

The Internet's Consequences for Politics

Henry Farrell (George Washington University)

September 13, 2011

Against Studying the Internet

How should political scientists study the Internet's influence on politics? Political science can surely help improve current public arguments about the Internet. Major public debates on the relationship between the Internet and politics now center around a few very general questions. Does the Internet exacerbate political polarization? Does the Internet empower ordinary citizens vis-a-vis political elites? Can the Internet help activists to topple dictators? The broad literature on these questions has tended, until quite recently, to polemics rather than substance. It has gathered information not to test hypotheses, but either to provide rhetorical support for grand, sweeping arguments or to bludgeon rival arguments into non-existence. While some participants in these broad debates have recently started to move towards more nuanced positions, there are remarkably few connections between these broader public arguments and the relevant literatures within political science. There is increasing need for the kinds of careful argument, hypothesis testing and elaboration of specific mechanisms that social science can provide.

However, political science is currently ill-equipped to answer these demands. A small group of political scientists work primarily on the Internet and related topics. However, this group has had surprising difficulty in connecting to a broader audience within the discipline, let alone more generally. The plausible reason is disciplinary conservatism - many social scientists, especially senior ones, still find it difficult to believe that the Internet is a matter for serious scholarly investigation. It is true that a few scholars working on more 'mainstream' topics have looked at this or that aspect of the Internet. For example, Milner [2006] has studied the relationship between democratization and Internet access, building on a previous statistical literature by sociologists [Hargittai, 1999] and others. Nickerson [2007] has examined whether email solicitations can improve turnout (he finds no evidence to support the hypothesis) . These findings, while interesting in their own right, remain largely disconnected from each other. However, the political science literature on the Internet does not cumulate in any very satisfying way. Hence, it is poorly suited to answering the larger and more interesting questions.

In short, while there is an urgent need for the kind of arguments that political scientists (and other social scientists) can provide, few political scientists are in a position to meet it. This may be starting to change. Some scholars are beginning to think seriously about how to disaggregate the Internet into a set of distinct (but overlapping) causal factors, and explore how these factors

differ from, or are similar to, those present in other forms of political communication. Hence, they are beginning to uncover specific *mechanisms* through which the Internet may affect politics, and to explore the form and consequences of these mechanisms using both qualitative and quantitative data. They are thus building on a previously under-exploited body of research, and extending it in important directions.

Thinking about the Internet in this way has some important implications. First - and most significantly - it suggests that one should not study the Internet *as such*. If the Internet operates through a bundle of different mechanisms with different implications, then one should obviously study these mechanisms (and, to the greatest extent possible, their interactions), rather than positing that the Internet is a single, undifferentiated causal force. Second, it provides the basis for actual comparisons between the Internet and other communication technologies. For example, one possible interpretation of the Internet's role in the Arab Spring is that it made communication among protesters and possible protesters faster, while giving a weaker boost, or no boost at all, to communications among state security forces. If this interpretation is correct, *ceteris paribus* one might expect to see similar consequences flowing from other communication technologies that gave potential dissenters an advantage of speed, such as e.g. the printing press in an earlier era. In short, by disaggregating the Internet, one can understand its consequences in comparative perspective, rather than thinking about it as a unique case. Third, it more easily accommodates interaction effects between the Internet and other forms of communication. It is effectively impossible to disentangle the Internet's consequences for real life politics from that of other communications technologies. Internet users are typically exposed to other forms of communication too, while politically relevant information may hop promiscuously back and forth between different online and offline media.

Adapting Henry Teune and Adam Przeworski's three decade-old prescription, we should stop thinking about the 'Internet' as a proper name, and instead start thinking of it as a bundle of mechanisms that we can in principle disentangle from each other. Doing this will not only allow us to compare the Internet with other communications technologies, but also to think systematically about how politics is likely to change as the Internet becomes ever more assimilated into everyday political activity. It will hence allow us to take advantage of a plethora of new data. While many significant sources of Internet data are not publicly available, what there is may potentially provide insight into important questions that political scientists and other social scientists have struggled with. If personal interactions on the Internet can be systematically captured in a way that offline transactions cannot, then we can for the first time begin to observe or reasonably infer e.g. informal communication flows, the dissemination of ideas (or, at least, somewhat crude proxies for same) across different social groups, and the actual network structures underlying communication. Doing this systematically will require political scientists to develop new tools and competences, and acknowledge vexing methodological problems. It will also require them to confront important ethical issues. Even so, it presents major opportunities.

The next section of this essay provides a brief history of debates over the Internet and politics.

Then, the article discusses three families of mechanisms through which the Internet may plausibly affect politics - *sorting* (the processes through which individuals come to identify and preferentially interact with those similar to themselves), *affecting the transaction costs of collective action*, and affecting *preference revelation* (the ease of discerning people's underlying preferences). This list of mechanisms is obviously illustrative rather than exhaustive - there are many other ways in which the Internet could plausibly affect politics. Next, the article asks, in two successive sections, how these mechanisms may help elucidate two important debates among political scientists. The first of these is a long-standing argument about whether the Internet does, or does not, exacerbate political polarization in the US. The second is an emerging argument about the reasons for the "Arab Spring" uprisings in Egypt, Tunisia and elsewhere. The concluding section speaks to the relationship between mechanisms and data, and lays out an agenda for future research.

Changing Debates Over the Internet and Politics

Debate over the political consequences of the Internet began not among political scientists, but among activists, politicians and law professors. As the Internet began to change during the 1990s from an obscure network connecting universities and research institutions, it became embroiled in political controversy. As the Internet snowballed, politicians were not sure whether to take credit for it or to deplore its potential excesses. What they were quite certain about was that it needed to be regulated.

Hence, early debates about the political consequences of the Internet saw it, for good or bad, as a threat to established political rules and norms. Legislators in the US worried that it would unleash a flood of child pornography, and introduced and passed the Communications Decency Act, which aimed to outlaw harmful content, and provide hosting companies with appropriate incentives to regulate areas that the state could not reach. This led to outrage among libertarian Internet activists, including John Perry Barlow, who authored a "Declaration of the Independence of Cyberspace" which declared that governments could not regulate the Internet, no matter how they tried, and that it would sweep away traditional forms of state tyranny Barlow [1996]. Many Internet activists believed that the Internet was inherently resistant to the state, because it had no obvious controlling nexus - according to one much-quoted adage, the Internet "interprets censorship as damage, and routes around it."

Related arguments took place over the Internet's international consequences. David Johnson and David Post claimed that traditional law was inappropriate to govern the Internet, since the Internet was not bounded by the physical borders of the system of sovereign states [Johnson and Post, 1995]. Peter Swire and Michael Froomkin argued that the Internet provided private actors with new opportunities vis-a-vis states Froomkin [1997], Swire [1998]. Other legal academics, such as Jack Goldsmith argued that the Internet presented no challenges to state sovereignty that could not be resolved through traditional legal principles such as conflict of law Goldsmith [1998], while Shanthi Kalathil and Taylor Boas claimed that the Internet helped authoritarian regimes rather than harming them [Kalathil and Boas, 2003].

Finally, some writers began to explore broader questions about the Internet's social consequences. Nicholas Negroponte [Negroponte, 1995] predicted that the Internet would lead to a world with much less social cohesion, as individuals stopped consuming mass-produced information from newspapers and televisions, instead reading personalized information sources (what he called the "Daily Me.") Discounting for Negroponte's somewhat naive techno-utopianism, this claim had clear political implications. In particular, it suggested that the Internet might lead to greater political polarization and extremism [DiMaggio et al., 2001].

Political science played a remarkably modest role in these early arguments. Only a few political scientists were interested in the topic. Bruce Bimber argued in a succession of articles (e.g. [Bimber, 1998]) against utopian claims that the Internet was about to precipitate a profound political transformation, suggesting that it would instead make the formation of political groups easier than it had been in the past. Lance Bennett saw the Internet as perhaps helping people to organize in 'uncivic cultures,' while Robert Putnam suggested (more or less as an aside) that it might further exacerbate the decline of American civil society. However, these political scientists found themselves having to engage with communications scholars, sympathetic sociologists (who themselves found it difficult to get traction in their field) and public intellectuals. Their departmental colleagues were resolutely uninterested.

The slack was taken up by legal academics. Larry Lessig's 1996 book, *Code and Other Laws of Cyberspace* [Lessig, 1999] was intended as a corrective to libertarian accounts of the Internet, most prominently that of Post and Johnson (also law professors). Lessig, a constitutional scholar, argued that computer code played an analogous role on the Internet to law in offline life - it provided a set of rules that shaped individual behavior. Hence, he suggested, major features of code should be subject to the same kinds of collective and democratically mandated decision making as were major laws. More generally, the Internet - with its plethora of different forms of social organization - enticed the attention of many legal academics, some of whom were unfortunately less interested in looking for general laws, than in finding convenient cases to showcase their various theoretical hobby-horses. This led to the brief flowering of a sub-discipline of 'Cyberlaw,' followed by a nearly equally rapid decline as the technology boom of the early 2000's evaporated.

