

MIT Open Access Articles

Towards Scalable Governance: Sensemaking and Cooperation in the Age of Social Media

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Rahwan, Iyad. "Towards Scalable Governance: Sensemaking and Cooperation in the Age of Social Media." *Philosophy & Technology* 30, 2 (November 2016): 161–178 © 2016 Springer Science+Business Media Dordrecht

As Published: <http://dx.doi.org/10.1007/s13347-016-0246-y>

Publisher: Springer-Verlag

Persistent URL: <http://hdl.handle.net/1721.1/110315>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



Towards Scalable Governance: Sensemaking and Cooperation in the Age of Social Media

AUTHOR ANONYMIZED AS PER JOURNAL POLICY

November 24, 2016

Abstract

Cybernetics, or self-governance of animal and machine, requires the ability to *sense* the world and to *act* on it in an appropriate manner. Likewise, self-governance of a human society requires *groups* of people to collectively sense and act on their environment. I argue that the evolution of political systems is characterized by a series of innovations that attempt to solve (among others) two ‘scalability’ problems: scaling up a group’s ability to make sense of an increasingly complex world, and to cooperate in increasingly larger groups. I then explore some recent efforts towards using the Internet and social media to provide alternative means for addressing these scalability challenges, under the banners of crowdsourcing and computer-supported argumentation. I present some lessons from those efforts about the limits of technology, and the research directions more likely to bear fruit.

1 Governance: Monoliths versus Societies

The field of *Cybernetics* is defined by one of its founders, MIT mathematician Norbert Wiener, as “the scientific study of control and communication in the animal and the machine” [Wiener et al., 1948]. Wiener, and other mid-twentieth century mathematicians and engineers who set the mathematical foundations for cybernetics in the 1940s and 1950s, were interested in engineering systems that can self-regulate. This work spun a variety of other disciplines, including modern Artificial Intelligence and control theory.

Thousands of years before Wiener resurrected and popularized the term, the earliest known mention of ‘cybernetics’ (from ‘kybernetike’, Greek for ‘governance’) was by Plato, as he contemplated how society can govern itself [Plato, 2015]. The parallel between the notion of governance in political science and cybernetics is not an accident. Both pertain the ability of an entity (an animal, a machine, or a society as a whole) to reflectively figure out how to regulate itself.

The analogy of society as a machine is a dangerous one, however. It is tempting to think of a complex system, such as society or the economy, as akin to a washing machine. A washing machine has well-understood behavior, can be reliably diagnosed, and faults can be addressed by fixing the malfunctioning component. But this analogy conceals the inherent unpredictability of large complex systems, which are made up of a multitude of interacting agents [Taleb, 2012]. If Chaos Theory has taught us anything, it is that even very simple systems can exhibit significant inherent unpredictability due to non-linear dynamics [Strogatz, 2014].

Luckily, nowadays neither scholars nor institutional decision makers are likely to fall prey of the washing machine fallacy, and the fact that societies are complex and largely unpredictable systems is widely accepted, even by laypeople. However, it is still very instructive to investigate the reasons that make that position fallacious, i.e., the key differences between the governance of a society and the regulation of a machine.

The misguided use of the ‘machine’ analogy to govern a system as complex as the economy (let’s call it the *washing machine fallacy*) has led to the collapse of many pure ‘command economies’ of communist regimes.¹ This is not to say, however, that social systems cannot self-regulate. Complex Systems science has taught us that the micro-level behaviors of individual agents can often lead to predictable, emergent, macro-level properties [Schelling, 2006, Holland, 2000]. Hence, the regulation of complex adaptive systems must (a) acknowledge the inherent unpredictability of various aspects of complex systems; (b) promote desirable and predictable emergent properties that *can* be predicted, by combining top-down control (e.g. setting interest rates, or punishing insider trading) with emergent behavior (e.g. competitive markets). Naturally, our ability to predict emergent properties of complex systems is still a subject of scientific inquiry, but we must do with what we have.

Another fundamental difference between machines and societies—the difference that this article is concerned with—can be revealed by investigating the basic building blocks of self-regulation. The process of regulating a machine or a society requires (among other things) two common fundamental abilities:

1. the ability to *sense* (and adapt one’s representation of) the environment; and
2. the ability to *act* on (or affect or control) this environment.

While a machine may regulate its behavior through sensors and control algorithms, a society has a fundamentally different set of mechanisms and constraints. Governance of a social system² is different from cybernetic control in that the processes of sensing and acting are distributed across groups of individuals. On one hand, sensing the environment becomes a problem of *distributed sensemaking* that incorporates information from a variety of sources and perspectives. On the other hand, acting becomes a problem of *cooperation* towards coordinated collective action, by a variety of actors with a multitude of potentially-conflicting goals (see Table 1).

	Machine / Single Organism	Society
Sense	Sensory input processing	Collective sensemaking problem
Act	Optimal action selection	Cooperation & negotiation problem

Table 1: Governance of monolithic versus social systems

This difference between governance of a monolithic system and a social system has a number of profound consequences. First, society’s ability to make sense of its environment relies on how well it can produce reliable information from a diverse set of agents. These agents have varying abilities to acquire information, for example by virtue of their physical location, position in their social network, or privileged access. They also have diverse abilities to process information through reasoning. As a result, the construction of knowledge requires processes by which group members can aggregate their information and reasoning in a manner (and speed) that produces timely informed action. I could not put it better than economist Friedrich von Hayek [Hayek, 1945]:

“[T]he problem of a rational economic order is determined precisely by the fact that the knowledge of the circumstances of which we must make use never exists in concentrated or integrated

¹See also the ambitious Project Cybersyn as an early example of cybernetic governance in 1970 Chile [Medina, 2011].

²In modern Political Science, *governance* is a broad term that refers to “all processes of governing, whether undertaken by a government, market or network, whether over a family, tribe, formal or informal organization or territory and whether through laws, norms, power or language” [Bevir, 2012]. As opposed to *government*, a formal body with the authority to make decisions in a given political system, governance includes all actors with the ability to influence the decision-making process, and encompasses not only how nation states are governed, but also corporations, universities, and so on.

form but solely as the dispersed bits of incomplete and frequently contradictory knowledge which all the separate individuals possess.

...

Or, to put it briefly, it is a problem of the utilization of knowledge which is not given to anyone in its totality.”

The second key aspect of governance of social systems is that action is not monolithic either—it must be undertaken by a variety of agents with diverse goals and interests. This creates a multitude of new challenges. One major challenge is *resource allocation*, since diverse agents may have claims on limited resources—such as land or labor—to achieve their individual goals. This requires mechanisms by which agents negotiate the use of such limited resources, including mechanisms for negotiating and agreeing on common goals—e.g. reducing inequality, maximizing social welfare, and so on.

