Enabling democratic decision-making with collective belief models and game theoretic analysis

Kshanti A. Greene, Joseph M. Kniss and George F. Luger
Department of Computer Science
University of New Mexico
kshanti, jmk, luger@cs.unm.edu

Abstract

We introduce a new approach to aggregating the beliefs and preferences of many individuals to form models for democratic decision-making. Traditional *social choice functions* used to aggregate beliefs and preferences attempt to find a *single* consensus model, but produce inconsistent results when a stalemate between opposing opinions occurs. Our approach combines the probabilistic beliefs of many individuals using Bayesian decision networks to form *collectives* such that the aggregate of each collective, or *collective belief* conforms to principles of rationality in social choice theory. We first extract the symbolic preference order from each individual's quantitative Bayesian beliefs. We then show that if a group of individuals share a preference order, their aggregate will uphold principles of rationality defined by social choice theorists. These groups will form the collectives, from which we can extract the Pareto optimal solutions. By representing the situation competitively as opposed to forcing cooperation, our approach identifies the situations for which no single consensus model exists, and returns a rational *set* of models and their corresponding solutions. Using an election simulation, we demonstrate that our approach predicts the behavior of a large group of decision-makers more accurately than a single consensus approach.

1 Introduction

We introduce a new approach to aggregating the beliefs and preferences of many individuals to form models for democratic decision-making. Computational models for social decision-making have the potential to change the decision-making paradigm in a society. For example, we see promise in leveraging the viral nature of popular social-networking tools to build computational models for collective decision-making that would enable individuals to have a direct influence in the social, economic, and political decisions that affect them. These collective intelligence models can communicate the diverse ideas, beliefs and preferences of individual stakeholders to decision-makers so they can be directly incorporated into policy. The potential of collective decision-making models to enable informed and thoughtful decision-making at societal scales motivates this work.

We combine the probabilistic beliefs of many individuals using Bayesian decision networks. As a decision-making tool, Bayesian networks provide a rich structure to represent the uncertainties, goals and multi-factor nature of real-life decision-making. However, social choice and Bayesian theorists state that it is not possible to combine, or *aggregate* arbitrary beliefs or preferences to form a consensus model that behaves rationally. For example, the Nobel-laureate economist Kenneth Arrow showed that attempting to find a "social preference" order given a ranked order of options from multiple individuals can result in a situation that does not conform to mathematical principles of rationality (**?**; **?**; **?**). Economists Hylland and Zeckhauser extended these findings for Bayesian decision-making by aggregating probabilistic beliefs and utilities.

Our approach discovers the set of ***collective belief models*** within a population of interest. In (**?**) and (**?**) we introduced our approach to separate individuals into similar groups or clusters based on their beliefs and preferences. In this paper, we define the "extent" to which individuals within a collective should agree. In particular, we show that if a group of individuals share a common preference order over the decision alternatives, that their aggregate will uphold the rationality properties defined by social choice theorists. In this paper, we describe and verify our process for enabling rational democratic decision-making, summarized below:

i Elicit beliefs and preferences from individuals

ii Compute the *expected* utility of a set of decision options using the Bayesian decision networks

iii Extract the symbolic preference order used in social choice theory from each individual's expected utility.

iv Form *collectives* from the groups who share a preference order.

v From these collectives we find their *collective belief*, which is the aggregate of the beliefs of the individuals in the collectives. We show that if a group of individuals share a preference order, their aggregate will uphold principles of rationality defined by social choice theorists.

vi Finally, we form a decision-making *game* in which each of the "players" is derived from one of the collective belief models. We then extract the Pareto optimal solution from the preference orders of the collectives.

Using a strategic election simulation, we show that our approach more accurately predicts the behavior of a large population of decision-makers than the single consensus approach. In fact, we show a situation in which the single consensus approach incorrectly predicts the outcome of a simulated election, while our approach accurately predicts the outcome and provides additional information about the behavior of voters that the single consensus approach does not. For instance, in a strategic situation we demonstrate how we can discover the Nash equilibrium solution.

