An Outline of Cognitive Democracy

Henry Farrell (George Washington University)

Cosma Shalizi (Carnegie Mellon University and The Santa Fe Institute)

Introduction

Democracy is uniquely fitted to help people with highly diverse perspectives to come together to solve problems collectively. Democracy can do this better than either markets and hierarchies, because it brings these diverse perceptions into direct contact with each other, allowing forms of learning that are unlikely either through the price mechanism of markets or the hierarchical arrangements of bureaucracy. Furthermore, democracy can, by experimenting, take advantage of novel forms of collective cognition that are facilitated by new media. Together, our arguments outline a cognitive approach to democracy.

Much of what we say is synthetic. On the normative side, we build both on the academic literature (Cohen (1986); Anderson (2007) on epistemic democracy; Knight and Johnson (2011) on the pragmatist case for radically egalitarian democracy; Rosenblum (2008) on partisanship; Landemore and Elster (2012) on *Collective Wisdom*), and on public intellectuals such as Cory Doctorow, Chris Hayes, Steven Berlin Johnson, Clay Shirky, Tom Slee and the late Aaron Swartz. We steal outrageously, in different ways, from Josiah Ober and Scott Page, and draw much more on Charles Lindblom than our explicit citations suggest. Much of the prior literature focuses on *epistemic* questions, about whether democracy can reliably determine the truth. Throughout, we emphasize the *cognitive* benefits of democracy, how it combines multiple mechanisms to build a powerful collective *process* for solving problems,

drawing on important results from cognitive science, sociology, machine learning and network theory.

We start by explaining one thing which social institutions *should* do: they find solutions to social problems, which means they have cognitive, information-processing goals. From this viewpoint, we discuss sophisticated arguments on behalf of markets (Hayek's "catallaxy") and hierarchy (Richard Thaler and Cass Sunstein's "libertarian paternalism"), and their inadequacies of these arguments. We then lay out our arguments for democracy, emphasizing how democratic procedures have problem-solving capacities that other social forms do not. The penultimate section discusses how democracy can learn from new forms of collective cognition on the Internet, treating these forms not as ideals to be approximated, but as imperfect experiments, whose successes *and failures* can teach us about the conditions for better decision making; this is part of a broader agenda for cross-disciplinary research involving computational scientists and democratic theorists.

Justifying Social Institutions

Three of the most wide-spread and enduring kinds of macro-institutions are *markets*, decentralized forms of economic exchange governed by prices; *hierarchies*, centralized forms of decision-making in which inferiors report to superiors and superiors issue authoritative orders to inferiors through more or less elaborate chains of command; and *democracies*, forms of collective decision making in which individuals are able both to argue with each other over proposals in a non-hierarchical setting and to vote over which proposals are accepted as collectively binding. (This list is not exhaustive.) What might justify such institutions?

The dominant tradition in political theory tends to evaluate macro-institutions such as politics, markets and hierarchies in terms of *justice* — whether institutions use procedures, or give results, that are just according to some reasonable normative criterion. Others, more cynically or more modestly, look to institutions' contributions to *stability* — whether they produce an acceptable level of social order, reducing violence and providing a modicum of

predictability. We start instead with a pragmatist question — whether these institutions are *useful* in helping us *solve difficult social problems.*[1]

Some political problems are simple — the solutions might not be easy to put into practice, but the problems are easy to analyze. But the most vexing problems are usually ones without any very obvious solutions. How do we change legal rules and social norms in order to mitigate the problems of global warming? How do we regulate financial markets so as to minimize the risk of new crises emerging, and limit the harm of those that happen? How do we best encourage the spread of human rights internationally?[2]

These problems all share two important features. First, they are *social*. That is, they are problems which involve the interaction of many human beings, with different interests, desires, needs and perspectives. Second, they are *complex* problems, in the sense that scholars of complexity understand the term. To borrow the definition of Page (2011, p. 25), they involve "*diverse* entities that interact in a *network* or *contact structure*" (italics in the original).[3] They are a result of behavior that is difficult to predict, so that consequences to changing behavior are extremely hard to map out in advance. Finding solutions is difficult, and even when we find one, it is hard to know whether it is better than others, let alone the best.

Macro-institutions will best be able to tackle these problems if they have two features. First, they should foster *a high degree of direct communication* between individuals with diverse viewpoints. This kind of intellectual diversity is crucial to finding good solutions to complex problems. Second, we argue that they should provide *relative equality* among affected actors in decision-making processes, so that socially or politically powerful groups do not block socially beneficial changes that conflict with their own particular interests.

---

[1] Justice and social order matter; useful institutions that were profound unjust or apt to create devastating social instability would be dubious to say the least. But justice and stability are not our concerns here.

[2] These problems have an international aspect — we bracket the thorny question of how to deal with them in the absence of anything that even faintly resembles global democracy. Pragmatists like Dewey have occasionally allowed themselves to hope that international politics was building closer connections between global problems and the publics that might set out to solve them. We are not especially optimistic ourselves.

[3] Much more could be said about the meaning of "complexity". In particular, it may later be useful to look at formal measures of the intrinsic complexity of problems in terms of the resources required to solve them (on "computational complexity" theory, see Moore and Mertens 2011), or the degree of behavioral flexibility of systems, such as interacting decision-makers (Badii and Politi, 1997; Shalizi *et al.*, 2004).

We base these contentions on arguments from work on collective problem solving and theories of political power. Both are clarified if we think of the possible solutions to a difficult problem as points on a landscape, where we seek the highest point[4]. Difficult problems present many peaks (solutions which are better and "higher" in value than the other points immediately adjacent to them). Such landscapes are *rugged* — they have some degree of organization, but are not so structured that simple algorithms can quickly find the best solution. There is no guarantee that any particular peak is *globally optimal* (i.e., the best solution across the entire landscape) rather than *locally optimal* (the best solution within a smaller subset of the landscape).

Solving a complex problem involves a search across this landscape for the best visible solutions. Individual agents have limited cognitive abilities, and (usually) limited knowledge of the landscape. Both of these make them likely to get stuck at local optima, which may be much worse than even other local peaks, let alone the global optimum. Less abstractly, people may settle for bad solutions, because they do not know better (they cannot perceive other, better solutions). It may even be, as McAfee argues (this volume) that there is no 'right' answer in any absolute sense, although it is hard to capture this possibility using the ideas that we invoke here.

Hong and Page (2004) use mathematical models to argue that *diversity of viewpoints* helps groups find better solutions. The intuition is that different individuals, when confronting a problem, "see" different landscapes — they organize the set of possible solutions in different ways, which are more or less useful in identifying good peaks. Very smart individuals (those with many mental tools) have better organized landscapes than less smart individuals, and so are less likely to get trapped at inferior local optima. However, at the group level, diversity of viewpoints matters a lot. Hong and Page find that "diversity trumps ability". Groups with high diversity between internal viewpoints are better able to identify optima than groups composed of much smarter individuals with more homogeneous viewpoints. By putting their

---

[4]This metaphor goes back to the work of the evolutionary geneticist Sewall Wright in the 1920s and 1930s (Provine, 1971).

diverse views together, the former are able to identify attractive solutions that would be invisible to more homogeneous groups.[5]

Hong and Page say little about how individuals share their diverse perspectives with each other, allowing them to build a common framework. However, their arguments suggest that actors' different points of view need to be *exposed directly* to each other, in order to identify the benefits and drawbacks of different points of view, the ways in which viewpoints can be combined to better advantage, and so on. Such claims are supported by work in sociology and elsewhere (e.g., Burt 1992).