But even when society has a common goal, a second major challenge to cooperative action is the *free rider* problem, which arises when the common goal conflicts with individual goals. Society as a whole may try to achieve some collective goal, such as sustaining a natural resource in order to ensure its the long-term availability. Yet, each individual agent may benefit from unilateral deviation from this goal, e.g. by over-exploiting the resource to his/her own advantage. This results in the so-called *Tragedy of the Commons* [Hardin, 1968], in which individuals acting independently and rationally according to each’s self-interest, bring about an outcome that is contrary to the best interests of the whole, by depleting some common resource. From the point of view of society as a whole, resolving the tension between the collective and individual interest is a constraint orthogonal to the challenge of identifying optimal collective actions.³

2 Evolution of Systems of Governance

I argue that the evolution of political systems is characterized by a series of innovations that attempt to *scale-up* both sensemaking and cooperation—subject to avoiding the ‘washing machine’ fallacy. That is, complex political structures arose because they provide a competitive advantage to groups that are able to better incorporate information from their members, and whose members are better able to cooperate effectively.

This conception of the development of political systems—as a solution to scalability of sense-making and cooperation—is consistent with the established wisdom in anthropology. Cultural anthropologists trace the evolution of political systems of governance from decentralized bands and tribes, to increasingly centralized chiefdoms, sovereign states and empires [Haviland et al., 2013]. For thousands of years, humans governed themselves through bands and tribes, whose members inhabited relatively simple environments and had simple technological artefacts (tools, weapons etc.). They were able to make sense of their environment by sharing information through language and social (cultural) learning [Henrich, 2015a, Boyd and Richerson, 2004, Mesoudi et al., 2006]. They were also able to cooperate (e.g. for hunting, warfare or farming) through a variety of deeply-rooted evolutionary mechanisms [Tomasello, 2009], most notably *kin selection*—helping others who share their genes [Hamilton, 1963], and *reciprocal altruism*—helping others who would later help them back [Trivers, 1971].

³It is worth noting that ultimately, individual humans are themselves not monoliths either. In cognitive science, there is significant theoretical and empirical support for the idea that individuals are themselves made up of societies of simpler, possibly competing agents [Minsky, 1988], and maintain internal coherence among competing goals and beliefs through a form of constraint satisfaction [Thagard, 2002] or interpersonal strategic game [Read, 2001]. However, the problems posed by such distributed governance have been effectively handled by natural evolution at the individual level, but remain very much a work-in-progress as we scale societal governance.

Over time, however, old institutions of sensemaking and cooperation became inadequate, as they cannot scale adequately to larger groups. In the face of inter-group competition, evolutionary pressure favored the emergence, and spread, of more complex social institutions to coordinate people's behaviors [Young, 2001]. On one hand, groups needed institutions that more efficiently aggregate information [Piketty, 1999]. Indeed, democratic institutions are often seen as a means through which large groups can make informed decisions without being individually informed. That is, people (the principal) delegate functions to the government (the agent) while using reliable information intermediaries (journalists, commentators etc.) to ensure reasoned choice without full information. As Lupia and McCubbins state [Lupia and McCubbins, 1998]:

“[D]ebate about democracy has always been a debate as to whether the pervasiveness of ignorance is going to lead to massive manipulation; we argue that the complexity of the world and the bounded rationality of the agents is actually matched by the huge ability to adapt of human beings and adequate democratic institutions.”

Likewise, centralized institutions became necessary for scaling up cooperation. For example, centralized sanctioning power is able to prevent subtle free-riding problems that undermine cooperation in larger groups [Gürerk et al., 2006, Baldassarri and Grossman, 2011, Sigmund et al., 2010]. Indeed, the founders of *social contract* theory, going back to Hobbes' Leviathan [Hobbes, 1651], posit that centralized government is legitimate precisely because it enables industrious people to cooperate through third-party enforcement of contracts among strangers.

In summary, as the world exploded in size and complexity, political institutions of sensemaking, cooperation and negotiation had to adapt in order to meet new scalability challenges. Modern political institutions, including the modern state, are a product of these evolutionary mechanisms of political development, which combine institutional innovation with learning. As Fukuyama puts it, “[s]ocieties are not trapped by their pasts and freely borrow ideas and institutions from each other” [Fukuyama, 2011].

3 New Enablers of Scalable Governance

In the previous sections, I argued that scaling up governance of society required (among others) the ability to scale up sensemaking and cooperation. This is because the ability of society to self-regulate depends on solving the complex challenges of collective sensing and collective action in the presence of distributed knowledge and diverse interests among political agents. I now turn to discussing the role of the Internet and social media in furthering the evolution of governance institutions, through the lens of sensemaking and cooperation.

Admittedly, much has been said about the topic of ‘Internet and governance’ [Lessig, 1998, Chadwick, 2006], and here I do not presume to cover it exhaustively. My perspective is particularly that of a technologist, rather than a political scientist, legal expert, or media studies scholar. Thus, I have a more humble, twofold objective: (a) to shed light on a handful of efforts to push sensemaking and cooperation to new scales, under the banners of argumentation-support and crowdsourcing; and (b) to present some lessons from those efforts about the limits of technology. My hope is to help guide future research agendas in these fields, specifically in relation to the goals of scalable governance.

3.1 Web-Enabled Scalable Sensemaking

By lowering the cost of information dissemination, communication technologies have always been seen as a means for revolutionizing sensemaking and cooperation in society. Even before the Internet, major

advances in communication technology were seen as powerful tools with the potential to transform governance. But some of these utopian visions have often appeared naive in hind sight. For example, in 1912, radio pioneer Guglielmo Marconi declared that “[t]he coming of the wireless era will make war impossible” [Narodny, 1912]. The Internet age is no different. The early transformative effects of the Internet has promised nothing short of revolutionizing democracy [Bryan et al., 2002, Hague and Loader, 1999]. However, this vision faced numerous early challenges, and as a result failed to fundamentally alter the public sphere or governance [Dahlberg, 2001]. Recently, commentators have gone to great lengths to diagnose the roots of the so-called cyber-utopian tendency [Morozov, 2011], and the economic forces that undermine it [Wu and Vietor, 2011].

However, that is not to say that the Internet has no potential to substantially scale up sensemaking. The canonical example is Wikipedia,⁴ the enormously successful encyclopaedia produced wholly by volunteer Internet users, and whose breadth is unmatched by any traditional encyclopaedia, in part because it provides explicit, reliable indicators of the goodwill and dedication of its editors [Goodwin, 2009]. Similar efforts have been very successful in scaling up the ability to produce collective knowledge about complex subjects, such as computer programming—for example *Stackoverflow*⁵ and mathematics—for example *Mathoverflow* [Tausczik et al., 2014]. The Web has also yielded platforms for sharing answers to broader, everyday questions, through so-called ‘Question and Answer’ sites like Quora⁶ and StackExchange.⁷ All of these systems have developed—often through discussion among the contributors themselves—policies for quality assurance. Some sites do this by allowing users to vote on the importance of various questions, in order to prioritize them. Another approach is to allow volunteers and users to flag inappropriate content, and enforce fairly sophisticated norms of accuracy and significance [Reagle, 2010]. Interestingly, the enforcement of complex norms has led, in the case of Wikipedia and other sites, to the emergence of power structures similar to those found in traditional political systems [Kittur et al., 2007].

