Why Bowl Alone When You Can Flashmob the Bowling Alley?: Implications of the Mobile Web for Online-Offline Reputation Systems.[*]

Dave Karpf, Political Science Department, University of Pennsylvania. davekarpf@gmail.com

"Location-sensing wireless organizers, wireless networks, and community supercomputing collectives all have one thing in common: *they enable people to act together in new ways and in situations where collective action was not possible before*."
Howard Rheingold, *Smart Mobs*

"The power of rating and reputation systems was not so much invented on the Net as 'revealed,' because rating and reputation systems are a natural and universal artifact of all forms of human cooperation. A reputation is really the collection of tags that are assigned to an individual or entity to reflect assessments of his or her competence." – John Clippinger, *A Crowd of One*

Howard Rheingold introduced the transformative potential of the mobile internet – "the internet untethered" – in his 2002 book, *Smart Mobs* [14]. Pondering the sight of high school students furiously texting back and forth as they wandered through Tokyo's Shibuya station, Rheingold speculated that the ubiquitous presence of these communications devices would have a host of transformative implications for individual habits, community-level interactions, and society as a whole. Removing the stationary computer access terminals necessary to reach the World Wide Web dissolves the distinction between "online" and "offline," and this in turn allows many of the lowered transaction costs associated with internet tools to be applied to areas of daily life that had previously been minimally impacted by the new communications regime.

It has been seven years since *Smart Mobs* was published. To put that in perspective, consider that the first web browser, Mosaic, was released in 1993. The elapsed time between Mosaic and Rheingold's research is equivalent to the elapsed time between the publication and today. In 2001, The Friendster social networking site had yet to launch, and Facebook's founder Mark Zuckerberg was still a high school student. Blogging was such a new platform that Rheingold had to refer to it as "weblogging" [14] and the political "netroots" were a tiny collection of anonymous, powerless partisan wordsmiths. Tagging systems, pioneered in 2005 by del.icio.us, were still years off in the distance. Most importantly, the mobile web, as accessible through the Apple's iPhone and Google's Android operating system, was still far beyond the horizon (the iPod was first released in October 2001). Rheingold's work stands the test of time remarkably well, but it is nonetheless well worth revisiting how the mobile web has changed in the intervening years. Particularly given that case research in the study of fields like distributed reputation systems has slowed to a crawl – still limited mostly to Resnick's excellent early research on eBay's "feedback forum"[13] – we would be well-advised to reexamine our understanding of this fast-growing field. Contemporary authors [2,3,10,17] identify reputation tracking as a necessary precondition for the development of large-scale online participatory communities, but they do so while offering limited discussion of the circumstances under which such systems can effectively function. Between the iPhone, Android, and other new entrants, the Mobile Web has now reached deep enough penetration in some segments of society for us to consider maturing examples of their unique applications.

This paper represents, in essence, a theory-building project, attempting to better understand the foundational elements of distributed reputation systems and the impact that new types and quantites of mobile-based data is likely to have on their range and application. It is divided into four sections. Section 1 asks "what is reputation," reviewing the existing literature on distributed reputation tracking and offering a generic model of reputation systems. Section 2 introduces the location-enabled Mobile Web, discussing the quantitative and qualitative shifts in "proxy data" that it makes available to potential reputation systems as the internet-through-iPhone reaches a critical density threshold. Section 3 then synthesizes the previous two sections into a brief thought experiment that indicates one likely application for actors attempting to engage in "offline" collective action. Section 4 presents a critical perspective on the field, delving into a host of likely implications both problematic and promising. Centrally, the thesis of this paper is that the Mobile Web represents a critical juncture in the utility of the internet to political and civic associations. By dissolving the barrier between online and offline and ushering in location-aware proxy data, the circumstances are ripe for a new set of applications that will likely allow offline participatory communities to enjoy the similar efficiencies to those enjoyed by online communities-of-interest.

1.1: What is Reputation?

An individual's reputation refers to a set of complex, context-specific community assessments based on past activities. By "complex" I mean that reputation crosses multiple spaces. An academic colleague could have a reputation as an excellent researcher, but narrow-minded in her field, providing excellent feedback but habitually late for meetings. By context-specific, I mean that the dimensions of one's reputation are variably accessed depending on the situation. The above academic colleague cannot receive a universal standardized reputation score, because habitual lateness matters more in committee assignments than in providing feedback on research. Likewise, one's business colleagues will hold reputational assessments of hows he acts in negotiations, while a very difficult set of connections will hold assessments about her capacities as, say, a chef. By community assessments, I mean that these assessments are embedded in a set of network connections. A business trip that doubles as a visit with old college friends will cross social network boundaries, and an individual will have different, equally valid, and potentially non-overlapping reputations in these two circles.