The second issue for collective problem solving is more obvious. Even when groups are able to identify good solutions (relatively high peaks in the solution landscape), they may not be able to reach them. In particular, actors who benefit from the status quo (or who would prefer less generally-beneficial solutions) may be able to use political and social power to block movement towards such peaks, and instead compel movement towards solutions that have lower social and greater particular benefits. Research on problem solving typically does not talk about differences in actors' interests, or in their ability to pursue them given e.g., bargaining power, level of rhetorical training and access to symbolic resources and so on. While different individuals initially perceive different aspects of the landscape, researchers assume that once they are able to communicate with each other, they will all agree on how to rank visible solutions. But actors may have diverse interests as well as diverse understandings of the world (and the two may indeed be systematically linked). They may even be working in such different landscapes, in terms of personal advantage, that one actor's peak is another's valley, and vice versa. Moreover, actors may differ in their ability to ensure that their interests are prosecuted. Recent work in political theory (Knight, 1992; Knight and Johnson, 2011), economics (Bowles and Naidu, 2008), political science (Hacker and Pierson, 2010) and sociology details how powerful actors can compel weaker ones to accept solutions that are to

---

[5]Something like this idea may be said to motivate the collective project that this chapter is a part of. This edited volume represents extended discussions between a variety of people with very different disciplinary backgrounds, and different views of the world.

the advantage of the former, but that have lower overall social benefits.

Here, relative equality of power can have important consequences. Individuals in settings with relatively equal power relations, are, *ceteris paribus* more likely to converge on solutions with broad social benefits, and less likely to converge on solutions that benefit smaller groups of individuals at the expense of the majority. Furthermore, equal power relations may not only make it easier to converge on "good" solutions when they have been found, but may stimulate the search for such solutions. Participating in the search for solutions and in decision-making demands resources (at a minimum, time), and if those resources are concentrated in a small set of actors, with similar interests and perspectives, the solutions they will find will be fewer and worse than if a wide variety of actors can also search.

Since we need macro-institutions to solve complex social problems, and their capacity to do this depends on bringing together people with different perspectives and sharing decision-making power relatively equally among those people, we will look at how the broad logics of different macro-institutions relate to those two criteria. Hierarchy, while a remarkable and enduring social form, suffers from dramatic informational flaws. In response, for decades now many scholars and policy makers have been devoted to pushing markets as *the* way to address social problems that are too complex to be solved by top-down authority. As we show below, markets, if they embody substantial power inequalities, and homogenize human relations, are unlikely to possess the virtues attributed to them, though they can have more particular benefits under specific circumstances. This prompts us to advocate democracy, not for the sake of justice or stability, but as a tool for solving complex problems.

Markets and Hierarchies for Problem-Solving

Both markets and hierarchies have their advocates among scholars and public intellectuals, who believe that one or the other (or both) are better ways of solving complex problems than democracy. Advocates of markets usually build on the ground-breaking work of Hayek

(1937, 1945), to argue this form of organization is better at eliciting information and putting it to good work than any more collective form. Advocacy of hierarchy is much older, and, perhaps for that reason, does not have any such unified tradition. However, Thaler and Sunstein (2008) have recently made a sophisticated case for the benefits of hierarchy. They advocate a combination of top-down mechanism design and institutions designed to guide choices rather than to constrain them — what they call libertarian paternalism — as a way to solve difficult social problems. Hayek's arguments are not the only case for markets, and Thaler and Sunstein's are not the only justification for hierarchy. They are, however, among the *best* such arguments, and hence provide a good initial way to test the respective power of markets, hierarchies and democracies to solve complex problems.

Markets

Hayek's account of the informational benefits of markets is ground-breaking (Hayek, 1937, 1945). Although it builds on the insights of others (particularly Michael Polanyi), it is arguably the first real effort to analyze how social institutions work as information processors. Hayek reasons as follows. Much of human knowledge (as Polanyi argues) is practical, and cannot be fully articulated ("tacit"). This knowledge is nonetheless crucial to economic life. Hence, if we are to allocate resources well, we must somehow gather this dispersed, fragmentary, informal knowledge, and make it useful.

Hayek is explicit that no one person can know all that is required to allocate resources properly, so there must be a *social* mechanism for such information processing. He considers three such mechanisms: central planning, planning by monopolistic industries, and decentralized planning by individuals. The first and second of these break down in the face of the need for vast amounts of tacit knowledge, which cannot be conveyed to any centralized authority. Centralized or semi-centralized planning are especially poor at dealing with the constant flows of major and minor changes through which an economy (or, as Hayek would prefer, a catallaxy or spontaneous order) approaches balance. To deal with such changes, we

need people to make the necessary decisions on the spot — but we also need a way to convey information about changes in the larger economic system to these people. The virtue of the price mechanism, for Hayek, is to compress diffuse, even tacit, knowledge about specific changes in specific circumstances into a single index, which can help individuals adapt to changes elsewhere. I do not need to grasp the intimate local knowledge of the farmer who sells me tomatoes in order to decide whether to buy them. The farmer needs to know the price of fertilizer, not how it is made, or what it could be used for other than tomatoes, or the other uses of the fertilizers' ingredients. The information that we need, to decide whether to buy tomatoes or to buy fertilizer, comes through prices, which may go up or down, depending on the aggregate action of many buyers or suppliers, each working with their own tacit understandings.

This insight is both crucial and beautiful[6], yet it has stark limits. It suggests that markets will be best at conveying a particular kind of information about a particular kind of underlying fact, i.e., the relative scarcity of different goods. As Stiglitz (2000) argues, market signals about relative scarcity are always distorted, because prices embed information about many other economic factors. More importantly, although information about relative scarcity surely helps markets approach some kind of balance, it is little help in solving more complicated social problems, which may depend not on allocating existing stocks of goods in a useful way, given people's dispersed local knowledge, so much as discovering new goods or new forms of allocation. More generally, Hayek's well-known detestation for projects with collective goals lead him systematically to discount the ways in which aggregate knowledge might work to solve collective rather than individual problems.[7]

---

[6]Imagine trying to discover whether a locally-grown tomato in Pittsburgh is better, from the point of view of greenhouse-gas emission, than one imported from Florida. After working out the differences in emissions from transport, one has to consider the emissions involved in growing the tomatoes in the first place, the emissions-cost of producing different fertilizers, farm machinery, etc., etc. The problem quickly becomes intractable — and this is before a consumer with limited funds must decide how much a ton of emitted carbon dioxide is worth to them. Let there be a price on greenhouse-gas emission, however, and the whole informational problem gets solved implicitly by ordinary market interactions.

[7]Hayek draws a sharp distinction between an empiricist liberalism which is associated with trial and error and the gradual evolution of robust institutions, and a collectivist liberalism which is doctrinaire and committed to the enforcement of a preconceived rationalist enterprise. He is uninterested in how democratic

This is unfortunate. To the extent that markets fulfill Hayek's criteria, and mediate all relevant interactions through the price mechanism, they foreclose other forms of intellectual exchange. In particular, Hayek's deprecation of "rationalism" leaves little place for reasoned discourse or the useful exchange of views. In Hayek's markets, people communicate only through prices. The advantage of prices, for Hayek, is that they inform individuals about what others want (or don't want), without requiring anyone to know anything about anyone else's plans or understandings. But there are many useful forms of knowledge that cannot readily be conveyed in this way.