Sensemaking is not only about understanding the present, but also about *anticipating* the future, including both the consequences of one’s own actions, as well as the natural evolution of the world. Recently we have witness various advances in scaling up collective anticipation, for example through the use of *prediction markets* [Arrow et al., 2008, Wolfers and Zitzewitz, 2004, Chen and Pennock, 2010]. In these markets, people trade predictions about future events, and the market price indicates what the crowd thinks the probability of the event is. Although these markets can succumb to so-called *speculative bubbles*, they allow for scalable collective anticipation of future events, which can in turn inform policy. Indeed, prediction markets are now used in various domains, from predicting political events, to selecting disruptive ideas in organizations.

So, the Internet seems to be doing well in terms of scaling up sensemaking in society, and allowing for the open sharing of knowledge, and the aggregation of both opinions and predictions. We seem to be on the promised path towards a deliberative democracy, in which knowledge is constructed by citizens, and in which information production is regulated through emergent, self-organized norms and quality-assurance mechanisms. Unfortunately, though, a number of caveats and bottlenecks remain. I discuss some of these below.

⁴<https://www.wikipedia.org/>

⁵<https://stackoverflow.com/>

⁶<https://www.quora.com/>

⁷<https://stackexchange.com/>

3.1.1 The Cognitive Bottleneck

One cognitive bottleneck to scalable sensemaking is the *cognitive overload* that people experience as they manage the torrent of information available online. Information competes for our limited attention [Weng et al., 2012]. As a result, individuals have to employ strategies that may impede the propagation of certain kinds of information, for example as people selectively allocate attention to easily discoverable content [Lerman, 2016].

Another caveat is that social knowledge construction may, in some cases, *impede* the performance of groups in collective intelligence tasks via *social influence* that undermines the diversity of opinions [Lorenz et al., 2011, Pentland, 2013]. A related effect is the so-called *social influence bias*, in which slight manipulations of initial opinions can have lasting effects on the collective opinions [Muchnik et al., 2013]. These effects not only diminish the diversity of opinions, but they can also be subject to long-term dynamics that lead to polarization [Conover et al., 2011].

These examples highlight the need for furthering our understanding of the limits of collective intelligence under the forces of social influence. Indeed, social influence and persuasion is often a good force [Mercier and Sperber, 2011], conducive of social learning and culture [Henrich, 2015b]. This is why our minds are especially attuned to social-proof. But it can sometimes backfire, raising the need for mechanisms for the detection and correction of misinformation. But whether social influence and social learning is a bug or a feature hinges on the nature of the task, and the structure and characteristics of the group.

3.1.2 The Temporal Bottleneck

A largely understudied collective sensemaking bottleneck is that current Web technology struggles with *time*. If an article in Wikipedia, or an answer to a question on Quora, remains incomplete or inaccurate, the consequences are fairly minor. However, in a natural disaster scenario, a political revolt, or an outright military conflict, time is the most scarce commodity, and information is perishable. A number of efforts have been successful in mobilizing highly committed volunteers to produce accurate maps of areas hit by violence or natural disasters in a timely manner [Meier, 2015, Popoola et al., 2013].

But there remains much to be done in order to scale this to larger teams and more complex situation awareness tasks [Mao et al., 2016], including better machine learning approaches to augment the ability of volunteers to vet information quickly [Gupta et al., 2014].

3.1.3 The Argumentation Bottleneck

Another major bottleneck in scaling up collective sensemaking is our inability to scale debate. In contested domains, various actors hold conflicting information, either due to selective information acquisition, or simply due to diversity of perspectives. We still lack a technology that can scale up sound argumentation and debate, even beyond a handful of discussants.

Argumentation is already common on discussion forums, micro-blogging sites, and news commentary. It is essential for civic participation, as it provides a mechanism for identifying pertinent issues and forming the collective agenda. However, Web technologies have so far been inadequate in supporting truly global collective argumentation. Most forums suffer from increased opinion polarization, spread of misinformed opinions, and wide-spread reasoning fallacies.

This raises the challenge of creating better tools and institutions to elevate online debate. Recently, a vision was articulated for using open data (semantic) technologies to create an infrastructure to support large-scale argumentation on the Web [Rahwan et al., 2007]. A number of efforts followed, towards building this technology, including the *Argument Interchange Format*, which provided the first common, exten-



Figure 1: The ArguBlogging widget for annotating Web content with arguments, facilitating global argumentation.

sible ontology of arguments [Chesñevar et al., 2006]. This led to early prototype implementations using Web ontology languages, highlighting its possible uses in argument annotation, querying, and classification [Rahwan, 2008, Rahwan et al., 2011]. This technology is now part of an emergent set of tools for “argumentative blogging” [Bex et al., 2013]. Figure 1 shows the ArguBlogging widget which allows people to annotate any Web content with their own arguments, which can be shared and cross-referenced openly. Instead of simply pressing a “Like” button, this technology enables a more nuanced, thoughtful, and targeted interaction with content.

Scalable argumentation is not only made difficult by the challenge of annotating and structuring arguments, but also in the fundamental limits on how these arguments can be aggregated. A substantial body of research exists on automating the aggregation of conflicting arguments through a process of automated reasoning [Rahwan et al., 2009b]. However, these techniques are not scalable, and are also not very good at incorporating uncertainty [Hunter, 2013]. There are also limits to our ability to use voting protocols for aggregating opinions about complex networks of conflicting arguments [Rahwan and Tohmé, 2010, Rahwan and Larson, 2008b]. Opinion aggregation can also succumb to strategic manipulation by individual agents [Rahwan and Larson, 2008a, Rahwan et al., 2009a, Awad et al., 2016, Rienstra et al., 2013, Sakama et al., 2015].

Although I have just surveyed attempts to scale-up debate, similar issues arise in the literature on judgment aggregation [List and Polak, 2010], which studies the aggregation of opinions over interconnected propositions. This includes the impossibility of aggregations that guarantee certain properties [Dietrich and List, 2007a], to limits to immunity from strategic manipulation [Dietrich and List, 2007b], to computational complexity [Endriss et al., 2012].

3.2 Web-Enabled Scalable Cooperation

Coined by Jeff Howe, the term *crowdsourcing* refers to the process of soliciting contributions from a large group of people—especially an online community—to achieve large-scale tasks [Howe, 2006]. We already

encountered crowdsourcing in the context of scalable sensemaking (see above). But crowdsourcing also captures the use of the Internet and social media to facilitate collective *action*.

Social media has already facilitated the mobilization of volunteers to map natural disasters in real-time, conduct large-scale search-and-rescue missions, and coordinate major political movements such as Occupy Wall Street or the Egyptian revolution. Again, there are many bottlenecks that prevent our ability to scale cooperation, and I discuss some of those below.

3.2.1 The Incentive Bottleneck

The *tragedy of the commons* occurs when multiple individuals, acting rationally in their own self-interest, will ultimately deplete a common good to the detriment of everybody, themselves included. To quote Nobel laureate Elinor Ostrom, “locally evolved institutional arrangements governed by stable communities and buffered from outside forces have sustained resources successfully for centuries” [Dietz et al., 2003]. Yet, some of the most important problems in modern society, such as pollution and food resource depletion, arise from our inability to scale-up our cooperative instincts beyond small groups.




Figure 2: Localizing externalities to the peers (yellow) of individuals A and B incentivizes their respective peers to use peer influence to encourage/discourage the behavior causing the positive/negative externality, respectively.

A very effective solution to this problem is to use *social network incentives*. In recent work, Yoeli et al [Yoeli et al., 2013] leveraged the role of *indirect reciprocity* [Nowak and Sigmund, 1998], which is to say that people exhibit cooperative behavior if this behavior is observable, since the gains in *reputation* increases the likelihood that they will be recipients of cooperation in the future. The result was a tripling of cooperation rates, in terms of signing up to energy efficiency programs. The same insight underlies *online reputation systems* [Jøsang et al., 2007].

Another alternative is to try to scale-up *direct reciprocity*. Instead of taxing individuals who pollute, or subsidising individuals who install solar panels, it may be possible to leverage peer-to-peer reciprocity at scale. One effort introduced a new approach to explored this, both using theory and preliminary experiments [Mani et al., 2013]. The idea is to reward the peers of an individual, rather than the individual himself, for demonstrating the desired behavior (see Figure 2). Analytically, it was shown that this mechanism can require a lower budget to operate than direct subsidies, even when accounting for the social cost of peer pressure.

While I focused here on addressing social dilemmas and avoiding free-riding behavior, incentives also play an important role when individuals need to resolve conflicts over the allocation of resources. This is the subject of study in the fields of *computational social choice* [Moulin et al., 2016], which explores the aggregation of societal preferences and fair allocation of resources. Here, again, we face issues of strategic manipulation [Faliszewski and Procaccia, 2010], impossibility of various kinds of fair resource


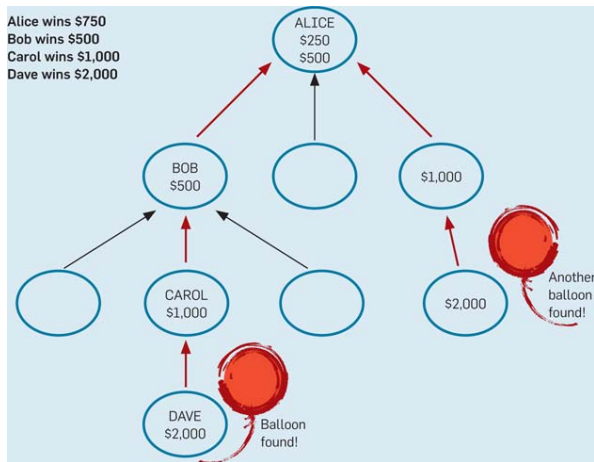


Figure 3: **DARPA Red Balloon Challenge.** (Left) Recursive incentive scheme. (Right) A recruitment cascade.

allocations [Procaccia, 2013], and so on. These are all problems that researchers in the algorithmic game theory community are exploring in depth [Nisan et al., 2007, Chen et al., 2013, Parkes and Wellman, 2015].

3.2.2 The Time-Space Bottleneck

Another bottleneck to large-scale cooperation has been the ability to get thousands or millions of people to collaborate on a task that spans a very large geographical area. Some recent examples highlight how the Internet and social media are changing this.

Various *citizen science* efforts have shown that it is possible to use the crowd to collect data over very large geographies, producing large, sometimes longitudinal data sets. Examples range from ecological data collection about bird species by bird watchers [Dickinson et al., 2010], to large scale archaeological searches through satellite imagery [Lin et al., 2014].

The *DARPA Network Challenge* (a.k.a. *Red Balloon Challenge*) [Tang et al., 2011] required competing teams to locate ten tethered weather balloons dispersed at random locations all over the continental United States. The winning team, based at MIT, won by locating all balloons in under 9 hours. The team used a novel incentive scheme to kick start a recruitment cascade of thousands of people within 48 hours (see Figure 3).

Analysis of the incentives and recruitment dynamics in the challenge [Pickard et al., 2011] suggested that the incentive mechanism theoretically guaranteed to lead to large recruitment cascades under a wide range of conditions. The data supported this finding, and revealed that the speed and branching of the diffusion (cascade size, branching factor, tree depth) reached levels above what is normally observed in viral propagation schemes such as chain letters [Iribarren and Moro, 2009]. Moreover, people actively recruited peers at greater geographical distance than expected by models of the geographical distribution of social ties [Liben-Nowell et al., 2005]. Subsequent analysis of Twitter data revealed that people with very high numbers of followers fail to sustain cascades despite the initial burst, providing further evidence of the role of incentives in maximizing cumulative recruitment. With this effort, we now understand various parameters of *time-critical social mobilization* in the Internet age for the first time.

A subsequent study used high-resolution, data-driven simulations of recruitment dynamics to demonstrate that the outcome of the Red Balloon Challenge is plausible without the presence of mass media, but lies at the limit of human capability, and requires all parameters to be at their empirical extremes

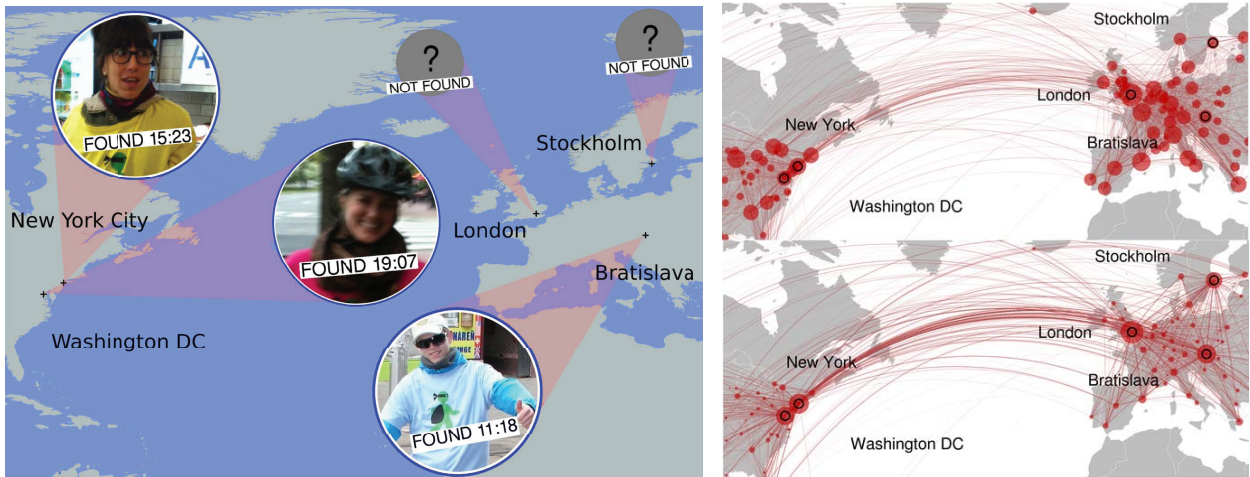


Figure 4: **State Department's Tag Challenge.** (Left) By mobilizing social media, we found these individuals and photographed them within 12 hours of release of their mug shots. (Right) We ran computer simulations to investigate how social mobilization has occurred; if people recruited friends at random (upper panel), this would result in a particular distribution of visits to different cities around the world; we estimate that people actually used a greedy heuristic, recruiting one in three friends in a targeted manner specifically in order to get closer to the target cities (lower panel).

[Rutherford et al., 2013a]. It also identified the role of population density and information overload in obscuring the target information. These observations are crucial if we are to rely on social media to disaster response or civic action.

In a similar example, crowdsourcing was used to win the *U.S. State Department's Tag Challenge*, which required teams to locate and photograph five 'thieves' (actors) in the US and Europe, based only on a mug shot released at 8:00am.⁸ The targets were visible for 12 hours, and followed normal itineraries around Stockholm, London, Bratislava, New York City and Washington D.C. Unlike the Red Balloon challenge, the targets were mobile, and they could blend into the crowd, raising additional complexity. Using a custom-built mobilization phone app, the team located 3 of the 5 suspects by mobilizing thousands using social media [Rahwan et al., 2013]. Analysis of the IP logs and Twitter data reveals fascinating spatio-temporal dynamics of the search [Rutherford et al., 2013b]. Not only was there evidence of geographical convergence towards target cities; it was also possible to recover the 'heuristic search algorithm' executed by the crowd's recruitment. Other heuristics were also identified in similar urban search challenges [Alstott et al., 2014].

These examples are powerful proofs-of-concept. They show that, in principle, it is possible to recruit hundreds of thousands or millions of participants to conduct difficult search tasks over very large geographies, sometimes in time-critical situations. More work needs to be done, however, to adapt these techniques to other domains. We also need to understand the conditions under which they fail.

3.2.3 The Complexity Bottleneck

One pertinent theoretical challenge is developing a model of time-critical team formation across a broader variety of domains and tasks. This would allow going beyond finding balloons and people, to more complex

⁸See: *Six Degrees of Mobilization*. The Economist, 2013. <http://www.economist.com/node/21560977>




Figure 5: Virtual board of our crowdsourcing platform, which allowed thousands of people to simultaneously collaborate on assembling shredded document for the DARPA Shredder Challenge.

tasks, such as disaster management. It is important to study the role of team assembly mechanisms in more complex settings, and to further elucidate the role of small time-scales in forming and maintaining these temporary coalitions. Another important direction is exploring the effect of various assumptions about the local knowledge, preferences, and capabilities of potential participants. I believe that a computational perspective is immensely powerful in understanding and designing these ‘social algorithms.’

A number of efforts are underway towards scaling up the complexity of large-scale, Internet-enabled cooperation [Doan et al., 2011]. These include various platforms for crowdsourcing of complex workflows [Kittur et al., 2011, Kulkarni et al., 2012, Zhang et al., 2012, Kulkarni et al., 2011]. Much further work is needed, however, before these types of complex workflows can enable cooperation in real-world dynamic scenarios [Bernstein et al., 2012].

I should also emphasize that scalable governance will always have its limits, because there are limits to what is knowable and implementable at any given point in time. When evaluating a policy or a potential solution to a cooperation problem, often the full extent of the outcome will only become apparent over long time scales.

3.2.4 The Sabotage Bottleneck

A crowdsourcing platform was also developed to solve the \$50,000 DARPA Shredder Challenge [DARPA, 2011]. Contestants were asked to solve five puzzles of increasing difficulty, each the result of a document being processed through different commercial shredders. The team managed to recruit thousands of individuals to collaboratively assemble the puzzle in real-time (Figure 5 shows the virtual canvas). Recent analysis revealed thousands of erroneous contributions and a number of large-scale attacks [Stefanovitch et al., 2014]. While the crowd was able to self-organize to recover from errors, they are (i) unable to contain malicious behavior, and (ii) the attacks displayed persistence over the subsequent participants, manifested in decreased participation and reduced problem solving efficiency. These results raise caution in the application of crowd-sourced problem solving for sensitive tasks involving financial markets and national security.

With any powerful technology come ethical considerations about the *risks* it entails, and a need for

safeguards. A technology capable of mobilizing thousands or millions of people deserves such attention. The risk inherent in the ability to mobilize people was recognized as far back as 1919, when United States Supreme Court justice Oliver Wendell Holmes identified the inherent power of any individual “*falsely shouting fire in a crowded theater*” to cause panic and even a deadly stampede. Attack is also a major challenge to the reliability of online reputation systems, necessitating defensive measures to ensure reputation acts as honest signals [Hoffman et al., 2009].

When it comes to large-scale cooperation via crowdsourcing, incentives can not only drive recruitment and action, but also verification. Some recent progress has been made on devising incentives for verification [Naroditskiy et al., 2012], but there is much to be done.

4 Summary

I argued that governance of social systems is similar to cybernetic regulation of machines, but with two fundamental differences: (a) sensing is distributed among agents with diverse and conflicting views on the present and future facts and goals; (b) action is distributed among agents with diverse and conflicting interests. Social systems are also complex and exhibit inherent unpredictability, which limits our ability to self-regulate society. This means that scaling up governance requires—among other things—an ability to scale up collective sensemaking, as well as collective cooperation and negotiation.

I then argued that scalable governance has already been happening over time, as evident in the historical development of political institutions for sensemaking (via information aggregation) and cooperation (via centralized monitoring and sanctioning).

The Internet and social media promise to usher a new major transition in scaling up sensemaking, cooperation and negotiation. But despite progress to date, we face a number of challenges, ranging from the way social media propagates and amplifies misinformation and bias, to our inability to scale argumentation and debate, to the tradeoffs between scalable cooperation and policing of open systems. I discussed many of these challenges, largely through the lens of my own experience, but also incorporating important work in many other fields of inquiry. The list of issues identified is, therefore, representative rather than exhaustive.

I hope the article helps frame a broad agenda towards building technology that scales-up civic participation in the age of hyper connectivity and big data. This agenda should focus on incentive schemes, Web technologies, and mobile applications for large-scale mobilization, crowdsourcing and sensemaking in social networks spanning theoretical, experimental and applications research.

References

- [Alstott et al., 2014] Alstott, J., Madnick, S., Velu, C., and Sánchez, A. (2014). Homophily and the speed of social mobilization: the effect of acquired and ascribed traits. *PloS one*, 9(4):e95140.
- [Arrow et al., 2008] Arrow, K. J., Forsythe, R., Gorham, M., Hahn, R., Hanson, R., Ledyard, J. O., Levmore, S., Litan, R., Milgrom, P., Nelson, F. D., et al. (2008). The promise of prediction markets. *SCIENCE-NEW YORK THEN WASHINGTON*-, 320(5878):877.
- [Awad et al., 2016] Awad, E., Booth, R., Tohme, F., and Rahwan, I. (2016). Judgement aggregation in multi-agent argumentation. *Journal of Logic and Computation*, (in press).
- [Baldassarri and Grossman, 2011] Baldassarri, D. and Grossman, G. (2011). Centralized sanctioning and legitimate authority promote cooperation in humans. *Proceedings of the National Academy of Sciences*, 108(27):11023–11027.
- [Bernstein et al., 2012] Bernstein, A., Klein, M., and Malone, T. W. (2012). Programming the global brain. *Communications of the ACM*, 55(5):41–43.
- [Bevir, 2012] Bevir, M. (2012). *Governance: A very short introduction*, volume 333. Oxford University Press.

- [Bex et al., 2013] Bex, F., Lawrence, J., Snaith, M., and Reed, C. (2013). Implementing the argument web. *Communications of the ACM*, 56(10):66–73.
- [Boyd and Richerson, 2004] Boyd, R. and Richerson, P. J. (2004). *The origin and evolution of cultures*. Oxford University Press.
- [Bryan et al., 2002] Bryan, C., Tambini, D., and Tsagarousianou, R. (2002). *Cyberdemocracy: Technology, cities and civic networks*. Routledge.
- [Chadwick, 2006] Chadwick, A. (2006). *Internet politics: States, citizens, and new communication technologies*. Oxford University Press, USA.
- [Chen et al., 2013] Chen, Y., Lai, J. K., Parkes, D. C., and Procaccia, A. D. (2013). Truth, justice, and cake cutting. *Games and Economic Behavior*, 77(1):284–297.
- [Chen and Pennock, 2010] Chen, Y. and Pennock, D. M. (2010). Designing markets for prediction. *AI Magazine*, 31(4):42–52.
- [Chesñevar et al., 2006] Chesñevar, C. I., McGinnis, J., Modgil, S., Rahwan, I., Reed, C., Simari, G., South, M., Vreeswijk, G., and Willmott, S. (2006). Towards an argument interchange format. *The Knowledge Engineering Review*, 21(4):293–316.
- [Conover et al., 2011] Conover, M., Ratkiewicz, J., Francisco, M. R., Gonçalves, B., Menczer, F., and Flammini, A. (2011). Political polarization on twitter. *ICWSM*, 133:89–96.
- [Dahlberg, 2001] Dahlberg, L. (2001). The internet and democratic discourse: Exploring the prospects of online deliberative forums extending the public sphere. *Information, Communication & Society*, 4(4):615–633.
- [DARPA, 2011] DARPA (2011). Darpa shredder challenge. <http://archive.darpa.mil/shredderchallenge>. Accessed: 2013-12-26.
- [Dickinson et al., 2010] Dickinson, J. L., Zuckerberg, B., and Bonter, D. N. (2010). Citizen science as an ecological research tool: challenges and benefits. *Annual review of ecology, evolution, and systematics*, 41:149–172.
- [Dietrich and List, 2007a] Dietrich, F. and List, C. (2007a). Arrows theorem in judgment aggregation. *Social Choice and Welfare*, 29(1):19–33.
- [Dietrich and List, 2007b] Dietrich, F. and List, C. (2007b). Strategy-proof judgment aggregation. *Economics and Philosophy*, 23(03):269–300.
- [Dietz et al., 2003] Dietz, T., Ostrom, E., and Stern, P. C. (2003). The struggle to govern the commons. *science*, 302(5652):1907–1912.
- [Doan et al., 2011] Doan, A., Ramakrishnan, R., and Halevy, A. Y. (2011). Crowdsourcing systems on the world-wide web. *Communications of the ACM*, 54(4):86–96.
- [Endriss et al., 2012] Endriss, U., Grandi, U., and Porello, D. (2012). Complexity of judgment aggregation. *Journal of Artificial Intelligence Research*, 45:481–514.
- [Faliszewski and Procaccia, 2010] Faliszewski, P. and Procaccia, A. D. (2010). Ai’s war on manipulation: Are we winning? *AI Magazine*, 31(4):53–64.
- [Fukuyama, 2011] Fukuyama, F. (2011). *The origins of political order: from prehuman times to the French Revolution*. Profile books.
- [Goodwin, 2009] Goodwin, J. (2009). The authority of wikipedia. In *Argument cultures: Proceedings of the Ontario Society for the Study of Argumentation conference*. Ontario Society for the Study of Argumentation.
- [Gupta et al., 2014] Gupta, A., Kumaraguru, P., Castillo, C., and Meier, P. (2014). Tweetcred: Real-time credibility assessment of content on twitter. In *Social Informatics*, pages 228–243. Springer.
- [Güererk et al., 2006] Güererk, Ö., Irlenbusch, B., and Rockenbach, B. (2006). The competitive advantage of sanctioning institutions. *Science*, 312(5770):108–111.
- [Hague and Loader, 1999] Hague, B. N. and Loader, B. (1999). *Digital democracy: Discourse and decision making in the information age*. Psychology Press.
- [Hamilton, 1963] Hamilton, W. D. (1963). The evolution of altruistic behavior. *American naturalist*, pages 354–356.
- [Hardin, 1968] Hardin, G. (1968). The tragedy of the commons. *science*, 162(3859):1243–1248.
- [Haviland et al., 2013] Haviland, W., Prins, H., McBride, B., and Walrath, D. (2013). *Cultural anthropology: the human challenge*. Cengage Learning.
- [Hayek, 1945] Hayek, F. A. (1945). The use of knowledge in society. *The American Economic Review*, 35(4):519–530.
- [Henrich, 2015a] Henrich, J. (2015a). *The Secret of Our Success: How Culture Is Driving Human Evolution, Domesticating Our Species, and Making Us Smarter*. Princeton University Press.

- [Henrich, 2015b] Henrich, J. (2015b). *The secret of our success: how culture is driving human evolution, domesticating our species, and making us smarter*. Princeton University Press.
- [Hobbes, 1651] Hobbes, T. (1651). *Leviathan, or, the Matter, Forme, and Power of a Common-Wealth Ecclesiasticall and Civill*.
- [Hoffman et al., 2009] Hoffman, K., Zage, D., and Nita-Rotaru, C. (2009). A survey of attack and defense techniques for reputation systems. *ACM Computing Surveys (CSUR)*, 42(1):1.
- [Holland, 2000] Holland, J. H. (2000). *Emergence: From chaos to order*. Oxford University Press.
- [Howe, 2006] Howe, J. (2006). The rise of crowdsourcing. *Wired magazine*, 14(6):1–4.
- [Hunter, 2013] Hunter, A. (2013). Modelling uncertainty in persuasion. In *International Conference on Scalable Uncertainty Management*, pages 57–70. Springer.
- [Iribarren and Moro, 2009] Iribarren, J. and Moro, E. (2009). Impact of human activity patterns on the dynamics of information diffusion. *Physical Review Letters*, 103(3).
- [Jøsang et al., 2007] Jøsang, A., Ismail, R., and Boyd, C. (2007). A survey of trust and reputation systems for online service provision. *Decision support systems*, 43(2):618–644.
- [Kittur et al., 2007] Kittur, A., Chi, E., Pendleton, B. A., Suh, B., and Mytkowicz, T. (2007). Power of the few vs. wisdom of the crowd: Wikipedia and the rise of the bourgeoisie. *World wide web*, 1(2):19.
- [Kittur et al., 2011] Kittur, A., Smus, B., Khamkar, S., and Kraut, R. E. (2011). Crowdforge: Crowdsourcing complex work. In *Proceedings of the 24th annual ACM symposium on User interface software and technology*, pages 43–52. ACM.
- [Kulkarni et al., 2012] Kulkarni, A., Can, M., and Hartmann, B. (2012). Collaboratively crowdsourcing workflows with turkomatic. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work*, pages 1003–1012. ACM.
- [Kulkarni et al., 2011] Kulkarni, A. P., Can, M., and Hartmann, B. (2011). Turkomatic: automatic recursive task and workflow design for mechanical turk. In *CHI’11 Extended Abstracts on Human Factors in Computing Systems*, pages 2053–2058. ACM.
- [Lerman, 2016] Lerman, K. (2016). Information is not a virus, and other consequences of human cognitive limits. *Future Internet*, 8(2):21.
- [Lessig, 1998] Lessig, L. (1998). Open code and open societies: values of internet governance. *Chi.-Kent L. Rev.*, 74:1405.
- [Liben-Nowell et al., 2005] Liben-Nowell, D., Novak, J., Kumar, R., Raghavan, P., and Tomkins, A. (2005). Geographic routing in social networks. *Proceedings of the National Academy of Sciences*, 102(33):11623–11628.
- [Lin et al., 2014] Lin, A. Y.-M., Huynh, A., Lanckriet, G., and Barrington, L. (2014). Crowdsourcing the unknown: The satellite search for genghis khan. *PLoS one*, 9(12):e114046.
- [List and Polak, 2010] List, C. and Polak, B. (2010). Introduction to judgment aggregation. *Journal of economic theory*, 145(2):441–466.
- [Lorenz et al., 2011] Lorenz, J., Rauhut, H., Schweitzer, F., and Helbing, D. (2011). How social influence can undermine the wisdom of crowd effect. *Proceedings of the National Academy of Sciences*, 108(22):9020–9025.
- [Lupia and McCubbins, 1998] Lupia, A. and McCubbins, M. D. (1998). *The democratic dilemma: Can citizens learn what they need to know?* Cambridge University Press.
- [Mani et al., 2013] Mani, A., Rahwan, I., and Pentland, A. (2013). Inducing peer pressure to promote cooperation. *Scientific reports*, 3.
- [Mao et al., 2016] Mao, A., Mason, W., Suri, S., and Watts, D. J. (2016). An experimental study of team size and performance on a complex task. *PLoS one*, 11(4):e0153048.
- [Medina, 2011] Medina, E. (2011). *Cybernetic revolutionaries: technology and politics in Allende’s Chile*. MIT Press.
- [Meier, 2015] Meier, P. (2015). Digital humanitarians. *How Big Data Is Changing the Face of Humanitarian Response*. CRC.
- [Mercier and Sperber, 2011] Mercier, H. and Sperber, D. (2011). Why do humans reason? arguments for an argumentative theory. *Behavioral and brain sciences*, 34(02):57–74.
- [Mesoudi et al., 2006] Mesoudi, A., Whiten, A., and Laland, K. N. (2006). Towards a unified science of cultural evolution. *Behavioral and Brain Sciences*, 29(04):329–347.
- [Minsky, 1988] Minsky, M. (1988). *Society of mind*. Simon and Schuster.
- [Morozov, 2011] Morozov, E. (2011). *The Net Delusion: How not to liberate the world*. Penguin.

- [Moulin et al., 2016] Moulin, H., Brandt, F., Conitzer, V., Endriss, U., Lang, J., and Procaccia, A. D. (2016). *Handbook of Computational Social Choice*. Cambridge University Press.
- [Muchnik et al., 2013] Muchnik, L., Aral, S., and Taylor, S. J. (2013). Social influence bias: A randomized experiment. *Science*, 341(6146):647–651.
- [Naroditskiy et al., 2012] Naroditskiy, V., Rahwan, I., Cebrian, M., and Jennings, N. (2012). Verification in referral-based crowdsourcing. *PLoS ONE*, 7(10):e45924.
- [Narodny, 1912] Narodny, I. (1912). Marconi’s plans for the world. *Technical World Magazine*, pages 145–150.
- [Nisan et al., 2007] Nisan, N., Roughgarden, T., Tardos, E., and Vazirani, V. V. (2007). *Algorithmic game theory*, volume 1. Cambridge University Press Cambridge.
- [Nowak and Sigmund, 1998] Nowak, M. A. and Sigmund, K. (1998). Evolution of indirect reciprocity by image scoring. *Nature*, 393(6685):573–577.
- [Parkes and Wellman, 2015] Parkes, D. C. and Wellman, M. P. (2015). Economic reasoning and artificial intelligence. *Science*, 349(6245):267–272.
- [Pentland, 2013] Pentland, A. S. (2013). Beyond the echo chamber. *Harvard Business Review*, 91(11):80–+.
- [Pickard et al., 2011] Pickard, G., Pan, W., Rahwan, I., Cebrian, M., Crane, R., Madan, A., and Pentland, A. (2011). Time-critical social mobilization. *Science*, 334(6055):509–512.
- [Piketty, 1999] Piketty, T. (1999). The information-aggregation approach to political institutions. *European Economic Review*, 43(4):791–800.
- [Plato, 2015] Plato (2015). *First Alcibiades*. CreateSpace Independent Publishing Platform.
- [Popoola et al., 2013] Popoola, A., Krasnoshtan, D., Toth, A.-P., Naroditskiy, V., Castillo, C., Meier, P., and Rahwan, I. (2013). Information verification during natural disasters. In *Proceedings of the 22nd international conference on World Wide Web companion*, pages 1029–1032. International World Wide Web Conferences Steering Committee.
- [Procaccia, 2013] Procaccia, A. D. (2013). Cake cutting: not just child’s play. *Communications of the ACM*, 56(7):78–87.
- [Rahwan, 2008] Rahwan, I. (2008). Mass argumentation and the Semantic Web. *Journal of Web Semantics*, 6(1):29–37.
- [Rahwan et al., 2011] Rahwan, I., Banihashemi, B., Reed, C., Walton, D., and Abdallah, S. (2011). Representing and classifying arguments on the semantic web. *The Knowledge Engineering Review*, 26(04):487–511.
- [Rahwan et al., 2013] Rahwan, I., Dsouza, S., Rutherford, A., Naroditskiy, V., McInerney, J., Venanzi, M., Jennings, N. R., and Cebrian, M. (2013). Global manhunt pushes the limits of social mobilization. *IEEE Computer*, 46(4).
- [Rahwan and Larson, 2008a] Rahwan, I. and Larson, K. (2008a). Mechanism design for abstract argumentation. In Padgham, L., Parkes, D., Mueller, J., and Parsons, S., editors, *7th International Joint Conference on Autonomous Agents & Multi Agent Systems, AAMAS’2008, Estoril, Portugal*, pages 1031–1038.
- [Rahwan and Larson, 2008b] Rahwan, I. and Larson, K. (2008b). Pareto optimality in abstract argumentation. In Fox, D. and Gomes, C., editors, *Proceedings of the 23rd AAAI Conference on Artificial Intelligence (AAAI-2008)*, Menlo Park CA, USA.
- [Rahwan et al., 2009a] Rahwan, I., Larson, K., and Tohmé, F. (2009a). A characterisation of strategy-proofness for grounded argumentation semantics. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 251–256.
- [Rahwan et al., 2009b] Rahwan, I., Simari, G. R., and van Benthem, J. (2009b). *Argumentation in artificial intelligence*, volume 47. Springer.
- [Rahwan and Tohmé, 2010] Rahwan, I. and Tohmé, F. (2010). Collective argument evaluation as judgement aggregation. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 1-Volume 1*, pages 417–424. International Foundation for Autonomous Agents and Multiagent Systems.
- [Rahwan et al., 2007] Rahwan, I., Zablith, F., and Reed, C. (2007). Laying the foundations for a world wide argument web. *Artificial Intelligence*, 171(10–15):897–921.
- [Read, 2001] Read, D. (2001). Intrapersonal dilemmas. *Human Relations*, 54(8):1093–1117.
- [Reagle, 2010] Reagle, J. M. (2010). *Good faith collaboration: The culture of Wikipedia*. MIT Press.
- [Rienstra et al., 2013] Rienstra, T., Thimm, M., and Oren, N. (2013). Opponent models with uncertainty for strategic argumentation. In *IJCAI*.

- [Rutherford et al., 2013a] Rutherford, A., Cebrian, M., Dsouza, S., Moro, E., Pentland, A., and Rahwan, I. (2013a). The limits of social mobilization. *PNAS*, 110(16):6281–6286.
- [Rutherford et al., 2013b] Rutherford, A., Cebrian, M., Rahwan, I., Dsouza, S., McInerney, J., Naroditskiy, V., Venanzi, M., Jennings, N., deLara, J., Wahlstedt, E., and Miller, S. (2013b). Targeted social mobilisation in a global manhunt. *PLOS ONE*, (in press).
- [Sakama et al., 2015] Sakama, C., Caminada, M., and Herzig, A. (2015). A formal account of dishonesty. *Logic Journal of IGPL*, 23(2):259–294.
- [Schelling, 2006] Schelling, T. C. (2006). *Micromotives and macrobehavior*. WW Norton & Company.
- [Sigmund et al., 2010] Sigmund, K., De Silva, H., Traulsen, A., and Hauert, C. (2010). Social learning promotes institutions for governing the commons. *Nature*, 466(7308):861–863.
- [Stefanovitch et al., 2014] Stefanovitch, N., Alshamsi, A., Cebrian, M., and Rahwan, I. (2014). Error and attack tolerance of collective problem solving: The darpa shredder challenge. *EPJ Data Science*, 3(13).
- [Strogatz, 2014] Strogatz, S. H. (2014). *Nonlinear dynamics and chaos: with applications to physics, biology, chemistry, and engineering*. Westview press.
- [Taleb, 2012] Taleb, N. N. (2012). *Antifragile: Things that gain from disorder*, volume 3. Random House Incorporated.
- [Tang et al., 2011] Tang, J., Cebrian, M., Giacobe, N., Kim, H., Kim, T., and Wickert, D. (2011). Reflecting on the DARPA Red Balloon Challenge. *Communications of the ACM*, 54(4):78–85.
- [Tausczik et al., 2014] Tausczik, Y. R., Kittur, A., and Kraut, R. E. (2014). Collaborative problem solving: a study of mathoverflow. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing*, pages 355–367. ACM.
- [Thagard, 2002] Thagard, P. (2002). *Coherence in thought and action*. MIT press.
- [Tomasello, 2009] Tomasello, M. (2009). *Why we cooperate*. MIT press.
- [Trivers, 1971] Trivers, R. L. (1971). The evolution of reciprocal altruism. *Quarterly review of biology*, pages 35–57.
- [Weng et al., 2012] Weng, L., Flammini, A., Vespignani, A., and Menczer, F. (2012). Competition among memes in a world with limited attention. *Scientific reports*, 2.
- [Wiener et al., 1948] Wiener, N. et al. (1948). *Cybernetics*. Hermann Paris.
- [Wolfers and Zitzewitz, 2004] Wolfers, J. and Zitzewitz, E. (2004). Prediction markets. *The Journal of Economic Perspectives*, 18(2):107–126.
- [Wu and Vietor, 2011] Wu, T. and Vietor, M. (2011). *The master switch: The rise and fall of information empires*. Vintage Books.
- [Yoeli et al., 2013] Yoeli, E., Hoffman, M., Rand, D. G., and Nowak, M. A. (2013). Powering up with indirect reciprocity in a large-scale field experiment. *Proceedings of the National Academy of Sciences*, 110(Supplement 2):10424–10429.
- [Young, 2001] Young, H. P. (2001). *Individual strategy and social structure: An evolutionary theory of institutions*. Princeton University Press.
- [Zhang et al., 2012] Zhang, H., Law, E., Miller, R., Gajos, K., Parkes, D., and Horvitz, E. (2012). Human computation tasks with global constraints. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 217–226. ACM.