

LA-UR- 09-00324

*Approved for public release;
distribution is unlimited.*

Title: Revisiting the Age of Enlightenment from a Collective Decision Making Systems Perspective

Author(s): Marko Rodriguez, T-5
Jennifer Watkins, IAT-1

Intended for: IEEE Intelligent Systems (Special Issue on Governance)



Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the Los Alamos National Security, LLC for the National Nuclear Security Administration of the U.S. Department of Energy under contract DE-AC52-06NA25396. By acceptance of this article, the publisher recognizes that the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

TITLE: Revisiting the Age of Enlightenment from a Collective Decision Making Systems Perspective

AUTHORS: Marko A. Rodriguez and Jennifer H. Watkins (both LANL)

ABSTRACT: Many of today's most valued principles of governance were developed during the eighteenth century's Age of Enlightenment. Twenty-first century democratic nations may benefit from revisiting the idealized decision making systems developed during the Enlightenment and reframing them within the socio-technical context of the Information Age. This article explores Thomas Paine's (English: 1737--1809) requirements of representation, Adam Smith's (Scottish: 1723--1790) self-interested actors, and Marquis de Condorcet's (French: 1743--1794) optimal decision making groups.

Revisiting the Age of Enlightenment from a Collective Decision Making Systems Perspective

Marko A. Rodriguez and Jennifer H. Watkins

Los Alamos National Laboratory

Los Alamos, NM 87501

Abstract—Many of today's most valued principles of governance were developed during the eighteenth century's Age of Enlightenment. Twenty-first century democratic nations may benefit from revisiting the idealized decision making systems developed during the Enlightenment and reframing them within the socio-technical context of the Information Age. This article explores Thomas Paine's (English: 1737–1809) requirements of representation, Adam Smith's (Scottish: 1723–1790) self-interested actors, and Marquis de Condorcet's (French: 1743–1794) optimal decision making groups.

Index Terms—collective decision making, e-participation, e-democracy, computational social choice theory

I. INTRODUCTION

Eighteenth century Europe is referred to as The Age of Enlightenment, a period when prominent philosophers began to question traditional forms of authority and power and the moral standards that supported these forms. The radical ideas of the Enlightenment helped to shape various revolutions across Europe as well as contribute to the formalization of the governing structure that would determine the course of a new nation: the United States of America. It was during this time that many of the values we hold in high esteem today were set in their present form. However, the present implementation of these values in decision making can benefit from a revamping that makes greater use of the technical advances made in the present Information Age. Such a revitalization of the modern decision making infrastructure can yield a greater precision and as such, better embody the ideals, not simply the mechanism, of these original thinkers. To move in this direction, the principle of citizen representation as articulated by Thomas Paine (English: 1737–1809) and the principle of competitive actors for the common good as articulated by Adam Smith (Scottish: 1723–1790) are considered from a techno-social, collective decision making systems perspective. Moreover, the validity of these ideals can be understood within the mathematical formulations of Marquis de Condorcet's (French: 1743–1794) requirements for optimal decision making.

M.A. Rodriguez is with T-5/Center for Non-Linear Studies, Los Alamos National Laboratory, Los Alamos, NM 87501 USA e-mail: marko@lanl.gov.

Jennifer H. Watkins is with the International and Applied Technology Group, Los Alamos National Laboratory, Los Alamos, NM 87501 USA e-mail: jhw@lanl.gov.

This research was conducted by the Collective Decision Making Systems (CDMS) project at the Los Alamos National Laboratory (<http://cdms.lanl.gov>).

Manuscript received April 19, 2009; revised January 11, 2009.

II. CONDORCET JURY THEOREM: ENSURING OPTIMAL DECISION MAKING

Marquis de Condorcet ardently supported equal rights and free and universal public education. These ideals underpinned his investigations in mathematics, philosophy, and political science, making his work as significant today as it was then. Condorcet was one of the first to widely apply mathematics to the nascent field of political science. Of his mathematical investigations into the nature of politics, one of his most famous results is the Condorcet statement and its associated theorem. In his 1785 *Essai sur l'Application de l'Analyse aux Probabilités des Decisions prises à la Pluralité des Voix* (english translation *Essay on the Application of Analysis to the Probability of Majority Decisions*), Condorcet states that when a group of “enlightened” decision makers chooses between two options under a majority rule, then as the size of the decision making population tends toward infinity, it becomes a certainty that the best choice is rendered [1]. The first statistical proof of this statement is the Condorcet jury theorem expressed as follows. Imagine there exists n independent decision makers and each decision maker has a probability $p \in [0, 1]$ of choosing the best of two options in a decision. If $p > 0.5$, meaning that each individual decision maker is enlightened, and as $n \rightarrow \infty$, the probability of a majority vote outcome rendering the best decision approaches certainty at 1.0. This is known as the “light side” of the Condorcet jury theorem. The “dark side” of the theorem states that if $p < 0.5$ and as $n \rightarrow \infty$, the probability of a majority vote outcome rendering the best decision approaches 0.0. Figure 1 plots the relationship between p and n , where a 100% probability of the group rendering the best decision is white and a 0% probability is black.