The paper is organized as follows. Section **??** summarizes the rationality principles for aggregating beliefs and preferences defined by social choice and Bayesian theorists and corresponding *impossibility* results. Section **??** presents our belief aggregation approach for democratic decision making and verifies that our collective belief models produce rational results. Section **??** describes an election simulation that compares our collective choice function to a traditional social choice function. Finally we conclude in Section **??** with a brief discussion of our continued and future research.

2 Background

Our approach connects two significant fields of study; Bayesian decision-making and social choice theory. In this section, we give a very brief description of Bayesian networks, and then summarize several decades of research that investigated the "irrationality" of combining, or aggregating preferences and beliefs to form one consensus model. We begin with a summary of Arrow's "impossibility" theorem (**?**; **?**; **?**). We then update similar findings by early Bayesian theorists (**?**) using a Bayesian network.

2.1 Bayesian Networks

We base our research on a framework that is well-studied in Artificial Intelligence. *Bayesian networks*, also known as *belief networks*, are a form of graphical model that integrate the concepts of graph theory and probabilistic reasoning (**?**). These networks define dependencies between variables that can represent causality, implication or correlation. In a typical Bayesian network, random variables are represented by nodes and conditional relationships by directed edges. A variable X is *conditioned on* all of its parents, described by the expression $P(X|Pa_x)$, where Pa_x is the set of parents of X. Inference is the recursive process of computing the posterior probability of a variable from probability distributions passed from a variable's parents and children (**?**).

Bayesian networks can be extended to address decision problems using *influence diagrams* (**?**), also known as Bayesian decision networks. In addition to nodes representing random variables (or *chance* nodes), influence diagrams contain *decision* nodes, representing a set of decision alternatives; and *utility* nodes, representing the value associated with an outcome.

Bayesian belief aggregation is the process of combining probability estimates from multiple human or software agents to form a consensus network (**?**; **?**). Belief aggregation typically uses an *opinion pool* function to form an aggregate distribution from multiple beliefs. An opinion pool function is an example of a *social choice* function in social choice theory.

2.2 Rationality of Aggregation

Our approach splits a population of individuals with divergent beliefs into groups of similar beliefs. Our objective is to define a splitting function such that the aggregate of beliefs within each group will have rational behavior. In this section we summarize the mathematical principles of rationality with respect to combining the preferences, beliefs and utilities of multiple individuals, as defined by theorists in the fields of voting, decision and probability theories. We then summarize well-known counter examples that demonstrate the failure of aggregation methods. In Section **??** we will present our belief aggregation approach and demonstrate that it produces groups of individuals whose aggregates upholds the rationality principles in the same situations.

Arrow's Axioms We first discuss rationality when individuals and groups consider their preferences deterministically and symbolically. In other words they supply a precise preference order over a set of options that is not based on the uncertainty of outcomes or their perceived risk. Voting theory, the study of combining votes or preferences to select an outcome, has found that combining preferences using a ranked order over the options results in irrational behavior. In particular, the economist Kenneth Arrow developed a theorem that states that there is no rational way to aggregate a number of arbitrary votes when there are three or more options to choose from (**?**). In summary, the social choice function breaks a transitive assumption that states if A is preferred to B by a majority of voters, and B is preferred to C by a majority, then A should be preferred to C. A simple example shows that transitivity cannot hold unless there is a dictatorship (one person's vote is the rule).

We now summarize the properties of rationality defined by Arrow and other researchers with respect to aggregating preferences. First, we provide the notation used to describe the preference relationships between two alternative options x and y. A combination of these relations forms a preference order over a set of three or more options.

