Networked Deliberation in Peer-Produced Collective Action

Edward L. Platt

May 3, 2018

Contents

Chapter 1

Introduction

The past several decades have seen the emergence of a networked information economy (Benkler, 2006). This new type of economy is characterized by fast, cheap, and bidirectional global communication, thanks to the internet, resulting in an increase in knowledge work where creativity has become the limiting resource in many endeavors. While the internet and the networked information economy have rightly been identified to hold great potential for civic participation, the past several decades have seen a drastic decline in voter turnout, political knowledge, participation in service organizations, and other indicators of civic engagement in the United States (Putnam, 2000; Ackerman & Fishkin, 2002). Civic engagement is crucial for the development of the social capital and trust that enable collective action, which in turn allows communities to advocate for their interests and build sustainable regimes to manage resources (Ostrom, 2000). Why has the potential of the internet to enable civic participation and collective action gone largely unfulfilled? What can be learned from relatively successful internet-enabled collaborations? And what questions still need to be answered to best harness the internet for collective action?

Specifically, I will examine what is know about the many steps between large-scale communication and large scale collective action. Wikipedia and other forms of commons-based peer production serve as a key and relatively successful example (Benkler, 2002). These examples suggest that the internet can enable collective action at a global scale, if the massive potential for communication and production can be combined with a means of filtering and accreditation. Fundamentally, this is a problem of (very) large scale decision-making. I review the literature traditional collective action, commons-based peer production, collective intelligence, and social choice theory to identify common themes relevant to internet-enabled decision-making for collective action.

Large-scale decision-making is difficult and emphasis typically shifts from discourse and deliberation to voting as groups grow in size (Ackerman & Fishkin, 2002; González-Bailón, 2010). But the ability of voting to achieve cooperation is limited by paradoxes and impossibility results from social choice theory (Brandt et al., 2012; Arrow, 1950; Condorcet, 1785). Instead of abandoning discourse and deliberation, in this prelim I examine the importance of these modes of participation in commons-based peer production and collective action. I also emphasize that the limitations of voting can potentially be overcome by the ability of discourse and deliberation to build consensus, trust, and cooperation (Habermas, 1964; Ostrom, 2000; Geiger, 2009). I suggest that techniques for large-scale deliberation could create more effective large-scale decision-making and more effective use of the internet for collective action.

In addition to my literature review, I propose an experiment to evaluate one approach to large-scale deliberation. Across different fields, the importance of small, overlapping groups emerges repeatedly as important for building trust, propagating and filtering information, and creating non-hierarchical structure (Ostrom, 2010; Benkler, 2006; Geiger 2009; Putnam, 2000; Freeman, 1972; Gray et al., 2015). However, there has not been a systematic, quantitative study of how the network structure of overlapping groups impacts their ability to enable large scale cohesion and discourse. I propose networked deliberation as a way to combine small groups into very large scale deliberations, and describe an experiment to test its effectiveness. I transfer quantitative measures from social choice theory to track the development of consensus over the course of deliberation and quantitatively compare the effectiveness of different networks. Elinor Ostrom wrote that “a core goal of public policy should be to facilitate the development of institutions that bring out the best in humans (2010).” In this prelim, I ask how information and communication technology can do the same.

Chapter 2 reviews theories of collective action and collective intelligence. Chapter 3 reviews case studies of large scale collective actions, both online and off. Chapter 4 describes my proposed research.

Chapter 2

Collective Action & Intelligence

Individuals who have few resources can join together to achieve common goals by engaging in collective action. However, to achieve their goals, those individuals must be able to collaborate effectively. Collaboration requires both that individuals can agree on which actions they should take, and that they can coordinate those actions effectively. Both collective decision-making and coordination have been studied in the in the context of collective intelligence. This chapter reviews work on both collective action and collective intelligence, and discusses how the latter can inform the former.

No unified theory or framework of collective action exists (Ostrom, 2000). For the purpose of this review, I will adopt a framework consisting of four stages: ideation, deliberation, decision-making, and execution (Figure 2.1). These stages are roughly sequential, but occur simultaneously in varying combinations over the course of actual collective action. The ideation stage consists of generating ideas for possible actions, as occurs in commons-based peer production. Deliberation consists of discussing problems and proposed solutions, as occurs in social learning and the discourse of the public sphere. In decision-making, group members aggregate their individual preferences into a social preference for the entire group, as formalized in social choice theory. Finally, in the execution stage, once a decision has been agreed on, group members act, either complying with the group's decision or, if some dissent and conflict remain, defecting.

2.1 Collective Action

Much of the existing literature on collective action focuses on two questions. First, what are the conditions that lead to cooperation? And second, when is cooperation effective?

2.1.1 Cooperation

Collective action depends on the cooperation of group members, including contributions of resources. Olson's zero contribution thesis argues that cooperation in public goods settings is only possible in small groups or through coercion (Olson, 1965). When individuals can free-ride on the contributions of others without fear of punishment, the situation resembles a multiplayer prisoner's dilemma. A rational egoist has an equilibrium solution of zero contribution and should never cooperate willingly. But willing cooperation does sometimes occur in the world. Experiments show that people often behave as *conditional cooperators*, and *willing punishers* (Ostrom, 2000). Conditional cooperators are trusting when they expect that trust to be reciprocated. Willing punishers choose to punish norm violations even when the costs outweigh the benefits.

Why does cooperation occur? Real individuals are not always motivated by rational egoism. As Yochai Benkler writes, "there exist ranges of human experience in which the presence of monetary rewards is inversely related to the presence of other, social-psychological rewards (2002)." He cites examples including marriage, sex, and gift exchange. Ostrom (2000) proposes that the human deontic problem-solving system (reasoning about norms, guilt, and shame) has enabled us to adopt cooperative behaviors despite their apparent short-term irrationality. Ironically, these behaviors may be more rational than they seem when taking a long-term, large-scale perspective. In evolutionary models of a repeated prisoner's dilemma, communities of trusting individuals can out-perform less trusting individuals (Axelrod, 1997a). Ostrom suggests that the deontic problem-solving system allows members of a society to learn seemingly irrational cultural norms such as trust, when those norms allow the society to thrive in the long run.

2.1.2 Collaboration

While collaboration allows groups to apply their resources to problems that individuals cannot solve on their own, some of those resources must be spent on communication and coordination, contributing to *process loss* (Steiner, 1972), possibly limiting the usefulness of collaboration. Both organizational psychologists and economists have investigated when the benefits of collaboration outweigh the costs.

In a review of organizational psychology research, Hill (1982) concluded that groups typically performed better than individuals, while taking longer to find a solution. However, except in a few cases, statistical aggregates or high skilled individuals performed at least as well as groups. Exceptions were for difficult and complex tasks like crossword puzzles which no individual could solve on their own, as has also been seen in agent-based and theoretical results (Hong & Page, 2004). Hill attributed benefits of group work to process gain from social learning (see Section 2.6), but concluded that talented individuals will often out-perform committees when those committees have low-performing individuals. However, this conclusion ignores that difficult, complex problems may be a more realistic and more important scenario.

Informed by case studies of shared resources such as municipal water supplies, police departments, forests, and fisheries, Elinor Ostrom (2000; 2010) proposed the following set of design patterns for long-surviving, self-organized common pool resource governance regimes.

Boundaries. Clear boundaries delineating who is and is not a member of a group enable the development of within-group reciprocity and trust.

Rules for appropriation and provision. Rules reduce uncertainty by defining required contributions of inputs and the conditions for resource use. Rules are congruent with local conditions. Such rules are not necessarily formal.

Collective Choice Arrangements. When all individuals affected by rules are able to participate in creating them, rules are both better adapted to local conditions and have more legitimacy in the eyes of the group.

Monitoring of Users and Resources. Users select individuals to monitor both resources and users (including the monitors themselves) within the regime.

Graduated Sanctions. Sanctions for users who violate rules and norms begin light, but become more severe with repeated misconduct. Sanctions help to publicly reinforce norms as well as the trust that they will be followed. Furthermore, sanctions allow past transgressors to demonstrate cooperation and rebuild trust.

Conflict Resolution Mechanisms. Conflicts over interpretations of rules inevitably arise. Fast and effective means of resolving those conflicts help maintain trust.

Minimal Recognition of Rights. External authorities recognize the rights of the group. Without such recognition, an appeal to authority could be used as a threat to destroy the group.

Nested Enterprises. When resources exist at multiple scales, they are governed by nested organizations, each adapted to a particular scale.

Ostrom also identifies the following common anti-patterns that pose a threat to sustained collective action by creating unpredictability, eroding trust, or creating rules at odds with local conditions.

1. Migration,

2. Government imposed rules,

3. Rapid changes in technology,

4. Transmission failure,

5. Reliance on external aid,
6. External aid disconnected from local knowledge and institutions,
7. Corruption,
8. Conflict between regimes.

2.1.3 Types of Collaboration

An analysis of a process as complex as collective action could potentially focus on any number of factors. Hill (1982) suggests a two-dimensional framework based on *interpersonal context* and *performance evaluation*. Interpersonal context describes how individuals interact. Isolated individuals have no interaction, but their outputs might be aggregated, e.g., Galton's ox (1907). Coacting individuals interact freely and synchronously, as in a typical, face-to-face meeting. Dispersed interactions occur asynchronously, as is more typical of online collaborations (Benkler, 2002). Performance can be evaluated based on each individual, the best member's output, the union of all members output (e.g., brainstorming), or a single collaborative output.

Ostrom (2000) also identified several factors that might be relevant to the functioning of a specific collective regime.

1. Type of production and allocation functions,
2. Predictability of resource flows,
3. Relative scarcity of the good,
4. Size of the group,
5. Heterogeneity of the group,
6. Dependence of the group on the good,
7. Common understanding of the group,
8. Size of the total collective benefit,
9. Marginal contribution by one person to the collective good,
10. Size of the temptation to free ride,
11. Loss to cooperators when others do not cooperate,
12. Having a choice of participating or not,
13. Presence of leadership,
14. Past experience and level of social capital,
15. Autonomy to make binding rules.

So far, I have focused on collective action at a macro level. In order to understand the role of the internet in large-scale collective action, it is necessary to use a finer-grained analysis. Different stages of the collective action process (Figure 2.1) tend to be studied in different fields. Ideation is studied in organizational psychology, deliberation in sociology and political science, decision-making in economics and political science, and action execution in sociology, anthropology, public policy and epistemology.

This prelim focuses specifically on deliberation, so I expand my analysis of the deliberation stage to include both predictions of which actions will lead to which outcomes and individual preferences over those outcomes. Predictions can be influenced by information and reasoning, whether that information is gained through social interaction or gained through non-social means such as a search engine. Preferences can only be influenced through social, human-to-human *discourse* (Habermas, 1964; Geiger, 2009), for example, by creating identification with a community, emotional contagion, empathy, and so on (discourse will be described in more detail in Section 2.4.2). For example, when friends discuss which toppings to order on a pizza, the strength of the friendship may influence which toppings they are willing to consider, but it is unlikely to influence their prediction of what will arrive if they place an order for pineapple. These two effects of deliberation are particularly important to distinguish in an online setting, where interactions can range from non-discursive to highly discursive.

2.2 The Economics of Peer-Production

The internet has enabled near-instantaneous, many-to-many communication at a previously unknown scale. These affordances of the internet are promising not just for enabling larger collaborations, but because they enable collaborations that follow fundamentally different economics. Yochai Benkler (2002, 2006) labels this new form of economic production *commons-based peer production*, in contrast to more traditional forms of organization and production: firms, markets, and states (Coase, 1937; Ostrom, 2010). The new form is "commons-based" for two reasons. First, because outputs are contributed to the knowledge commons (a public good). Second, despite the non-rivalrous nature of the knowledge commons, there is incentive for free-riding and social loafing, analogous to the "tragedy of the commons" that occurs for common goods. The form is peer-produced because it is non-hierarchical: tasks are self-selected within a group of peers, rather than assigned from a superior to a subordinate. Benkler (2002) discusses examples including: NASA click workers, Wikipedia, the Google search engine, and the web forums Slashdot and Kuro5hin. Commons-based peer production faces the same motivational and organizational challenges as other forms of collective action, but offers both new resolutions to those challenges as well as novel advantages.

In addition to the usual motivations for non-rational-egoist behavior, commons-based peer production offers a unique solution to the problem of motivation: decomposability (Benkler, 2002). When tasks can be broken up into small pieces,

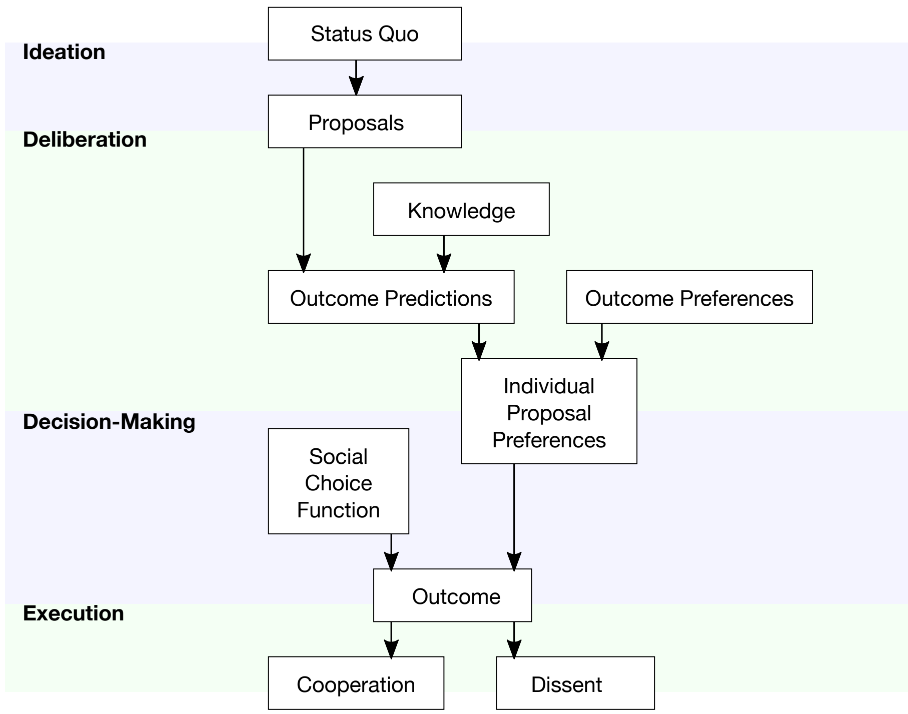

Figure 2.1: Collective action framework.

each of those pieces requires smaller amounts of motivation relative to tasks performed in a traditional firm. Furthermore, when those sub-tasks can be shared with a very large group, it is easier for individuals to identify the sub-tasks best matched to their particular level and type of motivation. Decomposability makes it possible for dispersed teams to collaborate on a problem (Hill, 1982). Dispersed interactions create additional coordination challenges: accreditation (quality control) and integration of outputs (Benkler, 2002). There are many possible approaches to accreditation and integration: hierarchical, norms-based, or aggregation/averaging, for example. Or, as in the case of peer review, these tasks can also be peer-produced. Peer-production of accreditation and integration is a key component of Benkler's formulation of commons-based peer production. Using Hill's classification (1982), commons-based peer production tasks use a group product evaluation criterion. Note that the dispersed "individuals" above could also be small, coacting sub-groups, allowing for a combination of dispersed and co-acting collaboration.

The need for accreditation and integration in peer-production suggests that it is not entirely non-rivalrous: if two incompatible outputs are produced by different sub-groups, at most one can be integrated into the groups final output. So while the ideation stage of peer-production is non-rivalrous, the decision-making stage is rivalrous. However, when the process of integration and accreditation is itself peer-produced, the final output is a hybrid: a rivalrous good produced through a non-rivalrous process. The dispersed decision-making process of accreditation and integration is a social learning process (discussed in more detail in Section 2.6). Commons-based peer production is thus an example of collective intelligence that combines non-rivalrous peer production with social-learning-based accreditation and integration (Figure 2.2).

Benkler (2002) identifies the economic conditions that favor peer production. Specifically, he notes that peering must be more efficient than market-based solutions (Coase, 1937) and that the cost of enforcing contract or property rights must outweigh the benefit (Demsetz, 1974). When do these conditions occur? Benkler attributes the benefits of peer production partly to larger groups having access to more resources, but primarily to allocation gains and information gains. He argues that, with the low cost of communication and storage, creativity is often the scarce resource, and that creative talent is variable and task-specific. The variability of creative work makes it difficult to allocate the right people to the right task. Peer production overcomes this difficulty by letting individuals self-assign. This benefit is amplified in the absence of property rights, which would otherwise limit individuals to working on tasks entirely within one firm. Benkler writes, "The widely distributed model of information production will better identify who is the best person to produce a specific component of a project, all abilities and availability to work on the specific module within a specific time frame considered." The lack of property rights also allows any combination of collaborators to form a sub-group to work on a sub-task, allowing increased functional diversity, which Hong and Page (2004) have shown is important for teams collaborating to solve difficult problems.

The relative value of peer production also depends on the specifics of the task

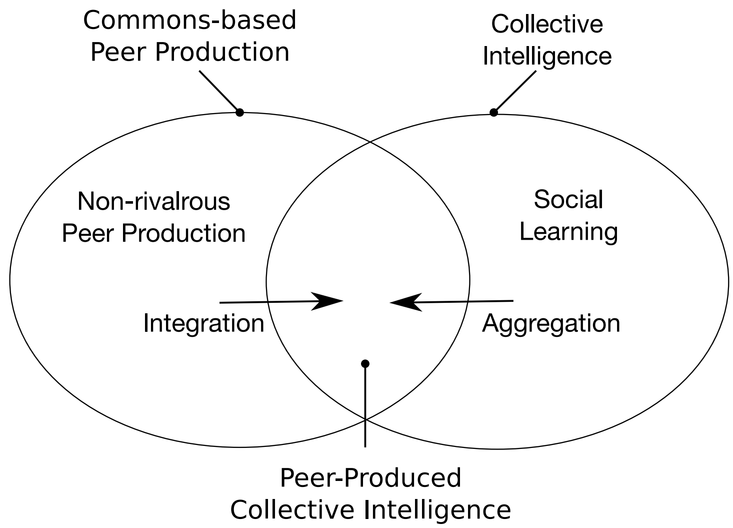

Figure 2.2: Commons-based peer production relies on social learning for accreditation and integration. Collective intelligence relies on peer production for ideation and preference aggregation.

being performed. Benkler writes, "peer production is limited not by the total cost or complexity of a project, but by its modularity, granularity, and the cost of integration (2002)." In Section 2.7, I will discuss insights on task complexity from collective intelligence and social learning, and how they might help inform internet-enabled collective action. Similarly, the cost of integration depends both on organizational networks and on networks of trust relationships. Hierarchical organizations have bottlenecks (i.e., managers and information brokers) who increase the cost of integrating ideas from different parts of the network. Given the need for organizational networks to reflect task structure (Conway, 1968), the more interdependent a task is, the more interdependent the organizational network will need to be. Interdependence creates complexity, raising the cost of integration. This cost can be mitigated in the presence of trust, specifically by the ability to predict the behavior of other members, as reflected by many of Ostrom's (2000) design patterns for self-organized collective action.

2.3 Social Choice Theory

One of the most fundamental challenges of collective action is the need to choose a single course of action when group members are in disagreement. This problem has been addressed by philosophers and mathematicians throughout history, and more recently by social choice theory, a subfield within microeconomics. Social choice theory focuses on the how groups choose between alternatives/candidates/proposals. In this paper, I will use the term "proposal."

When members of a group disagree on an issue, one approach is to put it to a vote. In certain contexts, voting can be quite effective. Assuming a binary choice with a correct answer, and decision-makers who independently choose the correct answer with probability p, the Condorcet jury theorem (Condorcet, 1785) states that when $p > 1/2$, the probability of the majority choosing correctly approaches 1 as the number of decision-makers grows. While a promising result, the assumptions made by this theorem are somewhat artificial. Many important decisions are "wicked problems," having no single correct answer (Hill, 1982; Anderson, 2006). Also, decision-makers are seldom independent (Hong & Page, 2009), and rarely homogeneous (Anderson 2006; Hong & Page 2004). I will return to these criticisms in later sections. A third difficulty forms the central focus of social choice theory: when more than two alternatives exist, there are many conflicting ways to compute a winner.

Simple majority and plurality voting are two straightforward voting methods which can give contradictory results. In plurality voting, the proposal with the most votes wins, even if it does not have a majority. In other words, it is possible for the majority of voters to disapprove of the winner. However, requiring a majority creates its own problems. When there are three proposals and no majority exists, supporters of the two less popular proposals can join together behind one of them allowing it to win, despite another being more highly preferred (Brandt et al., 2012).

More sophisticated approaches consider pairwise contests between proposals. For example, a Condorcet winner is a proposal that receives a majority of votes in pairwise contests against all other proposals (Condorcet, 1785; Brandt et al., 2012). It is difficult to argue against choosing the Condorcet winner, when such an alternative exists, but there is no guarantee that it will. The Condorcet paradox (Condorcet, 1785; Brandt et al., 2012) states that even when individual preferences are transitive, group preferences determined by pairwise contests can be intransitive. For example, if there are three alternatives (Rock, Paper, Scissors) it is possible that, in pairwise votes, Rock beats Scissors, Scissors beats Paper, and Paper beats Rock. In this example, none of the alternatives beat both of their competitors in pairwise votes. Similarly, while any choice between a finite set of proposals can be decomposed into a series of binary choices (i.e., option A or not, option B or not, etc.), such decompositions are path dependent, with the outcome depending on the order in which proposals are considered.

Voting methods can be formalized as *social welfare functions* in order to evaluate and compare them systematically (Arrow 1950; Brandt et al., 2012). Each voter i's preferences form a total order $<_i$ over all proposals. Social choice theory typically considers voting systems which depend only on the set of all voter preferences $\{<_i : i \in V\}$. This set is the *preference profile* (or just *profile*). The purpose of a voting system is to map the profile into a single social preference representing the group. Ties are allowed, but in order to avoid the Condorcet paradox, the social preference must be at least weakly transitive. The resulting social preference is thus represented by a weak order, and the voting system by a social welfare function mapping the profile to the social preference. Similarly, if only the first-place winner is needed, a *social choice function* maps each profile into a single winner (or set of winners in the case of a tie).

Using the above formalism, Arrow (1950) demonstrated that any social welfare function will be, in some sense, unfair. Specifically, Arrow's impossibility theorem states that no social welfare function can satisfy all of the the following three fairness criteria:

Weak Pareto efficiency. If A is preferred to B in all individual preferences, then A is preferred to B in the social preference.

Independence of irrelevant alternatives. Changing the position of any C within any individual preference will not change the order of A and B in the social preference.

No dictator. Group preference is not entirely determined by one member of the group.

Social choice theory has thus focused primarily on identifying desirable properties of voting systems and cataloguing the properties satisfied by various systems.

2.3.1 Voting Systems

Social choice theorists have studied many voting systems (Brandt et al., 2012). I will describe a few which are relevant to this prelim.

Condorcet method (Condorcet, 1785). A majority vote is conducted for each pair of proposals. If any proposal wins against all others, it is the winner Such a winner is not guaranteed to exist.

Dodgson's Method (Dodgson, 1876). An extension of the Condorcet method. For each proposal, the Dodgson score is the total number of adjacent pairs in the profile that need to be re-ordered to make that proposal a Condorcet winner. If a Condorcet winner exists, its Dodgson score is 0 and it wins. Calculating the Dodgson winner is NP-hard (Elkind & Slinko, 2016).

Borda count. (de Borda, 1781). Each proposal receives a score: the total number of proposals less popular than it over all individuals. This score is its Borda count, and the proposal with the highest Borda count is the winner. This winner can be found in polynomial time.

Copeland rule (Copeland, 1781). Majority votes are held between each pair of proposals. A proposal's Copeland score is the number of pairwise elections won minus the number lost. The proposal with the highest score wins.

Ranked Pairs / Tideman method. (Tideman, 1987). Pairwise majority contests are held. The results are ranked in order of descending margin. In this order, the winner of each comparison is placed above the loser in the social preference, unless doing so would create an intransitive cycle. The proposal with the highest social welfare value wins. Tideman can be interpreted as a greedy approximation of minimizing the dissatisfaction of voters with the social preference, and can be found in polynomial time (Meskanen & Nurmi, 2006).

2.3.2 Spatial Models

A noteworthy class of social welfare functions represents proposals and voter preferences as points in a mathematical space, and are thus called *spatial models*. One benefit of this representation is that positions can be compared using mathematically-defined distance measures.

The individual preferences within social choice theory can be compared using several distance measures. These measures can be used to define *distance-rationalizable* voting systems, which seek to minimize the winner's distance from some ideal position (Elkind & Slinko, 2016). For example, Dodgson's method minimizes the distance to a profile with a Condorcet winner.

The Kendall tau metric (Kendall, 1938; Brandt et al., 2012) defines the distance between two total orders (i.e., individual preferences) as the number of pairs for which the two orders disagree.

The Spearman correlation (Spearman, 1904) of two total orders is given by the linear (Pearson) correlation between the lists of rank positions. While Kendall tau takes only order into account, Spearman also takes into account the magnitude of rank difference.

An alternative approach to spatial modeling fixes voter preferences and seeks a dominant strategy for selecting a proposal with the highest chance of winning. Such a model could, for example, represent a political candidate choosing a platform based on their perception of voter preferences and other candidates' positions. In the simplest models, voter preferences are represented as a continuous variable in one dimension (e.g., liberal-conservative). Davis et al. (1970) consider both single-dimensional and multi-dimensional models. For single-peaked preferences in one dimension, the dominant strategy is the position of the median voter. For preferences with more than one peak, it is possible that no dominant strategy exists. For spatial models with multiple peaks or multiple dimensions, the best strategy can exhibit path dependence, depending on which order proposals are defined. Davis et al. also consider the influence of voter abstention on strategy, assuming that alternatives may lose votes to abstention for two reasons. First, voters may abstain if they are *indifferent* between two outcomes. Second, voters may abstain if they feel *alienated* from the political process and unable to influence the outcome. Under these assumptions, moderate positions can be effective at shifting extreme positions toward the center. It is also possible for an alternative to win over another with higher support if support for the majority position is more dispersed in preference space. In practice, voters' ideal points in preference space are latent but can be inferred from behavioral signals such as past votes (Clinton & Rivers, 2004).

Returning to the problem of aggregating preferences, the median voter can be useful even in multiple dimensions. The median can be generalized to multiple dimensions as the point which minimizes the sum of distances to all other points. If pairwise distances between voters form a *median graph*, then a series of 3-way votes can efficiently find the median (Goel & Lee, 2012; 2016). Furthermore, the median is either the Condorcet winner (if it exists) or a good approximation. Goel and Lee (2012) propose an efficient method to find an approximation of the generalized median. The Goel-Lee method chooses a triad of 3 voters at random, then asks each to vote between the top preferences of the other two. All three are then required to vote for winning proposal in the next time they vote. The population converges to the generalized median in $O(n \log^2 n)$ triads.

2.4 Discourse, Deliberation, and Democracy

While the paradoxes of social choice theory seem to bode ill for participatory decision-making at large scales, there is more to decision-making than voting. Voter preferences are not fixed, but influenced by public opinion and discourse. The ability to change preference profiles through deliberation suggests the possibility of resolving conflicts prior to the decision-making stage. The ability of the internet to enable discourse and deliberation at a large scale suggests it might make new forms of decision-making, collective intelligence, and therefore collective action, possible. Discourse, deliberation, and voting are distinct but intimately related components of participatory decision-making. Group members share information and opinions through discourse, persuade through deliberation, and aggregate preferences through votes.

2.4.1 Democracy

Democracy is delineated from other forms of collective action by universal participation without the need for actions to be universally supported. It relies on two institutions: votes and talk (Anderson, 2006). Talk includes both the discourse of the public sphere as well as the deliberation of formal and informal bodies. Votes are necessary both to identify winning proposals to give them legitimacy.

In contrast to authoritarian governance, there are several arguments for democracy. Centralization limits the ability of decision-makers to use widely-dispersed information, while democracy enables collective intelligence (Hayek, 1945). Procedural arguments favor democracy as more fair, while epistemic arguments favor it as more effective (Anderson, 2006). Anderson gives three conditions necessary for democracy to be an option: 1. A matter of public interest, 2. Necessity for joint action, and 3. Reliance on the law. Because cooperation is required despite a lack of universal agreement, democracy relies on some level of coercion of dissenting group members through the law.

Talk (discourse and deliberation) plays several roles in democracy. At the ideation and deliberation stages of collective action, the role of talk is to generate and evaluate ideas. This process consists of discourse within and communication between "multiple, cross-cutting organizations" (Anderson, 2006). In "strong democracy," participation is emphasized over voting (González-Bailón, 2010). By doing so, conflict resolution is shifted from coercion of dissenters at the execution stage to building agreement at the deliberation stage. For democracy to function, members do not need to agree with the group's decision, but they do need to cooperate. Dissent, by prompting deliberation, can secure the cooperation of dissenting individuals through "mutual accommodation" (Anderson, 2006).

2.4.2 Discourse and the Networked Public Sphere

As public discourse moves from older broadcast media and onto the internet, its context has changed. Yochai Benkler coined the term networked public sphere (Benkler, 2006) to refer to the unique affordances the internet offers for public discourse, building on sociological concept of the public sphere (Habermas, 1964). The public sphere refers to the social communication conducted by private individuals when they assemble into a public body, for example: newspapers, radio, and television.

The networked public sphere is notable for the affordances offered by the internet: scale, speed, decentralization, and low cost of speech (Benkler, 2002; 2006). The internet operates at a global scale, allows instantaneous publication, is difficult for any single entity to control, and allows the entire public the opportunity to speak without large investments of capital. However, the internet also faces challenges as a medium for pubic discourse, including: fragmentation, emergent centralization, authoritarian filtering, replacing traditional media's role as a watchdog, and the digital divide.

The wide participation enabled by the large scale of the internet also creates the potential for fragmentation of the public sphere into isolated communities (Geiger, 2009; Benkler, 2006). There have been two common solutions for integrating the discourse of very large groups of people online. The first is algorithmic. For example, the Google search engine aggregates data from sources across the internet to gauge the quality of content (Benkler, 2006). However, algorithmic integration is insufficient for creating a public sphere because it is not discursive: running an algorithm is not a social communicative act between humans (Geiger, 2009). The second resolution is the use of technology to enable discursive interactions between large groups, such as reputation systems, peer review, and structured posting privileges (Benkler, 2006).

Discursive integration of peer-produced content depends crucially on the ability of the internet to amplify salient information to a wide audience while filtering out irrelevant information. When the amount of information online is far greater than any one person can sift through, how can this amplication and filtering happen without resorting to non-discursive algorithmic solutions? The answer is again, decentralization. By participating in small topical and organizational groups, individuals have venues to send and receive information relevant to that group and anyone in it. And by participating in multiple such groups, information can quickly spread through the communities where it is relevant. For example, information about a new species of insect might travel from a university email list, to a forum for insect enthusiasts, to a blog for exterminators. This is what Benkler (2006) refers to as a "filtering and transmission backbone." Geiger (2009) describes the process: "Filtering, accreditation, synthesis, and salience are created through a system of peer review by information affinity groups, topical or interest based. These groups filter the observations and opinions of an enormous range of people and transmit those that pass local peer review to broader groups and ultimately to the polity more broadly." Such systems allow filtering and synthesis through trial-and-error rather than through careful planning; through evolution rather than engineering. Such systems are particularly applicable when errors can recovered from relatively easily.

The structure and protocols of the internet were designed to be decentralized (Baran, 1964), and while necessary for a decentralized sociotechnical system, decentralized architecture and protocols are not sufficient. Centralization can emerge within a decentralized framework due to social, organizational, and economic processes. In fact, the internet exhibits centralization on the levels of router connections and web links (Albert et al., 2000) as well as market share (Noam, 2003). However, this centralization differs from the centralization of traditional media in important ways. Small clusters of sites exceed the activity predicted by power-law distributions (Pennock, 2002). These overlapping topical and organizational clusters create the "filtering and transmission backbone" (Benkler, 2006). This backbone is bidirectional, allowing utterances from peripheral parts of the network to quickly reach central sites for amplification. This backbone also contains redundant paths not typical of more traditional media, allowing utterances many opportunities for both filtering and amplification. As long as communication is redundant and bidirectional, emergent centralization on the internet is not at odds with public discourse.

The existence of overlapping topical, organizational, and geographic clusters plays an important role in civic participation and discourse (Putnam, 2000; Anderson, 2006; Geiger, 2009), in part by counteracting fragmentation and centralization (Benkler, 2006). When individuals from different clusters interact, there is a potential for a clash of culture and norms. Norris and Inglehart (2009) describe four possible scenarios for cross-cultural communication. In the L.A. effect, a dominant cluster exports its own cluster to others. In the Bangalore effect, cultures from different clusters mix and coexist. In the Taliban effect, conflicting cultures polarize against each other. Finally, in the firewall effect, different cultures simply ignore each other despite the ability to communicate. While the L.A. and Bangalore effects are conducive to a networked public sphere, the others are not. For internet-enabled collective action to be possible on a large scale, technology must conflict resolution through communication, rather than through polarization or isolation.

2.4.3 Deliberation

Deliberation, in contrast to voting, is social communication intended to change the preferences of the participants (González-Bailón et al., 2010). This might occur through social learning (Hill, 1982) or through changes in preferences created

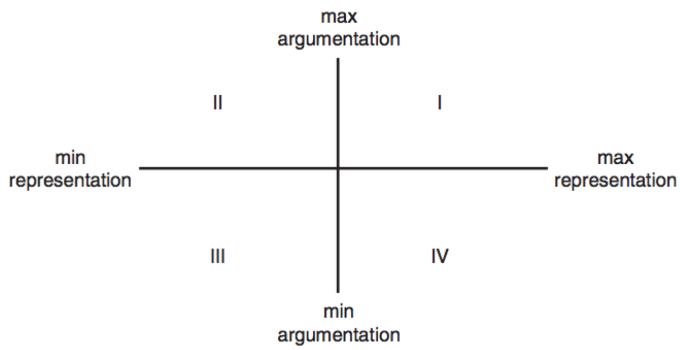

Figure 2.3: Types of deliberation. Adapted from (Ackerman & Fishkin, 2002).

by the act of deliberation itself (Anderson, 2006). In other words, deliberation can alter preferences towards actions by altering the predicted outcomes of those actions, or by changing the preferences of the outcomes themselves (Figure 2).

In "strong democracy" there is universal participation in deliberation, not just in voting (González-Bailón et al., 2010). Deliberation can be classified on two axes: representation, and argumentation (Figure 2.3) (Ackerman & Fishkin, 2002). An editorial is limited in representation and argumentation. A televised debate is limited in representation, but high on argumentation. A national poll is high on representation but low on argumentation. "Constitutional moments" when the population is widely debating a topic fall into the category of high representation and high argumentation, but have traditionally been infrequent and short-lived. Deliberative democracy requires both representation and argumentation, but the two have been difficult to achieve simultaneously, at least with traditional media. Examples of internet-based deliberation, such as Wikipedia (Benkler, 2006) and Slashdot (González-Bailón et al., 2010; Benkler, 2006) have been able to achieve both.

Deliberative democracy appeals to procedural arguments for democracy. However, the epistemic merits of deliberative democracy depend on context. Deliberating groups do well on "eureka problems," in which solutions can be easily recognized once they are found, but less well on wicked problems (Gigone & Hastie, 1997). When groups self-organize into homogeneous groups, deliberation can increase polarization (Schkade et al., 2007) thereby decreasing functional diversity, which is necessary for improved team performance (Hong & Page 2004). Effective deliberation requires careful attention to both task type and the human tendency towards homophily.

2.4.4 Consensus

There is no consensus on the definition of "consensus." DeTar (2013) classifies several forms of consensus decision-making according to whether they have open membership, egalitarian, formal procedures, binding decisions.

Corporate (nemawashi). Consensus-building informally before formal decision-making. This version can be considered a type of deliberation.

Scientific. A preponderance of agreement among researchers.

Standards. A non-binding agreement between stakeholders.

Consociationalism / Factional. Underrepresented factions are given representation in a formal decision-making body. All formal decisions must be passed by unanimous consent.

Mob. Emergent, unstructured coordination, such as flocking or rioting.

Assembly. Formal body with open membership. All members may or may not have veto power.

Affinity. Small-scale informal decision-making in a closed group. All members have veto power.

Studies of consensus in social and organizational psychology have typically focused on affinity consensus, categorically excluding formal process (Gentry, 1982). Findings suggest that affinity consensus requires more time than voting, but produces better solutions (Nemiroff & King, 1975). When trained in the following affinity consensus principles, group members are more likely to generate new, emergent solutions combining the resources of many group members (Hall & Williams, 1970).

- Avoid arguing for your own position.
- Avoid win-lose stalemates in discussion.
- Avoid changing your mind only in order to avoid conflict and to reach agreement and harmony.
- Avoid conflict-reducing techniques such as majority vote, averaging, bargaining, coin flipping and the like.
- View differences of opinion as both natural and helpful rather than as a hindrance in decision-making.
- View initial agreement as suspect.

These findings suggest that epistemic arguments for democracy based on collective intelligence might better apply to consensus.

While "consensus" connotes agreement, many groups using affinity consensus do not require unanimous agreement. Quakers, for instance, use consensus to seek the "Spirit of God" and the "sense of the meeting" (Gentry, 1982). So consensus implies recognition of the conclusion reached by a group, but not necessarily agreement with that conclusion.

2.5 Conflict and Influence

If deliberation can be used to reduce conflict before decision-making, it is necessary to understand the social processes of conflict and influence. In representative democracy, focus is shifted from discourse to voting, Enabling scalability at the cost of civic participation and the risk of the Condorcet paradox. Internet-enabled peer production has been able to overcome scaling difficulties while maintaining (or growing) discourse and participation. And as seen in the previous section, conflict can be a sign of functional diversity that enable emergent solutions. But what are the elements of large-scale collective actions that are able to resolve conflict?

The successful resolution of conflict often requires trust and cooperation. In the simple example of the prisoner's dilemma, the dominant strategy for a single round game is to defect rather than cooperate, despite this strategy yielding the outcome with lowest social welfare. However, in an evolutionary model with repeated games, communities can develop norms regarding trust, allowing them to cooperate, resolve conflict, and achieve higher social welfare (Axelrod, 1997a; Ostrom, 2000). In this context, "a norm exists in a given social setting to the extent that individuals usually act in a certain way and are often punished when seen not to be acting in this way (Axelrod, 1997a)." Many of Ostrom's design principles for self-organized collective action concern trust and norms, such as clear group boundaries (to identify who the norms apply to) and graduated sanctions (to reinforce norms) (Ostrom, 2000). In addition to norms, communities can exhibit metanorms: norms regarding the enforcement of norms. When metanorms exist to punish individuals for failing to punish transgressions, genetic algorithm simulations suggest that communities can converge to a state where transgressions are rare (Axelrod, 1997a). Internal enforcement, i.e., guilt and shame, also play an important role in influencing behavior. Guilt relies on an internalized sense of allowed behavior, while shame relies on "social proof", the visibility of actions to others (Axelrod 1997a).

While dominance hierarchies are not compatible with collective action, there are non-hierarchical forms of dominance which might be. Using "leveling mechanisms" such as public opinion, ridicule, and disobedience, some small human societies are effective at "keeping aggressive and dominating individuals in check" (Boehm et al., 1993). Leaders in such societies are often described as being "a first among equals." It is not that dominance is non-existent in such societies, it is that it is non-hierarchical. *Reverse dominance hierarchies* use dominance to reduce inequality. Similarly, in "heterarchies" individuals are simultaneously subordinate and superordinate to many others, enabling the use of dominance for conflict resolution without creating a single hierarchy of power (Sharp, 1958; Tonkinson, 1988; Crumley, 1995).

2.6 Social Learning

A crucial part of collective intelligence is the ability of group members to learn from each other. In social learning, individuals receive different information and are each able to communicate with some subset of others. There are many types of social learning, depending on the task being learned, the type of interaction between group members, the network structure of who communicates with whom. Social learning has been extensively modeled both mathematically and using agent-based models.

In the simplest learning problems, agents receive a *generated signal* (Hong & Page, 2009): a measurement of some true value with independent random noise. For example, Galton's ox (Galton, 1907) could be fall into this category if one assumes that errors in guessing are independent and effectively random. The Condorcet jury theorem also falls into this category (Condorcet, 1785). Such problems can be solved by repeated averaging of information with neighbors (DeGroot, 1974). This approach converges to the true value as long as the communication is regular: all individuals have influence on the same number of neighbors (Golub & Jackson, 2012). When individuals have different influence, the process experiences "persuasion bias" and the noise of the most influential individuals is interpreted as signal. The speed of convergence also depends on the communication structure, specifically the second eigenvector of the interaction matrix, which quantifies the existence of bottlenecks in the network (Golub & Jackson, 2012), with higher connectivity producing faster convergence (Zollman, 2010). In the somewhat more realistic "bounded confidence" model, agents average opinions with neighbors only if those opinions differ from their own by less than a "tolerance parameter" (according to a spatial model). In this model, moderately tolerant groups benefit from higher connectivity, while low-tolerance groups benefit from central individuals who can act as bridges (Zollman, 2012). While simple social learning tasks such as these are relatively well-understood, they are not particularly good models for real-world problems. Most problems in life are complex.

Complex social learning problems are composed of interdependent sub-problems. Interdependence results in many locally optimal solutions when sub-problems are solved in isolation. There may be more sub-problems than any individual agent can consider. In this case, agents are in an *intepreted signal* regime, and must take a *perspective* by focusing on some subset of the full problem and use *heuristics* to make decisions with limited knowledge (Hong & Page, 2009).

Complex social learning problems can be modeled as optimization problems. Agents search a space of parameters for one that optimizes an objective function. In "rugged landscapes," the objective function has many local maxima, making search by gradient ascent impractical. Such landscapes can be created using the NK model (Kauffman, 1987). In this model, states have N dimensions or loci, with each dimension having a finite number of possible values (e.g., 0 or 1 in the binary case). Associated with each locus is a function F_i depending on that locus and K others. The F_i are randomly assigned values in $[0, 1]$ for each possible value of locus i and its K neighbors. An objective function is created by averaging all of the F_i. By varying K, the ruggedness of the objective function can be tuned.

Agent-based models of complex social learning problems have addressed the role of network structure and social learning strategies in complex problems. Agents can employ a range of behavioral strategies in a social learning setting. In conformity-based strategies (Mason & Watts, 2012), an agent adopts the solution most popular among its neighbors. In best-neighbor strategies (Lazer & Friedman, 2007; Grim et al. 2013), an agent evaluates the solutions of their neighbors and adopts the best one. Models exhibit a tradeoff between exploration of new solutions and exploitation of known solutions depending on both network structure (Lazer & Friedman, 2007) and social learning strategy (Barkoczi & Galesic, 2016). Such a tradeoff is a characteristic of real-world complex problems (March, 1991). A similar tradeoff exists for performance and efficiency (Platt & Romero, 2018). Sparsely connected networks result in parallel problem solving and greater functional diversity (Lazer & Friedman, 2007). Conformity-based strategies also enable greater functional diversity by allowing popular but sub-optimal solutions to persist (Barkoczi & Galesic, 2016). Combinations of network structure and learning strategy which balance exploration and exploitation yield the best long-term performance (Barkoczi & Galesic, 2016).

In a real-world setting, the importance of network connectivity can be interpreted as the importance of intersecting and cross-cutting social groups (Anderson, 2006; Benkler 2006; Putnam, 2000). The importance of social learning strategy has implications for the effectiveness of discourse decision-making, which explores many new and existing solutions, versus voting which singles out the best existing solution. However, it is important to note that the social component of collective action is not simply a matter of exchanging information, but also a matter of building trust and mutual accommodation (Anderson, 2006), which are not reflected in current agent-based models.

2.7 Networks, Teams, and Tasks

2.7.1 Decentralization and Non-hierarchicalism

The potential benefits of internet-enabled collective action stem largely from allowing groups to grow in size without requiring centralization. "Decentralization" has long been a buzzword when discussing the internet, but the term has taken on several distinct meanings. It is necessary to consider precisely in what sense internet-enabled collective action is decentralized and what the affordances of that type of decentralization are.

The internet grew out of research on decentralized communication networks (Baran, 1964). The goal of this early research was to introduce redundancies so that the network could not be disabled by nuclear attacks on central hubs. Baran described three architectures: centralized, decentralized, and distributed (Figure 2.4). In Baran's taxonomy, centralized networks have a single central

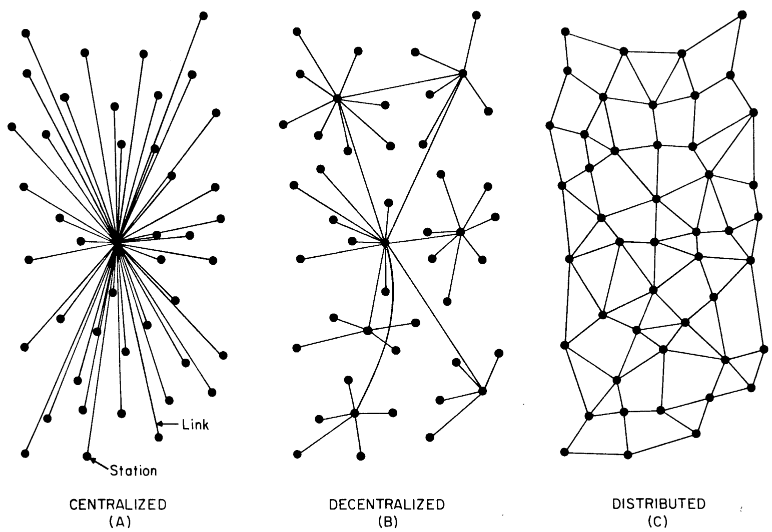

Figure 2.4: Centralized, decentralized, and distributed network topologies. Adopted from (Baran, 1964).

hub that intermediates between all other nodes in the network (for example, the local telephone networks of the time). Decentralized networks are then any network that does not have a single central hub, but might still have a division between endpoints and hubs. In the extreme, distributed networks (e.g., mesh networks) have all nodes play equivalent roles, acting both as end-points and intermediaries. In this taxonomy, decentralized networks exhibit two distinct properties. First, there is more than one intermediary node. Second, endpoints can be connected by multiple redundant paths, passing through different intermediaries. For example, a corporate organizational hierarchy exhibits the first type of centralization, but not the second: middle-managers exist as additional intermediaries, but there is exactly one chain of command from the CEO down to each employee. So it is important to distinguish not just between centralized and decentralized networks, but also between hierarchical and non-hierarchical networks.

To further complicate matters, the term "hierarchy" is used in multiple ways. Crumley (1995) distinguishes between scalar and control hierarchies. Elements in a scalar hierarchy are composed of sub-components, which are themselves composed of sub-sub-components, and so on (e.g., state, city, address). Control hierarchies on the other hand, involve dominance relationships between components; each relationship is composed of one subordinate and one superordinate element. So the divisions, departments, teams, and individuals of a corporate firm form a scalar hierarchy, while the CEO, regional managers, local managers, and individual contributors form a control hierarchy. It is worth emphasizing that the former is possible without the latter, as in the case of self-managed organizations (Gray et al., 2015) and heterarchies (Crumley, 1995).

The existence of multiple types of decentralization implies that a nuanced approach is necessary to study decentralized collective action. Sociologists and network theorists have developed many techniques for discussing decentralization more precisely. The presence of structural holes in communities allows them to be divided into smaller sub-communities (Burt, 2009; González-Bailón & Wang, 2016). Those structural holes can in turn be spanned by liaisons/bridges (Rappoport, 1963). Granovetter (1974) proposed that weak ties, individuals with a weak connection and no common friends, play an important role as bridges. He formalized a bridge as a connection that disconnects the entire network if broken, and an n-local bridge as a connection between two nodes whose next shortest connecting path is length n. Liaisons, bridges, and weak-ties have the potential to act as brokers and gatekeepers of information (González-Bailón & Wang, 2016), a form of social capital (Putnam, 2000). The existence of multiple paths in a decentralized network allows for transmission of conflicting information and attitudes. Balance theory (Rappoport, 1963) suggests that in social networks, when two ties reflect conflicting attitudes, only one will survive, placing constraints on the types of networks that can remain stable.

Focusing on how to achieve non-hierarchical decentralization, DeTar (2013) describes several means:

Replication. Elements are copied and exist independently. For example, the free/open-source Wordpress blogging software can be considered decentralized in the sense that anyone can run a copy independently on their own server.

Federation. In a federated system, such as email or the Mastodon social network, all elements are independent but interoperable. Different email servers can have different owners, run different software, and exist in different locations, but still communicate via a shared protocol.

Subsidiarity. Systems providing different functionality exist independently. For example, embedding a YouTube video on a Wordpress blog that allows users to log in through their Mastodon account.

The above examples raise another point: decentralization in one domain typically relies on centralization in another (DeTar, 2013; Wilcox-O'Hearn, 2001). For example, the internet's interconnection network is decentralized by relying on a centralized protocol. This centralization/decentralization relationship suggests that some form of centralization is unavoidable. Due to this constraint, Zooko's triangle (Wilcox-O'Hearn, 2001) conjectures that any network protocol can be, at best, two of the following: decentralized, secure, human-readable.

Although less common than centralized, or hierarchical-decentralized organizations, non-hierarchical decentralized regimes can be found in the world today. Elinor Ostrom (2010) has studied several such organizations, including municipal water and police departments, dubbing this type of organization polycentric. In contrast to the assumption that non-hierarchical organization necessarily results in chaos and disorder, Ostrom found that under certain conditions, polycentric governance can take advantage of economies of scale while avoiding diseconomies. She attributed these advantages to three mechanisms: 1. small-to-medium-sized cities are often more effective at monitoring their own performance, 2. citizens who are dissatisfied can vote with their feet by moving into another group, and 3. local incorporated communities can contract with larger producers. Ostrom also noted the role of polycentric governance structure in promoting cooperative behavior. Rational choice models, often used to explain non-cooperation, do not take communication or individual agency into account. As such, Ostrom proposes the following conditions for cooperation within polycentric governance:

1. Communication feasible with all participants,
2. Reputations are known,
3. High marginal per capita return,
4. Ability to enter/exit,
5. Longer time horizon,
6. Agreed-upon sanctioning.

Based on observations of polycentric governance regimes, Ostrom summarizes: "complexity is not the same as chaos in regard to metropolitan governance." In this prelim, I investigate some of the questions necessary to extend Ostrom's observation to internet-enabled collective action.

2.7.2 Task Structure

While social network structure is well-studied, the structure of tasks themselves is less well studied. It may not be obvious that tasks can have structure or what that means. As in social learning, one dimension of task structure is interdependence: how effectively can sub-parts of a problem be solved independently? Low-interdependence and task granularity, are key component of commons-based peer production (Benkler, 2002; Kittur & Kraut, 2008). The granularity of a problem relative to the number of problem solvers also influences whether it is best modeled as a generated signal or interpreted signal problem (Hong & Page, 2009). When a problem can be considered on many dimensions, problem solvers are more likely to focus on different subsets up the problem, as in interpreted signal problems.

Centola and Macy (2007) discussed how the role of network structure is modulated by task structure. They distinguish between simple contagion (information diffusion) and complex contagion. In the former, only a single exposure is necessary for information to propagate, while in the later, multiple exposures are necessary. Such might be the case when individuals are looking to others for social signals about the relative risk or safety of adopting an innovation. They note that long ties are unable to spread complex contagions, despite their effectiveness at propagating information such as job opportunities (Granovetter, 1974).

In the organizational psychology literature, Hill (1982) identifies several classes of group tasks:

Learning Examples include: memorization, concept-attainment/classification.

Mastery/creativity. Examples include: learning the meaning of words and identifying synonyms/antonyms.

Abstract problem solving. Examples include: completing partially written poems, finding efficient travel routes connecting multiple points.

Brainstorming. Examples include: thinking of proposals to ensure schools perform well as enrollment increases.

Complex problem solving. Examples include: a manager incorporating conflicting suggestions on ways to improve the work process.

These task structures appear to have implications for how groups collaborate. In any analysis of collective action, and in the design of internet-based tools for collective action, it will be necessary to be aware of the nature of the task being approached and how it differs from any existing examples.

2.7.3 Organizational Networks

Organizations and firms with formally defined roles and relationships provide a context in which network structure is at least partially visible and can be deliberately manipulated. One particularly interesting observation regarding organizational networks is Conway's law (Conway, 1968) which says that the organizational structure of complex systems reflects the social structure of the organization that built them (Figure 2.5). This observation stems from the fact that when sub-tasks are assigned to groups, groups that are disconnected from each other will not be able to solve sub-tasks that are highly interdependent. In Conway's words, "Given any design team organization, there is a class of design alternatives which cannot be effectively pursued by such an organization because the necessary communication paths do not exist." For collective action, the implication is that large interdependent problems are not solvable without a way for large groups to coordinate and cooperate.

This prelim is largely concerned with the potential for non-hierarchical, decentralized collaboration enabled by the internet, but this should absolutely not be confused with structurelessness. In a classic essay, feminist scholar Jo Freeman wrote about the dangers of (apparent) structurelessness (Freeman, 1972). Freeman argues that all groups have structure, whether it is formal or hidden. Hidden structure can be used to mask power and make participation closed to the greater group, both of which can limit the ability of groups to take full advantage of their resources. Regarding the scaling of communication requirements with group size, Freeman writes, "This inevitably limits group participants to

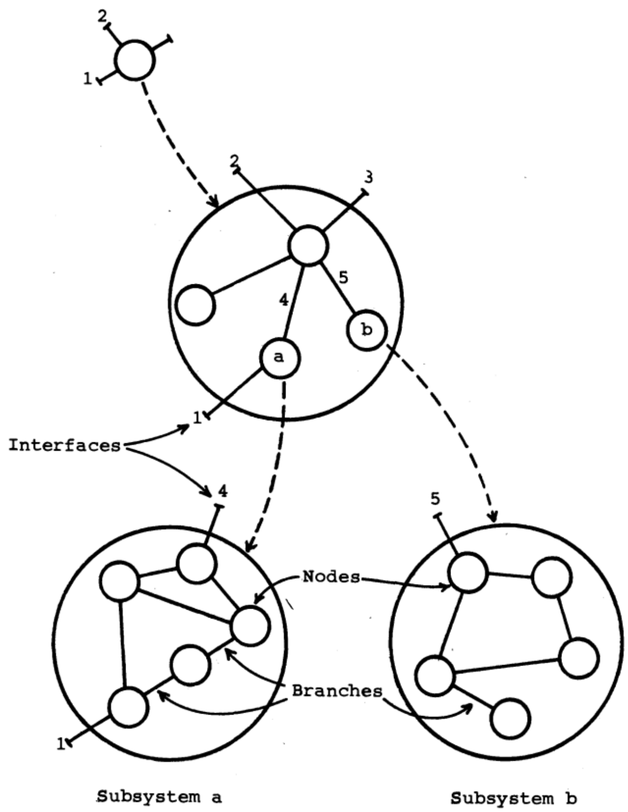

Figure 2.5: Conway's law: "There is a very close relationship between the structure of a system and the structure of the organization that designed it." Adapted from (Conway, 1968).

about five, or excludes some from some of the decisions. Successful groups can be as large as 10 or 15, but only when they are in fact composed of several smaller subgroups which perform specific parts of the task, and whose members overlap with each other so that knowledge of what the different subgroups are doing can be passed around easily." While Freeman seems to make the common assumption that structure implies control hierarchy, I argue in this prelim that through concepts like heterarchy (Crumley, 1995), explicit structure can exist in large-scale, non-hierarchical, decentralized collective action.

2.7.4 Team Composition

The success of teams can also be influenced by the personal characteristics of the team members. The Diversity Trumps Ability theorem (Hong & Page, 2004) describes the conditions under which a diverse group of individuals can outperform a homogeneous group of higher-performing individuals. Specifically: 1. agents must be intelligent, 2. the problem must be difficult such that no individual can solve it on their own, 3. the agents must be diverse, and 4. the best agent must be unique. Hong and Page also used an agent-based model to verify that a random sample of agents can outperform a team of the best agents. The social dynamics of teams must also be taken into account. Hill (1982) noted that interactions between team members were influenced by social roles such as gender. Finally, levels of inequality can influence a team's ability to work together. Nishi et al. (2005) found that when simulated wealth levels were visible in a prisoner's dilemma game, participants wealthy participants took advantage of the less wealthy, leading to a rich-get-richer dynamic, lowering both cooperation and overall group performance. Team success is influenced by many individual characteristics of the team members, and even by the visibility of those characteristics.

Chapter 3

Case Studies

3.1 Anthropological Studies

Anthropologists have identified and studied several human societies that operate cooperatively, including the Yir Yoront (Sharpe, 1958) and the Mardujarra (Tonkinson, 1988), both among the indigenous peoples of Australia. One common feature of such cooperative peoples is mutual dependence. For example, when scarce resources and unpredictable weather make cooperation necessary for survival. Mutual dependence can be formed by two-way flows of obligations (Tonkinson, 1988). Or mutual dependence can be the result of complex non-hierarchical structure, as Sharpe (1958) describes: "A [Yir Yoront] man has no dealing with another man (or woman, either) on exactly equal terms. And where each is at the same time in relatively weak positions and in an equal number of relatively strong positions, no one can be either absolutely strong or absolutely weak. A hierarchy of a pyramidal or inverted-Y type to include all the men in the system is impossible." In the Yir Yoront, these structures are based on geographically dispersed kinship obligations. Just as in commons-based peer production, cooperation is achieved through dispersed collaboration and overlapping sub-groups (Benkler, 2002), forming a heterarchy (Crumley, 1995). These non-hierarchical structures allow conflict resolution through reverse dominance hierarchies (Boehm et al., 1993) based on norms and deontic decision-making, just as observed in Western cooperative regimes (Ostrom, 2000).

3.2 Experimental Studies

Small scale networked social learning experiments have shown that networked groups can solve very difficult problems, but the relationship between network structure, task-type, and group performance is still not fully understood. In a series of lab experiments, Kearns (2012) has found that very difficult problems, such as graph coloring and unanimous consensus, can be solved by networked groups. In graph coloring tasks, the performance of low-connectivity cycle networks was improved by adding shortcuts to increase connectivity, while the same changes made it more difficult for groups to solve a consensus task. Mason & Watts (2012) similarly found that high-connectivity, efficient networks were preferable for experimental groups of problem solvers seeking to maximize an NK model objective function. In one notable exception, Kearns (2012) found that self-organized groups tended to perform poorly in networked bargaining experiments, and that these networks were differentiated from others by highly skewed degree distributions.

A novel form of large-scale survey, the wiki survey, has been demonstrated to be effective at eliciting both preferences and new ideas (Salganik & Levy, 2015). Participants are presented with pairs of proposals, and prompted to either choose one or submit a third alternative. The granular, dispersed nature of the task is more amenable to peer production (Benkler, 2002) and adapts to collect as much or as little information as participants are willing to provide. Salganik and Levy (2015) partnered with New York City and the OECD to conduct trials of approximately 1,500 participants and 30,000 pairwise comparisons each. Over the course of the surveys, the number of unique proposals increased between 5 and 10 times. New proposals included new ideas, but also re-framings of existing ideas. These findings suggest a combination of social learning and discursive processes. The wiki survey is a primary inspiration for the experiment in this prelim.

3.3 Wikipedia

The persistent record of activity on Wikipedia and other wikis has made it possible for researchers to analyze the evolution of editor behavior over long timescales. On Wikipedia, several researchers have noted changes in the community behavior starting around 2004 and ending around 2007, the so-called "golden era" (Keegan & Fiesler, 2017; Kittur et al., 2007, Forte & bruckman, 2008). In an early study of Wikipedia, Viégas & Wattenberg (2004) found many of Ostroms (2000) design principles for self-organized cooperation. Policies such as "neutral point of view (NPOV) are visibly recorded and actively discussed by any interested contributors. History logs allow for monitoring. The ability to revert harmful edits means no permanent damage can be done, making graduated sanctions feasible. A later study (Viégas et al., 2007) found that criteria had become stricter, e.g., for "featured" articles, and that automated tools were performing some of the quality control and accreditation tasks originally performed by humans. However, the same study noted that editors had developed templates allowing quality control issues to be flagged and addressed at different times by different people, enabling dispersed collaboration. These developments are signs of a change in the nature of decentralized work on Wikipedia.

Based on interviews, Forte & Bruckman (2008) describe a shift from individual decision-making to deliberative and committee-based decision-making, and conclude that Wikipedia has become less centralized. Similarly, Kittur et al. (2007) found a shift from direct work on articles to indirect work on policy and anti-vandalism. Keegan & Fiesler (2017) examined the edit history of Wikipedias policy pages, finding that after 2007, edits to policy pages declined in favor of discussion on policy talk pages. Kittur et al. (2007) found that an increasing fraction of edits were made by newer users, but that content was primarily created by long-standing users. Similarly, in a study of Wikia wikis, Shaw and Hill (2014) found that it became more difficult for new users to become editors, that existing admins used their powers more, and that new users were more likely to have their edits reverted. These studies indicate ways in which wikis become both more and less centralized over time. Overall, Wikipedia has responded to an influx of users by creating a more decentralized structure based on committees, while keeping power somewhat centralized among earlier editors through less policy flexibility, a high bar for creating new content, and higher inequality between formal admin positions and non-admin editors.

Wikipedia has also been studied to gain a better understanding of the factors that influence teamwork at a large scale. Kittur & Kraut (2008) quantified implicit and explicit coordination (gini coefficient of editors edit counts and number of talk page edits, respectively), finding that both types were helpful in the early stages of an article, and that larger teams relied on implicit coordination to see any benefit from their greater resources. Romero et al. (2015) found that higher status articles also relied on more implicit coordination and that crowded articles, having many editors relative to the size, exhibit more explicit coordination. Kittur et al. (2009) also studied the role of task interdependence, for example: improving coverage by adding content (low-interdependence) or improving readability by synthesizing many existing contributions (high-interdependence). They found that articles with larger number of editors showed the benefit primarily in low-interdependence tasks, again concluding that coordination is necessary for larger groups to take advantage of their additional resources. Larger group can bring not just more time, but also more diversity and experience to a project. Robert & Romero (2015) found that more diverse and experienced teams were better able to take advantage of large group size. Platt & Romero (2018) found that when editors of a WikiProject interact with a fewer coeditors, those projects both have more high-quality articles and improve article quality more quickly. In summary, Wikipedia appears to be able to take advantage of its large number of contributors, in part, because those editors self-organize into small, diverse sub-groups who use implicit coordination to reduce the necessity for explicit coordination.

3.4 Email Lists and Forums

Even before the web, the internet enabled online communities in the form of email lists and forums. Franco et al. (1995) studied the course of a “flame war” (a series of angry replies) on an email list. The event began when a private reply containing criticism was assumed to be public and the recipient responded publicly. Franco et al. found that high prestige users were more likely to receive replies, and that their messages could change the tone of the conversation. In this example, the messages transitioned from divisive to unifying, allowing the community to publicly develop and enforce its principles. Forums, whether one Usenet or the web, offer similar functionality to email lists but sometimes enable additional features. For example, the web forum Slashdot has formal moderation positions, as well functionality for about 90

3.5 Organizations and Social Movements

Both organizations and social movements have started using internet-based platforms to enable large-scale cooperation. Buurtzorg Nederland, a non-profit Dutch home-care provider employs 8000 nurses divided into 700 self-managed teams (Gray et al., 2015). Burtzorg uses a custom web application to enable coordination within and between teams and achieves higher patient satisfaction for the same expenditure as other firms. In the political sphere, the German Pirate party used a delegative democracy tool called LiquidFeedback from 2011 to 2014 (King et al., 2015). LiquidFeedback allows individuals to either vote directly on proposals or delegate their vote to a proxy (who might then delegate to another person). Using the voting and delegation records, King et al. were able to evaluate the evolution of political power over time, finding that power inequality increased but that "super-voters" with high power tended to vote with the majority of other voters, seldom using their power to influence results. On a smaller scale, DeTar (2013) created the InterTwinkles tool for affinity consensus decision-making in small groups (e.g., housing cooperatives), finding that the tools were often used for dispersed ideation and deliberation between face-to-face meetings. DeTar also noted the importance of technical architecture that supports the creation of ad-hoc groups within a community. Finally, González-Bailón and Wang (2016) studied the Twitter networks of the Indignados movement in Spain and the Occupy movement in the US. They found that after the face-to-face events ended, their latent networks remained active online. They also identified brokers that bridged the two movements, but noted that only about 10

Chapter 4

Research Proposal

As suggested in the previous chapters, large-scale communication is necessary for large-scale collective action, but it is not sufficient. Large-scale decision-making is also necessary. This research proposal attempts to address how large-scale, non-hierarchical decision-making can be implemented, particularly with the help of internet-based technology. Traditionally, large-scale collective action has been organized through coercive dominance hierarchies, e.g., states and firms. Findings from studies of collective intelligence and social learning suggest that non-hierarchical decision-making can avoid informational bottlenecks and biases found in hierarchical organizational structures (Hayek, 1945; Benkler, 2006; Jackson & Golub, 2012). Even in democratically-run groups relying on voting, coercion is necessary to ensure the compliance of dissenting members, and due to the Condorcet paradox (Condorcet, 1785) and Arrows impossibility theorem (Arrow, 1950), winners may not even represent the groups preferences very well.

Without coercion, the only way to ensure the members of a group comply decisions is to ensure that decisions have broad support (Ostrom, 2000). Broad support is often referred to as "consensus, which has become in imprecise term with multiple meanings, so I will prefer "concurrence" when referring to the state of general agreement. Discourse and deliberation can identify points of common agreement and even strengthen concurrence through mutual accommodation (Anderson, 2006; Habermas, 1964; Geiger, 2009). But effective deliberation has traditionally been difficult at large scales (Ackerman & Fishkin, 2002; González-Bailón et al., 2010). Both offline (Putnam, 2000) and online (Benkler, 2006) deliberative communities exhibit small, overlapping sub-groups, creating a "filtering and transmission backbone" (Benkler, 2006). These non-hierarchical structures (Crumley, 1995) allow for the benefits of small group dynamics while enabling ideas to quickly propagate from group to group. I hypothesize that the specific network structure of these sub-groups will influence how effectively large-scale deliberations can generate broad agreement, and propose an experiment to test this hypothesis.

4.1 Related Work

Motivated by falling levels of civic participation (Putnam, 2000) scholars have proposed several designs for facilitating deliberation, both online and off. Face-to-face examples include deliberative polling (Fishkin et al. 2000), citizens juries (Crosby, 1995), planning cells (Dienel & Renn, 1995), consensus conferences (Sclove, 1995), and deliberation days (Ackerman & Fishkin, 2002). These highly-structured approaches generally involve pre-deliberation and post-deliberation surveys combined with a combination of reading material, lectures, small group discussions, and large plenary assemblies. While potentially effective at changing opinions, Schkade et al. (2007) found that when individuals were organized into homogeneous groups, structured deliberation led to increased polarization, so any structured deliberation scheme must be mindful of the human tendency towards homophily. Such face-to-face events are also, by nature, difficult to conduct at a large scale.

Many internet-based systems have been developed to aid decision-making, with some of them enabling various levels of deliberation. Wiki surveys (Salganik & Levy, 2015) allow survey respondents to submit new proposals and consider proposals from other respondents. The InterTwinkles suite of tools for consensus decision-making includes a range of tools for deliberation and voting (DeTar, 2013). Loomio (Jackson & Kuehn, 2016) combines a forum with a flexible interface for consensus decision-making.

4.1.1 Deliberation and Preferences

Deliberation plays several important roles in democratic decision-making. It is a social learning process, allowing members of a group to exchange information. Social learning allows decision-makers to learn of options they had not yet considered and helps them re-evaluate the efficacy of known options at producing their desired outcomes. Deliberation is also a discursive process, potentially changing a decision-makers desired outcomes through social interaction (Habermas, 1964; Anderson, 2006), and by building trust that enables cooperation (Ostrom, 2000; Axelrod, 1997a).

The ability to change individual preferences through deliberation suggests a way to avoid the complications of Arrows impossibility theorem, and to create concurrence within large collective-action groups. However, methods are necessary to determine if and when deliberation actually achieves this goal. I will rely on quantitative measures from social choice theory (described in Section 4.2) to measure how concurrence changes over the deliberative process. These measures have typically been used as formal justificiations for distance-rationalizable voting systems, but I propose using them to track the progress of deliberation in multiple scenarios. It is worth noting that these measures are constructed from first principles to measure similarity/distance between preferences and voting profiles. For the purposes of this proposal, I will assume they are valid for this purpose. An experimental verification could be novel, but is outside the scope of this project.

RQ1. Do consensus profile distances from social choice theory change over the course of deliberation?

As described above, deliberation can potentially change preferences in multiple ways. As social learning, deliberation can change preferences for actions without changing preferences for outcomes, by providing decision-makers with information that allows them to revise their predictions of the efficacy of different actions. Conversely, as discourse, deliberation can change decision-makers preferences over outcomes, by building trust and group identity.

RQ2. In deliberation, how much are changes in preferences over actions due to changes in preferences over outcome?

4.1.2 Networked Deliberation

As groups grow in size, deliberation often becomes too cumbersome and focus is shifted to voting (Ackerman & Fishkin, 2002; González-Bailón, 2010). One possible remedy is to divide large groups into smaller sub-groups.

RQ3. How does the change in concurrence over the course of deliberation depend on group size?

Dividing a group into sub-groups allows for the benefits of small-group dynamics, but does not allow information to flow between groups. However, by shuffling groups and repeating the process, it becomes possible for information to diffuse throughout the entire network. The decision-makers and groups form an affiliation network, and I refer to this scheme as *networked deliberation. Members might participate in multiple groups simultaneously, or sequentially.* This experiment will focus on the sequential case, allowing re-evaluation of preferences between stages. A single stage of the process can also be used to compare the behavior of small groups to large groups.

Def 1. Sequential networked deliberation: N voters are partitioned into groups of size M and allowed to deliberate. The process is repeated T times with new partitions each time.

Network structure has been shown to influence social learning (Lazer & Friedman, 2007; Grim et al., 2013; Barkoczi & Galesic, 2016), suggesting the importance of understanding the role of network structure in networked deliberation. In networked social learning, the lengths of shortest paths between nodes plays an important role in determining how fast information can travel through a network and the level of functional diversity present in the population (Barkoczi & Galesic, 2016). The shorter the shortest paths connecting typical nodes, the more efficient, a network is considered.

RQ4a. Does the efficiency of the network topology influence the level of concurrence reached in a networked deliberation?

RQ4b. Does the efficiency of the network topology influence the number of stages necessary to reach concurrence in networked deliberation?

4.2 Methods

4.2.1 Quantifying Consensus

In order to address the research questions above, it is necessary to define a measure of concurrence. Within social choice theory, consensus can be defined by a consensus class, a set of preference profiles meeting one of several consensus criteria (Elkind & Slinko, 2016). Examples include strong unanimity (of rank orders), unanimity (of winners), majority (existence of majority), Condorcet (existence of Condorcet winner), and transitivity (of social preference). On the most restrictive end, in strong unanimity and unanimity, all decision-makers prefer the same alternative. In contrast, transitivity only requires that the social preferences induced by individual preferences are transitive. Distance-rationalizable voting systems are defined by projecting the preference profile onto the nearest consensus profile according to a suitable distance. For example, Dodgson's method (Dodgson, 1876; Brandt et al., 2012) defines a distance based on swapping adjacent entries in individual preference profiles. However, for most voting systems, finding the consensus profile that minimizes distance is NP-Hard (Elkind & Slinko, 2016). Social choice theory also provides measures of distances between two individual preference rankings. These measures can be extended to create measures of concurrence for entire preference profiles that can be calculated efficiently.

One possibility is to measure the distance from strong unanimity. I propose two measures: the mean Kendal tau metric and the mean Spearman correlation. The Kendal tau metric is the number fraction of pairwise contests that have different results between two profiles (Kendall, 1938). By averaging this measure over all pairs of members, an overall concurrence measure can be calculated for the group. This measure does not incorporate information about win/loss margins, only whether an alternative wins or loses. When the margin is important, I propose instead, the mean Spearman correlation. The Spearman correlation is the linear coefficient of correlation between the rank orders of alternatives for two voters (Spearman, 1904). As with Kendall tau, the Spearman correlation can be averaged over all pairs of members to determine a measure of concurrence for the entire group. It must be noted that representing concurrence with a single number will inevitably ignore important information. Neither of these measures distinguishes a small number of voters disagreeing on many comparisons from a large number of voters disagreeing on a few. The Spearman correlation is also unable to distinguish these from situations in which a small number of voters disagree by a very large margin.

While unanimity might be preferred, it may sometimes be more realistic to focus on transitive consensus. This consensus class requires only that the social preferences between pairs of alternatives are transitive. Such preference profiles guarantee that a Condorcet winner will exist, bringing many different voting methods into agreement on which should win.

To measure the distance from transitivity, I propose the number of ranked pair violations. When performing a vote by the Tideman ranked pair system, the relative order of two alternatives in the social preference is determined by the winner of a pairwise vote, except in the special case that contests with a higher margin have already constrained their order due to transitivity. In the latter case, the winner of the pairwise contest may have a lower position in the social preference. If this situation occurs, the social preference is intransitive and there may not be a Condorcet winner. Each such pair represents a contest violating the social ranking, so by counting the fraction of such pairs, we can quantify the strength of intransitivity in the groups preferences. This measure has the property that a Condorcet winner is guaranteed when the value is 0.

4.2.2 Network Topologies

This experiment requires two network topologies: an efficient topology with short path lengths and an inefficient topology with long path lengths. Both topologies should have the same group size and have relatively little overlap between different groups. The networks are affiliation networks between groups and participants, but can be projected into co-affiliation networks between participants. I will work directly with the co-affiliation projections, in which each group is represented as a fully connected clique of participant nodes. Let N be the number of participants, M be the number of participants in a group, and D be the number of groups a participant belongs to (its degree in the affiliation network).

For the efficient network, I will use a randomized network equivalent to randomly assigning participants to groups in each stage. The nodes are randomly divided into N/M partitions of size M and edges are added between all pairs belonging to the same partition. This process is repeated D times.

To achieve a network with long paths while retaining small overlap between groups, I propose a topology using local subsets of prime residue classes. The prime residue classes guarantee small overlap, while the local subsets ensure long paths. The topology is defined as follows. Let p_i be the ith prime number. Let each node be labeled by an integer in $[0, N-1]$. At step i, divide the nodes into p_i partitions by their remainder modulo p_i. Next, divide each partition into sub-partitions of size M, starting with the lowest M node labels, then the next lowest M, and so on. Repeat for i in $[0, D-1]$. Note that the sub-partitions of size M connect nodes at most $M * p_i$ apart in label, preventing the creation of any shortcut edges. The number of nodes may not be an exact multiple of M, so the final groups may be truncated. However, the network will be a sub-graph of a network that does have a multiple of M nodes, so the structure will not be fundamentally changed, and these edge effects should become negligible as the number of nodes increases.

A typical numerically simulated efficient network with N = 150, M = 8, D = 3, has diameter 3 and mean shortest path of approximately 2. For the same parameters, the inefficient network has diameter 7 and mean shortest path 3.17.

4.2.3 Experiment

I propose developing a platform for networked deliberation and conducting an experimental trial in collaboration with a civic organization. The trial will follow the example of the wiki survey case studies conducted in collaboration with the NYC Mayors Office and the OECD (Salganik & Levy, 2015). Members or stakeholders of the organization will be invited to use the online platform to deliberate on proposals for an issue relevant to the organization.

The platform will consist of two components: preference elicitation and deliberation. The preference elicitation component will ask participants to rank order existing proposals, as well as to suggest any additional proposals. Proposals will be ranked graphically using a drag-and-drop interface for ease of use. Participants will have the option of leaving any number of proposals unranked, which will be considered tied for the lowest place. Participants will also have the option of writing free-form text arguing for or against each proposal. Additionally, participants will be able to rank order their desired outcomes independently from the proposals, providing the information necessary to address RQ3.

The deliberation component will function similarly to a typical online discussion forum. There will be one discussion thread per proposal. Participants will be able to add replies to each proposal and a second level of replies to those replies. Top-level replies will appear chronologically.

Participants will be divided into three groups. One of the most well-studied scales for group size is Dunbars number (Dunbar, 1992), approximately 150 people. As the goal of this experiment is to evaluate processes for large-scale deliberation, each group will be at least as large as Dunbars number. The three groups will be as follows:

Group 1 (control). This group will engage in deliberation between all group members, with sub-group structure. There will be only one stage of deliberation throughout the entire trial.

Group 2 (efficient network). Members of this group will be divided into approximately 19 subgroups of size 8. The trial will be divided into three stages, and members will be reassigned to different subgroups at each stage according to the efficient (randomized) network structure.

Group 3 (inefficient network). Members of this group will be divided into approximately 19 subgroups of size 8. The trial will be divided into three stages, and members will be reassigned to different subgroups at each stage according to the inefficient (local residue class) network structure.

The trial will proceed according to the following stages:

Pre-deliberation. Participants are assigned to a group and assigned an id within that group which will determine their sub-groups at each subsequent stage. Participants rank-order existing proposals and suggest any alternative proposals.

Deliberation (3 stages). Participants are shown the deliberation forum, with one thread for each proposal. They are able to see posts and replies only from other members in their current sub-group (as determined by their groups network structure and the current stage of deliberation). Posts from previous stages remain visible, but cannot be replied to. Each stage of deliberation will last 1 week. After each stage, participants are allowed to change their rank-order preferences and are assigned to a new sub-group for the next round of deliberation.

Post-deliberation and debriefing. Participants have a final chance to revise their rank order preferences. Participants are also asked for any feedback on how the deliberation process has helped them develop their opinions and understand other viewpoints.

4.3 Analysis

Using each measure of consensus in the Methods section, the concurrence of each group will be tracked over the course of deliberation and differences between various network topologies will be reported. The winning proposals according to various voting methods will also be reported for each group and at each point in the deliberation process.

The concurrence beore and after the first stage of deliberation will be used to address RQ1. If any significant change is detected for any of the groups it will verify that it is possible for concurrence to change over the course of deliberation. During the first stage, the 19 low-efficiency and 19 high-efficiency sub-groups will effectively be 38 independent deliberating groups, because the network topology only comes into play at reassignment, after the first stage. These groups will be used to test whether the magnitude and direction of changes in concurrence are consistent.

In addition, concurrence in proposal preference will be compared to concurrence in outcome preference to address RQ2. A change in concurrence over outcome preferences suggests that the discursive component of deliberation played a role in changing preferences. The relative magnitude of change in outcome and proposal preferences can be compared to determine how significant any discursive component is relative to non-discursive knowledge transfer.

Similarly, the change in concurrence over stage 1 will be compared between the 38 experimental sub-groups and the large control group to address the effects of group size for RQ3.

Finally, changes over the course of multiple stages will be examined to address RQ4. A significant departure between the high and low-efficiency groups would suggest a relationship between network structure and the effectiveness of deliberation.

4.4 Discussion

The proposed experiment has several parameters, including the number of members in each group, the number of members in each sub-group, and the number of stages. The group size has been chosen to be above Dunbars number, which represents one scale separating small group behavior from large group behavior. It is possible that there are scales beyond Dunbars number which define different regimes for social behavior, but none have been identified. Multiple stages of deliberation are needed to allow participants to belong to multiple sub-groups, but too many stages would be cumbersome for participants. I have chosen 3 stages to balance these two requirements. Similarly, the group size of 8 has been chosen such that the mean shortest path of the efficient network is lower than the number of stages, while that of the inefficient network is higher. Over the course of 3 stages, ideas from any individual could feasibly reach most others in the efficient network, but not for the inefficient network.

There are several possible alternative methods for studying the effect of network structure on deliberation. There are many platforms online where deliberation occurs every day. While a customized platform offers the most potential, existing platforms could also be analyzed. In threaded forums, e.g. Slashdot (González-Bailón, 2010), sub-threads could be analyzed as sub-groups and bots could be used to elicit preferences throughout the deliberation. This method would not be able to determine preferences for all deliberators or at all stages, but could still produce enough information to compare different network structures, to the extent that natural variation exists. In-person events, similar to deliberation days (Ackerman & Fishkin, 2002), could be used to divide participants into sub-groups according to various network structures for a day of deliberation on a particular topic. An in-person approach has the benefit of keeping participants attention over the course of the experiment, with the drawback of requiring physical space and transportation, making large-scale participation more difficult.

Acknowledgements

Thanks to Daniel M. Romero, Ceren Budak, and Scott E. Page for their time, guidance, and helpful feedback. And thanks to Jeffrey W. Lockhart for many insightful discussions.

References

Abelson, R. P. (1964). Mathematical models of the distribution of attitudes under controversy. *Contributions to mathematical psychology*, *14*, 1–160.

Acemoglu, D., Dahleh, M. A., Lobel, I., & Ozdaglar, A. (2011). Bayesian learning in social networks. *The Review of Economic Studies*, *78*(4), 1201–1236.

Ackerman, B., & Fishkin, J. S. (2002). Deliberation day. *Journal of Political Philosophy*, *10*(2), 129–152.

Adamic, L. A., & Glance, N. (2005). The political blogosphere and the 2004 US election: divided they blog. In *Proceedings of the 3rd international workshop on Link discovery* (pp. 36–43). ACM.

Ali, I., Cook, W. D., & Kress, M. (1986). On the minimum violations ranking of a tournament. *Management Science*, *32*(6), 660–672.

Anderson, E. (2006). The epistemology of democracy. *Episteme*, *3*(1-2), 8–22.

Arrow, K. J. (1950). A difficulty in the concept of social welfare. *Journal of political economy*, *58*(4), 328–346.

Arrow, K. J. (1974). *The limits of organization.* WW Norton & Company.

Arrow, K. J. (2012). *Social choice and individual values* (Vol. 12). Yale university press.

Asch, S. E. (1955). Opinions and social pressure. *Scientific American*, *193*(5), 31–35.

Axelrod, R. M. (1997a). An Evolutionary Approah to Norms. In *The complexity of cooperation: Agent-based models of competition and collaboration.* Princeton University Press.

Axelrod, R. M. (1997b). The Evolution of Strategies in the Iterated Prisoners Dilemma. In *The complexity of cooperation: Agent-based models of competition and collaboration.* Princeton University Press.

Barabasi, A.-L. (2003). Linked: How everything is connected to everything else and what it means.

Barabsi, A.-L., & Albert, R. (1999). Emergence of scaling in random networks. *Science*, *286*(5439), 509–512.

Baran, P., & others. (1964). On distributed communications. *Volumes I-XI, RAND Corporation Research Documents, August*, 637–648.

Barkoczi, D., & Galesic, M. (2016). Social learning strategies modify the effect of network structure on group performance. *Nature communications*, *7*.

Benkler, Y. (2002). Coase's Penguin, or, Linux and" The Nature of the Firm". *Yale law journal*, 369–446.

Benkler, Y. (2006). *The wealth of networks: how social production transforms markets and freedom.* New Haven [Conn.]: Yale University Press.

Boehm, C., Barclay, H. B., Dentan, R. K., Dupre, M.-C., Hill, J. D., Kent, S., ... Rayner, S. (1993). Egalitarian behavior and reverse dominance hierarchy [and comments and reply]. *Current Anthropology*, *34*(3), 227–254.

Borge-Holthoefer, J., Perra, N., Gonalves, B., Gonzlez-Bailn, S., Arenas, A., Moreno, Y., & Vespignani, A. (2016). The dynamics of information-driven coordination phenomena: A transfer entropy analysis. *Science advances*, *2*(4), e1501158.

Brandt, F., Conitzer, V., & Endriss, U. (2012). Computational social choice. *Multiagent systems*, 213–283.

Brooks Jr, F. P. (1995). *The Mythical Man-Month: Essays on Software Engineering, Anniversary Edition, 2/E.* Pearson Education India.

Burt, R. S. (2009). *Structural holes: The social structure of competition.* Harvard university press.

Butler, C. L., & Rothstein, A. (2007). *On conflict and consensus: A handbook on formal consensus decisionmaking.* Creative Commons.

Cartwright, D., Harary, F., & Norman, R. (1965). Structural models: An introduction to the theory of directed graphs. *New York*.

Centola, D., & Macy, M. (2007). Complex contagions and the weakness of long ties. *American journal of Sociology*, *113*(3), 702–734.

Cialdini, R. B., & Cialdini, R. B. (2007). *Influence: The psychology of persuasion.* Collins New York.

Clinton, J., Jackman, S., & Rivers, D. (2004). The statistical analysis of roll call data. *American Political Science Review*, *98*(2), 355–370.

Coase, R. H. (1937). The nature of the firm. *economica*, *4*(16), 386–405.

Cohen, J., Cohen, P., West, S. G., & Aiken, L. S. (2013). *Applied multiple regression/correlation analysis for the behavioral sciences.* Routledge.

Condorcet, M. d. (1785). Essay on the Application of Analysis to the Probability of Majority Decisions. *Paris: Imprimerie Royale*.

Conway, M. E. (1968). How do committees invent. *Datamation*, *14*(4), 28–31.

Copeland, A. H. (1951). *A reasonable social welfare function* (Tech. Rep.). mimeo, 1951. University of Michigan.

Crosby, N. (1995). Citizens juries: One solution for difficult environmental questions. In *Fairness and competence in citizen participation* (pp. 157–174). Springer.

Crumley, C. L. (1995). Heterarchy and the analysis of complex societies. *Archeological Papers of the American Anthropological Association*, *6*(1), 1–5.

Davis, O. A., Hinich, M. J., & Ordeshook, P. C. (1970). An expository development of a mathematical model of the electoral process. *American Political Science Review*, *64*(2), 426–448.

de Borda, J.-C. (1781). Memoire sur les elections au scrutin, historie de lacademie royale des sciences. *Paris, France*.

DeGroot, M. H. (1974). Reaching a consensus. *Journal of the American Statistical Association*, *69*(345), 118–121.

DeMarzo, P. M., Vayanos, D., & Zwiebel, J. (2003). Persuasion bias, social influence, and unidimensional opinions. *The Quarterly Journal of Economics*, *118*(3), 909–968.

Demsetz, H. (1974). Toward a theory of property rights. In *Classic Papers in Natural Resource Economics* (pp. 163–177). Springer.

DeTar, C. (2013). *InterTwinkles: online tools for non-hierarchical, consensus-oriented decision making* (Unpublished doctoral dissertation). Massachusetts Institute of Technology.

Dewey, J. (1940). *Creative democracy: The task before us.* GP Putnam's Sons New York.

Dienel, P. C., & Renn, O. (1995). Planning cells: A gate to fractal mediation. In *Fairness and competence in citizen participation* (pp. 117–140). Springer.

Dodgson, C. (1876). A method of taking votes on more than two issues. *The theory of committees and elections*.

Downs, A. (1957). An economic theory of political action in a democracy. *Journal of political economy*, *65*(2), 135–150.

Dunbar, R. I. (1992). Neocortex size as a constraint on group size in primates. *Journal of Human Evolution*, *22*(6), 469–493.

Elkind, E., & Slinko, A. (2016). Rationalizations of voting rules. *Handbook of Computational Social Choice*.

English, H. B., & English, A. C. (1958). A comprehensive dictionary of psychological and psychoanalytical terms: A guide to usage.

Fishkin, J. S. (1997). *The voice of the people: Public opinion and democracy.* Yale university press.

Fishkin, J. S., Luskin, R. C., & Jowell, R. (2000). Deliberative polling and public consultation. *Parliamentary affairs*, *53*(4), 657–666.

Fjermestad, J., & Hiltz, S. R. (1998). An assessment of group support systems experimental research: methodology and results. *Journal of Management Information Systems*, *15*(3), 7–149.

Forte, A., & Bruckman, A. (2008). Scaling consensus: Increasing decentralization in Wikipedia governance. In *Hawaii International Conference on System Sciences, Proceedings of the 41st Annual* (pp. 157–157). IEEE.

Franco, V., Piirto, R., Hu, H.-Y., Lewenstein, B., Underwood, R., & Vidal, N. (1995). Anatomy of a flame: conflict and community building on the Internet. *IEEE Technology and Society Magazine*, *14*(2), 12–21.

Freeman, J. (1972). The tyranny of structurelessness. *Berkeley Journal of Sociology*, *17*, 151–164.

French Jr, J. R. (1956). A formal theory of social power. *Psychological review*, *63*(3), 181.

Friendkin, N., & Johnsen, E. (1999). Social influence networks and opinion change. *Adv Group Proc*, *16*, 1–29.

Galton, F. (1907). Vox populi (The wisdom of crowds). *Nature*, *75*(7), 450–451.

Geiger, R. S. (2009). Does Habermas understand the Internet? The algorithmic construction of the blogo/public sphere.

Gentry, M. E. (1982). Consensus as a form of decision making. *J. Soc. & Soc. Welfare*, *9*, 233.

Gigone, D., & Hastie, R. (1997). Proper analysis of the accuracy of group judgments. *Psychological Bulletin*, *121*(1), 149.

Goel, A., & Lee, D. (2012). Triadic consensus. In *Proc. 8th Workshop Internet and Network Econ* (pp. 434–447). Springer.

Goel, A., & Lee, D. T. (2016). Towards large-scale deliberative decision-making: Small groups and the importance of triads. In *Proceedings of the 2016 ACM Conference on Economics and Computation* (pp. 287–303). ACM.

Golub, B., & Jackson, M. O. (2010). Naive learning in social networks and the wisdom of crowds. *American Economic Journal: Microeconomics*, *2*(1), 112–149.

Gonzalez-Bailon, S., Kaltenbrunner, A., & Banchs, R. E. (2010). The structure of political discussion networks: a model for the analysis of online deliberation. *Journal of Information Technology*, *25*(2), 230–243.

Gonzlez-Bailn, S., & Wang, N. (2016). Networked discontent: The anatomy of protest campaigns in social media. *Social networks*, *44*, 95–104.

Gonzlez-Bailn, S., Wang, N., & Borge-Holthoefer, J. (2014). The emergence of roles in large-scale networks of communication. *EPJ Data Science*, *3*(1), 32.

Granovetter, M. (1978). Threshold models of collective behavior. *American journal of sociology*, 1420–1443.

Granovetter, M. S. (1973). The strength of weak ties. *American journal of sociology*, 1360–1380.

Gray, B. H., Sarnak, D. O., & Burgers, J. S. (2015). Home care by self-governing nursing teams: the Netherlands Buurtzorg model. *The Commonwealth Fund*.

Grim, P., Singer, D. J., Fisher, S., Bramson, A., Berger, W. J., Reade, C., . . . Sales, A. (2013). Scientific networks on data landscapes: question difficulty, epistemic success, and convergence. *Episteme*, *10*(4), 441–464.

Habermas, J. (2006). Political communication in media society: Does democracy still enjoy an epistemic dimension? The impact of normative theory on empirical research. *Communication theory*, *16*(4), 411–426.

Habermas, J., Lennox, S., & Lennox, F. (1974). The public sphere: An encyclopedia article (1964). *New German Critique*(3), 49–55.

Hacker, K. L., & van Dijk, J. (2000). *Digital democracy: Issues of theory and practice*. Sage.

Hall, J., & Williams, M. S. (1970). Group dynamics training and improved decision making. *The Journal of Applied Behavioral Science*, *6*(1), 39–68.

Hayek, F. A. (1945). The use of knowledge in society. *The American economic review*, *35*(4), 519–530.

Heider, F. (1946). Attitudes and cognitive organization. *The Journal of psychology*, *21*(1), 107–112.

Hill, G. W. (1982). Group versus individual performance: Are N+ 1 heads better than one? *Psychological bulletin*, *91*(3), 517.

Hoebel, E. A. (1958). *Man in primitive world: an introduction to anthropology.* (Tech. Rep.).

Holland, J. H. (1992). *Adaptation in natural and artificial systems: an introductory analysis with applications to biology, control, and artificial intelligence*. MIT press.

Hong, L., & Page, S. (2009). Interpreted and generated signals. *Journal of Economic Theory*, *144*(5), 2174–2196.

Hong, L., & Page, S. E. (2004). Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proceedings of the National Academy of Sciences of the United States of America*, *101*(46), 16385–16389.

Jackson, S. K., & Kuehn, K. M. (2016). Open Source, Social Activism and" Necessary Trade-offs" in the Digital Enclosure: A Case Study of Platform Co-operative, Loomio. org. *tripleC: Communication, Capitalism & Critique. Open Access Journal for a Global Sustainable Information Society*, *14*(2), 413–427.

Kauffman, S., & Levin, S. (1987). Towards a general theory of adaptive walks on rugged landscapes. *Journal of theoretical Biology*, *128*(1), 11–45.

Kearns, M. (2012). Experiments in social computation. *Communications of the ACM*, *55*(10), 56–67.

Keegan, B., & Fiesler, C. (2017). The Evolution and Consequences of Peer Producing Wikipedia's Rules. In *ICWSM.*

Kendall, M. G. (1938). A new measure of rank correlation. *Biometrika*, *30*(1/2), 81–93.

Kittur, A., Chi, E., Pendleton, B. A., Suh, B., & 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, A., & Kraut, R. E. (2008). Harnessing the wisdom of crowds in wikipedia: quality through coordination. In *Proceedings of the 2008 ACM conference on Computer supported cooperative work* (pp. 37–46). ACM.

Kittur, A., Lee, B., & Kraut, R. E. (2009). Coordination in collective intelligence: the role of team structure and task interdependence. In *Proceedings of the SIGCHI conference on human factors in computing systems* (pp. 1495–1504). ACM.

Kittur, A., Suh, B., Pendleton, B. A., & Chi, E. H. (2007). He says, she says: conflict and coordination in Wikipedia. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 453–462). ACM.

Kling, C. C., Kunegis, J., Hartmann, H., Strohmaier, M., & Staab, S. (2015). Voting Behaviour and Power in Online Democracy: A Study of Liquid-

Feedback in Germany's Pirate Party. In *Ninth International AAAI Conference on Web and Social Media.*

Kumar, R., Novak, J., Raghavan, P., & Tomkins, A. (2005). On the bursty evolution of blogspace. *World Wide Web*, *8*(2), 159–178.

Landa, D., & Meirowitz, A. (2009). Game theory, information, and deliberative democracy. *American journal of political science*, *53*(2), 427–444.

Landau, H. G. (1951). On dominance relations and the structure of animal societies: I. Effect of inherent characteristics. *The bulletin of mathematical biophysics*, *13*(1), 1–19.

Lazer, D., & Friedman, A. (2007). The network structure of exploration and exploitation. *Administrative Science Quarterly*, *52*(4), 667–694.

Lee, D. T., Goel, A., Aitamurto, T., & Landemore, H. (2014). Crowdsourcing for participatory democracies: Efficient elicitation of social choice functions. In *Second AAAI Conference on Human Computation and Crowdsourcing.*

Lessig, L. (2009). *Code: And other laws of cyberspace.* ReadHowYouWant.com.

Leuf, B., & Cunningham, W. (2001). The Wiki way: quick collaboration on the Web.

Maltz, T. (1996). Customary law & power in Internet communities. *Journal of Computer-Mediated Communication*, *2*(1), JCMC213.

March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization science*, *2*(1), 71–87.

March, J. G., & Simon, H. A. (1958). Organizations.

Mason, W., & Watts, D. J. (2012). Collaborative learning in networks. *Proceedings of the National Academy of Sciences*, *109*(3), 764–769.

Meskanen, T., & Nurmi, H. (2006). Distance from consensus: A theme and variations. In *Mathematics and democracy* (pp. 117–132). Springer.

Nemiroff, P. M., & King, D. C. (1975). Group decision-making performance as influenced by consensus and self-orientation. *Human Relations*, *28*(1), 1–21.

Newcomb, T. M. (1953). An approach to the study of communicative acts. *Psychological review*, *60*(6), 393.

Newcomb, T. M. (1956). The prediction of interpersonal attraction. *American psychologist*, *11*(11), 575.

Nishi, A., Shirado, H., Rand, D. G., & Christakis, N. A. (2015). Inequality and visibility of wealth in experimental social networks. *Nature*, *526*(7573), 426.

Noam, E. M., Hay, D., Baye, M. R., & Morgan, J. (2003). The Internet: Still wide open and competitive? *OII Internet Issue Brief No. 1*. Retrieved from `https://dx.doi.org/10.2139/ssrn.1308496`

Norris, P., & Inglehart, R. (2009). *Cosmopolitan communications: Cultural diversity in a globalized world.* Cambridge University Press.

Olson, M. (1965). *Logic of collective action: Public goods and the theory of groups.* Harvard University Press.

Ostrom, E. (2000). Collective action and the evolution of social norms. *Journal of economic perspectives*, *14*(3), 137–158.

Ostrom, E. (2010). Beyond markets and states: polycentric governance of complex economic systems. *Transnational Corporations Review*, *2*(2), 1–12.

Ostrom, V., Tiebout, C. M., & Warren, R. (1999). The organization of government in metropolitan areas: A theoretical inquiry. In *Polycentricity and local public economies: Readings from the workshop in political theory and policy analysis* (pp. 31–51). University of Michigan Press.

Pennock, D. M., Flake, G. W., Lawrence, S., Glover, E. J., & Giles, C. L. (2002). Winners don't take all: Characterizing the competition for links on the web. *Proceedings of the national academy of sciences*, *99*(8), 5207–5211.

Platt, E. L., & Romero, D. M. (2018). Network structure, efficiency, and performance in WikiProjects..

Poster, M. (1997). Cyberdemocracy: Internet and the public sphere. *Internet culture*, 201–218.

Putnam, R. D. (2000). Bowling alone: Americas declining social capital. In *Culture and politics* (pp. 223–234). Springer.

Rapoport, A. (1963). Mathematical models of social interaction.

Redfield, R. (1956). How human society operates. *Man, Culture, and Society*, 345–371.

Reinelt, G. (1985). *The linear ordering problem: algorithms and applications* (Vol. 8). Heldermann.

Rheingold, H. (2000). *The virtual community: Homesteading on the electronic frontier*. MIT press.

Rheingold, H. (2007). *Smart mobs: The next social revolution*. Basic books.

Robert, L., & Romero, D. M. (2015). Crowd size, diversity and performance. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (pp. 1379–1382). ACM.

Romero, D. M., Huttenlocher, D., & Kleinberg, J. M. (2015). Coordination and Efficiency in Decentralized Collaboration. In *ICWSM* (pp. 367–376).

Salganik, M. J., & Levy, K. E. (2015). Wiki surveys: Open and quantifiable social data collection. *PloS one*, *10*(5), e0123483.

Schelling, T. C. (2006). *Micromotives and macrobehavior* ([New ed.] with a new preface and the Nobel Lecture ed.). New York: Norton.

Schkade, D., Sunstein, C. R., & Hastie, R. (2007). What happened on deliberation day. *Cal. L. Rev.*, *95*, 915.

Sclove, R. (1995). *Democracy and technology*. Guilford Press.

Sharp, L. (1958). *People without politics*. Bobbs-Merrill.

Shaw, A., & Hill, B. M. (2014). Laboratories of oligarchy? How the iron law extends to peer production. *Journal of Communication*, *64*(2), 215–238.

Spearman, C. (1904). The proof and measurement of association between two things. *The American journal of psychology*, *15*(1), 72–101.

Steiner, I. (1972). Group processes and group productivity. *New York: Academic*.

Stern, H. (1990). A continuum of paired comparisons models. *Biometrika*, *77*(2), 265–273.

Stewart, G. L. (2006). A meta-analytic review of relationships between team design features and team performance. *Journal of management*, *32*(1), 29–55.

Tideman, T. N. (1987). Independence of clones as a criterion for voting rules. *Social Choice and Welfare*, *4*(3), 185–206.

Tonkinson, R. (1988). Egalitarianism and inequality in a Western Desert culture. In *Anthropological Forum* (Vol. 5, pp. 545–558). Taylor & Francis.

Turner, C., & Martin, E. (1985). *Surveying subjective phenomena* (Vol. 2). Russell Sage Foundation.

Tyler, M. C. (2004). 115 and counting: The state of ODR 2004. In *Melissa Conley Tyler, Ethan Katsh, and Daewon Choi, editors, Proceedings of the Third Annual Forum on Online Dispute Resolution.*

Vigas, F. B., Wattenberg, M., & Dave, K. (2004). Studying cooperation and conflict between authors with history flow visualizations. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 575–582). ACM.

Vigas, F. B., Wattenberg, M., & McKeon, M. M. (2007). The hidden order of Wikipedia. In *International conference on Online communities and social computing* (pp. 445–454). Springer.

Weber, M. (1978). *Economy and society: An outline of interpretive sociology* (Vol. 1). Univ of California Press.

Wilcox-OHearn, Z. (2001). *Names: Distributed, secure, human-readable: Choose two.* Retrieved from `https://web.archive.org/web/20011020191610/http://zooko.com/distnames.html`

Zollman, K. J. (2012). Social network structure and the achievement of consensus. *Politics, Philosophy & Economics*, *11*(1), 26–44.

Zollman, K. J. S. (2010). Social structure and the effects of conformity. *Synthese*, *172*(3), 317–340.