The Condorcet jury theorem is one of the original formal justifications for the application of democratic principles to government. Moreover, it verifies the intuition that it is important to utilize a large group of enlightened decision makers. Democracies do not rely on a single decision maker, but instead use senates, parliaments, and referendums to add increase the size of the voting population in order to better ensure optimal decisions are rendered. Additionally, it is a democratic ideal to include those who will be affected by the decision in the decision making process. Whether for diversity of opinion or ideals, including many people in decision-making is already a goal of democracy. The necessity of enlightened decision makers is apparent in the lengths to

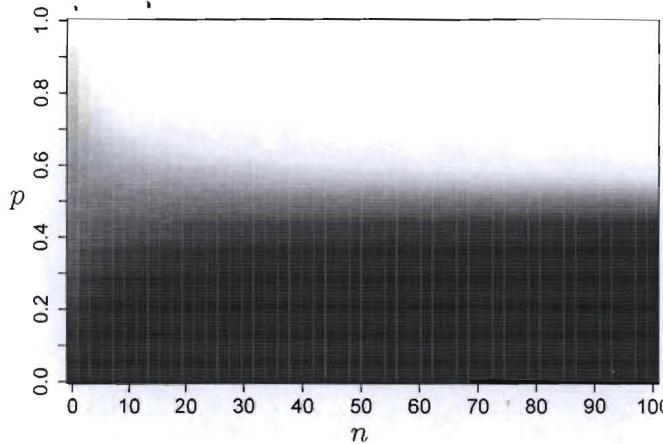


Fig. 1. The relationship between $p \in [0, 1]$ and $n \in (1, 2, \dots, 100)$ according to the Condorcet jury theorem model. Darker values represent a low probability of a majority vote rendering the best decision and the lighter values represent a high probability of a majority vote rendering the best decision.

which democratic societies go to choose leaders. Elections are opportunities to compare candidates' expertise, skills, and other leadership qualities that serve as indicators of one's ability to make optimal decisions in the future. It is hoped that these decisions are voted on by an informed (or enlightened) electorate.

While the Condorcet jury theorem does not reveal any startling conditions for a successful democracy, it does distill the necessary conditions to two variables (under simple assumptions). If a decision making group has a large n and a $p > 0.5$, then the group is increasing its chances of optimal decision making. Unfortunately, the theorem does not suggest a means to achieve these conditions. Technology-enabled social algorithms provide an means by which to achieve the conditions of the "light" side of the Condorcet jury theorem and thus, achieve optimal decision making. Given the advances in information systems, communication technologies, and social algorithms, it may be possible to build and deploy a decision making infrastructure that solicits the decision making influence of those individuals that are more likely than not to choose the best decision for a particular problem [2]. This article presents two such algorithms that show promise as mechanisms for the future of societal-scale, collective decision making. One algorithm exaggerates Thomas Paine's requirements of representation in order to accurately simulate the behavior of a large decision making population ($n \rightarrow \infty$), and the other makes use of Adam Smith's market philosophy to induce participation by the enlightened within that population ($p \rightarrow 1$). Both algorithms together, utilize the Condorcet jury theorem to the society's advantage.

III. DYNAMICALLY DISTRIBUTED DEMOCRACY: SIMULATING A LARGE DECISION MAKING POPULATION

Thomas Paine was originally born in England, but in his middle years, he relocated to America due to the recommendation of Benjamin Franklin. It was in America, in the time leading up to the American Revolution, where his enlightened

ideals were well received. In 1776, the same year as the authoring of the Declaration of Independence, Thomas Paine wrote a widely distributed pamphlet entitled *Common Sense* which outlined the values of a democratic regime [3]. This pamphlet discussed the equality of man and the necessity for all those at stake, to partake in the decision making processes of the group. When populations are small "some convenient tree will afford them a State house", but as the population increases the necessity for representatives who act on behalf of their constituents becomes the necessary consequence. Moreover, and being the central tenet of political representation, it is important that representatives "act in the same manner as the whole body would act were they present". In other words, it is necessary to simulate the behavior of a large group even if only a subset of that group can actively participate.