Notation The following notation represents the possible preference order relationships between two options x and y (**?**).

x, y: alternative options; i, j: individual people

i. xP_iy: Person i strictly prefers x to y

ii. xI_iy: Person i is indifferent to x and y (doesn't prefer either)

iii. xPy: The society S prefers x to y

iv. xIy: S is indifferent to x and y

v. xRy: S prefers x to y or is indifferent to them

We illustrate these preference order relations with a simple example. Suppose person A prefers vanilla ice cream to chocolate, and chocolate to strawberry. Their pairwise preference orders would be vP_ac and cP_as. Putting them together for all flavors results in the preference order vP_acP_as. Suppose another individual B prefers vanilla to chocolate, but is indifferent between chocolate and strawberry. B's preference order could be vP_bcI_bs or vP_bsI_bc. Since both

Given a society of interest, S and a social choice function:

1. *Completeness (CP)*: Social choice function returns an order that includes all relevant alternatives
2. *Transitivity (TP)*: if S prefers A to B and B to C then S prefers A to C (also replace “prefers” to “is indifferent to”)
3. *Pareto optimality (should be at least weekly Pareto optimal)*:
   (a) Weak Pareto principle (WP). For all x and y, if xP_iy for all i, then xPy:
   (b) Strong Pareto principle (SP). For all x and y, if xR_iy for all i, and xP_iy for some i, then xPy:
4. *Non-dictatorship (NDP)*. There is no dictator. Individual i is a dictator if, $\forall x$ and $\forall y$, $xP_iy \rightarrow xPy$.

Figure 1: Rationality properties defined by Arrow and other social choice theorists.

Group	Vanilla	Chocolate	Strawberry
X	1	2	3
Y	2	3	1
Z	3	1	2

Table 1: Table showing the preference orders for ice cream for three groups (1 is the most preferred flavor).

individuals prefer vanilla to chocolate, we could make a generalization that states $vPcRs$. Note that we could not say $vPsRc$ because this would contradict A’s preference.

There are several properties for rationality discussed by social choice theorists (?; ?; ?). The properties shown in Figure **??** summarize those addressed in this paper. A simple example (?) demonstrates a situation in which a social choice function fails to conform to these properties without relaxing another property (such as non-dictatorship). Consider Table **??** that shows three separate group’s preferences for ice cream, where a rank of 1 indicates that the flavor is the top choice for the individual. When attempting to find a consensus preference order that combines all the groups given their ice cream preferences, one compares each pair of flavors in the following manner using a majority vote:

- Two out of the three individuals prefer vanilla (v) to chocolate (c), so vPc
- Two out of the three individuals prefer chocolate to strawberry (s), so cPs
- Given the above preferences, if vPc and cPs, then vPs should hold, however, we see that the majority actually prefers strawberry to vanilla (sPv)

Thus, the transitivity principle cannot hold without relaxing another principle.

Bayesian Rationality We now discuss rationality in the presence of uncertainty. In this case, probabilities are used to represent an individual’s (or group’s) belief in the likelihood of an event given its causal factors, while utilities indicate the value of an outcome given the inherent uncertainty. It is natural that the properties defined above for rationality in

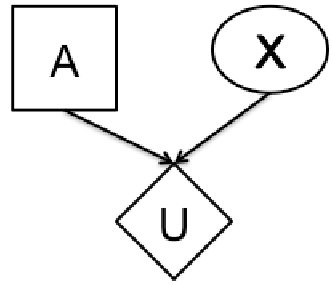

Figure 2: Bayesian decision network for a simple decision involving two possible actions in A, a binary variable X and utilities U.

the non-deterministic case also hold in probabilistic belief aggregation.

Hylland and Zeckhauser (?) show that aggregating probabilistic beliefs and utilities can result in an aggregate that breaks the Pareto optimality rule. Their results apply to the Bayesian decision network in Figure **??**. The authors find the aggregate of the expected utility $U(A)$ by first finding the aggregates of the agents’ beliefs $P(X)$ and conditional utilities $U(A|X)$ and then finding $U_0(A) = U(A|X)P(X)$.