People may learn something about those understandings as a *by-product* of market interactions. In John Stuart Mill's description:

> But the economical advantages of commerce are surpassed in importance by those of its effects which are intellectual and moral. It is hardly possible to overrate the value, in the present low state of human improvement, of placing human beings in contact with persons dissimilar to themselves, and with modes of thought and action unlike those with which they are familiar. Commerce is now what war once was, the principal source of this contact. (Mill, 1909, §III.17.1)

When this learning happens, however, it is incidental, a side effect of buying and selling to best advantage. As markets grow depersonalized, they lead to less of this salutary contact between different modes of life. The resurgence of globalization; the creation of an Internet where people buy and sell with those they will only ever known as account names; the replacement of local mores with global standards; all these provide enormous efficiency gains and allow information about supply and demand to flow more smoothly. Yet they also undermine the Millian benefits of commerce, making it less likely that individuals with different points of view will be exposed to each other's perspectives. More tentatively, markets may themselves have a homogenizing impact on differences between individuals and across soci-

---

deliberation might fold processes of experimentalism within itself. See, e.g., Hayek (1958), and, for an internal critique, Kerstenetzky (2000).

eties, increasing peace but reducing diversity (Hirschman, 1992). Sociologists such as Meyer *et al.* (1977) find evidence of increased cultural and social convergence across nations, as a result of exposure to common market and political forces.

Finally, it is unclear whether markets tend to reduce or reinforce power inequalities in modern democracies. It is almost certainly true, as Marx argues, that the spread of markets helped undermine such historical forms of domination as feudalism. It is *not* clear that markets continue to do so. On the one hand, market participation gives people some ability (presuming equal market access, etc.) to break away from abusive relationships. On the other, markets give greater voice and choice to those with more money; if money talks in politics, it shouts across the agora. Nor is this limited to the marketplace. The market facilitates and fosters asymmetries of wealth which in turn may translate, directly or indirectly, into asymmetries of political influence (Lindblom, 1977, 1982). Untrammeled markets are associated with gross income inequalities, which in turn infect politics with a variety of pathologies. Markets are at best indifferent levelers of unequal power relations.

Hierarchy

Authoritative hierarchy is endorsed (if not presumed) by what are, historically, the oldest, most wide-spread and most enduring traditions of political philosophy. In modern times, especially within the last century, it has come to seem normatively objectionable, but it remains an extraordinary political achievement. States with clear, accountable hierarchies can achieve vast and intricate projects, and businesses use hierarchies to coordinate highly complex chains of production and distribution.[8] Accordingly, it is not without its defenses even now. In particular, the recent influential book of Thaler and Sunstein (2008) is a sustained brief for hierarchy in solving complex problems.

Thaler and Sunstein argue that "choice architects", people who have "responsibility for

---

[8] "Thus bridges are built; harbours open'd; ramparts rais'd; canals form'd; fleets equip'd; and armies disciplin'd every where, by the care of government, which, tho' compos'd of men subject to all human infirmities, becomes, by one of the finest and most subtle inventions imaginable, a composition, which is, in some measure, exempted from all these infirmities." — Hume (1739, book III, part II, sect. vii).

organizing the context in which people make decisions," can design institutions so as to spur people to take better choices rather than worse ones. Thaler and Sunstein are self-consciously paternalist, claiming that flawed thinking consistently stops people from doing what is in their best interests. However, they also find direct control of people's choices morally opprobrious. Libertarian paternalism seeks to guide but not eliminate choice, so that the easiest option is the "best" choice, the one people *would* make if they only had enough attention and discipline. It provides paternalistic guidance through libertarian means, shaping choice contexts to make it more likely that individuals will make the right choices rather than the wrong ones.

This is, in Thaler and Sunstein's words, a politics of "nudging" choices rather than dictating them. Although Thaler and Sunstein do not put it this way, it is also a plea for the benefits of hierarchy in organizations, especially in government. Thaler and Sunstein's "choice architects" are hierarchical superiors, specifically empowered to create broad schemes that will shape the choices of many other individuals. Their power to do this does not flow from, e.g., accountability to those whose choices get shaped. Instead, it flows from positions of authority within firm or government, which allow them authoritatively to craft pension contribution schemes within firms, environmental policy within the government, and so on.

Thaler and Sunstein's recommendations have outraged libertarians, who believe that a nudge is just a well-aimed shove — that individuals' freedom will be reduced nearly as much by Thaler and Sunstein's choice architecture, as it would be by direct coercion. We are also unenthusiastic about libertarian paternalism, but for different reasons. While we do not talk, here, about coercion, we have no particular normative objection to it, provided that it is proportionate, directed towards legitimate ends, and constrained by well-functioning democratic controls. Instead, we worry that the kinds of hierarchy that Thaler and Sunstein presume actively inhibit the unconstrained exchange of views that is essential to solving complex problems.

There are, after all, reasons why bureaucracies have few modern defenders. Hierarchies need power asymmetries to work. Inferiors take orders from superiors, in a chain of command.

This is good for pushing orders *down* the chain, but notoriously poor at transmitting useful information *up*, especially information superiors did not anticipate wanting. Furthermore, as scholars from Max Weber on have emphasized, bureaucracies systematically encourage a culture of conformity in order to increase predictability and static efficiency.

Thaler and Sunstein discuss ordinary people's bad choices at length. However, they have remarkably little to say about how it is that the architects housed atop the hierarchy can figure out better choices on these individuals' behalf. Sometimes, Thaler and Sunstein suggest that choice architects can rely on introspection: "Libertarian paternalists would like to set the default by asking what reflective employees in Janet's position would actually want." At other times, they imply that choice architects can use experimental techniques; this is hard for any sort of *social* problem, about how people interact with each other[9], and becomes intractable precisely with *complex* problems, where the set of candidate solutions is too vast for systematic exploration. Finally, Thaler and Sunstein sometimes argue that choice architects can use results from the social sciences to find optima.

One mechanism of information gathering that they systematically ignore is active feedback from citizens. Although they argue in passing that signals from choice architects can help guide *consumers*, e.g., giving information about the content of food, or by shaping online interactions to ensure that people are exposed to others' points of view, they have no place for feedback from the individuals whose choices are being manipulated to help guide the choice architects, let alone to constrain them. As Mettler (2011) has pointed out, Thaler and Sunstein depict citizens as passive consumers, who need to be guided to the desired outcomes, rather than active participants in politics.

This also means that Thaler and Sunstein's proposals don't take advantage of diversity. Choice architects, located within hierarchies which tend to promote conformity, are likely to have a much more limited range of ways of understanding problems than the population whose choices they are seeking to structure. In Page's terms, these architects may be very

---

[9]Since it becomes hard to separate "treatment" and "control" groups in an experiment.

"able", but will also be very similar to each other in background and training, so that *as a group* they will see a far more limited set of possibilities than a group of randomly selected members of the population (who are likely to have less sophisticated but far more diverse heuristics). Cultural homogeneity among hierarchical elites helps create policy disasters (the "best and brightest" problem). Direct involvement of a wider selection of actors with more diverse heuristics would alleviate this problem.

However, precisely because choice architects rely on hierarchical power to create their architectures, they will have difficulty in eliciting feedback, even if they want to. Inequalities of power notoriously dampen real exchanges of viewpoints. Hierarchical inferiors within organizations worry about contradicting their bosses. Ordinary members of the public are uncomfortable when asked to contradict experts or officials. Work on group decision making (including Sunstein 2003) is full of examples of how perceived power inequalities lead less powerful actors either to stay silent, or merely to affirm the views of the more powerful, even when they have independently valuable perspectives or knowledge.

In short, libertarian paternalism is flawed, not because it restricts peoples' choices, but because it makes heroic assumptions about choice architects' ability to figure out what the choices should be, and blocks the architects' channels for learning better. Libertarian paternalism may still have value where people likely do want, e.g., to save more or take more exercise, but face commitment problems, or when other actors have an incentive to misinform these people or to structure their choices in perverse ways in the absence of a "good" default. However, it will be far less useful, or even actively pernicious, in complex situations, where many actors with different interests make interdependent choices. Thaler and Sunstein are far more convincing when they discuss how to encourage people to choose appropriate pension schemes[10] than when they suggest that environmental problems are the "outcome of a global choice architecture" that could be usefully rejiggered through voluntaristic mechanisms.

---

[10]Though it is hard, today, to be quite as sanguine as they were about encouraging workers to invest more in the stock market.

Democracy as a way to solve complex problems

Is democracy better at identifying solutions to complex problems? Many — even on the left — doubt this. We believe that they are wrong, and that democracy offers a better way of solving complex problems. Since, as we've argued, power asymmetries inhibit problem-solving, democracy has a large advantage over both markets and technocratic hierarchy. The fundamental democratic commitment is to equality of power over political decision making. Real democracies do not attain this any more than real markets deliver perfect competition, or real hierarchies deliver an abstractly benevolent interpretation of rules. But just as a commitment to markets is a commitment to improving competition, and a commitment to hierarchy (in its positive aspects) is a commitment to greater disinterestedness, a commitment to *democratic* improvements is a commitment to making power relations more equal.

Besides equality, the other requirement for solving complex social problems is exposing diverse points of view to each other; we saw that both markets and hierarchies are bad at this. But democracy is good fostering such engagement and rendering it productive, because, along with voting, it involves *debate*. In competitive elections and in more intimate discussions, democratic actors argue over which proposals are better or worse, exposing their different perspectives to each other. It is for this reason that, as Knight and Johnson (2011, p. 151) put it, "democratic decision processes make better use of the distributed knowledge that exists in a society than do their rivals" such as market coordination or judicial decision making. It is the combination of voting and debate that makes it work.

And yet, democratic debate looks ugly: it is partisan, rancorous and vexatious, and people seem to never change their minds. Not least for this reason, the dominant tradition of American liberalism actually has considerable distaste for the less genteel aspects of democracy. The early 20th century Progressives and their modern heirs deplore partisanship and political rivalry, instead preferring technocracy, moderation and deliberation (Rosenblum, 2008). Reinforcing this, earlier, epistemic, accounts of the benefits of democracy have relied on results such as Condorcet's Jury Theorem, which says that as long as people are even

slightly more likely to be right than wrong, and do not seek to influence each other, the more people who vote on a particular question, the more likely it is that they will arrive at the correct answer. Such approaches assume assumed that people judge best when they judge in isolation from each other. This in turn implies that democracy would work better still if debate could be suppressed — that debate is not just vulgar, it is a weakness.

Such reactions miss two crucial points. The first is that, as Knight and Johnson (2011) say, politics is a response to the problem of diversity. Actors with differing — indeed conflicting — interests and perceptions find that their fates are bound together, and that they must make the best of this. We are lucky when genuine conflicts over things which really matter to people lead only to ugly words. Second, and more subtly, as Knight and Johnson argue, politics can also seek to *harness* diversity so as to generate useful knowledge. This does not require impartial deliberation; rather, partial and self-interested debate can have practical benefits. Knight and Johnson suggest that approaches based on diversity (e.g., Page 2007 and Anderson (2007)) are better guides here than the earlier epistemic accounts.

We agree, yet Knight and Johnson do not explain *how* clashes between different actors with different viewpoints result in better decision making. Without some account of the mechanisms engaged, it could be that conflict between perspectives results in worse rather than better problem solving. To make a good case for democracy, we not only need to bring diverse points of view to the table, but show that the *specific ways* in which they are exposed to each other have beneficial consequences for problem solving.

Important micro-level work speaks directly to this, however. Mercier and Sperber (2011) advance an "argumentative" account of reasoning, on which reasoning is not intended to reach right answers, but rather to evaluate the weaknesses of others' arguments and come up with good arguments to support one's own position. This explains both why confirmation bias and motivated reasoning are rife, and why the quality of argument is significantly better when actors engage in real debates. Experimentally, individual performance when reasoning in non-argumentative settings is "abysmal," but is "good" in argumentative settings. This, in turn,

means that groups are typically better in solving problems than is the best individual within the group.[11] Indeed, where there is both diversity of opinion and competition, confirmation bias can have positive effects in pushing people to evaluate and improve their arguments.

A separate line of research in experimental social psychology (Nemeth *et al.*, 2004; Nemeth and Ormiston, 2007; Nemeth, 2012) indicates that problem-solving groups produce more solutions, which outsiders assess as better and more innovative, when they contain persistent dissenting minorities, and are encouraged to engage in, rather than refrain from, mutual criticism. (Mercer (2000) reports similar effects in school-children.) This makes a great deal of sense from Mercier and Sperber's perspective. It also provides more reasons to value the "impure" kinds of dissent described by Tommie Shelby in his contribution to this volume.

This research provides micro-level evidence that political argument will improve problem solving, even if we are skeptical about human beings' ability to abstract away from their specific circumstances and interests. Neither a commitment to deliberation in its strongest sense, nor even standard rationality, is required for argument to help solve problems.

This has clear implications for democracy, which forces actors with very different perspectives to engage with each others' viewpoints, as pointed out long ago by Lindblom (1965). Even the most homogeneous-seeming societies contain great diversity of opinion and of interest (the two are typically related) within them. In a democracy, no single set of interests or perspectives is likely to prevail on its own. Thus, much of the time, political actors have to build coalitions with others holding dissimilar views, which requires engagement between these views. This need to form coalitions will also encourage political actors to find options which are broadly rather than narrowly beneficial. Likewise, actors have to publicly contend with others holding opposed perspectives in order to persuade uncommitted others to favor their position, rather than another. As new issues arise, actors have to persuade even their old allies of how their shared perspectives should be reinterpreted anew.

---

[11]Several decades of work in experimental psychology indicates that groups are, indeed, better at problem-solving than the best individuals within the group (Laughlin, 2011). This tradition does not, however, seem to have given much attention to the role of intra-group debate and argument.

More generally, many of the features of democracy that skeptical liberals deplore are actually of considerable benefit. Mercier and Sperber's work provides micro-foundations for arguments about the benefits of political contention, such as Knight and Johnson's and (differently) John Stuart Mill's, and of arguments for the benefits of partisanship, such as Rosenblum (2008)'s sympathetic critique and reconstruction of Mill. Their findings suggest that the confirmation bias that political advocates have are subject to can have crucial benefits, so long as it is tempered by the ability to evaluate good arguments in context.

Other work suggests that the macro-structures of democracies too can aid problem-solving. Lazer and Friedman (2007) suggest that networks which isolate individuals (so that they do not converge too quickly on the same solution) do the same kind of work that confirmation bias does in Mercier and Sperber's studies — it preserves diversity and encourages actors to keep exploring solutions that may not have immediate payoffs.[12]

This work offers a justification for the organization of democratic life around political parties. Party politics tends to organize debate into intense clusters of argument among people (partisans for the one or the other party) who agree in broad outline about how to solve problems, but who disagree vigorously about the specifics. Links between these clusters are sparser than links within them, and are usually mediated by competition. One might see each cluster as engaged in exploring the space of possibilities around a particular set of solutions, maintaining some limited awareness of other searches being performed within other clusters, and sometimes discreetly borrowing from them in order to improve competitiveness, but nonetheless preserving an essential level of diversity (cf. Huckfeldt *et al.* 2004). Such very general considerations do not justify any *specific* partisan arrangement, as there may be better (or worse) arrangements available. What it does is highlight how party organization and party competition can have benefits that are hard or impossible to match in a less clustered and more homogeneous social setting. Specifically, it shows how partisan arrangements can

---

[12]Broadly congruent results have come from experiments on learning and problem-solving in controlled networks of human subjects in the laboratory (Mason *et al.*, 2008; Judd *et al.*, 2010; Mason and Watts, 2012). Using multiple semi-isolated sub-populations ("islands") is a common trick in evolutionary optimization, precisely to prevent premature convergence on local optima (Mitchell, 1996).

be better at solving complex problems than non-partisan institutions, because they better preserve and better harness diversity.