One of the conclusions of the Condorcet jury theorem is that, assuming the other conditions of the theorem are met, a larger decision making population is better than a smaller one. To explore the implications of achieving this large n , consider the various forms of government. Assuming a $p > 0.5$, the Condorcet jury theorem would hold that a representative democracy would be more likely to make optimal decisions than a dictatorship and that a direct democracy would be more likely to make optimal decisions than the representative democracy. In practice, the desire for optimal decision making is tempered by the tremendous burden that constant voting would impinge on citizens (not to mention the logistical problems such a voting system would incur). For this reason, representation is a necessity. However, if the representative body votes as its represented constituents would have voted, then representative democracy and direct democracy are equal. The problem then is to develop a social algorithm that accurately simulates the decision making behavior of the whole population without demanding the burdensome amount of time this would require.

Before presenting such a solution, it is important to define a collective decision under total participation (i.e. a direct democracy with full participation). Assuming a two-option majority rule, an individual's judgement can be placed along a continuum between the two options such that the "political tendency" of decision maker i can be denoted $x_i \in [0, 1]$. For example, given United States politics, a political tendency of 0 represents a fully Republican perspective, a tendency of 1 represents a fully Democratic perspective, and a tendency of 0.5 denotes a moderate. Given this definition, there are two ways to quantify the group. One way is to calculate the average tendency of all decision makers. That is $d^{\text{tend}} = \frac{1}{n} \sum_{i=0}^{i=n} x_i$, where $d^{\text{tend}} \in [0, 1]$ is the collective tendency of the group. Given a uniform distribution of political tendency within x , the global tendency approaches 0.5 as the size of the group increases towards infinity. The other way to quantify the group is to require that the individual's tendency be reduced to a binary option (i.e. a two option vote). If a decision maker has a political tendency that is less than 0.5, then they will vote 0. If they have a tendency equal to 0.5 then a fair coin toss will determine their vote. Otherwise, for a tendency greater than 0.5, they will vote 1. This majority wins vote is denoted d^{vote} ,

where $d^{\text{vote}} \in \{0, 1\}$.

Imagine a direct democracy in the purest sense, where a raise of hands or a shout of voices is replaced by an Internet architecture and a sophisticated error- and fraud-proof ballot system. All citizens have the potential to vote on any decisions they wish; if they could not vote on a particular decision for whatever reason, they could abstain from participating. In practice, not every decision will be voted on by all n citizens. Citizens will be constrained by time pressures to only participate in those votes in which they are most informed or most passionate. If we assume that all citizens have a tendency, whether they vote or not, how would the collective tendency and collective vote change as citizen participation waned?

Let $d_{100}^{\text{tend}} \in [0, 1]$ and $d_{100}^{\text{vote}} \in \{0, 1\}$ denote the tendency and vote given by 100% participation. Let $d_k^{\text{tend}} \in [0, 1]$ and $d_k^{\text{vote}} \in \{0, 1\}$ denote the collective tendency and vote if only k -percent of the population participated. The error in the calculation of the collective tendency for k -percent participation is calculated as

$$e^{\text{tend}} = |d_{100}^{\text{tend}} - d_k^{\text{tend}}| \in [0, 1].$$

The further away the active voters' collective tendency is from the population's collective tendency, the higher the error. The gray line in Figure 2 plots the relationship between k and e^{tend} . As citizen participation wanes, the ability for the remaining, active participants to reflect the tendency of the whole becomes more difficult. Next, the error in the collective vote is calculated as the proportion of voting outcomes that are different than what a fully participating population would have voted and is denoted e^{vote} . The gray line in Figure 3 plots the relationship between k and e^{vote} . As participation wanes, the proportion of decisions that differ from what would have occurred given full participation decreases. As with collective tendency, a small active voter population is unable to replicate the behavior of the whole.