The importance of reputation to collective action and public goods provision has been well established in the literature. Game theorist Robert Axelrod identifies the importance of reputation in mitigating a host of prisoners' dilemma-type social challenges, including collective action and the provision of public goods. Axelrod finds that the "Tit for Tat" strategy in pure strategic games, in which actor A's move at time T+1 is determined by actor B's move at time T, is a robust long-term strategy when the prisoners' dilemma is played over multiple iterations [1]. This "shadow of the future" is the formal equivalent of reputation between actors in a network. Dennis Chong provides an illustrative discussion of reputation's impact on the American Civil Rights Movement, noting that it is an essential motivating element that compels actors to participate in high-risk tactics that would otherwise face high levels of free ridership.[5] Likewise, Mancur Olson's opus work*, The Logic of Collective Action*, posits that large groups and small

[*] Paper Presented at the WebSci'09 Conference, March 19, 2009, Athens, Greece. Dave Karpf is a Research Fellow at the Miller Center for Public Affairs and a Doctoral Candidate in Political Science at the University of Pennsylvania. The author would like to thank Professors Rogers Smith, Jack Nagel, and John Lapinski, as well as Lokman Tsui, Rasmus Kleis Nielsen, Daniel Kreiss, David Faris and two anonymous reviewers for their valuable feedback.

groups are fundamentally different from one another because individual contributions in small groups are significant enough to be tracked, recognized, rewarded and punished. It is when groups attempt to scale up to larger size that Olson suggests we should expect underprovision of public goods.[12]

1.2: Distributed Reputation Systems

If reputation is a collection of complex, context-specific community assessments, a distributed reputation system can be best understood as a technologically-mediated approximation of such community-embedded assessments. As famously articulated by the classic *New Yorker* cartoon, initial forays into the internet faced the challenge of anonymity: "on the internet, nobody knows you're a dog."[15] Reputation systems replace internet anonymity with, at a minimum, internet pseudonymity. This allow members of a community to provide evaluations, assessments, tags, or other indicators of positive or negative contributions. Systems as diffuse as eBay's "feedback forum," Slashdot's "mojo" system, Amazon's recommendation structure, Wikipedia's administrator structure and google's PageRank algorithm all fundamentally share a few core attributes that allow their users to filter for preferred content, reward good behavior/contributions and sanction negative behavior/contributions. Note that some of these systems track the reputations of *individual users* while others track the reputations of *ideas*, *products*, or *institutions*. Some researchers have separated such systems into divergent categories,[9] but for the purposes of this paper, we gain more leverage by examining their commonalities. In all of these cases, the value of the reputation system can be reduced to three features: *proxy quality, set size,* and *algorithm*.