A second wave of literature began in the mid-2000s. As US politicians began to move beyond simple campaign websites and the like [Gibson and Ward, 2000, Bimber and Davis, 2003], and use the Internet actively to raise money and organize supporters, Americanist political scientists began to pay more attention to it. The surprise insurgent campaign of Howard Dean, which came close to upsetting the race for the Democratic presidential nomination, highlighted the importance of social media. Dean supporters relied extensively on social tools such as MeetUp in order to organize themselves. Dean's early success suggested that the transaction costs of political activity had been radically lowered [Schmitt, 2003], while his campaign's rapid collapse suggested that the new technologies had their limits too [Johnson, 2005]. Matthew Hindman argued that Internet fundraising was crucial to Dean's success, and that his campaign also illustrated an important disparity between liberals and conservatives - the former were far readier to organize online than

the latter [Hindman, 2005]. Dean’s campaign proved a harbinger for both the 2004 presidential campaign (where John Kerry’s success in Internet fundraising helped counter George W. Bush’s greater support among more traditional donors), and 2008, where Barack Obama successfully married traditional campaigning structures with a highly successful online organizing and fundraising operation [Wilcox, 2008].

Questions about new models of fundraising and organization went hand in hand with inquiry into citizen engagement. The work of Robert Putnam [Putnam, 2001] and Theda Skocpol [Skocpol, 2004] raised important questions about the decline of civic engagement in American political life, even if their diagnoses disagreed in their particulars. Putnam worried about the decline of traditional non-political civil society organizations which had spanned social divisions, while Skocpol called for a resurgence of vigorous contention among organizations that were capable of connecting ordinary citizens to greater political battles. The Internet fed directly into these debates. Pippa Norris and her colleagues had previously worried that a “digital divide” might develop between those who had sufficient skills to master these new technologies and to use them to participate in politics and those who did not (Norris [2001]; see more recently, Schlozman et al. [2010]). This debate spilled out into wider arguments about economic development, and the question of whether countries in the developing world would fall behind because of poor infrastructure, or create their own models of knowledge-based growth [Boas et al., 2005], but also took on a new significance in the US and other advanced industrialized democracies as social media began to take off. Blogs and related forms of communication (Twitter, YouTube, to a lesser extent Facebook), provided individuals with the potential opportunity to communicate with large numbers of other people, at little or no cost.

Some (especially bloggers themselves) argued that this dramatic increase in ease of access would lead to the unseating of traditional elites and the democratization of public debate. Others [Shirky, 2003, Farrell and Drezner, 2008] pointed out that these forms of media carried their own inherent inequalities - the distribution of blogs’ hyperlinks and readership numbers was heavily skewed, so that a small number of elite blogs received the lion’s share of attention, and the vast majority of blogs had very few readers. The distribution of Twitter authors’ followers and YouTube video views is similarly skewed. As Matthew Hindman [2010] has pointed out, the more successful political bloggers also tended overwhelmingly to be white, male, and with advanced degrees from well-regarded academic institutions. As time has passed, the larger and more successful blogs have increasingly either become absorbed within mainstream media outlets (which are expanding onto the WWW), or have sought to become commercial competitors to these media outlets. Even so, community blogs, in which large numbers of likeminded people interact, remain an attractive model for some kinds of democratic theory [Benkler and Shaw, 2010]

As the Internet’s novelty has diminished, broad debates over the general implications of the Internet for democracy have given rise to a set of more specific inquiries into particular ways in which the Internet might shape politics. It is increasingly possible to break down broad questions (such as the consequences of the Internet for citizenship) into specific lines of inquiry regarding e.g. effects on political knowledge, on political participation, and exposure to different opinions [Bruce,

forthcoming]. A thriving comparative literature is starting to examine how the Internet shapes democratic politics in a variety of national contexts [Chadwick, 2006, Vaccari, 2008, Gibson and McAllister, 2011, Gibson et al., 2003, Tkach-Kawasaki, 2003]. Easier access to good data is furthermore transforming the field. Thus, for example, Hindman [2010] is able to use readership data in order better to gauge the political visibility of different online sources. Gentzkow and Shapiro [2010] use a massive dataset to measure the exposure of Americans to ideologically diverse online voices, while Lawrence et al. [2010] use extensive survey data to identify patterns of blog readership. While none of this data provides a fully satisfactory understanding of the *causal* relations that are at work, it surely helps eliminate unsubstantiated arguments.

Identifying Plausible Causal Mechanisms

Even though social scientists are making real progress in understanding the Internet, their impact on public debate is limited. Broader arguments about the Internet are still dominated by public intellectuals (some of whom are excellent; some rather less so), legal academics such as Larry Lessig, Jonathan Zittrain and Cass Sunstein, and a smattering of sociologists and communications scholars. Some theories from political science filters through into debates indirectly, through legal scholars such as Sunstein who maintain an interest in empirical research, but much is lost. Thus, for example, US public debates about the unrest in Iran in 2009 were dominated by claims which were belied both by the available data, and by the existing literature on social movements and contentious politics [Aday et al., 2010].

Political scientists can best contribute to these debates by showing how their way of thinking about the world can shed light on major arguments over the relationship between the Internet and politics. In part, this will require changing the questions. For example, the question of whether the Internet helps or hurts democracy has provoked much controversy over the last decade (see e.g. the controversy between Shirky [2009] and Morozov [2011b]). However, it is more or less unanswerable as posed, because it proposes no specific theory as to the connection between the Internet and democracy. Hence, those who believe that the Internet helps democracy can point to one set of plausible relationships, while those who believe that it hurts the Internet can point to another. Moreover, even those who agree on the broad question may differ dramatically on specifics, depending on which more particular theory they adopt. For example, Clay Shirky argues that policy makers who believe that the Internet helps democracy by lowering the costs of collective action, may make a major mistake by pushing to fund democracy activists in non-democratic countries. Their actions may politicize the Internet so that non-democratic regimes seek to block access, undermining the possibility that the Internet will build an active civil society and hence a more sustainable set of supporting institutions for political freedom [Shirky, 2011].

It is here that political science can make an important contribution. Although Shirky does not use academic language, he is clearly arguing that those who adopt a monocausal logic, in which there is a simple relationship between new technologies and political outcomes, may make very serious mistakes. Even more sophisticated observers, who assume relatively invariant causal relationships

may err badly. Instead of positing monocausal relationships, or even more sophisticated predictive explanations, we may be better off if we instead try to focus on the different *mechanisms* that might intervene between forms of communication such as the Internet, and final political outcomes. Here, our initial ambition is not to predict, but to explain. As Jon Elster describes mechanisms, they are:

frequently occurring and easily recognizable causal patterns that are triggered under generally unknown conditions or with indeterminate consequences [Elster, 2007] p.36.

He gives the example of childrens' exposure to an alcoholic environment - some children may respond by becoming alcoholics themselves, others by eschewing alcohol. Both of these involve mechanisms, and while we cannot easily say in advance which mechanism will pertain in a given circumstance, we can *ex post* explain both, quite different reactions as causally linked to the initial condition. The weakness of mechanisms-based explanations is that they are indeterminate. Their advantage is that they can help us to identify patterns which, even if they do not readily allow us to predict outcomes, at least help us to map out the territory of causal relationships, by identifying frequently occurring causal patterns, and distinguishing them from other such patterns [Elster, 2007].

Looking for mechanisms is a useful first step towards finding broader *laws* - causal relationships that are triggered under more or less predictable conditions. If one can discover the circumstances under which one mechanism rather than another is likely to apply, one has made progress towards identifying laws, albeit weak ones. But in order to do this, one must first map out plausible mechanisms, and think, where possible, about the ways in which they might plausibly interact with each other.

To illustrate how this might work, I look at three families of plausible causal mechanisms, which can be teased out of the existing literature. This list is not intended to be exhaustive - there are surely other important mechanisms that could be invoked to explain the relationship between the Internet and politics. Nor, however, is it arbitrary; together, they can help us understand the consequences of the Internet for politics across quite different areas of interest. More specifically, in the two succeeding sections, I examine their relevance to two major debates over the political consequences of the Internet - arguments over whether the Internet may pose a particular threat to autocratic regimes, and arguments over whether the Internet is leading to increased political polarization in the US (and elsewhere).

First - there are mechanisms which link the Internet to political outcomes via their consequences for the costs of *collective action*. If the Internet lowers (or, less likely, increases) the costs of certain kinds of collective action by e.g. making it cheaper to communicate with others, providing the means for decentralized action etc, it may plausibly affect a variety of political outcomes. Actors who now find it cheaper or easier to organize than they had previously may be better able to pursue their goals than previously, especially when those goals conflict with the goals of other actors, who are not so advantaged. Shirky [2009] provides one version of this argument, which explicitly theorizes

the difficulty of collective action in Coasean terms [Coase, 1960]. He argues that the Internet has lowered the costs of collective action. The result is that collective activities which used to require central coordination and hierarchy can now be carried out through much looser forms of coordination. Benkler [2006] identifies this kind of collective coordination as a generalizable form of production, which does not really fit in the traditional dichotomy between market and state. He suggests that ‘social sharing’ will often be more efficient and have lower transaction costs than market-based forms of exchange.