In the situation the authors describe, it appears that the aggregate prefers a_2 over a_1. However, each individual actually prefers a_1 to a_2:

First individual:
$U_1(A = a_1) = 1.0 * 0.75 - 1.2 * 0.25 = 0.45$
$U_1(A = a_2) = 0 * 0.75 + 0 * 0.25 = 0.0$

Second individual:
$U_2(A = a_1) = -1.2 * 0.25 + 1.0 * 0.75 = 0.45$
$U_2(A = a_2) = 0 * 0.25 + 0 * 0.75 = 0.0$

Aggregate:
$U_0(A = a_1) = -.1 * 0.5 - .1 * 0.5 = -.1$
$U_0(A = a_2) = 0$

Thus, the aggregation result breaks Pareto optimality. (?) extend these findings and define a preference ordering among expected utilities, such that option o_i is strictly preferred to o_2 if and only if the expected utility of o_1 is greater than that of o_2. We extend this preference ordering in Definition **??**.

3 Approach

Traditional *social choice functions* used to aggregate beliefs and preferences attempt to find a consensus for a whole population, but fail to produce rational results when a stalemate between opposing opinions occurs. By representing the situation competitively as opposed to forcing cooperation, our approach identifies the situations for which no single consensus exists, and returns a *set* of solutions from which a game-theoretic solution can be found. This section describes our approach to enable *rational* democratic decision making and verify that our *collective choice function* upholds the principles for rationality defined by social choice theorists.

3.1 Forming Collectives

We now discuss how we overcome the irrationality results described in Section **??**. In this section we define the properties of our collective belief models and show that they maintain rational behavior in the previous examples that break rationality principles. Our initial approach (**?**) was to cluster a population into groups depending on their "similarity." In this paper we present a more rigorous definition for the properties of clusters in order to overcome the theoretical limitations.

Definitions The following definitions will be used to describe and defend our approach.

Definition 1. *Rank Order: A rank order is a partially ordered set such that, given a set of options* O *containing* n *options* $\mathrm{o_1..o_n}$*, a ranked order* R *over* O *is an order such an item* $\mathrm{r_i}$ *is preferred to (or indifferent to)* $\mathrm{r_j}$ *if and only if* $\mathrm{r_i}$ *is before* $\mathrm{r_j}$ *in the order.*

We now provide a mathematical definition of the term "collective." We first consider the term's definitions in the English language. The Merriam Webster definition states that a collective is "marked by similarity among or with the members of a group" and "involving all members of a group as distinct from its individuals." (m-w.com) The second definition implies that there is some generalization about the group that can be stated without referring to the individual group members.

We are interested in a more rigorous definition of a collective that captures the implications of these English language definitions in a mathematical or set theoretic form. Since we have not found one in the literature, we will form our own definition of a collective as follows. A collective is a group such that a specific generalization of the group holds for all members of the group. In set theory, this is simply a subset, with the generalization being the property that defines the subset.

Definition 2. *Collective: A collective* C *w.r.t. a property* A *is a subset of a population* P *(*$\mathrm{C} \subseteq \mathrm{P}$*) s.t.* A *holds for all members of* C*. If* A *holds for an individual* $\mathrm{p} \subseteq \mathrm{P}$ *then* p *is a member of* C*.*

Arrow's theorem shows that there is no aggregate in general that represents the group based on the equal participation of the group (**?**; **?**; **?**). However, if a group of individuals has an identical rank order of their preferences, then we can make a generalization about that group based on this rank order. We illustrate with an example: if $O = \{A, B, C\}$ and $R_i = BCA$ and $R_j = BCA$, then $R = BCA$. Thus, if a group of individuals G shares a rank order R over a set of discrete valued options O then R can define a collective C. We now we state the following definition for a rank order shared by a group of individuals.

Definition 3. *Rank Order Collective: If* R *is a rank order over a set of options* O*, and* C *is a collective s.t.* $\forall \mathrm{c_i} \in \mathrm{C}$*, the rank order of* $\mathrm{c_i}$ *is* R*, then* C *is a* rank order collective*. In this case,* A *as defined in Definition* **??** *is specified to be* R.