- Proxy Quality. While small-group reputation assessments are derived from complex interpersonal reactions, distributed online reputation systems must identify some form of proxy data as a stand-in for reputation. Google's PageRank uses the hyperlinked structure of the web as such a proxy, organizing search results according to those sites that receive the most hyperlinks, particularly frequent hyperlinks from high-traffic locations. Hindman points out that this has the result of skewing the system in favor of a small number of major sites.[8] Such a skew, however, solves the otherwise enormous search problem by giving searchers more popular, and thus relatively more trustworthy, sources of information. Slashdot and community blogs like DailyKos award "mojo points" to those comments or user-contributions deemed helpful, funny, or informative by their peers. This lets users filter results in a crowded discussion thread and also creates the basis for *superusers* – active community members, well-respected by their peers, who are given additional moderation or posting privileges. EBay's feedback forum invites buyers and sellers to rate their interactions, and these ratings aggregate into a permanent "feedback profile. All reputation systems rely on some form of proxy data, and where no good proxies can be identified or appropriate data is unavailable, reputation systems fail to operate.
- Set Size. Some types of proxy data are more abundant than others. The smaller the set of assessments, the easier it is to manipulate a distributed recommendation system, subverting it to malicious or counter-productive ends. Assessments can be obtained through various types of data, some of which are more costly to produce than others. These data types can be arranged on a spectrum of data availability, as demonstrated below in figure 1. Google and Amazon rely on the "data stream" for their proxy data. This is data that is added to the web as a byproduct of web activity. Clickstreams, hyperlinks, and overlapping "preferred lists" all query databases of existing web activity, drawing upon information that accrues as a costless byproduct of web surfing or participation. Data stream proxies are sometimes protected by proprietary firewalls, but otherwise are by far the most abundant of all data sources, because the data will accrue unless users take specific steps to erase their tracks. A minimally costly form of proxy data is produced through "integrated user experience," proxies. Comment rating systems on community blogs, for instance, allow users to click a button and score the contribution they have just read positively or negatively. Digg, Reddit, and StumbleUpon all invite users to similarly recommend interesting content as they navigate the web. This requires additional activity, but it is a near-seamless form of engagement, embedded into their web-surfing experience, and thus the additional transaction costs are minimal. At the higher end of the spectrum are "user feedback forms" such as those used by eBay, RateMyProfessor, and Meetup. These forms invite users to go online after a transaction or event has occurred and provide a rating. Because the activity of submitting a form is temporally separated from the transaction or event itself, the transaction costs are higher. Though in some of these cases, feedback forms can be somewhat heavily used – Resnick reports greater than 50% participation rates in the eBay feedback forum – the more costly it is for users to contribute proxy data, the more limited the available set will be.[13]
- The Algorithm. Algorithms can be simple or complex, proprietary or open, numeric or qualitative. In all cases, however, the purpose of the algorithm is to convert the set of proxy data into a *useful* product. "Useful" is used intentionally here because distributed reputation ratings, like reputation itself, are context-specific. In each and every case, the designers of a distributed reputation system have a set of activities they want to promote and discourage, and they design their algorithm to convert the raw proxy data into a system that rewards good behavior, products, activities, or ideas while sanctioning bad behavior, counter-productive activity, or lower-quality products or ideas.

Distributed reputation systems as a class, then, can be understood as converting a set of proxies through an algorithm into a useful rating system that supports the goals and purpose of the online community or network. It stands to reason, then, that improvements in either the quality or quantity of available proxy data would allow designers to craft new and better algorithms, applying distributed reputation tracking to a wider range of activities. With that background, let us now turn to the shifts in available data brought about by the internet-through-iPhone, colloquially known as the "mobile web."