The possible political benefits of lower transaction costs are clear. *Ceteris paribus*, actors who face lower costs to collective action will be more likely to mobilize in pursuit of their interests. However, some writers caution that the lowering of transaction costs may have unexpected long term consequences. Schmitt [2003] extrapolates from Benkler’s work to provide an early (and surprisingly overlooked) analysis of the long term consequences of low transaction cost politics. He suggests that it will lead to transactional politics, in which “low barriers to entry mean low barriers to exit,” and it will be difficult to maintain long term loyalty (Ammori [2005] provides a case study which offers some empirical support for this proposition). Morozov [2011b] claims that actors will be motivated to engage in cheap and ineffective but showy forms of politics, such as joining Facebook groups, rather than more expensive or risky forms of collective action. While he does not provide an explicit theoretical framework as Shirky does, one plausible interpretation of his claim is that the Internet lowers the cost of purely expressive political action more than it does the cost of actual physical protest (where one may still be hurt or killed), hence making it more tempting to engage in decorative but useless kinds of political action. Gladwell [2010] argues that the Internet is far more likely to create ‘weak ties’ than the strong ties that social movement theorists argue are the bedrock of costly political action, and hence will not have the beneficial impact that Shirky claims.

Second, there are mechanisms which link the Internet to political outcomes via *homophily*. The literature on homophily is far better developed in sociology than in political science (see McPherson et al. [2001] for a helpful overview) - the term describes the propensity of individuals who are similar on some meaningful dimension to form clusters with each other. There are a variety of ways in which the Internet makes it more likely that individuals with shared views or preferences will cluster together. Again, Benkler [2006] provides a useful description of one plausible mechanism. Individuals with relatively unusual interests have traditionally found it difficult to identify others with those interests, especially if they live in small communities. The Internet makes it far more likely that individuals with unusual interests will find each other, both because it vastly expands the set of actors with whom one can meaningfully interact, and because it provides effective means of both searching across that set of actors, and signalling to others within that set. This may have negative as well as possible social consequences. Shirky [2009] points to how the Internet has allowed anorexics to find each other, and create an online community that may bolster their distorted self-images. But homophily may also occur more indirectly. Individuals may, for example, converge around a common source of online information that is attractive given their shared interests, and

only cluster together as a secondary consequence of this shared interest.

As both Benkler and Shirky observe, homophily may be an important precursor to the kinds of collective action observed above. However, homophily may also plausibly shape individual perceptions and preferences in ways that have nothing to do with transaction costs. It might, for example reinforce intergroup boundaries. It may have consequences for knowledge, so that actors know more about others who share relevant characteristics, and less about others with different characteristics than they would if their social connections were formed randomly. To the extent that contact between diverse actors is associated with collective problem solving [Page, 2008], societies with high tendencies towards homophily may be poorer at solving complex problems. Equally, repeated affirmation of commonly held beliefs within a homophilous group may lead group members to feel more assured in those beliefs.

Third, the Internet may very plausibly affect the likelihood of *preference falsification* [Kuran, 1997]. As Timur Kuran argues, individuals will have incentives to conceal their true preferences in a wide variety of social situations. However, there are situations where preference falsification is socially endemic. For example, it is risky, even in mildly authoritarian regimes, to state that one would prefer to live in a different political order. Thus, most people in such systems either passively acquiesce to the regime or actively demonstrate their loyalty, even if they secretly loathe the regime. This means that people will lack information about others' true preferences, and may in turn be reluctant to display their own true preferences. Hence, even if a regime is nearly universally loathed, it may still be relatively secure, as long as it is able to control public culture so as to prevent people from expressing their true preferences.

However, preference falsification may also work in less dramatic ways. Mutz [2006] finds that Americans are far less likely to engage in overt political activity when they have politically heterogeneous social networks (i.e. networks that include both Democrats and Republicans) than when they have homogenous social networks. As Mutz notes, one plausible explanation of this finding is that Americans in heterogeneous networks have greater incentive to conceal their actual political preferences in order to avoid social discord, such as when e.g. their left-leaning neighbor discovers that they are in fact strong conservatives, and decides not to invite their child to little Norah's birthday party.

It is notable that these mechanisms may reinforce each other; however, they do so in ways which are partially contingent on circumstances. When people join together in homophilous groupings, this may make collective action more likely. Collective action, in turn, may mean that homophily is more likely to change people's beliefs - when people actively work together, they may become more likely to identify with each other closely. Homophily may also interact with preference falsification. The costs to revealing one's true beliefs are likely to be lower when one associates primarily with others who are likely to share those beliefs. Very obviously, preference falsification may affect the costs of collective action - potential dissidents living in oppressive regimes are less likely to engage in action against the regime when they believe that others support the regime. Under other circumstances though, reducing the degree of preference falsification may *decrease* the likelihood

of collective action. When people believe that their actions will reveal their political preferences, they may be less likely to engage in these actions if they believe that others disapprove.

Thus, disentangling the ways in which these mechanisms work is a crucial first step towards understanding the empirical consequences of the Internet. In the next two sections, I illustrate this, by first examining how these mechanisms may be brought to bear on debates over the Internet and the spread of democracy to non-democratic regimes, and then applying them to arguments over whether the Internet is increasing political polarization in the US.

Is the Internet Helping to Polarize American Politics?

American politics is becoming increasingly polarized between left and right. The causes for this polarization, and the extent to which this is an elite or mass phenomenon are the topic of vigorous debate among political scientists [Fiorina and Abrams, 2008, McCarty et al., 2006]. It is clear that the Internet does not explain polarization, if for no other reason than that polarization substantially precedes the widespread use of the Internet. Nonetheless, there is vigorous public argument about whether or not the Internet may substantially exacerbate polarization.

This debate springs from an earlier argument between Negroponte [1995] and Sunstein [2002] over the consequences of the internet for individual consumption of information. Negroponte claimed that the Internet would allow people to replace broad interest newspapers with a “Daily Me” precisely tailored to their needs. Sunstein responded that this, far from heralding a libertarian utopia, would weaken democracy by reducing contact between people with dissimilar points of view, and increasing political polarization. Here, Sunstein claimed that experiments showed that people who discuss political issues with others who share their views end up adopting more extreme positions than they had before. This could be a result of them being persuaded by other group members, adjusting their beliefs to conform with those of the group, or becoming more certain in their beliefs because of confirmation by others. Sunstein also argued in this, and in later work that the Internet was likely to create a fragmented communication market that would have especially strong polarizing effects. This might lead to cascades of bad information, the spread self-reinforcing political extremism and sundry other horrors.

These arguments - and the counter-arguments of those disagreeing with Sunstein - can usefully be recomposed in terms of the mechanisms described above. Most obviously, Sunstein’s major claims invoke arguments about homophily. Sunstein argues that the Internet tends both to bring like-minded people together, and to make them more like-minded than they previously were. However, they also have implications for collective action. For example, Mutz [2006] argues, there is a trade-off between cross-cutting deliberation and involvement with politics, suggesting that greater polarization may go hand-in-hand with higher political engagement. Finally, the Internet’s consequences for preference falsification may dampen polarization but also make individuals less likely to engage in overt politics.

The Internet makes it easy for individuals to express their political views on the Internet whether through blogging, Facebook, Twitter, or a myriad of other social technologies. Many of these

social technologies are difficult to research - the important data is proprietary. However, there is emphatic evidence that politically engaged content producers tend to cluster together with others who share their political orientation. In a classic study, Adamic and Glance [2005] demonstrate that US political bloggers tend to belong to either a left-leaning or right-leaning cluster, and to link overwhelmingly to those who share their political beliefs rather than those who on the other side. Hargittai et al. [2008] find similar dynamics among a more select group of bloggers. Conover et al. [2011] show similar clustering on Twitter, at least in ‘retweets’ (messages which simply repeat and pass on another message written by another Twitter user).

This supports the hypothesis that the Internet encourages clustering among politically active content producers. However, it says little about the possible *consequences* of clustering. Sunstein’s argument is two-fold - that the Internet encourages like-minded people to find each other, and that when they find each other, they will reinforce each others’ beliefs so that they become inward looking and extreme. While the evidence clearly supports the first claim, it is insufficient to support the second. In part, this reflects fundamental methodological problems - it is impossible to distinguish between simple homophily (individuals clustering together because they are similar) and social influence (individuals becoming more like each other because they are in contact with each other) in any generally satisfactory way using standard data [Shalizi and Thomas, 2011]. In part, this reflects genuine ambiguity in research results. Although Hargittai et al. [2008] find a high degree of political clustering, they also carry out content analysis which suggests that genuine political debate may take place across clusters, and argue that there is no evidence of increasing polarization over time. Conover et al. [2011] find that even if retweets are clustered, ‘mentions’ where one Twitter user is referred to by another in order to attract her attention are not. Moreover, because much Twitter conversation is organized via unregulated hashtags which can be used by anyone, readers of partisan hashtags are exposed “to content from users on both sides of the political spectrum.” (p.94) However, it is likely that much of this exchange is intended to provoke rather than to persuade (arguments that are articulated in 140 characters are likely in any event to do little to persuade the ideologically opposed) [Yardi and boyd, 2010].