We will now map Bayesian outcomes to the rank order concept. Each variable X_i in a Bayesian network has a number of possible values ($\{x_{i1}, x_{i2}, ..., x_{in}\}$), where n is the arity of X_i. The posterior probability of a variable in a Bayesian network being a specific value (x_{ij}) is derived through inference.

Given the posterior probabilities of the values of a variable, there will be an order from *most likely* to *least likely* for these possible values. For example, given a binary variable X, with a probability distribution $P(X = T) = 0.25, P(X = F) = 0.75$, the order of values is FT. This order is analagous to the rank order in Definition **??**. In the case of a Bayesian decision network, the result of inference will be is a set of *expected* utilities for the possible decision options. The decision options can be ranked by order of *highest expected utility* to *lowest expected utility*, or *best* to *worst* option. Given a Bayesian network, the rank order can be determined for an arbitrary variable or decision.

Definition 4. *Bayesian Rank Order: A Bayesian rank order* B *with respect to a variable* X *in a Bayesian network is the rank order of the posterior probabilities* $(\mathrm{P(X = x_1), P(X = x_2), ..., P(X = x_n)})$ *of the values of* X*. In a Bayesian decision network,* B *is the rank order of the expected utilities* $(\mathrm{U(o_1), U(o_2), ..., U(o_n)})$ *of the options of* D*, where* $\mathrm{U(o_i)}$ *is the expected utility of decision option* $\mathrm{o_i}$.

Before we define a collective in terms of the Bayesian rank order, we show that an aggregate of $m > 1$ equivalent Bayesian rank orders R always results in R. We illustrate this by showing that the mean of $m > 1$ arbitrary sets of n ordered values in $\mathbb{R}$ results in an ordered set of values.

Theorem 1. *The sum of* $\mathrm{m} > 1$ *sets of* n *real valued (*$\mathrm{r} \in \mathbb{R}$*), ordered numbers will result in an ordered set.*

Proof 1. *We prove this by contradiction. Imagine that we have two sets of* n *cups (*C_1 *and* C_2*) that are ordered from left to right such that a cup on the left has* $\geq$ *the liquid as the cup to its right. Now consider that one adds the liquid in each of* C_2*'s cups to* C_1*'s cup. So cup* c_{2i} *is added to cup* c_{1i}*. The amount of liquid in the cups in* C_2 *will remain in order. If this were not the case, then cup* c_{2j} *would have had to have had more in it than cup* c_{2i} *(where* i *is farther left than* j*). However, this would mean that* C_2 *was not in order, which would be a contradiction.*

Since the arithmetic mean will be found by normalizing the combined amounts in each cup by the same amount (k), finding the mean will not affect the rank order. The *collective belief* of a collective is then the aggregate of the beliefs that define the *Bayesian rank order* for the collective.

Definition 5. *Collective Belief: If* B *is a Bayesian rank order over the probabilities of a variable* X *(or the utilities of a decision* D*) and* C *is a collective s.t.* $\forall \mathrm{c_i} \in \mathrm{C}$*, the Bayesian rank order of* $\mathrm{c_i}$ *is* B*, then the* collective belief β *of* C *is the aggregate of the* k *probability distributions (or expected utilities) supplied by the members of* C.

Definition 6. *Collective Belief Model: A collective belief model is a Bayesian belief network derived from aggregating the structure and parameters that resulted in the collective belief* β *for a collective* C*.*

A collective belief model can be formed for each collective that emerges from a population by aggregating the prior probability distributions supplied by the individuals in C for all ancestors and descendants of X (or D). The following definitions define the *set* of rank order collectives for a population P.

Definition 7. *Partition: A partition* T *of a population* P *is a set of collectives such that all individuals in* P *are in exactly one collective* $\mathrm{C_j}$*. A partition will be either weak or strong as defined below.*

Definition 8. *Strong Partition: A strong partition* $\mathrm{T_s}$ *of a population* P *is a partition such that each individual can be assigned to one and only one collective* $\mathrm{C_j}$*. In other words, there is a unique partition* $\mathrm{T_s}$ *for the population* P.