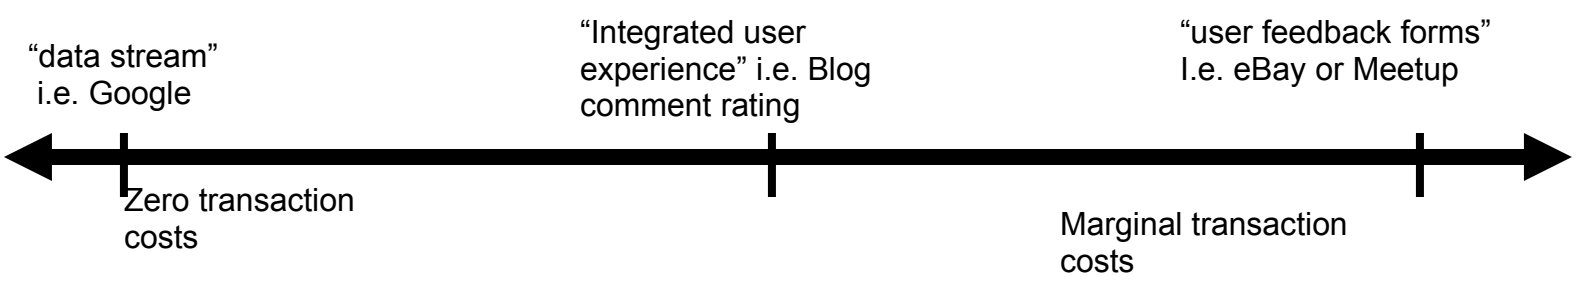

Figure 1: Continuum of data availability

2: Implications of the Mobile Web for Reputation Systems

In a recent article for *Wired* magazine, Mathew Honan offered the following summary: "simply put, location changes everything. This one input – our coordinates – has the potential to change all the outputs. Where we shop, who we talk to, what we read, what we search for, where we go – they all change once we merge location and the Web."[9] Indeed, with over 10 million iPhones already in circulation and the recent release of Google's Android operating system, we are poised to witness a massive expansion of this new version of the mobile web. There is a meaningful distinction to be drawn between the mobile web, as it is experienced today, and the mobile interfaces that Rheingold was discussing in *Smart Mobs*. Rheingold was primarily discussing the growth of SMS (Short Message Service), commonly known as text messaging. While SMS played an intriguing role in the "People Power II" Philippines revolt that led to the overthrow of a government [14] and has been used to organize "flashmobs" that virally spread an agreed-upon time and location, then arrive to engage in some form of collective action, the range of text-based applications has remained rather limited. SMS simply does not allow for the transmission of very *much* information. Apple capitalized on this limitation, as well as the limited breadth of phone-based web browsers, in it's initial marketing campaign. One of the first commercials announcing the iPhone to consumers intoned, "This is not a watered down version of the internet. Or the mobile version of the internet. Or the kinda-sorta-looks-like-the-internet internet. It's just the internet. On your phone." [16] The iPhone's data plan allows users to connect to wireless access points where available, and to access the internet through AT&T's broadband data network where wireless is unavailable. From the user's perspective, this means that web pages, google maps, wikipedia, and other major internet sites are now ubiquitously accessible.

The iPhone and its competitors affect reputation systems in at least two important ways. First is the flood of location-aware data that they make available, leading to a dramatic shift in the *types* of proxy data that can potentially be accessed by potential organizations or communities. Consider Yelp.com, for instance, which launched in 2005 as a sort of "Zagat ratings by the masses, for the masses," combining elements of Wikipedia, community blogs, social networks, and the Yellow Pages, all with a helping of Google Maps on the side. The site currently has over 3.3 million registered users and has expanded from a San Francisco start-up to include both major and mid-sized US cities.[1] Visitors to the site can browse through user-generated reviews of restaurants, shopping, nightlife, beauty and spas, or 15 other categories. Registered users can write their own reviews and give 1-5 star ratings of any location they visit. They can also rate each other's reviews, giving them credit for being "useful," "funny," or "cool," or flagging them as inappropriate content. Registered users are also invited to fill out profile information and network with each other. The most active reviewers on the site are invited to be members of the "Elite '08," a superuser-designation that includes invitations to local "thank-you" mixers.

As reporter Saul Hansell from the *New York Times* puts it, "Most people aren't drawn to write a witty review of the scrambled eggs at the local diner simply to get their ego stroked. But enough people find it rewarding to turn Yelp into one of the richest repositories of local reviews on the Web." [7] Though its launch predates the iPhone's market entry, Yelp's recent growth outside of tech-friendly cities like San Francisco and New York can be traced to the growth of mobile accessibility. Yelp provides a location-aware variation on Wikipedia's production model. Whereas Wikipedia serves as a repository for internet-based hobbyists to add their knowledge, Yelp serves as a repository for local wisdom. As reviews become accessible and increasingly writeable on-the-go, Yelp becomes an exponentially more valuable tool. So long as location-aware internet services are tethered to laptop or desktop portals, their utility is circumscribed. Yelp's iPhone application lets users access this repository of local wisdom at all times, encouraging higher participation rates and the accumulation of more assessments. While long-established sites like YellowPages.com and CitySearch.com attempt to provide top-down directors of available services, Yelp invites a growing participatory to do the legwork behind local search. Location-awareness increases the utility of Yelp's services, in turn broadening the size and scope of the service's applicability.

At least as important as the addition of new *types* of proxy data is the shift in data *abundance* as internet access comes untethered from stationary desktop or laptop portals and the line between online and offline is rendered utterly porous. It is perhaps too early to engage in detailed survey research on how the iPhone and its ilk change usage patterns, but a few things seem certain at this juncture. The gap between "integrated user experience" and "user feedback forms" represented in figure 1 is, almost by definition, reduced as the internet becomes omnipresent. If attendees to a meetup.com event can enter their ratings of the meeting *while it is in process*, then the costs of rating offline events approximate the costs of rating a blog comment. If "Yelpers" can rate the service at a restaurant while they are enduring an especially long wait, they are more likely to fill the extra time by submitting a low rating. It stands to reason that, at least to the extent that contributions to websites like epinions.com, yelp.com, or ratemyprofessor.com are limited by the marginal costs of participation, the wide diffusion of the mobile web will lead to massive increases in the production and availability of such assessments.[2]

The spread of this next technological wave does not itself predetermine the algorithms or uses that will result, and as previously asserted, reputation algorithms are designed to augment the goals of specific communities or organizations. All that we can be certain of at this juncture is that the introduction of new types of data and the increase in availability of other types of data will expand the space in which meaningful reputation algorithms can be developed. If we posit that effective reputation systems require both the quantity and quality of proxy data to be above some minimum threshold, then it follows that expanding the range of both, as the mobile web is doing, will make feasible a wide array of new applications.