Related dynamics characterize the *mass readership* of online political information. Prior [2007] examines the consequences of cable and satellite television, and argues that it leads to a twofold sorting process. First, it sorts between those who are interested in politics and those who are not. Those with little interest in politics might previously have been exposed to some political information (e.g. listening to the news while waiting to hear about baseball scores). Now, they can avoid political information entirely if they want to (by e.g. consuming sports channels on cable). Those who do have an interest in politics, in contrast, can now consume far greater amounts of information. However, they tend to be overwhelmingly partisan. The result - for those who dislike partisanship - is a world in which the best lack all conviction and the worst are full of passionate intensity. Political independents - who tend to be apathetic - stop consuming political information, so that the tastes of partisans increasingly come to dominate the consumption of political news.

Lawrence et al. [2010] find that similar patterns reproduce themselves among consumers of

political blogs. Those who read these blogs tend to be far more interested in politics than those who do not. They also tend to be far more partisan, and far more ideologically coherent in their views. Indeed, consumers of political blogs are roughly as polarized (in terms of their attitudes to a set of hot-button political issues) as US Senators, and far more polarized than e.g. viewers of Fox News or NBC. Finally, readers of political blogs tend to be ideologically one-sided in their reading habits - readers of left wing blogs typically do not read right wing blogs, and vice versa. However, their survey evidence tells us nothing about the direction of causation. While it shows a clear correlation between the consumption of partisan information and intensely held ideological and political preferences, it cannot tell us whether this is because individuals who were already intensely ideological and partisan are more likely to search out skewed information sources, or whether skewed information sources make people more ideological than they otherwise would be. This reflects broader problems in distinguishing between media choice and media effects [Bennett and Iyengar, 2008, 2010].

Even so, it would be a mistake to take readers of political blogs as a good proxy for the general population. They are almost certainly a highly unrepresentative and self-selecting minority. At the micro level, Wojcieszak and Mutz [2009] examine online discussion forums, and find that individuals are more likely to encounter dissenting political views on in non-political online forums than in explicitly political ones. Munson and Resnick [2011] find that non-political blogs account for roughly 25% of political postings, and that these posts received slightly more comments than non-political posts at the same blogs. At the macro-level, Gentzkow and Shapiro [2010] examine evidence from a much larger sample of the population who read mainstream news sources rather than sites (such as blogs) which account for a relatively small share of online media consumption), and compare people's exposure to opposing political views via the Internet to exposure via traditional news sources, and exposure via personal contacts. They find no evidence to support the Sunstein hypothesis that the Internet is leading to increased polarization among the general population. Instead, both the Internet and traditional media expose people quite frequently to other points of view, and both do so much better than people's offline encounters via personal relationships. Even intense partisans are regularly exposed to other points of view online. Extreme sites are much more frequently observed than extreme users, so that most Internet users visit a varied selection of websites. "Visitors to the most conservative sites are typically more likely to visit nytimes.com in the same month than the average Internet user or the average visitor to Yahoo! News."

All these results draw on various kinds of data, and are necessarily fragmentary. There is ample scope for future research. These caveats acknowledged, these results taken together may tentatively suggest that a more diffuse version of Prior's sorting process is taking place. On the one hand there is a broad population, which includes many partisans, but which is exposed to a broad variety of online sources with different ideological perspectives. On the other, there is a more specific, and highly politically aware sub-group, which preferentially seeks out partisan information via blogs and other means, and which is less likely to be exposed to dissenting opinions. While the latter group is far smaller than the former, it is likely to be more involved in politics, and

hence more influential than its size would suggest. Hence, the Sunstein hypothesis is more likely to apply to the small group of highly political aware people who e.g. read blogs, than the wider population described by Gentzkow and Shapiro. However, even there, there is no very good data that might allow one to measure the effects of cross-cutting and in-group exposure within either of these populations. It is entirely possible that cross-cutting exposure will do nothing to change the political opinions e.g. of strong partisans, or that the effects of exposure to like minded people will be limited. More field experiments - which might e.g. change individuals' browsing patterns and investigate whether this results in changes in attitudes - might help us see whether differential exposure affects political opinions. More micro-level work along the lines of Munson and Resnick [2010], which identifies differences between individuals who prefer information that affirms their initial priors, and individuals who prefer some degree of challenging information, would also help clarify whether different mechanisms work on different people.

The degree of cross-cutting discussion has plausible implications for collective action. Most simply, it will be easier to organize collective action among like-minded people with a similar view of politics than among those who disagree. Indeed, 'netroots' bloggers justify their focus on talking with the likeminded by arguing they are not interested in engaging in debate to convince others, but instead in organizing effective political action [Lawrence et al., 2010]. Yet there may also be more subtle effects. One of the most interesting recent research findings is Mutz [2006]'s identification of a trade-off between cross-cutting social contacts and political action. Her finding is straightforward. The more that Americans have cross-cutting social connections (e.g. connections with neighbors or friends who identify with a different political party), the less likely they are to engage in political activity. This suggests a troubling political dilemma. On the one hand, there is ample reason from democratic theory to favor cross-cutting discussion [Chambers, 2003, delli Carpini et al., 2004]. On the other, there is also excellent reason to encourage democratic participation [Macedo, 2005]. Any trade-off between the two poses a vexing problem for democratic theorists [Thompson, 2008].

Empirical evidence on the relationship between cross-cutting exposure and participation is only beginning to emerge. Lawrence et al. [2010] find no evidence that blog readers' exposure to competing views has any consequences for participation. The difference in participation between readers of leftwing blogs and readers of *both* left and rightwing blogs in their sample is statistically insignificant. In contrast, however, readers of right wing blogs participate in politics at a significantly lower rate than either of the above (although they participate more than do non-blog readers). This suggests that other factors than exposure to dissonant points of view explain variation in blog readers' propensity to political action. One plausible explanation is provided by [Benkler and Shaw, 2010], who find that prominent left-wing blogs are more likely than their right wing counterparts to exhort readers to donate money, write letters, or engage in other forms of political activity. It would be interesting both to expand this research to a wider set of websites, and to see whether this has changed over the last few years, as e.g. the Tea Party movement has sought to connect online and offline forms of political action. It would also be interesting to see how or whether this pattern generalizes when applied to a larger and more representative audience than political blog

readers.

Finally, preference falsification has direct relevance both for the phenomena described by Mutz and for polarization more generally. The most plausible interpretation of the effects that Mutz observes is that they stem directly from individuals' willingness to falsify their political preferences in order to lubricate social relations. Kuran [1997] makes it clear that preference falsification can result from pervasive social pressures in democracies just as much as from fear of persecution in oppressive states (although most of his discussion of preference falsification in democracies concerns support for affirmative action, a policy which he evidently detests, and argues is a false preference equilibrium). If one holds minority political opinions within one's social, familial or work community, one may frequently be tempted to dissimulate about one's true preferences, to avoid social friction. More generally, preference falsification is likely to mean that the spectrum of *publicly expressed* and expressible opinion and political preferences is substantially narrower than the spectrum of *privately held* opinion and political preferences. It is those whose political preferences are furthest from the mode who are likely to bear the highest cost for revealing their true preferences.

Preference revelation may cut in different directions. On the one hand, the Internet may plausibly lower the costs of expressing one's true preferences under many circumstances. The spectrum of opinion among people on the Internet is obviously likely to be wider on the Internet than in most geographic communities, making it easier for individuals with minority preferences to find others who either share, or are tolerant of their own true beliefs. They may also find it easier to 'try out' different identities or ways of interfacing with the world, especially when they can do so anonymously. Here, the Internet is a radically expanded (and less sexist) version of the eighteenth coffee-house culture identified by Sennett [1977] in which men could escape from the confines of their domestic identity. Research examining the relationship between minorities' exposure to reinforcing points of view via the Internet or other means, and political efficacy, could have quite interesting results.

On the other, the Internet may also dramatically raise the costs of some unpopular forms of political action. For example, California keeps records on those who make financial contributions to support ballot initiatives. This information was once relatively difficult to find, providing donors with effective privacy. Now that it is machine-readable it can be combined with maps so as to show who contributed to specific causes within your neighborhood. Eightmaps.com provides maps showing the name, approximate location and (where the information is available) address of contributors to California's controversial Proposition Eight anti-gay marriage initiative. Mutz's findings suggest that the ready availability of such information via the Internet will dissuade people from donating to controversial causes in future, even if they do not receive the death threats that some donors to Proposition Eight received. When the transaction costs of finding out information on one's neighbors' political proclivities was high, people were unlikely to seek out this information without good reason. This is no longer true.