Definition 9. *Weak Partition: A weak partition* $\mathrm{T_w}$ *of a population* P *is a partition such that each individual could fit into multiple collectives, but is assigned to only one.*

Finally, our *collective* choice function is defined as follows.

Definition 10. *Collective Choice Function: Given a population* P *and a set of options* O*, over which each individual* i *in* P *provides a rank order* $\mathrm{R_i}$*:*

i Separate a population P *into* m *groups such that each individual in group* $\mathrm{G_j}$ *has a rank order* $\mathrm{R_j}$ *over the options O.* ($\mathrm{R_i} = \mathrm{R_j}$ *iff* $\mathrm{i} \in \mathrm{G_j}$).

ii Each group $\mathrm{G_j}$ *becomes a collective* $\mathrm{C_j}$ *defined by its collective rank order* $\mathrm{R_j}$*. The union of each* C_j *is a partition.*

iii In a Bayesian environment, use an opinion pool function (e.g. arithmetic or geometric mean) to compute the aggregate of C_j*'s probability distributions or utilities.*

We will form our rank order collectives based on the pairwise preference order relations defined in the Section **??**. In particular we are interested in the following *society* relations, which are generalizations about a society, population, or in our case, a collective (C): xPy, xIy and xRy. A preference order among more than two options can be made up of any combination of the R, P and I relations. The *rank order* for each rank order collective is a preference order formed in this manner.

Rationality of Collectives We now present a sketch of how our *collective* choice function upholds the properties defined in Figure **??** by creating a solution for each collective, such that the properties hold within each collective. The social choice functions described in Section **??** attempt to find a *single* consensus, but fail to produce rational results when a stalemate occurs (Arrow's findings) and when the beliefs are in opposition (Hylland and Zeckhauser's findings). Instead of failing, our approach identifies the situations for which no single consensus exists, returning a *set* of solutions composed of a consensus for each collective. We briefly discuss each rationality principle defined in Figure **??** in the context of our collective choice function.

1. *Completeness (CP)*: Our collective choice function will return a *set* of solutions $S = \cup_{j=1}^{m} R_j$ where R_j is the rank order of the rank order collective C_j that emerges from a population P.
2. *Transitivity (TP)*: Since there is only one rank order per collective over a set of options O, and a rank order upholds the properties of a partially ordered set, then transitivity holds by the properties of a partially ordered set.
3. *Weak Pareto principle (WP)*: The weak Pareto principle is upheld implicitly by our definition of *rank order collective*. For all individuals $i \in C$, if the rank order for $i = R$ then the rank order for $C = R$. Thus, for all options x and y, if xP_iy then xPy for C. (**?**) also state this finding when there is consensus among the preference order of expected utilities.
4. *Non-dictatorship (NDIP)*: The rank order R of a collective happens to be the same as an arbitrary individual i in the collective C. However, this is not a dictatorship since the order of the remainder of the collective ($C' = C - i$) would persist even if individual i changed his belief.

3.2 Applying Game Theoretics

The next step in our social decision-making process is to form a game to enable us to find the Pareto optimal solutions for a partition derived from an arbitrary population. In some cases we can form a strategic game that may enable us to find Nash equilibrium, minimax and other game-theoretic solutions. In our games, each "player" is defined by the collective rank order or collective belief discovered by our collective choice function. Essentially, each player represents the shared beliefs and preferences of its collective.