This also helps us think more clearly about the *possibility conditions* for highly successful problem solving in democracies. We summarize these in the most cursory fashion (we hope to expand on this in further work). First, in contrast to existing epistemic accounts, cognitive accounts suggest that individuals need to be able to *expose* their different points of view to each other, rather than the polling of individuals in strict isolation from each other required by Condorcet's Jury Theorem. This goes hand-in-hand with a different account of problem solving — rather than asking whether people can determine whether a given decision will be correct or incorrect, as Condorcet does, it asks when individuals will be able to discover hitherto-unperceived solutions within a complex landscape. Second, individuals need to be at least "weak learners" in the terms of statistical learning theory (Schapire and Freund, 2012). Individuals who are fundamentally obtuse, profoundly blinded by ideology, or whimsically perverse will detract from collective learning rather than help it. Third and related, as Mercier and Sperber suggest, cognitive democracy requires that individuals participating in democratic argument have some core commitment to the truth, even if they disagree strongly about what the truth is. People need not be as disinterested as Gardner (this volume) would like them to be, but neither should they be so warped by self-interest that they cannot see the truth, or allow themselves to care for it. Fourth, even if people disagree on how to solve a problem, they agree on what the problems *are* that need to be solved in the first place, and have some minimal common empirical standards (see also McAfee, this volume).

Clearly, these conditions are falsified in our everyday political experience. When these conditions are not met, cognitive democracy will work less well. In the extreme, it may not work at all. Even so, where actual conditions even vaguely approximate the ideal, we contend that democracy will be better able to solve complex problems than either markets or hierarchy. First, democracy embodies a commitment to political equality that the other two macro-institutions do not. Clearly, actual democracies achieve political equality more or less

imperfectly. Yet if we are right, the better a democracy is at achieving political equality, the better it will be, *ceteris paribus*, at solving complex problems. Second, democratic *argument*, which people use either to ally with or to attack those with other points of view, is better suited to exposing different perspectives to each other, and hence capturing the benefits of diversity, than either markets or hierarchies. Notably, we do not make heroic claims about people's ability to deliberate in some context that is free from faction and self-interest. Instead, even under realistic accounts of how people argue, and some minimal degree of commitment to the truth, democratic argument will have cognitive benefits, and indeed can transform private vices (confirmation bias) into public virtues (the preservation of cognitive diversity)[13]. Democratic structures — such as political parties — that are often deplored turn out to have important cognitive advantages.

Democratic experimentalism and the Internet

We have no reason to think that actually-existing democratic structures are as good as they could be, or even close. If nothing else, designing institutions is, itself, a highly complex problem, where even the most able decision-makers have little ability to foresee the consequences of their actions. Even when an institution works well at one time, it does so in a context of other institutions and social and physical conditions, which are all constantly changing. Institutional design and reform, then, is always a matter of more or less ambitious "piecemeal social experiments", to use the phrase of Popper (1960). As emphasized by Popper, and independently by Knight and Johnson, one of the strengths of democracy is its ability to make, monitor, and learn from such experiments.[14] (Knight and Johnson

---

[13]This resonates with Karl Popper's insistence (Popper, 1960, 1969) that, to the extent science is rational and objective, it is not because individual scientists are disinterested, rational, etc. — he knew perfectly well that individual scientists are often pig-headed and blinkered — but because of the way the social organization of scientific communities channels scientists' ambition and contentiousness. The reliability of science is an emergent property of scientific institutions, not of scientists.

[14]Bureaucracies can do experiments, such as field trials of new policies, or "A/B" tests of new procedures, now quite common with Internet companies. (See, e.g., the discussion of such experiments in Pfeffer and Sutton 2006.) Power hierarchies, however, are big obstacles to experimenting with options which would upset those power relations, or threaten the interests of those high in the hierarchy. Market-based selection of variants (explored by Nelson and Winter 1982) also has serious limits (see, e.g., Blume and Easley 2006).

particularly emphasize the difficulty markets have in this task.) Democracies can, in fact, experiment with their own arrangements.

For several reasons, the rise of the Internet makes this an especially good time for experimenting with democratic structures themselves. The means available for communication and information-processing are constrain the possibilities for collective decision-making. (Bureaucracy was not an option in the Old Stone Age, nor representative democracy without something like cheap printing.) We do not yet *know* the possibilities of Internet-mediated communication for gathering dispersed knowledge, for generating new knowledge, for complex problem-solving, or for collective decision-making, but we really ought to find out.

In fact, we are already starting to find out. People are building systems to accomplish all of these tasks, in narrower or broader domains, for their own reasons. Wikipedia is, of course, a famous example of allowing lots of more-or-less anonymous people to concentrate dispersed information about an immense range of subjects, and to do so both cheaply and reasonably reliably[15]. Crucially, it is not unique. News-sharing sites like Digg, Reddit, etc., are ways of focusing collective attention and filtering vast quantities of information. Sites like StackExchange have become a vital part of programming practice, because they encourage the sharing of know-how about programming, with the same system spreading to many other technical domains. The knowledge being aggregated through such systems is not *tacit*, rather it is articulated and discursive, but it was dispersed and is now shared. Similar systems are even being used to develop new knowledge. One mode of this is open-source software development, but it is also being used in experiments like the Polymath Project for doing original mathematics collaboratively[16].

There are, after all, many reasons why there are no markets in alternative institutions; to begin with, even if such a market could get started, it would be a prime candidate for efficiency-destroying network externalities, leading at best to monopolistic competition. (Cf. the advice in Shapiro and Varian 1998 to businesses about manipulating standards-setting processes.)

[15]Most of the *content* of Wikipedia comes from a large number of users each of whom makes a substantial contribution or contributions to a very small number of articles (Swartz, 2006). Wikipedia readers are, however, not nearly as diverse as one might like.

[16]For an enthusiastic and intelligent account of ways in which the Internet might be used to enhance the practice of science, see Nielsen (2011). (We cannot adequately explore, here, how scientific disciplines fit into our account of institutions and democratic processes.)

More modestly, there are the ubiquitous mailing lists, discussion forums, etc., where people with similar interests discuss them, on basically all topics of interest to people with enough resources to get on-line. These are, largely inadvertently, experiments in developing collective understandings, or at least shared and structured disagreements, about these topics.

All such systems have to face tricky architectural problems (Shalizi, 2008; Farrell and Schwartzberg, 2008). They need ways of of making findings (or claims) accessible, of keeping discussion productive, and so forth. Often, participants are otherwise strangers to each other, which is at the least suggestive of the problems of trust and motivation which will face efforts to make mass democracy more participative. This opens up an immense space of possible institutional designs, which is still very poorly understood — but almost certainly presents a rugged search landscape, with an immense number of local maxima and no very obvious path to the true peaks. One of the great aspects of the current moment, for cognitive democracy, is that it has become (comparatively) very cheap and easy for such experiments to be made online, so that this design space can be explored.

Even the online ventures which do not work are informative. They range from poorly-designed sites which never attract (or actively repel) a user base, or produce much of value, to online groupings which are very successful in their own terms, but are, cognitively, full of fail, such as thriving communities dedicated to conspiracy theories. These are not just random, isolated eccentrics, but highly structured networks engaged in sharing and developing *very bad* ideas. (See, for instance, Bell *et al.* (2006) on the networks of those who share delusions that their minds are being controlled by outside forces.) If we want to understand what makes successful online institutions work, and perhaps even draw lessons for institutional design more generally, it will help tremendously to contrast the successes with such failures.[17]

The other way in which this helps learning is that all these experiments are leaving incredibly detailed records. People who use these sites or systems leave detailed, machine-

---

[17]By "successful," we mean successful in solving problems. Here, then, pragmatism provides a standard for evaluation, differing from e.g. the more content neutral assessment of online engagement that Fung and Shkabatur offer (this volume).

accessible traces of their interactions with each other, even ones which *tell us about what they were thinking.* This is an unprecedented flood of detail about experiments with collective cognition, and indeed with all kinds of institutions, and about how well they served various functions. Not only could we begin to just *observe* successes and failures, but we can probe the mechanisms behind those outcomes.

This points, we think, to a very clear constructive agenda. To exaggerate a little, it is to see whether the Internet enables modern democracies to make as much use of their citizens' minds as did Ober (2008)'s Athens. We want to learn from existing online ventures in collective cognition and decision-making. We want to treat these ventures are, more or less, spontaneous experiments[18], and compare the success and failures (and *partial* successes and failures) to learn about institutional mechanisms which work well at harnessing the cognitive diversity of large numbers of people who do not know each other well (or at all), and meet under conditions of relative equality, not hierarchy. If this succeeds, what we learn from this will provide the basis for experimenting with re-designing democratic institutions themselves.

We have, implicitly, been viewing institutions through the lens of information-processing. To be explicit, the human actions and interactions which instantiate an institution also implement abstract computations (Hutchins, 1995). Especially when designing institutions for collective cognition and decision-making, it is important to understand them as computational processes. This brings us to our concluding suggestions about some of the ways social science and computer science can help each other.

Hong and Page's work provides a particularly clear, if stylized, model of how diverse individual perspectives or heuristics can combine for better problem-solving[19]. This observation is highly familiar in machine learning, where the large and rapidly-growing class of "ensem-

---

[18]Obviously, the institutions people volunteer to participate in on-line will depend on their pre-existing characteristics, and it would be naive to ignore this. We cannot here go into strategies for causal inference in the face of such endogenous selection bias, which is pretty much inescapable in social networks (Shalizi and Thomas, 2011). *Deliberate* experimentation with online institutional arrangements is attractive, if it could be done effectively and ethically (cf. Salganik *et al.* 2006; Salganik and Watts 2008; Bakshy *et al.* 2012), but, as noted above, little is yet known about how to experiment on social *interactions.*

[19]The Hong and Page model can be seen as a formalization of the "disjointed incrementalism" described in Braybrooke and Lindblom (1963); Lindblom (1965); we plan to pursue this analogy elsewhere.

ble methods" work, explicitly, by combining multiple imperfect models, which helps only because the models are different (Domingos, 1999; Schapire and Freund, 2012) — in some cases it helps *exactly to the extent* that the models are different (Krogh and Vedelsby, 1995). Different ensemble techniques correspond to different assumptions about the capacities of individual learners, and how to combine or communicate their predictions. The latter are typically extremely simplistic, and understanding the possibilities of non-trivial organizations for learning seems like a crucial question for both machine learning and for social science.

Conclusions: Cognitive Democracy

Democracy, we have argued, has a capacity unmatched among other macro-institutions to actually experiment, and to make use of cognitive diversity in solving complex problems. To realize these potentials, democratic structures must themselves be shaped so that social interaction and cognitive function reinforce each other. But the cleverest institutional design in the world will not help unless the resources — material, social, cultural — needed for participation are actually broadly shared. This is not, or not just, about being nice or equitable; cognitive diversity is not something we can afford to waste.

We differ from Gardner's proposal (this volume) that we should seek to nurture disinterested professional standards, on the Internet and elsewhere. Our perspective points to the value, up to a point, of *interestedness* - that is of vigorous and even rowdy debate between individuals who have different perspectives and different interests but roughly equal ability to make their voice heard. Where Gardner wants professionalism, we look to a kind of rumbustious amateurism. That said, there are some points of congruence between our perspectives. As we argue above, if democracy is to thrive, it requires some level of agreement about the problems it confronts (even if it disagrees about how best to solve these problems). This agreement may be difficult to generate within democracy itself. We may need to turn e.g. to scientific advice, ideally produced by disinterested and non-political experts, to generate this agreement.

Our views have consequences for how we regard the political effects of the Internet. In some ways, we agree with e.g., Kahne (this volume) and others who see the Internet enabling more participatory politics. However, while we see the Internet as a space for experiments and learning, we think that the failures of collective cognition online are as important as the successes. We doubt that the Internet is inherently liberatory — what matters most for our purposes is that it makes collective relationships that would otherwise be opaque more transparent, and hence easier to experiment with and to learn from.

More broadly, we emphasize that both Internet and non-Internet forms of collective cognition will only be able to take full advantage of diversity in conditions of *political equality*. Without such equality, some voices will dominate, and others be diminished, or silenced. Like Allen and a number of other contributors to this volume, we do not think that the conditions of equality are satisfied by standard arguments for deliberation — argument can and will be driven by enlightened self-interest and will take a wide variety of forms. We cannot specify in advance the particular conditions under which equality can best be achieved — they can only be discovered through experimentation. However, the commitment to equality is crucial. Much techno-utopianism suggests, directly or indirectly, that new forms of technologically mediated communication can serve as a substitute for more traditional forms of political (and economic) inequality. Our claim is quite different: that the conditions under which technological utopianism will work properly, by providing valuable experimental input into politics, are conditions of equality (which may or may not be *made easier* to achieve through technological advances, but which are surely not an inevitable by-product of technological change). This implies that a genuine commitment to democracy and the liberatory potential of technology is a commitment to political radicalism. We embrace this.

References

Anderson, Elizabeth (2007). "The Epistemology of Democracy." *Episteme: A Journal of Social Epistemology*, **3**: 8–22.

Badii, Remo and Antonio Politi (1997). *Complexity: Hierarchical Structures and Scaling in Physics*. Cambridge, England: Cambridge University Press.

Bakshy, Eytan, Itamar Rosenn, Cameron Marlow and Lada Adamic (2012). "The Role of Social Networks in Information Diffusion." In *Proceedings of the 21st International World Wide Web Conference [WWW 2012]* (Michael Rabinovich and Steffen Staab, eds.), pp. 519–528. New York: ACM. URL `http://arxiv.org/abs/1201.4145`. doi:10.1145/2187836.2187907.

Bell, Vaughan, Cara Maiden, Antonio Muñoz-Solomando and Venu Reddy (2006). ""Mind Control" Experience on the Internet: Implications for the Psychiatric Diagnosis of Delusions." *Psychopathology*, **39**: 87–91. URL `http://arginine.spc.org/vaughan/Bell_et_al_2006_Preprint.pdf`.

Blume, Lawrence and David Easley (2006). "If You're so Smart, Why Aren't You Rich? Belief Selection in Complete and Incomplete Markets." *Econometrica*, **74**: 929–966. URL `http://www.santafe.edu/media/workingpapers/01-06-031.pdf`.

Bowles, Samuel and Suresh Naidu (2008). *Persistent Institutions*. Tech. Rep. 08-04-015, Santa Fe Institute. URL `http://tuvalu.santafe.edu/~bowles/PersistentInst.pdf`.

Braybrooke, David and Charles E. Lindblom (1963). *A Strategy of Decision: Policy Evaluation as a Social Process*. Glencoe, Illinois: The Free Press.

Burt, Ronald S. (1992). *Structural Holes: The Social Structure of Competition*. Cambridge, Massachusetts: Harvard University Press.

Cohen, Joshua (1986). "An Epistemic Conception of Democracy." *Ethics*, **26**: 26–83.

Domingos, Pedro (1999). "The Role of Occam's Razor in Knowledge Discovery." *Data Mining and Knowledge Discovery*, **3**: 409–425. URL `http://www.cs.washington.edu/homes/pedrod/papers/dmkd99.pdf`.

Farrell, Henry and Melissa Schwartzberg (2008). "Norms, Minorities, and Collective Choice Online." *Ethics and International Affairs*, **22(4)**. URL `http://www.carnegiecouncil.org/resources/journal/22_4/essays/002.html`.

Hacker, Jacob S. and Paul Pierson (2010). *Winner-Take-All Politics: How Washington Made the Rich Richer — And Turned Its Back on the Middle Class*. New York: Simon and Schuster.

Hayek, Friedrich A. (1937). "Economics and Knowledge." *Economica*, **4**: 33–54. URL `http://www.econlib.org/Library/NPDBooks/Thirlby/bcthLS3.html`. Reprinted in (Hayek, 1948, pp. 33–56).

— (1945). "The Use of Knowledge in Society." *American Economic Review*, **35**: 519–530. URL `http://www.econlib.org/Library/Essays/hykKnw1.html`. Reprinted in (Hayek, 1948, pp. 77–91).

— (1948). *Individualism and Economic Order*. Chicago: University of Chicago Press.

— (1958). "Freedom, Reason, and Tradition." *Ethics*, **64**: 229–245.

Hirschman, Albert O. (1992). *Rival Views of Market Society and Other Recent Essays*. Cambridge, Massachusetts: Harvard University Press.

Hong, Lu and Scott E. Page (2004). "Groups of diverse problem solvers can outperform groups of high-ability problem solvers." *Proceedings of the National Academy of Sciences*, **101**: 16385–16389. URL `http://www.cscs.umich.edu/~spage/pnas.pdf`.

Huckfeldt, Robert, Paul E. Johnson and John Sprague (2004). *Political Disagreement: The*

Survival of Diverse Opinions within Communication Networks. Cambridge, England: Cambridge University Press.

Hume, David (1739). *A Treatise of Human Nature: Being an Attempt to Introduce the Experimental Method of Reasoning into Moral Subjects*. London: John Noon. Reprint (Oxford: Clarendon Press, 1951) of original edition, with notes and analytical index.

Hutchins, Edwin (1995). *Cognition in the Wild*. Cambridge, Massachusetts: MIT Press.

Judd, Stephen, Michael Kearns and Yevgeniy Vorobeychik (2010). "Behavioral Dynamics and Influence in Networked Coloring and Consensus." *Proceedings of the National Academy of Sciences*, **107**: 14978–14982. doi:10.1073/pnas.1001280107.

Kerstenetzky, Celia Lessa (2000). "Hayek: The Evolutionary and the Evolutionist." *Rationality and Society*, **12**: 163–184. doi:10.1177/104346300012002002.

Knight, Jack (1992). *Institutions and Social Conflict*. Cambridge, England: Cambridge University Press.

Knight, Jack and James Johnson (2011). *The Priority of Democracy: Political Consequences of Pragmatism*. Princeton: Princeton University Press.

Krogh, Anders and Jesper Vedelsby (1995). "Neural Network Ensembles, Cross Validation, and Active Learning." In *Advances in Neural Information Processing Systems 7 [NIPS 1994]* (Gerald Tesauro and David Tourtetsky and Todd Leen, eds.), pp. 231–238. Cambridge, Massachusetts: MIT Press. URL `http://books.nips.cc/papers/files/nips07/0231.pdf`.

Landemore, Hélène and Jon Elster (eds.) (2012). *Collective Wisdom: Principles and Mechanisms*. Cambridge, England: Cambridge University Press.

Laughlin, Patrick R. (2011). *Group Problem Solving*. Princeton: Princeton University Press.

Lazer, David and Allan Friedman (2007). "The Network Structure of Exploration and Exploitation." *Administrative Science Quarterly*, **52**: 667–694. URL `http://www.hks.harvard.edu/davidlazer/files/papers/Lazer_Friedman_ASQ.pdf`.

Lindblom, Charles E. (1965). *The Intelligence of Democracy: Decision Making through Mutual Adjustment*. New York: Free Press.

— (1977). *Politics and Markets: The World's Political-Economic Systems*. New York: Basic Books, 1st edn.

— (1982). "The Market as Prison." *The Journal of Politics*, **44**: 324–336. URL `http://www.jstor.org/pss/2130588`.

Mason, Winter A., Andy Jones and Robert L. Goldstone (2008). "Propagation of Innovations in Networked Groups." *Journal of Experimental Psychology: General*, **137**: 427–433. doi:10.1037/a0012798.

Mason, Winter A. and Duncan J. Watts (2012). "Collaborative Learning in Networks." *Proceedings of the National Academy of Sciences*, **109**: 764–769. doi:10.1073/pnas.1110069108.

Mercer, Neil (2000). *Words and Minds: How We Use Language to Think Together*. London: Routledge.

Mercier, Hugo and Dan Sperber (2011). "Why do humans reason? Arguments for an argumentative theory." *Behavioral and Brain Sciences*, **34**: 57–111. doi:10.1017/S0140525X10000968.

Mettler, Suzanne (2011). *The Submerged State: How Invisible Government Policies Undermine American Democracy*. Chicago: University of Chicago Press.

Meyer, John W., John Boli, George M. Thomas and Francisco O. Ramirez (1977). "World Society and the Nation-State." *American Journal of Sociology*, **103**.

Mill, John Stuart (1909). *Principles of Political Economy with some of their Applications to Social Philosophy*. London: Longmans, Green and Co., seventh edn. URL `http://www.econlib.org/library/Mill/mlPCover.html`.

Mitchell, Melanie (1996). *An Introduction to Genetic Algorithms*. Cambridge, Massachusetts: MIT Press.

Moore, Cristopher and Stephan Mertens (2011). *The Nature of Computation*. Oxford: Oxford University Press.

Nelson, Richard R. and Sidney G. Winter (1982). *An Evolutionary Theory of Economic Change*. Cambridge, Massachusetts: Harvard University Press.

Nemeth, Charlan J., Bernard Personnaz, Marie Personnaz and Jack A. Goncalo (2004). "The Liberating Role of Conflict in Group Creativity: A Study in Two Countries." *European Journal of Social Psychology*, **34**: 365–374. URL `http://escholarship.org/uc/item/4k70n7v8`. doi:10.1002/ejsp.210.

Nemeth, Charlan Jeanne (2012). "Minority Influence Theory." In *Handbook of Theories in Social Psychology* (P. A. M. Van Lange and A. W. Kruglanski and Higgins. E. T., eds.), vol. 2, pp. 362–378. New York: Sage. URL `http://escholarship.org/uc/item/1pz676t7`.

Nemeth, Charlan Jeanne and Margaret Ormiston (2007). "Creative Idea Generation: Harmony versus Stimulation." *European Journal of Social Psychology*, **37**: 524–535. URL `http://escholarship.org/uc/item/72b658w9`. doi:10.1002/ejsp.373.

Nielsen, Michael (2011). *Reinventing Discovery: The New Era of Networked Science*. Princeton: Princeton University Press.

Ober, Josiah (2008). *Democracy and Knowledge: Innovation and Learning in Classical Athens*. Princeton, New Jersey: Princeton University Press.

Page, Scott E. (2007). *The Difference: How the Power of Diveristy Creates Better Groups, Firms, Schools, and Societies*. Princeton, New Jersey: Princeton University Press.

— (2011). *Diversity and Complexity*. Princeton: Princeton University Press.

Pfeffer, Jeffrey and Robert I. Sutton (2006). *Hard Facts, Dangerous Half-Truths and Total Nonsense: Profiting from Evidence-Based Management*. Boston: Harvard Business School Press.

Popper, Karl R. (1960). *The Poverty of Historicism*. London: Routledge, 2nd edn. First edition, 1957.

— (1969). *Conjectures and Refutations: The Growth of Scientific Knowledge*. London: Routledge, 3rd edn.

Provine, William B. (1971). *The Origins of Theoretical Population Genetics*. University of Chicago Press.

Rosenblum, Nancy (2008). *On the Side of the Angels: An Appreciation of Parties and Partisanship*. Princeton, New Jersey: Princeton University Press.

Salganik, Matthew J., Peter S. Dodds and Duncan J. Watts (2006). "Experimental study of inequality and unpredictability in an artificial cultural market." *Science*, **311**: 854–856. URL http://www.princeton.edu/~mjs3/musiclab.shtml.

Salganik, Matthew J. and Duncan J. Watts (2008). "Leading the Herd Astray: An Experimental Study of Self-Fulfilling Prophecies in an Artificial Cultural Market." *Social Psychological Quarterly*, **71**: 338–355. URL http://www.princeton.edu/~mjs3/salganik_watts08.pdf.

Schapire, Robert E. and Yoav Freund (2012). *Boosting: Foundations and Algorithms*. Cambridge, Massachusetss: MIT Press.

Shalizi, Cosma Rohilla (2008). "Social Media as Windows on the Social Life of the Mind." In *AAAI 2008 Spring Symposia: Social Information Processing*. URL `http://arxiv.org/abs/0710.4911`.

Shalizi, Cosma Rohilla, Kristina Lisa Klinkner and Robert Haslinger (2004). "Quantifying Self-Organization with Optimal Predictors." *Physical Review Letters*, **93**: 118701. URL `http://arxiv.org/abs/nlin.AO/0409024`. doi:10.1103/PhysRevLett.93.118701.

Shalizi, Cosma Rohilla and Andrew C. Thomas (2011). "Homophily and Contagion Are Generically Confounded in Observational Social Network Studies." *Sociological Methods and Research*, **40**: 211–239. URL `http://arxiv.org/abs/1004.4704`. doi:10.1177/0049124111404820.

Shapiro, Carl and Hal R. Varian (1998). *Information Rules: A Strategic Guide to the Network Economy*. Boston: Harvard Business School Press, 1st edn.

Stiglitz, Joseph E. (2000). "The Contributions of the Economics of Information to Twentieth Century Economics." *The Quarterly Journal of Economics*, **115**: 1441–1478. URL `http://www2.gsb.columbia.edu/faculty/jstiglitz/download/papers/2000_Contributions_of_Economics.pdf`. doi:10.1162/003355300555015.

Sunstein, Cass R. (2003). *Why Societies Need Dissent*. Cambridge, Massachusetts: Harvard University Press.

Swartz, Aaron (2006). "Who Writes Wikipedia?" URL `http://www.aaronsw.com/weblog/whowriteswikipedia`.

Thaler, Richard and Cass R. Sunstein (2008). *Nudge: Improving Decisions about Health, Wealth, and Happiness*. New Haven, Connecticut: Yale University Press.

Cognitive Democracy and the Internet

Henry Farrell (GWU)
Cosma Shalizi (Carnegie-Mellon/Santa Fe)

Remaking Democratic Argument

- A project with big, sweeping intentions

- Reground democratic theory in new work on complex problem solving.

- Test arguments with modeling (in progress) and empirical data from experiments in 'collective cognition' on the Internet (later).

- Think pragmatically about how to fold experimentalism into democratic practice.

Basic intuitions

1. That the biggest problems that democracies face are complex (in technical sense).
2. That the best way to find 'good' solutions to these problems is through maximizing diversity of perspective.
3. That democracy provides a more promising means to this than either markets or hierarchy.
4. That emerging results can be used to improve democracy's problem-solving capacity.

Complex Problems

- Result from the interaction of diverse entities.

- Tend to generate a *rugged solution space* with peaks and troughs.

- Simple algorithms (such as 'hill-climbing') weak at identifying good solutions.

- No guarantee that a local optimum is a global optimum, or even a *good* optimum.

- Useful approximation to political problems and likelihood of getting stuck at poor local optimum

Page-Hong Result

- Page-Hong find that in simulations of complex problem solving, diversity trumps ability.

- High ability problem solvers tend to employ similar heuristics.

- For complex problems, groups of high ability agents are outperformed by groups of (more diverse) less able problem solvers.

- Suggests that harnessing diversity of perspectives is key to complex problem solving.

Democracy and Diversity

- New approach to understanding democracy - in terms of its ability to harness diverse perspectives for problem solving.
- Builds on and extends Page-Hong results.
- Also - recent work by Sperber and Mercier on group problem solving and surprising usefulness of cognitive bias.
- Also - experiments by Mason and Watts on problem-solving by rival groups (rough proxy for partisanship).

Current Research I

- Applying arguments from mathematical evolutionary biology and information theory to institutional change.
- Intuition is that societies with greater equality will be better able to take advantage of internal diversity to reshape institutions in salutary ways.
- Work in progress (with Danielle Alan, IAS) using Athens-Sparta comparison as toy models to work out empirics.
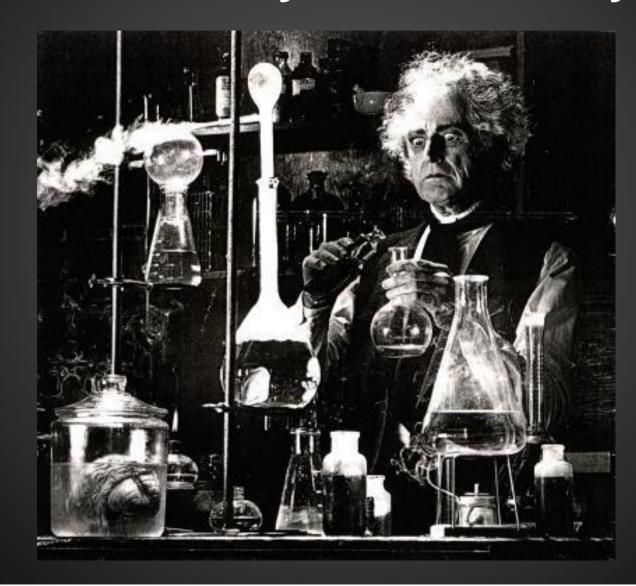
Current Research II

- Simulations of complex problem solving.
- Complex problems - linear functions of a high dimensional vector.
- Problem solvers - two classes of agents - elites (with higher ability), and non-elites (with lower).
- Try to match problems with varying coincidence (reflecting likely clashes of values between elites and non-elites)
- Question - how do different models work in solving general social problems?

Agent Problem-Solving

- Agents use a tree-based heuristic with few historical training cases.
- Are networked together and can learn from neighbours (so that we can study the consequences of network topology in later iterations).
- To be done - scale model, more complex solutions as well as problems, allow different communities within network to have different interests.

For the Future: The Internet as Laboratory of Democracy

Empirical Problems

- Lots of interesting data:
  - StackExchange
  - Wikipedia
- But very hard to figure out metrics for 'success' in problem solving.
- Will likely require a mixture of human and machine coding.
- Which will in turn require funding ...

Cognitive Democracy in Theory

New ideas (from cognitive psychology, complexity theory, network theory, machine learning) provide new insights into democracy.

Quite different e.g. from epistemic democracy approaches based on Condorcet etc - focus is on modeling collective learning and information processing rather than aggregation of pre-existing knowledge.

Basis for understanding cognitive benefits of democracy and other macro-institutions.

Cognitive Democracy in Practice

Obvious empirical fit with large-scale availability of data on social interaction and 'collective cognition' on the Internet.

No coincidence given machine learning connection.

Offer benefits for both research and practice.

But important difficulties remain, especially in measuring `success' in problem solving using existing data.