---

[1] This amounts to fewer "registered users" than, for instance, Meetup.com, but that is due to the ill-defined nature of that statistic. According to Alexa.com traffic rankings, Yelp is the 637th most-visited web site on the World Wide Web, while Meetup ranks 1,359th (www.alexa.com, accessed on November 8, 2008.). Since inactive registrations are never removed from a system, membership size and membership activity can be only loosely correlated.

[2] This is an important and unknown limiting factor. I am not suggesting that, if rating systems were somehow rendered costless, everyone would choose to provide such ratings. Personal tastes or preferences must be taken into account. As an analogue, lowering the costs of posting personal thoughts on the web through simple and intuitive blogging software packages has not resulted in the entire world blogging. Many people have no interest in using the web in this manner. Likewise, lowering the costs of offering feedback on meetings, events, businesses, or individuals will allow for a fuller realization of the demand curve for such activity, but we have no idea what such a demand curve will look like.

The following section offers one such hypothetical reputational application as a guiding example.

3: Reputation-in-Motion. Augmenting Offline Collective Action

In a 2008 interview with Natalie Foster, former Deputy Field Network Director for MoveOn.org, I raised the topic of distributed reputation systems. "That's the Holy Grail," she replied, "We always said in MoveOn that what we needed more than anything else was a way to record feedback on which house parties and meetings went well and which one's didn't."[3] Technically, MoveOn does offer such a feedback mechanism in the form of automated post-event e-mails that are sent to participants asking for their experience. Despite the organization's enormous size and scale, however, the reply rate to these e-mails is below the necessary threshold for providing systematically useful information. MoveOn has been limited in developing an algorithm because of where these assessments lie on the continuum of data abundance, and the low response rate that this entails. How might we expect the spread of the mobile web to affect political associations like MoveOn, then?

Consider the following: local members of a major internet-mediated political association hold a monthly (offline) meeting to discuss their latest activities in 2014. By this time, the mobile web has reached a high diffusion rate, at least through their range of represented economic classes– akin to the camera feature on mobile phones today. The organization has built a rather simple application that it encourages members to use during the meeting. After they have signed up for the meeting online, the app connects to a GoogleMap, helping them find the meeting space. It also lets them review the list of fellow attendees, reinforcing organizational network ties. The meeting agenda is also available on the app. During the meeting, members are invited to give "kudos" and "troll ratings" to each other for positive and negative contributions to the group. "great idea" or "well-facilitated," "obnoxious comment," "too jargony" or "unfriendly to new members." They are also invited to rate the various portions of the meeting from 1 to 5 stars. Due to heavy 1-star feedback in previous months, the boring newsletter editor's report has been dropped from the agenda and the meetings as a whole have become more participatory and engaging. Toward the end of the meeting, they break out into small groups and assign tasks for the coming month. The group has agreed to a public education campaign and petition drive, and members are identifying local farmer's markets where they can set up a table and talk with community members. Rather than writing the tasks on scratch paper, which would then inevitably get misplaced, group leaders enter assignments directly into the "notepad" feature of the application. These assignments are stored in a central database and remain accessible to group members over the course of the month. Members can note when the tasks have been completed, remind each other of upcoming deadlines, and give each other support and thanks for completed tasks. Members who routinely overcommit at meetings can, once presented with that pattern, by politely invited to be more realistic in their time commitments. Members who are particularly good "workhorses" are recognized and publicly thanked for their contributions.

All of these habits, sans-application, are existing activities in a civic association that I observed over a six month period in 2007. The introduction of the application would be designed to augment their existing activities, simplifying local coordination and communication efforts. The data provided by the application could then be captured by the larger internet-mediated organization to recognize, thank, and empower the local volunteers who are most respected and valued by their peers. The national organization could grant such top "kudos"-receivers with quasi-superuser status, inviting them to national conferences or events where they are given heightened governance responsibilities and are publicly thanked for their important contributions. Meanwhile, "troll-ratings" could be aggregated into a "steer-clear" list for the organization as a whole. Those members who speak loudly and antagonize their peers could be barred from receiving updates or routed away from leading teams or hosting distributed events. The "shadow of the future" which large civic associations routinely lose when moving from local scale volunteerism to national-scale volunteerism can be restored through a collection of simple, helpful applications that aggregate into a useful reputation algorithm.

Zack Exley, formerly of MoveOn and current President of the New Organizing Institute, has colorfully described traditional civic associations as being dominated by the "Tyranny of the Annoying:" *"The Tyranny of the Annoying stems from the fact that, except in times of extreme crisis, it is just not worth it for mature, serious people to put up with all the indignities that go along with taking and maintaining leadership o any political entity. This principle guarantees that every Elks Club, Union Local, DAR Chapter, or Democratic town committee will tend toward being controlled by annoying people –they are the ones with egos desperate to be fed by winning petty little power plays and plenty of time on their hands.*" [6] The internet has enabled the *non*-annoying to participate in *online* communities-of-interest for several years now, with distributed reputation systems playing an essential role in distinguishing the valued participants with the respect of their peers from the loudest and most obnoxious. The fictional application I describe above combines a number of existing features, none of which are particularly complex. The introduction of location-aware data, the ubiquity of mobile telephony, and the resultant slide of proxy assessments along the continuum of data availability make distributed reputation tracking a potential reality for *offline* associations, leading the "Tyranny of the Annoying" to be replaced by the "Wisdom of the Crowd."

4.Another Step Toward Panopticon? Dangers of Reputation's Expansion

The lowered transaction costs of the internet have always come with a dark side. From "digital divides" that exacerbate class-based and educational cleavages within society to newly-empowered terrorist and hate-group networks, from a world where privacy is increasingly sacrificed for expediency to shifting cultural practices that carry unintended consequences, we have plentiful reasons to avoid treating the mobile web as an unmitigated good. As the technology diffuses and academics, entrepreneurs, and government officials shape and influence it, there are four areas where skepticism is warranted and careful discussion is called for.

First, how will the mobile web affect the digital divide? Put another way, will access to "just the internet… on your phone" follow the trajectory of the Blackberry PDA or of the camera phone. Early in their diffusion rates, both of these represented a novel bundle of hardware features, only available to consumers of high-end mobile phones. Camera phones eventually attained near-market saturation, available on plenty of low-end, cheap models. Blackberries, meanwhile, have remained a high-end commodity, reaching deep market penetration online among white-collar workers, but rare in other segments of society. Rheingold notes that the United States lagged behind Japan and Sweden in the first version of the mobile internet because of divergent policy stances that created a poor environment for broad diffusion. National governments would be well-advised to look at the policy lessons from the SMS-based mobile internet and actively support the diffusion of this technology. Mobile web-based applications carry the implicit promise of augmenting our daily lives in a host of small-but-powerful ways. If they are only available to the wealthy, privileged few, they will come to represent a major source of class-based cleavage.

---

[3] Natalie Foster, personal correspondance. May 16, 2008.

Second, we should remain aware that lowered online transaction costs make *everything* easier, including a host of socially negative activities. Mathew Honan made note of this in his recent *Wired* article, offering an anecdote of noticing a woman taking photos with her iPhone in Golden Gate Park, then logging onto Flickr that night, identifying the photo, and searching through linked photos to determine where she lived. [9] The mobile web may be a boon to civic associations, but it will likely also make predatory internet stalking far more *efficient*. Likewise, if the mobile web is a boon to knitting enthusiasts or service organizations, how do we prevent it from similarly augmenting the activities of hate groups? These negative byproducts of technological diffusion are unavoidable, but they can also be mitigated. Rather than waiting for the next incarnation of the "Myspace sexual predator" scandal, policymakers and entrepreneurs should consider tools that can make the worst applications of information abundance somewhat less threatening.