Extracting the Pareto optimal solution A Pareto optimal solution is one in which no players can do better (have a higher utility) without another player doing worse (**?**). In our case, this means that for each option $o_i \in O$, o_i is a Pareto optimal solution if there are no other options o_j, for which all collectives provide a higher ranking than o_i. Given the rank orderings from our collectives we would like to extract the Pareto optimal solution from a strong partition T_s derived from an arbitrary population P. We first consider the situation in which all relations in a rank order are the strict preference (P), in other words the solution will be *weakly* Pareto optimal. We can show that in this case, the weak Pareto optimal solution S_{wp} is the union of each of the collective's most preferred option, corresponding to the first item r_{j1} in R_j, where R_j is the rank order for the collective C_j.

$$S_{wp} = \cup_{j=1}^{m} r_{j1} \tag{1}$$

Theorem 2. *The weak Pareto optimal solution* $\mathrm{S_{wp}}$ *for a partition* $\mathrm{T_s}$ *that uses only the P (strict preference) relation is composed of the union of the highest ranked option (*$\mathrm{r_{j1}}$*) from each collective* C_j *in* T_s*:*

Proof 2. *We prove this by contradiction. If* S_{wp}*, derived as in eq.* **??***, were not the weak Pareto solution, then there would be some option* $\mathrm{o_i} \notin \mathrm{S_{wp}}$ *that was preferred by someone to all other options* $\mathrm{o_j}$*. However, this cannot be the case since all options in* $\mathrm{S_{wp}}$ *are the most preferred options. Thus* $\mathrm{S_{wp}}$ *must be the weak Pareto solution.*

We discuss the algorithm to find the strong Pareto optmial solution in (**?**). We cannot extract a Pareto optimal solution from a weak partition T_w using this approach, because one cannot distinguish preference and indifference.

Applying the collective choice function We now revisit the examples in section **??** using our collective choice function. In the ice cream example illustrated by Table **??**, our collective choice function will result in three collectives X', Y' and Z', equivalent to the groups X, Y and Z. Assuming a strict preference ordering, we see that each of the ice cream choices is a Pareto optimal solution, since the union of each collective's preferred option is the entire set of options. In this stalemate situation our approach discovers a rational solution to this problem while the majority vote social choice function fails.

We next revisit the Bayesian example derived from Hylland and Zeckhauser's (**?**) findings discussed in Section **??**. The example showed that this social choice function resulted in a solution that was not Pareto optimal. The following table shows the Bayesian ranked orders for the individuals i_1 and i_2 based on their beliefs about $U(A|X)$ and $P(X)$. The last row in the table shows the ranked order results from the social choice function computed in Section **??**, represented by i_0.

Individual	$\mathbf{U(A\|X)}$	$\mathbf{P(X)}$	$\mathbf{U(A)}$
i_1	$a_1a_2a_2a_1$	FT	a_1a_2
i_2	$a_2a_1a_1a_2$	TF	a_1a_2
i_0	$a_2a_1a_2a_1$	$TForFT$	a_2a1

For the conditional utility $U(A|X)$, the rank order is composed of the rank order for A given that X is false, followed by the rank order for A given that X is true. For $P(X)$, a rank order of FT indicates that the individual believes that X is more likely to be false than true, while TF indicates the opposite. Our collective choice function will not aggregate different rank orders, therefore i_1 and i_2 would not be combined to find i_0. Instead, each individual would form its own collective, resulting in the collective rank order a_1a_2. Clearly, in this case the only Pareto optimal solution is a_1.

4 Experiments

Our vision is to utilize direct input from human contributors to form our collective belief models. In (**?**; **?**) we elicited opinions from real people using Amazon's Mechanical Turk (mturk.com) and surveymonkey.com. However, for the following experiment, we simulate human opinions in order to analyze the behavior of changes in parameters.

4.1 Election Simulation

We present the results of an election simulation in which there are three candidates, two sharing a majority of the vote within a few points of each other, while a third has a small minority. We will show that our collective choice function is a better predictor of the voting behavior of a population than a social choice function that finds a single consensus. In fact, the social choice function *incorrectly* predicts the outcome of the election in a contrived (but believable) situation. The simulation is strategic in that each individual may change her vote based on how she believes the remainder of the population will vote.