Thus, research on polarization and the Internet strongly suggests that the Internet facilitates homophily among at least one important sub-group of the US population - those who are politically

aware and have strong partisan and ideological leanings. It does not, however, indicate either that this leads to a breakdown of relations between those with different partisan identities. It is furthermore extremely difficult (and perhaps close to impossible in a non-experimental setting) to discern whether or not heightened clustering leads to political extremism, or whether, alternatively, those with unusually strong pre-existing political views are more likely to seek each other out. Among the broader population, the evidence seems to suggest that the Internet is associated with more exposure to alternative views than was possible previously, rather than less. Although increased polarization and intra-group contact may lead to higher levels of collective action within each group, other sources of variation (e.g. differences in the willingness of left-leaning and right-leaning online voices to ask their readers directly to get involved) may be more important. Finally, although the work of Mutz gives good initial reason to suspect that preference falsification is a key mediating factor between levels of online polarization and willingness to engage in political action, we need more research (using *inter alia* experimental designs) to determine the extent to which it is important, and (if important) the direction in which it cuts.

Does the Internet Help Spread Democracy?

Participants in public debate frequently claim that the spread of the Internet is necessarily tied to an increase in political freedom and the diffusion of democracy. Such claims are often diffuse and poorly supported, but are far from entirely ridiculous. Not only is the relevant literature emphatic on the need for authoritarian regimes to control communications [Hardin, 1990], but the behavior of these regimes, and of insurgents seeking to defeat them, provides ample evidence of their importance. Television and radio stations are usually heavily supervised and censored in authoritarian regimes, and are often the first target of rebels, since control of information can shape public beliefs and perceptions. However, the Internet turned out to be far more vulnerable to censorship than initially expected. In part, this was because its architecture of links has proven far less decentralized than libertarians believed - decentralized patterns of link formation mean that some nodes are far more important in the network than others, and removing or ‘censoring’ these nodes can degrade communication across the whole network [Albert et al., 2000]. In part, this is because the Internet depends on physical connections, which may be compromised to facilitate surveillance or censorship. Not only is Internet censorship possible - it has now become a “global norm” [Deibert et al., 2010].

More recent debates have focused not on the Internet’s inherent resistance to censorship, but on the positive ways in which it can (or cannot) make democracy more likely to replace autocracy. Statistics point to a correlation between Internet access and democracy. However, this correlation might plausibly be the result of democracies being more likely to allow widespread Internet access, rather than the Internet encouraging democracy [Milner, 2006]. [Howard, 2010] has recently deployed a different body of evidence to argue that the diffusion of the Internet is making democracy more likely across the Arab world. Specifically, he claims that the conjunction of several causal factors in a particular country, including most particularly the growth of an Internet based civil

society, is conducive to democratic transition.

Howard's arguments feed into a more micro-level debate about the ways in which the Internet does, or does not, provide tools that empower social movements. In a recent discussion in *Journal of Democracy*, Diamond [2010] p. 70 argues that even despite censorship, the Internet can allow citizens to "report news, expose wrongdoing, express opinions, mobilize protest, monitor elections, scrutinize government, deepen participation, and expand the horizons of freedom." This may allow dissidents to bring governments down. However, it may also help promote the liberalization of society, and the creation of a pluralistic public sphere, even before democratization occurs. Deibert and Rohozinski [2010] reply by pointing out that not only democracy activists use the Internet to mobilize and organize; so too do criminal enterprises, espionage networks and political and religious extremists. Morozov [2011a] argues that methods of Internet control are moving beyond technological measures to more subtle forms of social engineering aimed at disrupting networks of dissenters.

This discussion has bled over into interpretation of recent events, starting with the 'color revolutions' in the former Soviet bloc and in Lebanon, the 2009 social upheavals in Iran, and most recently and topically, the uprisings of the so-called Arab Spring (still ongoing at the time of writing). The Internet has been repeatedly described as a major precipitating factor in all of these. Good quantitative and qualitative data on these events is still extremely sparse, but suggests that the Internet's role has been more complicated than popular accounts might suggest.

There is evidence that the Internet played a significant role in previous upheavals, such as the 'color revolutions' in the former East bloc [McFaul, 2005] despite limited popular access [Dyczok, 2005]. However, the role of the Internet has likely been exaggerated in many accounts, as it has in later unrest in Moldova [Mungiu-Pippidi and Munteanu, 2009] where e.g. Twitter seems at most to have played an extremely limited role on the ground [Morozov, 2009]. Many public commentators claimed that the social unrest in Iran was a 'Twitter revolution,' despite the nearly complete lack of any evidence that Twitter was used by demonstrating Iranians to communicate among themselves, let alone to organize a revolution [Aday et al., 2010]. Most recent and most urgent is the nascent debate over the Internet's role in the Arab Spring - the process of political upheavals across Arab countries that began in Tunisia, led to the falls of the Egyptian and Libyan governments, and is still reverberating in Syria. No synthetic accounts have yet been published comparing the countries that saw major changes (Tunisia, Egypt, Libya perhaps Yemen), those that saw moderate political concessions to protesters (e.g. Jordan), those that saw significant but unsuccessful protests (e.g. Algeria, Bahrain), and those that saw no major protests at all (e.g. Saudi Arabia). Howard and Hussain [2011] and Lynch [2011] provide somewhat different accounts of change in Tunisia and Egypt.

Howard and Hussain [2011] finds it no surprise that Tunisia (with a very large population of Facebook users), and Egypt (which has more Internet users than any other country in the region, except Iran) saw greater civil disobedience and resulting pressure towards change. In his argument, popular frustration had not been translated into action until cellphones and the Internet became

available, helping to unite disparate grievances in a common agenda, and turning localized patches of discontent into a “structured movement with a collective consciousness about both shared rights and opportunities for action.” Social media were not the democracy movement themselves, but they provided the movement with necessary scaffolding and a means of organizing outside the control of the state. Lynch [2011], while agreeing with some of these arguments, is much more circumspect in his interpretation of the evidence. He suggests that “[w]hile protestors effectively used social media in their struggles, it is surprisingly difficult to demonstrate rigorously that these new media directly caused any of the outcomes with which they have been associated.” There is evidence strongly suggesting that new media empowered activists, but it is largely correlational - there is no smoking gun to demonstrate causation. Facebook and other organizing tools played a significant role - but only together with more traditional means of organizing, and more traditional media such as al-Jazeera. In short, Lynch argues that social media probably mattered, but only in conjunction with other media and organizing tools.

We need better data to resolve controversies over whether the Internet is essential or inessential to explaining the political changes in Egypt and Tunisia. Online data on its own is insufficient to measure empirical consequences. Both boyd [forthcoming] and Freelon [2011] show that there was substantial discussion of the Arab Spring protests on Twitter. However, both are careful to stick to their data and discuss their findings in terms of the dynamics of the global conversation around these events rather than seeking to demonstrate that online conversations in any way shaped events on the ground. Good data on the timing and place of protests as well as the timing and point of origin of online discussions might help clarify whether such a relationship exists. But even in the absence of such data, a closer focus on mechanisms can help us to clarify both the research questions we should ask, and the outcomes that may plausibly be associated with those mechanisms.

Access to Internet-based tools substantially lowered the costs of collective action for protestors in both Tunisia and Egypt. The *Kefaya* movement, which began in the early 2000s, used the Internet to organize itself in a semi-clandestine way before it began to publicize its cause Lynch [2011]. In a regime where unsanctioned political organizations were either co-opted or crushed, sometimes quite brutally, the Internet made it far cheaper for these individuals to coordinate collective action, and to attract others to demonstrations in spite of a harsh government censorship regime (at one point the government confiscated all copies of an independent newspaper carrying an advertisement for a *Kefaya* sponsored demonstration [Oweidat et al., 2008]). As the movement began to use Facebook groups more and more, it found it easier and cheaper to organize activities [Faris, 2008]. For sure, the Egyptian government adapted to the strategies of protesters, responding far more effectively to Facebook organized protests in 2007 than in 2006 [Lynch, 2011]. However, according to one one journalistic account of the protests of 2011, key Egyptian dissidents used Facebook not only to organize protests, but to direct authorities’ attentions to some protest marches, while communicating via more surreptitious channels about another demonstration, which turned out to be crucial [Levinson and Coker, 2011].

Even if the Internet lowers the costs of collective action, it may do so in ways that have un-

expected consequences. Both Lynch [2011] and Howard and Hussain [2011] caution that even if the protests were sufficient to destabilize the old regimes in Tunisia and Egypt, they may not be sufficient to build self-sustaining democracies in their places. Lynch is by far the more emphatic - he suggests that authoritarian Arab states may be able to adapt to these new challenges, as they have to previous ones. One plausible reason for pessimism is that, as both Johnson [2005] and Schmitt [2003] suggest, the very ease of Internet-based collective action may make it more difficult to build the more durable structures that can achieve long-term political goals. The failure of the Tahrir Square protesters to create a viable political party or movement is strongly suggestive - while they helped bring down a government, they have been unable to organize in the ways that would allow them to appeal to a broader political constituency. In contrast, the Muslim Brotherhood, which had a much stronger traditional organizational base appears likely to play a key role in any new political dispensation. More generally, Faris and Etling [2008] worry that Internet based forms of social organization may be strong enough to displace existing regimes, but structurally incapable of building viable replacements, resulting in the “formation of “thin” democracies with weak horizontal structures” (p.82) Citizen mobilization on its own is insufficient to create stable political systems.

Lynch [2011] argues that preference falsification played a significant role in the Tunisian uprising, but was far less important to recent events in Egypt. Before the outbreak of civil unrest, Tunisia was a highly cloistered society, in which individuals did not talk openly about politics for fear that the apparently ubiquitous secret police would hear about it. The result was a general dearth of information on people’s true political preferences. While the regime was likely highly unpopular, no-one in the general public had any good sense of quite how unpopular it was. When people did realize that a majority of their fellow citizens shared their dislike of the regime, this generated an informational cascade which led to the regime’s rapid demise. In Egypt, in contrast, the previously existing Kefiya movement had already demonstrated the regime’s unpopularity with a broad swathe of the population. This in turn suggested to observers that the Mubarak regime, even if it still had power, lacked popular support.

Even so, a focus on preference falsification helps focus attention on the importance of identity and signals in precipitating (or preventing) unrest. The identity of protesters is key - protesters who appear to be more “representative” of the general population provide more convincing signals of the privately-held preferences of the majority than unrepresentative ones [Lohmann, 1994]. This helps explain why Arab governments facing mass demonstrations were at pains to claim that the demonstrators constituted an unrepresentative minority and that the demonstrations were being fomented by outsiders. The indirect availability of evidence to the contrary (e.g. cellphone video of protests uploaded and then rebroadcast by satellite television stations) plausibly helped damage the government’s credibility. The Egyptian government’s decision to shut down the Internet likely backfired, since it suggested that the Egyptian government privately believed that the demonstrations were sufficiently representative of public opinion to present a real danger to the regime. Equally, protesters took actions that were intended to suggest that they represented the views of a major-

ity - Marc Lynch describes how they deliberately recruited “non-activists into the early protests, which sent a signal of a widespread societal consensus.” Here, social networking technologies had significant disadvantages - the organizers knew that the “demonstrations’ success would depend on the participation of ordinary Egyptians in working-class districts . . . where the Internet and Facebook aren’t as widely used,” and leafletted these areas to ensure wider participation [Levinson and Coker, 2011].

Hence, government, protesters and others seek to use signals to shape the beliefs of individual members of the public about what others believe. Some forms of signalling - such as e.g. the 2009 campaign to persuade Facebook users to have green bordered pages in solidarity with the protesters in Iran - have negligible political impact. They are costless, do not come from typical Iranians, and are unlikely to be seen by them. More costly signals are more likely to be effective [Kricheli et al., 2010]. Thus, the Internet’s most important effect on preference falsification may not be to send purely online signals, but to document the costly offline signals sent by others (e.g. by disseminating information about protests in which individuals take the risk of being hurt or killed), or, when it is suddenly censored as in Egypt, to allow individuals to make inferences about the private information held by the regime about its popularity. This dynamic may change, as regimes begin more systematically to investigate individuals’ online social networks, on the one hand making online signals of dissent more credible (because they are costly), and on the other making individuals less willing to pay them (again, because they are costly).

The evidence on the political consequences of homophily is scanty. However, it is sufficient to explode one frequent misconception among Western observers - that Arab online discourse is a hotbed of extremist rhetoric, which serves to recruit young Muslims to active or passive sympathy with terrorism. While there are extremist websites, the importance and visibility of these sites is grossly overestimated by many observers, and they plausibly alienate far more Muslims than they attract - the publicization of beheading videos by smaller groups have made it more difficult for al Qaeda to broaden its appeal to mainstream Muslims [Lynch, 2006]. Substantially fewer contributions to major online discussion boards supported the September 11 attacks than opposed them in their immediate aftermath [Abdulla, 2007]. The most comprehensive survey of the Arabic language blogosphere finds that there is “very little support for terrorism or violent jihad in the Arabic blogosphere and quite a lot of criticism.” [Etling, 2009].

However, there is insufficient evidence to decide whether homophily may play a more subtle role. What circumstantial evidence there is suggests a greater tendency towards heterophily among regime opponents before the regime is displaced - shared dislike of one’s government may prove a powerful solvent to many social barriers. In Iran, both critics and supporters of the regime took to the blogosphere in large numbers [Aday et al., 2010]. However, not only did they frequently link to each other, but the cleavage between them did not structure the blogosphere as a whole, which included a variety of partly intersecting clusters, some religious (but largely apolitical), others concerned with poetry and literature and so on (ibid). Egyptian political blogs tended to (a) be strongly critical of the Mubarak regime, and (b) be willing to emphasize their common opposition

to the regime rather than their specific differences (whether they were Islamists, secularists, or human rights advocates). There is evidence that members of the Muslim Brotherhood who blogged were more likely to be pluralist than their brothers who did not [Lynch, 2007] (it is an open question whether this was because bloggers were inherently more open to engagement, or because their views were changed as a result of their engagement with others on the Internet).

These tentative cross-group solidarities appear not to have survived the change in regime. Secularists are beleaguered in the new Egypt, and have found it difficult to build alliances with other groups. This has led them to turn in on themselves, sometimes with highly unfortunate consequences. One prominent Internet activist recently argued on her Facebook page that people “should not be upset” if judicial inaction on police killings led armed groups to carry out assassinations. Others have made their anger with their erstwhile religious collaborators public. The Muslim Brotherhood, for its part, has disassociated itself from secularist protesters, and pursued engagement with other religious groups and (more tacitly) with the interim government. While it has limited its public ambitions, it has made it clear that it wishes to see the reimposition of sharia law. Most of the Muslim Brotherhood bloggers have either left the organization (while remaining political active in other ways) or been expelled from it over the last year. Again, it is impossible to know whether their disaffection springs from innate characteristics (their behavior may reflect long-standing differences between their outlook on life and the modal Muslim Brotherhood member) or the results of exposure (being exposed to other points of view may have led them to become less like other members who were not exposed), or some combination of the two.

Thus, the evidence suffices to show that simplistic arguments about how the Internet affects group identification - whether they are claims that the Internet is radicalizing the Muslim world, or that the Internet has built a sustainable pro-democracy alliance between religious and non-religious parties - do not hold. However, it is insufficient, in the absence of more research, to provide a genuinely compelling account of what *is* happening. One plausible account - which is compatible with the available evidence, and could be properly tested by further research - is the following. In authoritarian regimes - where a variety of different groups are dissatisfied with the *status quo*, the Internet can help bring otherwise disparate groups together. By providing common organizing technologies, without at the same time requiring that groups create a rigid hierarchy with a unified party line, it can allow these groups to work together, and to build up ties with each other. However, should the regime collapse, these fragile ties will collapse too, as groups’ interests begin to clash more sharply, so that individuals have to choose the one or the other faction. Under these circumstances, it is plausible that the polarizing effects of the Internet may come into play, reinforcing already-existing group divisions.

Conclusions

Over the next decade, the relationship between the Internet and politics will become an increasingly important question for the discipline. Paradoxically, it is likely that there will be ever fewer scholars specializing in the Internet and politics (not that there were ever very many to begin

with). However, this will not be because political scientists will lose interest in the Internet and related technologies. Rather, it will be because these technologies have become so integrated into regular political interactions that it will be impossible to study e.g. the politics of fundraising, of election advertising, of political action, of public diplomacy or of social movements *without* paying close attention to Internet mediated forms of action or communication. As the Internet becomes politically normalized, it will be ever less appropriate to study it in isolation, but ever more important to think clearly, and carefully, about its relationship to politics.

Unbundling the Internet into discrete (yet sometimes mutually reinforcing or undermining) mechanisms, will help political scientists to accomplish this task, as the two previous sections demonstrate. It will better integrate political science with broader debates, by breaking down the assumptions behind these debates into more limited and specific arguments, that can be assayed against the empirical evidence. It will facilitate the transition towards a scholarly discussion in which the Internet is both ubiquitous and invisible, because its intermediating role is taken for granted. Finally, it will better allow the comparison of specific aspects of Internet technology (e.g. specific forms of social media), which may affect politics through this or that mechanism, with earlier technologies which affected politics through similar, or through different mechanisms.

This will also enable the discipline as a whole to face up to a new - and possibly transformative - challenge. The movement of political activity to the Internet is generating massive amounts of data, as individuals' political conversations, donations, and forms of political organizing leave electronic traces. Much of this data is difficult to reconstruct. For example, even although it is plausible that chain emails affect people's political beliefs, it is difficult to sample them in any satisfactory fashion, or, given problems of selection bias, to make satisfactory inferences [Liben-Nowell and Kleinberg, 2008] about how they spread [Golub and Jackson, 2010]. Other data is proprietary and difficult to access (e.g. Facebook relationships). Managing this kind of data, let alone analyzing it, creates a plethora of new problems [King, 2011, Stodden, 2010]. Finally, much of this data is sensitive, creating serious new ethical concerns about privacy.

Even so, the early work of computer scientists suggests that this plethora of data generates spectacular opportunities for political scientists and other social scientists to understand hitherto mysterious political relationships [Lazer et al., 2009]. For example, by using discrete fragments of text as markers, it is possible both to examine the dissemination and gradual transformation of ideas across social networks [Leskovec et al., 2009] and to make inferences as to the underlying shape of the network across which they travel [Gomez Rodriguez et al., 2010]. Automated analysis can allow researchers to identify the likely partisanship of authors via both textual analysis [Ahmed and Xing, 2010] and analysis of voting in social recommendation systems [Zhou et al., 2011]. Online experiments can allow scholars to distinguish between people who seek information confirming or contradicting their priors, and to see how these groups respond to treatments that highlight or organize contradictory information in different ways [Munson and Resnick, 2010]. More generally, online forms of social organization as different as Wikipedia, Digg and the Daily Kos can be treated as extraordinarily well-documented, and potentially comparable cases of social information

processing [Shalizi, 2007]. Data like this - as it becomes more broadly used - will make it far easier to get to grips with the actual empirical consequences of the mechanisms underlying collective action, political polarization and preference falsification, and a variety of other important mechanisms too.

Taking full advantage of this data will require a sea-change in social scientists' skill set, so that it encompasses computer programming, network analysis, large scale data analysis, and complex systems theory (or, at a pinch, the ability to work and communicate with other colleagues with skills in computer programming, network analysis, *und so weiter*). Yet it will also require greater theoretical sophistication. Transforming data into knowledge will require us to develop our conceptual vocabulary in dialogue with the data, so that we can both examine the usefulness of mechanisms described in the existing literature, and generate new ones that can be tested against the evidence. While the conversation about 'computational social science' is only in its beginning stages, the conversation about how computational social science and theory generation should work in tandem. If the argument about the Internet's consequences for politics helps structure our current knowledge better, and paves the way towards this broader debate, it will have done all that it should do, and more.

References

- Rasha A. Abdulla. Islam, jihad, and terrorism in post-9/11 arabic discussion boards. *Journal of Computer-Mediated Communication*, 12(3):1063–1081, 2007.
- Lada A. Adamic and Natalie Glance. The political blogosphere and the 2004 us election: Divided they blog. In *Proceedings of the 3rd international workshop on Link discovery*, pages 36–43. ACM, 2005.
- Sean Aday, Henry Farrell, Marc Lynch, John Sides, John Kelly, and Ethan Zuckerman. Blogs and bullets: New media in contentious politics. *US Institute of Peace*, 2010.
- Amr Ahmed and Eric P. Xing. Staying informed: Supervised and semi-supervised multi-view topical analysis of ideological perspective. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, pages 1140–1150. Association for Computational Linguistics, 2010.
- Réka. Albert, Hawoon Jeong, and Albert-László. Barabási. Error and attack tolerance of complex networks. *Nature*, 406(6794):378–382, 2000.
- Marvin Ammori. Shadow government: Private regulation, free speech, and lessons from the sinclair blogstorm. *Michigan Telecommunications and Technology Law Review*, 12:1, 2005.
- John Perry Barlow. A declaration of the independence of cyberspace. *Humanist*, 56:18–19, 1996.
- Yochai Benkler. *The Wealth of Networks: How Social Production Transforms Markets and Freedom*. Yale University Press, 2006.

- Yochai Benkler and Aaron Shaw. A tale of two blogospheres: Discursive practices on the left and right. *Berkman Center Research Publication No. 2010-6, Harvard Public Law Working Paper No. 10-33*, 2010.
- W. Lance Bennett and Shanto Iyengar. A new era of minimal effects? the changing foundations of political communication. *Journal of Communication*, 58(4):707–731, 2008.
- W. Lance Bennett and Shanto Iyengar. The shifting foundations of political communication: Responding to a defense of the media effects paradigm. *Journal of Communication*, 60(1):35–39, 2010.
- Bruce Bimber. The internet and political transformation: Populism, community, and accelerated pluralism. *Polity*, pages 133–160, 1998.
- Bruce A. Bimber and Richard Davis. *Campaigning Online: The Internet in US elections*. Oxford University Press, USA, 2003.
- Taylor Boas, Thad Dunning, and Jennifer Bussell. Will the digital revolution revolutionize development? drawing together the debate. *Studies in Comparative International Development (SCID)*, 40(2):95–110, 2005.
- danah boyd. The revolutions were tweeted: Information flows during the 2011 tunisian and egyptian revolutions. *International Journal of Communication*, forthcoming.
- Bimber Bruce. Digital media and citizenship. *Sage Handbook on Political Communication*, forthcoming.
- Andrew Chadwick. *Internet politics: States, Citizens, and New Communication Technologies*. Oxford University Press, USA, 2006.
- Simone Chambers. Deliberative democratic theory. *Annual Review of Political Science*, 6(1):307–326, 2003.
- Ronald H. Coase. The problem of social cost. *Economic Analysis of the Law*, pages 1–13, 1960.
- MD Conover, J. Ratkiewicz, M. Francisco, B. Gonçalves, A. Flammini, and F. Menczer. Political polarization on twitter. In *Proc. 5th Intl. Conference on Weblogs and Social Media*, 2011.
- R. Deibert, John G. Palfrey, Rafal Rohozinski, and Jonathan Zittrain. *Access Controlled: The Shaping of Power, Rights, and Rule in Cyberspace*. The MIT Press, 2010.
- Ronald Deibert and Rafal Rohozinski. Liberation vs. control: The future of cyberspace. *Journal of Democracy*, 21(4):43–57, 2010.
- Michael X. delli Carpini, Fay L. Cook, and Lawrence R. Jacobs. Public deliberation, discursive participation, and citizen engagement: A review of the empirical literature. *Annual Review of Political Science*, 7:315–344, 2004.

- Larry Diamond. Liberation technology. *Journal of Democracy*, 21(3):69–83, 2010.
- Paul DiMaggio, Eszter Hargittai, W. Russell Neuman, and John P. Robinson. Social implications of the internet. *Annual review of sociology*, pages 307–336, 2001.
- M. Dyczok. Breaking through the information blockade: Election and revolution in ukraine 2004. *Canadian Slavonic Papers*, 47(3-4), 2005.
- Jon Elster. *Explaining Social Behavior: More Nuts and Bolts for the Social Sciences*. Cambridge Univ Pr, 2007.
- Bruce Etling. *Mapping the Arabic Blogosphere: Politics, Culture, and Dissent*. Internet & Democracy Project, Berkman Center for Internet & Society, 2009.
- Robert Faris. Revolutions without revolutionaries? network theory, facebook, and the egyptian blogosphere. *Arab Media and Society*, 2008.
- Robert Faris and Bruce Etling. Madison and the smart mob: The promise and limitations of the internet for democracy. *Fletcher Forum for World Affairs*, 32:25, 2008.
- Henry Farrell and Daniel Drezner. The power and politics of blogs. *Public choice*, 134(1-2):15–30, 2008.
- Morris Fiorina and Samuel Abrams. Political polarization in the american public. *Annual Review of Political Science*, 11:563–588, 2008.
- Deen Freelon. The mena protests on twitter: Some empirical data. *dfreelon.org*, 2011.
- A. Michael Froomkin. The internet as a source of regulatory arbitrage. *Borders in cyberspace: information policy and the global information infrastructure*, pages 1462–543, 1997.
- Matthew Gentzkow and Jesse M. Shapiro. Ideological segregation online and offline, 2010.
- Rachel Gibson and Ian McAllister. Do online election campaigns win votes? the 2007 australian youtube election. *Political Communication*, 28(2):227–244, 2011.
- Rachel Gibson and Stephen Ward. A proposed methodology for studying the function and effectiveness of party and candidate web sites. *Social Science Computer Review*, 18(3):301, 2000.
- Rachel Gibson, Andrea Rommele, and Stephen Ward. German parties and internet campaigning in the 2002 federal election. *German Politics*, 12(1):79–108, 2003.
- Malcolm Gladwell. Small change: Why the revolution will not be tweeted. *The New Yorker*, 4: 42–49, 2010.
- Jack L. Goldsmith. Against cyberanarchy. *The University of Chicago Law Review*, pages 1199–1250, 1998.

- Benjamin Golub and Matthew O. Jackson. Using selection bias to explain the observed structure of internet diffusions. *Proceedings of the National Academy of Sciences*, 107(24):10833, 2010.
- Manuel Gomez Rodriguez, J. Leskovec, and Andreas Krause. Inferring networks of diffusion and influence. In *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1019–1028. ACM, 2010.
- Russell Hardin. The social evolution of cooperation. *The Limits of Rationality*, page 358, 1990.
- Eszter Hargittai. Weaving the western web: Explaining differences in internet connectivity among oecd countries. *Telecommunications Policy*, 23(10-11):701–718, 1999.
- Eszter Hargittai, Jason Gallo, and Matthew Kane. Cross-ideological discussions among conservative and liberal bloggers. *Public Choice*, 134(1):67–86, 2008.
- Matthew Hindman. The real lessons of howard dean: Reflections on the first digital campaign. *Perspectives on Politics*, 3(01):121–128, 2005.
- Matthew Hindman. *The Myth of Digital Democracy*. Princeton University Press, 2010.
- Philip Howard and Muzammil M. Hussain. The upheavals in egypt and tunisia: The role of digital media. *Journal of Democracy*, 22(3):35–48, 2011.
- Philip N. Howard. *The digital origins of dictatorship and democracy: information technology and political Islam*. Oxford University Press, USA, 2010.
- David R. Johnson and David Post. Law and borders—the rise of law in cyberspace. *Stanford Law Review*, 48:1367, 1995.
- Steven Berlin Johnson. *Two Ways to Emerge: And How to Tell the Difference Between Them*. <http://extremedemocracy.com/chapters/Chapter2005>.
- Shanthi. Kalathil and Taylor C. Boas. *Open Networks, Closed Regimes: The Impact of the Internet on Authoritarian Rule*. Carnegie Endowment, 2003.
- Gary King. Ensuring the data-rich future of the social sciences. *Science*, 331(6018):719, 2011.
- Ruth Kricheli, Yair Livne, and Beatriz Magaloni. Taking to the streets: Theory and evidence on protests under authoritarianism. *Social Science Research Network Working Paper Series*, 2010.
- Timur Kuran. *Private Truths, Public Lies: The Social Consequences of Preference Falsification*. Harvard University Press, 1997.
- Eric Lawrence, John Sides, and Henry Farrell. Self-segregation or deliberation? blog readership, participation, and polarization in american politics. *Perspectives on Politics*, 8(01):141–157, 2010.

- D. Lazer, A.S. Pentland, L. Adamic, S. Aral, A.L. Barabasi, D. Brewer, N. Christakis, N. Contractor, J. Fowler, M. Gutmann, T. Jebara, G. King, M. Macy, D. Roy, and M. Van Alstyne. Life in the network: the coming age of computational social science. *Science (New York, NY)*, 323(5915):721, 2009.
- Jure Leskovec, Lars Backstrom, and Jon Kleinberg. Meme-tracking and the dynamics of the news cycle. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 497–506. Citeseer, 2009.
- Larry Lessig. *Code and Other Laws of Cyberspace*. Basic Books, 1999.
- Charles Levinson and Margaret Coker. The secret rally that sparked an uprising. *The Wall Street Journal*, 2011.
- David Liben-Nowell and Jon Kleinberg. Tracing information flow on a global scale using internet chain-letter data. *Proceedings of the National Academy of Sciences*, 105(12):4633, 2008.
- Susanne Lohmann. The dynamics of informational cascades: The monday demonstrations in leipzig, east germany, 1989-91. *World Politics*, pages 42–101, 1994.
- Marc Lynch. Al-qaedas constructivist turn. *Praeger Security International*, 2006.
- Marc Lynch. Young brothers in cyberspace. *Middle East Report*, 37(245):26, 2007.
- Marc Lynch. After egypt: The limits and promise of online challenges to the authoritarian arab state. *Perspectives on Politics*, 9(2):301–310, 2011.
- Stephen Macedo. *Democracy at Risk: How Political Choices Undermine Citizen Participation and What We Can Do About It*. Brookings Institution, 2005.
- Nolan M. McCarty, Keith T. Poole, and Howard Rosenthal. *Polarized America: The Dance of Ideology and Unequal Resources*. MIT Press Cambridge, MA, 2006.
- M. McFaul. Transitions from postcommunism. *Journal of Democracy*, 16(3):5–19, 2005.
- Miller McPherson, Lynn Smith-Lovin, and James M. Cook. Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, pages 415–444, 2001.
- Helen V. Milner. The digital divide. *Comparative Political Studies*, 39(2):176, 2006.
- E. Morozov. Whither internet control? *Journal of Democracy*, 22(2):62–74, 2011a.
- Evgeny Morozov. More analysis of twitter’s role in moldova. *Net Effect*, 2009.
- Evgeny Morozov. *The Net Delusion: The Dark Side of Internet Freedom*. Public Affairs, 2011b.
- A. Mungiu-Pippidi and I. Munteanu. Moldovas “twitter revolution”. *Journal of Democracy*, 20(3):136–142, 2009.

- S.A. Munson and P. Resnick. The prevalence of political discourse in non-political blogs. *ICWSM 2011*, 2011.
- Sean A. Munson and Paul Resnick. Presenting diverse political opinions: how and how much. In *Proceedings of the 28th International Conference on Human Factors in Computing Systems*, pages 1457–1466. ACM, 2010.
- Diana C. Mutz. *Hearing the Other Side: Deliberative Versus Participatory Democracy*. Cambridge University Press, 2006.
- Nicholas Negroponte. *Being Digital*. Alfred A. Knopf, 1995.
- David W. Nickerson. Does email boost turnout? *Quarterly Journal of Political Science*, 2(4): 369–79, 2007.
- Pippa Norris. Digital divide: Civic engagement, information poverty, and the internet worldwide. *Nova York*, 2001.
- Nadia Oweidat, Cheryl Benard, Dale Stahl, Walid Kildani, Edward O’Connell, and Audra Grant. *The Kefaya movement: a Case Study of a Grassroots Reform Initiative*, volume 778. Rand Corp, 2008.
- Scott E. Page. *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*. Princeton University Press, 2008.
- Markus Prior. *Post-Broadcast Democracy: How Media Choice Increases Inequality in Political Involvement and Polarizes Elections*. Cambridge Univ Pr, 2007.
- Robert D. Putnam. *Bowling alone: The Collapse and Revival of American Community*. Simon and Schuster, 2001.
- Kay L. Schlozman, Sidney Verba, and Henry E. Brady. Weapon of the strong? participatory inequality and the internet. *Perspectives on Politics*, 8(2):487–509, 2010.
- Mark Schmitt. Dean’s penguin: Or, technology and the nature of political interaction. *The Decembrist*, 2003.
- Richard Sennett. *The fall of public man*. CUP Archive, 1977.
- Cosma Shalizi and Andrew Thomas. Homophily and contagion are generically confounded in observational social network studies. *Sociological Methods & Research*, 40(2):211, 2011.
- Cosma Rohilla Shalizi. Social media as windows on the social life of the mind. In URL <http://arxiv.org/abs/0710.4911>. Forthcoming in the proceedings of the 2008 AAAI symposium on Social Information Processing, 2007.
- C. Shirky. The political power of social media. *Foreign Affairs*, 90(1):28–41, 2011.

- Clay Shirky. Power laws, weblogs, and inequality. *Clay Shirky's Writings About the Internet*, 2003.
- Clay Shirky. *Here Comes Everybody: The Power of Organizing Without Organizations*. Penguin Group USA, 2009.
- Theda Skocpol. *Diminished Democracy: From Membership to Management in American Civic Life*, volume 8. University of Oklahoma Press, 2004.
- Victoria Stodden. The scientific method in practice: Reproducibility in the computational sciences. *Social Science Research Network Working Paper Series*, 2010.
- Cass R. Sunstein. *Republic.com*. Princeton Univ Pr, 2002.
- Peter Swire. Of elephants, mice, and privacy: International choice of law and the internet. In *International Law*, volume 32, page 991. HeinOnline, 1998.
- Dennis F. Thompson. Deliberative democratic theory and empirical political science. *Annual Review of Political Science*, 11:497–520, 2008.
- Leslie M. Tkach-Kawasaki. Politics@ japan. *Party Politics*, 9(1):105, 2003.
- Cristian Vaccari. Surfing to the elysée: The internet in the 2007 french elections. *French Politics*, 6(1):1–22, 2008.
- Clyde Wilcox. Internet fundraising in 2008: A new model? In *The Forum*, volume 6, page 6. bepress, 2008.
- Magda Wojcieszak and Diana Mutz. Online groups and political discourse: Do online discussion spaces facilitate exposure to political disagreement? *Journal of Communication*, 59(1):40–56, 2009.
- Sarita Yardi and danah boyd. Dynamic debates: An analysis of group polarization over time on twitter. *Bulletin of Science, Technology & Society*, 30(5):316, 2010.
- Xiaodan. Zhou, Paul Resnick, and Qiaozhu Mei. Classifying the political leaning of news articles and users from user votes. 2011.