Third, and related, is the rapidly changing privacy environment. Fears from a decade ago that web-based credit card transactions would lead to widespread identity theft were overstated. But a seemingly lesser demon has proved far more pervasive. As the "data streams" of the internet lead to a costless aggregation of online information, the privacy environment changes in unexpected and unfortunate ways. Under previous information regimes, there was an implicit assumption of privacy in most personal communications, simply because it would be prohibitively costly to track and aggregate such information. The expansion of online storage capacity and value derived from "clickstream" data means that there is now a strong incentive for companies and governments to mine previously-unwieldy amounts of information about our private lives. Here geolocation is a dangerous social ill because it dramatically reduces internet anonymity. Both government policy choice and social education programs will be required to help mitigate this issue, as many citizens are unlikely to grasp how seemingly harmless Code can be aggregated into an essentially panoptic environment.

Fourth and finally, the mobile web is likely to continue to shift social practices and norms. Rheingold detailed a few of these cultural shifts in his book, discussing teenagers who constantly multitask and divide their attention between the person in front of them and the person texting them. These social practices lie outside the scope of my field, and thus I cannot comment authoritatively upon them. But, though changing technology does not necessitate any *particular* changing social practice, technological change does pave the way to some sort of change or another. There is a great need to study and understand these changing social practices, and indeed such research can hopefully inform public decision-making and help shape emerging norms and mores for the better.

5.Conclusion

We stand today at a critical point in the diffusion of the mobile web. As the internet comes untethered from computer and laptop terminals, the lines between online and offline become overwhelmingly blurry. This leads to new types of data, and floods of data types that previously proved scarce. The purpose of this paper has been to outline some of these changes and speculate upon a few of their likely applications, with a particular eye to the civic and political associations who attempt to mobilize collective action or otherwise ensure the provision of public goods. It has been largely US-focused, because that constitutes my field of study, but I have sought to generalize findings such that other researchers can apply them to other fields. What we can say for sure at this point is that the mobile web is approaching a tipping point in many parts of society. This will lead to a host of novel applications, and both positive and negative consequences. Moving forward, web science scholars should play an actively role in seeking to better understand these trends and inform emerging government policy.

Works Cited

[1] Axelrod, Robert. 1984. *The Evolution of Cooperation*. New York: Basic Books.

[2] Benkler, Yochai. 2006. *The Wealth of Networks.* New Haven, CT: Yale University Press.

[3] Bruns, Axel. 2008. *Blogs, Wikipedia, Second Life, and Beyond: From Production to Produsage*. New York: Peter Lang

[4] Clippinger, John. 2007. *A Crowd of One*. New York: Public Affairs

[5] Chong, Dennis. 1991. *Collective Action and the Civil Rights Movement*. Chicago: University of Chicago Press.

[6] Exley, Zack. 2007. "Can the Internets Make Me President?" blog post.
http://www.neworganizing.com/wiki/index.php/Can_the_Internets_make_me_president%3F

[7] Hansell, Saul. 2008. "Why Yelp Works," *New York Times*, May 12, 2008. http://bits.blogs.nytimes.com/2008/05/12/why-yelp-works/

[8] Hindman, Matthew. 2008. *The Myth of Digital Democracy*. Princeton, NJ: Princeton University Press.

[9] Honan, Mathew. February 2009. "I Am Here." *Wired* Magazine.

[10] Lev-On, Azi and Russell Hardin. 2007. "Internet-based Collaborations and their Political Significance," *Journal of Information Technology and Politics*. Volume 4, Issue 2.

[11] Masum, Hassan and Yi-Cheng Zhang. 2004. "Manifesto for the Reputation Society." *First Monday*. Volume 9, Number 7.

[12] Olson, Mancur. 1971. *The Logic of Collective Action, 2nd Edition*. Cambridge, MA: Harvard University Press.

[13] Resnick, Paul, Richard Zeckhauser, John Swanson, and Kate Lockwood. 2006. "The value of reputation on eBay: A controlled experiment." *Experimental Economics*. Issue 9, pgs 79-101. Economic Science Association.

[14] Rheingold, Howard. 2002. *Smart Mobs: The Next Social Revolution*. New York: Basic Books.

[15] Steiner, Peter. 1993. "On the Internet, Nobody knows you're a dog." Cartoon. *The New Yorker*. (Vol 69, no. 20) Pg 61.

[16] Username: mazmohad. June 6, 2007. "iPhone Commercials" Posted to YouTube. Accessed February 24, 2009.
http://www.youtube.com/watch?v=6IZMr-ZfoE4

[17] Weinberger, David. 2007. *Everything is Miscellaneous*. New York: Holt Paperbacks.