Given a simulated population, our simulation computes each individual's expected utility for each candidate using a simple decision network, shown in Figure **??**. Each individual will provide a *conditional utility* for each candidate that represents the utility of a candidate *given* the candidate wins. In other words, the projected utility of each candidate *if* that candidate won. In this simulation the social choice functions utilize the *expected* utility of a candidate, which depends on the individual's belief in that candidate winning (oval labeled *Win*). For instance, an individual may prefer one candidate, but believe that the likelihood of him winning is low, therefore the expected utility of the individual's second favorite candidate may actually be higher.

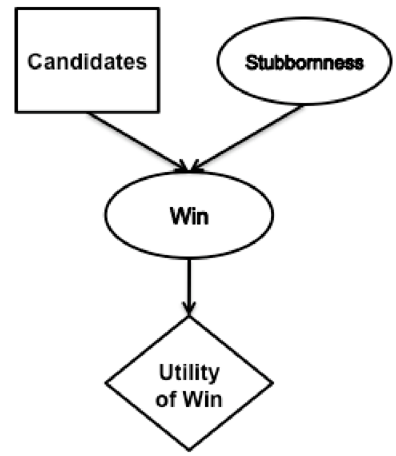

Figure 3: A Bayesian decision network for an election simulation in which the rectangle represents the decision options (candidates), the oval labeled *Win* represents the belief that each of these candidates will win and the diamond represents the conditional utility of the winner. *Win* is also dependent on *Stubbornness*, which is the probability that an individual will be resistant to changing beliefs.

Initialization: The simulation is initialized with a population P of individuals and a set of candidates O. Each simulated individual will provide: $U(O|W)$, a conditional utility for each candidate given the candidate wins; $P(W = T|O)$, an initial belief in each candidate winning; and $P(S)$, a probability of being stubborn, where stubbornness = 1.0 means that the individual is completely resistant to changing her belief. The simulation then repeats a "polling" process until *convergence* occurs. We define convergence as a situation in which no individual has switched votes for a specified number of repetitions.
Simulation Parameters:

- *weights*: contains a weight for each candidate to indicate the proportion of the population that initially prefers each candidate
- N, the population size
- C, the number of repetitions with no switches that is required for convergence

Each of the N individuals is randomly assigned a preferred candidate based on the weights parameter. Each individual will set the conditional utilities for each candidate based on her preference. The candidates are on a range from "left" to "right," such that an individual that prefers the leftmost or rightmost candidate will give second preference to the middle candidate (not the opposing candidate). An individual that prefers the middle candidate has an equal probability of leaning to the left or right for their second most preferred candidate. The stubbornness value for each individual is assigned randomly between 0.5 and 1.0. Each individual initializes her *Win* beliefs based on the candidate's utility and her stubbornness. Every individual will consider her preferred candidate more likely, but a more stubborn individual will overwhelmingly favor her preferred candidate.

Each individual will then compute her expected utility for each candidate $o \in O$ using the following formula:

$$U(o) = U(o|w) * P(w|o) \tag{2}$$

The individual will then set her initial "vote" to the candidate with the highest expected utility. The initial *vote count* is the count of votes to each candidate for the whole population.

Social choice functions We are comparing a *single consensus* social choice function and our *collective* choice function. The single choice function is an average of all individuals' expected utility for each candidate:

$$\forall o \in O, U_0(o) = \frac{1}{n} \sum_{i=1}^{n} U_i(o) \tag{3}$$

Where O is the set of candidates and n is the population size. Our collective choice function will find the expected utility for each candidate, *for each collective C_j*:

$$\forall C_j, \forall c \in C, U_j(o) = \frac{1}{k_j} \sum_{i=1}^{k_j} U_i(o) \tag{4}$$

Where k_j is the number of individuals in collective C_j.

Simulation run: At each time-step of the simulation, each individual will update her beliefs based on the *vote count*, which is provided to all individuals, and her stubbornness. She will then update her expected utility and select the candidate with the highest expected utility as her "vote". After all individuals have updated their belief, a new *vote count* is determined.

We ran the simulation several times, changing the initialization parameters to observe the final results after convergence. First, we were interested in determining if the single consensus social choice function accurately predicts the final vote count. