying, among other things, the nature of the stimulus (name versus photo), item format (open- versus-closed ended, open-ended followed by closed-ended; closed-ended versus probability scoring), the discouragement versus encouragement of DKs. Analyses reveal little hidden knowledge.

A Model of Information Contagion

People are more likely to share attitudinally congenial information they are exposed to. Some of that information is correct, some incorrect. Depending on the social group, people may also be more likely to share 'scandalous' stories for their implicit social value, than their retractions. For instance, a hoax story that Internet Explorer users have significantly lower IQ than Firefox or Chrome users was shared an astonishing million plus times over Twitter. Its retraction was shared less than a tenth as many times. I test the two hypotheses using data from Twitter over three specific scandals. I estimate ideology of a node as a function of ideology of those s/he follows. A model of contagion is being computed, and data collection is underway for including a wider range of stories.

Affect in Politics

The signal finding in political behavior over time is that most people are clueless about politics. Hence only a few reason in ideological terms. Hence weak ideological polarization among the masses (if true) is not all that unusual. A more plausible scenario is one of 'affective polarization', where people come to hold increasingly negative feelings about opposing partisans. I am currently investigating[2] how changing news media (tone, content), increased possibility to self-select into media based on style and content, and changing nature of political campaigns, are increasing partisan affect among the mass public. I explore this by exploiting a unique survey that carries 'social distance' measures, and show that Americans have become 'affectively' polarized, i.e. they now carry increasingly sharp negative feelings towards the out partisan group. Using exogenous variation in campaign ads that stems from nature of the electoral system ('battleground states' versus non-battleground states, cost of buying ads across different DMAs, etc.), we show that the media campaigns are culpable in this rise in partisan antipathy.

In the second strand of research, I move beyond the extant explanations about how news media affect political behavior - agenda-setting, and priming - and into affect. I find some evidence that news media affect a range of political behavioral intentions via the affect (conditional on information, and a consequence of style) they produce.

[2]Jointly with Shanto Iyengar and Yph Lelkes

References

Billig, M. and H. Tajfel. 1973. "Social categorization and similarity in intergroup behaviour." *European Journal of Social Psychology* 3(1):27–52.

Lang, A. 1995. "Defining audio/video redundancy from a limited-capacity information processing perspective." *Communication Research* 22(1):86.

Luskin, R.C. 1990. "Explaining political sophistication." *Political Behavior* 12(4):331–361.

Luskin, R.C., J.S. Fishkin and R. Jowell. 2002. "Considered opinions: Deliberative polling in Britain." *British Journal of Political Science* 32(03):455–487.

Prior, M. 2008. "Visual Political Knowledge." *Unpublished Manuscript* .

Prior, M. and A. Lupia. 2008. "Money, time, and political knowledge: Distinguishing quick recall and political learning skills." *American Journal of Political Science* 52(1):169.