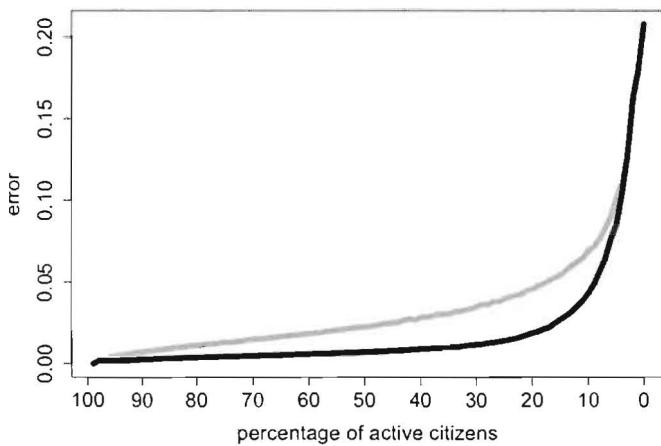


Fig. 2. The relationship between k and e^{tend} for direct democracy (gray line) and dynamically distributed democracy (black line). The plot provides the average error over a simulation that was run with 1000 artificially generated networks composed of 1000 citizens each. The preferential attachment, network growth algorithm was used to generate a degree distribution that is reflective of typical social networks "in the wild" (i.e. scale-free properties). Moreover, links between citizens of opposing political tendencies are less likely to link than like citizens.

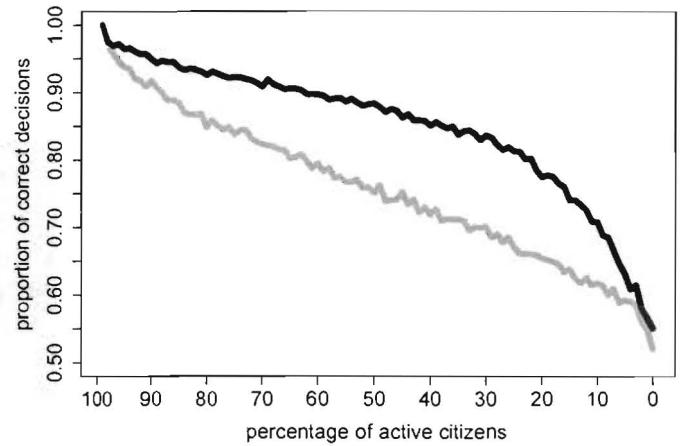


Fig. 3. The relationship between k and e^{vote} for direct democracy (gray line) and dynamically distributed democracy (black line). The plot provides the proportion of identical, correct votes over a simulation that was run with 1000 artificially generated networks composed of 1000 citizens each.

Dynamically distributed democracy (DDD) is an algorithm that provides a means by which a subset of the population can simulate the decision making results of the whole population [4]. As such, the algorithm reflects the primary tenet of representation as originally outlined by Thomas Paine and quantifiably valued by Marquis de Condorcet. The argument for the use of DDD as mechanism for representation goes as follows. Not everyone in a population needs to vote as others in that same population more than likely have a nearly identical political tendency. What does need to be recorded is the frequency of that sentiment in the population. If an active voter is similar in tendency to 10 non-participating citizens, then the active voter's ballot can be weighted by 10 to reflect the tendencies of the non-participating citizens. DDD accomplishes this weighting through a similarity- or trust-based social network that is used to direct voting power to active voters so as to mitigate the error incurred by waning citizen participation.

As previously stated, let $\mathbf{x} \in [0, 1]^n$ denote the political tendency of each citizen in this population, where \mathbf{x}_i is the tendency of citizen i and, for the purpose of simulation, is determined from a uniform distribution. Assume that every citizen in a population of n citizens uses some social network-based system to create links to those individuals that they believe reflect their tendency the best. In practice, these links may point to a close friend, a relative, or some public figure whose political tendencies resonate with the individual. Let $\mathbf{A} \in [0, 1]^{n \times n}$ denote the link matrix representing the network, where the weight of an edge, for the purpose of simulation, is denoted

$$\mathbf{A}_{i,j} = \begin{cases} 1 - |\mathbf{x}_i - \mathbf{x}_j| & \text{if link exists} \\ 0 & \text{otherwise.} \end{cases}$$

In words, if two citizen's are identical in their political tendency, then the strength of the link is 1.0. If their tendencies are completely opposing, then their trust (and the strength of the link) is 0.0. Note that an assortativity parameter [5] is used to bias the connections in the network towards citizens

with similar tendencies. The assumption here is that given a system of this nature, it is more likely for citizens to create links to similar-minded individuals than to those who opinions are quite different. It is the purpose of the social network to capture these links. The resultant link matrix \mathbf{A} is then normalized to be row stochastic in order to generate a probability distribution over the weights of the outgoing edges of a citizen. Figure 4 presents an example of an $n = 100$ artificially generated trust-based social network, where red denotes a tendency of 0.0 (Republican), purple a tendency of 0.5 (moderate), and blue a tendency of 1.0 (Democrat).

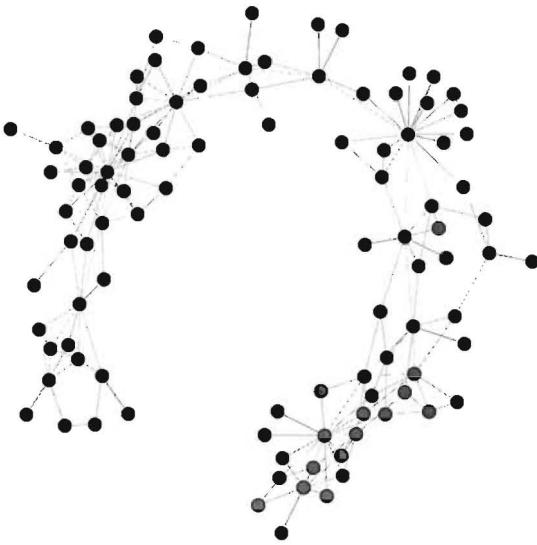


Fig. 4. A visualization of the network of trust links between citizens. Each citizen's color denotes their "political tendency", where full red is 0, full blue is 1, and purple is 0.5. The layout algorithm chosen is the Fruchterman-Reingold layout.

Given this social network infrastructure, it is possible to better ensure that the collective sentiment is appropriately represented through a weighting of the active, participating population. Given that the set of active voters can change from decision to decision, the algorithm gets its name: dynamically distributed democracy. Every citizen, active or not, is initially provided with $\frac{1}{n}$ "vote power" and this is represented in the vector $\bar{y} \in \mathbb{R}_+^n$, such that the total amount of vote power in the population is 1. Let $y \in \mathbb{R}_+^n$ denote the total amount of vote power that has flowed to each citizen over the course of the iterative algorithm. Finally, $\mathbf{a} \in \{0, 1\}^n$ denotes whether citizen i is participating ($a_i = 1$) in the current decision making process or not ($a_i = 0$). The values of \mathbf{a} are biased by an unfair coin that has probability k of making the citizen an active participant and $1 - k$ of making the citizen inactive. The iterative DDD algorithm is presented below, where \circ denote entry-wise multiplication and $\epsilon \approx 1$.

In words, active citizens serve as vote power "sinks" in that once they receive vote power, from themselves or from a neighbor in the network, they do not pass it on. Inactive citizens serve as vote power "sources" in that being inactive, they propagate their vote power over the network links to their neighbors iteratively until all (or ϵ) vote power has

```

 $y \leftarrow 0$ 
while  $\sum_{i=0}^{i < n} y_i < \epsilon$  do
     $y \leftarrow y + (\bar{y} \circ \mathbf{a})$ 
     $\bar{y} \leftarrow \bar{y} \circ (1 - \mathbf{a})$ 
     $\bar{y} \leftarrow \mathbf{A} \bar{y}$ 
end

```

reached active citizens. At which point, the tendency in the active population is defined as $\delta^{\text{tend}} = \mathbf{x} \cdot y$. Figure 2 plots the error incurred using DDD (black line), where k denotes the percentage of actively participating citizens. The error is defined as

$$e^{\text{tend}} = |d_{100}^{\text{tend}} - \delta^{\text{tend}}| \in [0, 1].$$

Next, the collective vote, δ_k^{vote} , is determined by a weighted majority as dictated by the vote power accumulated by active participants. Figure 3 plots the proportion of votes that are different from what a fully participating group would have rendered using DDD (black line). In essence, if a citizen, for any reason, is unable to participate in a decision making process, then they may abstain from participating knowing that the underlying social network will accurately distribute their vote power to their neighbor or neighbor's neighbor. The DDD algorithm is better able to simulate full active participation even as citizens abstain from participation.

Thomas Paine outlines that representatives should maintain "fidelity to the public" and believes this is accomplished through frequent elections [3]. The utilization of an Internet-based social network system afford repeated "elections" in the form of citizens creating outgoing links to other citizens as they please, when they please, and to whom they please. That is, citizens can dynamically choose representatives who need not be picked for a handful of potential representatives and, upon faltering in their ability to represent a citizen, can immediately have an incoming edge retracted from them. Such an architecture turns the representative's status from that of elected public official to that of self-intentioned citizen.

While many countries have political institutions that are set up according to a left, right, and moderate agenda, the individual perspectives of a citizen may be more complex. In many cases, the complexities of a citizen's political tendencies may only be amenable to a multi-dimensional representation. In a multi-relational trust-based social network, the links are augmented with labels in order to denote the type of trust one citizen has for another. In this way, voting power propagates over the links in a manner that is biased to the domain of the decision. For example, citizen i may trust citizen j in the domain of "education" but not in the domain of "health care". The design of such a system has been articulated in [6]. While a mechanism for dynamically distributing vote power amongst a set of active decision makers within a population has been presented, the means by which ballots are posed has not. Research into societal-scale, decision support systems is presented in [7].

With the Internet, supporting Web technologies, and DDD, it is possible to dynamically determine a representative-layer of government that more accurately reflects a full direct

democracy. In this respect, the larger population helps to ensure, according to the Condorcet jury theorem, that the decisions are either definitely right or definitely wrong. Other technologies could be utilized to induce only those that are more likely than not to choose the optimal decision to actively participate.

IV. PREDICTION MARKETS: INCENTIVIZING AN ENLIGHTENED MAJORITY

Adam Smith was a Scottish moral and economic philosopher who is best known for his two most famous works entitled *The Theory of Moral Sentiments* (1759) and *An Inquiry into the Nature and Causes of the Wealth of Nations* (1776). It is in the latter work that Adam Smith outlines the benefits of utilizing market mechanisms to accurately determine, in a decentralized fashion, the value of goods and production requirements. It is articulated that, in some cases, the greatest contribution to the common good can be derived not by the cooperation, but instead by the competition of those agents involved in the market. When an agent pursues “his own interest he frequently promotes that of the society more effectually than when he really intends to promote it” [8]. Market mechanisms can be used beyond the determination of commodity prices and can be generally applied to information aggregation and ultimately, to optimal collective decision making.

It is no easy task to ensure that a decision making population contains an enlightened majority. The means to gauge whether a particular individual will be an optimal decision maker fall on proxy measures such as IQ and expertise. If these measures are to successfully determine whether a decision maker has $p > 0.5$, they must be carefully considered for applicability to the decision at hand. Different decisions require different knowledge. Furthermore, in democratic governments, it is an ideal to put as few strictures on political participation as possible (e.g., women’s suffrage is more democratic). Therefore, the goal is to provide a mechanism that calls for the fewest *a priori* restrictions on participation and assumptions about the knowledge required. What is needed then is a self-selection mechanism that incentivizes those who have knowledge regarding the problem and are confident in their knowledge and discourages participation from others without forbidding it. A prediction market is such a mechanism.

Prediction markets reward individuals for buying low and selling high, thus encouraging those who believe they know which way the market will move to contribute their information in the form of the price at which they purchase and sell shares. A prediction market differs from commodity markets (such as the New York Stock Exchange) in that stocks represent future events. For example, given the market question “Will prediction markets be used in U.S. government?”, shares of stocks in a “yes” outcome and in a “no” outcome are purchased and sold on the market. These outcomes can be placed on a continuum such that “no” is 0 and “yes” is 1. Then, similar to the previous section, an individual’s judgement is a point along the continuum between the two outcomes such that the “prediction” of decision maker i can be denoted $x_i \in [0, 1]$. The price at which a share is bought or sold denotes the

collective prediction of the decision makers. Let the number of shares of “0” stock on the market be denoted s_0 and the number of “1” stocks on the market be s_1 . There are many ways to determine stock prices, such as through a traditional continuous double-auction format, or using a market maker format such as Hanson’s logarithmic market scoring rule [9] which states that the current price of outcome “0” is

$$c_0 = \frac{e^{s_0}}{e^{s_0} + e^{s_1}},$$

where e is the base of the natural logarithm. Similarly, the current price of outcome “1” is $c_1 = \frac{e^{s_1}}{e^{s_0} + e^{s_1}}$ or simply $1 - c_0$.

Suppose there exists a prediction market in a future event with two possible outcomes represented by two stocks denoted “0” and “1”. The outcome, \mathbf{o} , is determined by finding the mean of a finite collection of environmental signals, each between 0 or 1. This environment could be considered the voting population in a general election between a Republican (“0”) and a Democrat (“1”) or any scenario where the outcome with the greater representation in the environment is considered the optimal decision. Also, suppose there exists a set of n participants in the market, hereafter referred to as traders. A trader has access to some of the factors in the environment. The mean of these accessible signals is the trader’s prediction, \mathbf{x}_i . Trader i ’s knowledge is denoted p_i and is equal to the proportion of signals in the environment to which the the trader has access. Trader i has $p_i < 0.5$ if they have access to fewer than half of the environmental signals. Similarly, trader i has $p_i > 0.5$ if they have access to more than half of the environmental signals. For the simulation, $\mathbf{p} \in [0, 1]^n$ is normally distributed around p according to a binomial variance. Thus, the average trader knowledge refers to the mean of \mathbf{p} across all traders rather than to a homogenous collective of traders all with equivalent p . This alternative Condorcet jury theorem is proven in [10] and yields a slower convergence than that depicted in Figure 1. However, for the sake of a market simulation, where there exist a heterogenous set of traders, it is necessary to assume such a distribution.

Given this framework, the purpose of the simulation is to demonstrate that decision making can remain optimal even as the average knowledge of the traders decreases (i.e. as p decreases). The market price, c_0 , is compared to the actual outcome \mathbf{o} , such that a more optimal decision is one with the lowest absolute difference. More specifically, the error of the market prediction is defined as

$$e^{\text{pred}} = |c_0 - \mathbf{o}|.$$

Traders interact with the market in turn. During a trader’s turn, they compare their belief, \mathbf{x}_i , to the current market value, c_0 . If $\mathbf{x}_i > c_0$, the trader will buy shares of “0” stock and, in the same instance, sell shares of “1” stock. If traders do not consider their amount of knowledge, then the participation by the traders is homogenous. That is, traders lacking knowledge will behave the same as traders with knowledge. The number of shares that a trader will buy, is determined by the homogenous decision algorithm,

$$s^{\text{homo}} = \sigma(|\mathbf{x}_i - c_0|),$$

where $\sigma \in [0, 1]$ serves as a smoothing factor to eliminate large swings in the market value. Note that the value of σ is the same for all traders in the homogenous group. When trader's do not consider their knowledge, they participate in the market by purchasing enough stocks to move the market price a σ proportion between the current price and the trader's prediction. Each time shares are traded, the traded stock's share count, s , is incremented if the trader bought or decremented if the trader sold. Thus, as shares are purchased and sold, c_0 fluctuates.

The gray line in Figure 5 plots the relationship between the mean of p and e^{pred} for trader's not considering their knowledge. As the average trader knowledge decreases, the poorly informed agents continue to exert as much power in moving the price in the market as the well-informed traders, causing market error to increase.

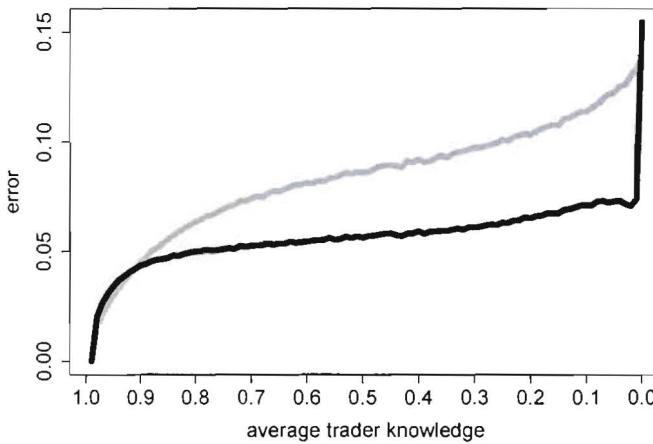


Fig. 5. The relationship between p and e^{pred} for homogenous participation (gray line) with $\sigma = 0.5$ and self-selecting heterogeneous participation (black line). The plot provides the average error over a simulation of 1000 traders over 1000 runs. The space between the lines represents accuracy that is gained by allowing trader's to participate with the number of shares they can buy or sell at one time varying based on their knowledge.

A prediction market provides incentives for contributing predictions that are strongly backed by knowledge. The higher the demand for a stock, the higher its price. Therefore, simply following the crowd, without personal knowledge that this prediction is accurate, comes with a significant financial risk. Additionally, money is earned by buying a stock when the price is low and selling it when the price is high. The likelihood that a trader can predict that a stock will increase in value is based on their knowledge of the environment. In a market, the choice of when to participate, which stocks to buy and sell, and how many is left the each trader. Individuals select their own extent of involvement with the knowledge that poor choices result in a financial loss. To simulate the effect of an incentive system on trader behavior, a heterogeneous decision algorithm is used:

$$s^{\text{hetero}} = p_i (|x_i - c_0|),$$

where p_i is the knowledge of the trader. Trading decisions are identical to those of the homogenous decision algorithm except that the number of shares bought or sold is dependent

upon the trader's knowledge. Trader's with perfect knowledge buy or sell the number of shares necessary to move the market price to their prediction and all other traders buy or sell some fraction thereof.

When trader's bias their participation in the market based on how confident they are in their prediction, the average trader knowledge has less of an effect on the market error. The black line in Figure 5 plots the relationship between p and e^{pred} for trader's that allow the degree of knowledge to effect the number of shares they will buy or sell. As the average trader knowledge decreases, the poorly informed traders (those with $p_i < 0.5$), buy fewer shares and so exert less power in moving the market price than well informed traders (those with $p_i > 0.5$). Therefore, the market error remains low, even as the proportion of traders with $p_i < 0.5$ increases.

Not only do prediction markets amplify the abilities of those with $p_i > 0.5$, they also increase the abilities of individuals to contribute useful information. Because prediction markets and their associated incentives operate through time, an additional benefit of a market mechanism is that it motivates the traders to become more knowledgeable or informed through time. Traders have a monetary incentive to seek out new information, especially if the environment is changing, and update their positions in the market accordingly. This further improves the likelihood that a collective will have $p > 0.5$ by the time the market closes.

V. CONCLUSION

According to the Condorcet jury theorem, the key to optimal decision making lies in the use of a large decision making collective, where individuals act independently and with each having a probability of being correct that exceeds 50%. As articulated by Thomas Paine, the involvement of a large collective of decision makers is burdensome and logically challenging. Therefore, social algorithms are needed to simulate the benefits of a large collective when only a subset participates. The solution presented in this article is dynamically distributed democracy (DDD). DDD uses a social network to automatically forward the decisions of those who don't participate in the form of additional weighting to those who do. Using DDD, even as participation wanes, the accuracy of the representation remains high. The stipulation that the collective be more than 50% likely to be correct is difficult to translate to real-world decision making situations simply because there is no way of knowing which decision makers meet this condition. Social algorithms are needed to incentivize those with $p > 0.5$ of choosing the optimal answer to participate without *a priori* restricting participation. This article presented a market environment as a solution. Prediction markets emphasize both self-selection such that there are disincentives for participating without confidence in your solution and competition such that individuals are likely to act independently. Markets allow participants to choose when, if, and the extent of their participation. It was shown that in a market, even as the probability of being correct decreases, the accuracy of the market remains high.

As has been demonstrated, technology-enabled social algorithms have the potential to mold collections of decision

makers to the statistical ideals that yield optimal decisions. However, the promise of social algorithms is even greater. They have the potential to allow for the values of democratic nations to be more perfectly executed. For example, DDD supports the implementation of direct democracies allowing for greater governance by the people and prediction markets allow an unrestricted populace to participate without degrading the integrity of the decision making. The future of democracy revolves around developing such social algorithms and their Internet-based implementations so as to not only produce optimal decisions, but do so in accordance with the highest values of democratic societies.

REFERENCES

- [1] M. de Condorcet, "Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix," Paris, France, 1785.
- [2] J. H. Watkins and M. A. Rodriguez, *Evolution of the Web in Artificial Intelligence Environments*, ser. Studies in Computational Intelligence. Berlin, DE: Springer-Verlag, 2008, ch. A Survey of Web-Based Collective Decision Making Systems, pp. 245–279. [Online]. Available: <http://repositories.cdlib.org/hcs/WorkingPapers2/JHW2007-1/>
- [3] T. Paine, "Common sense," American colonies, 1776.
- [4] M. A. Rodriguez and D. J. Steinbock, "A social network for societal-scale decision-making systems," in *Proceedings of the North American Association for Computational Social and Organizational Science Conference*, Pittsburgh, PA, USA, 2004. [Online]. Available: <http://arxiv.org/abs/cs.CY/0412047>
- [5] M. Newman, "Assortative mixing in networks," *Physical Review Letters*, vol. 89, no. 20, 2002.
- [6] M. A. Rodriguez, "Social decision making with multi-relational networks and grammar-based particle swarms," in *Proceedings of the Hawaii International Conference on Systems Science*. Waikoloa, Hawaii: IEEE Computer Society, January 2007, pp. 39–49. [Online]. Available: <http://arxiv.org/abs/cs.CY/0609034>
- [7] M. Turrof, S. R. Hiltz, H.-K. Cho, Z. Li, and Y. Wang, "Social decision support system (SDSS)," in *Proceedings of the Hawaii International Conference on Systems Science Hawaii International Conference on Systems Science*. Waikoloa, Hawaii: IEEE Computer Society, January 2002, pp. 81–90.
- [8] A. Smith, *An Inquiry into the Nature and Causes of the Wealth of Nations*. London, England: W. Strahan and T. Cadell, Londres, 1776.
- [9] R. Hanson, "Logarithmic market scoring rules for modular combinatorial information aggregation," *Journal of Prediction Markets*, vol. 1, no. 1, pp. 3–15, February 2007.
- [10] G. Owen, B. Grofman, and S. L. Feld, "Proving a distribution-free generalization of the Condorcet jury theorem," *Mathematical Social Sciences*, vol. 17, no. 1, pp. 1–16, February 1989.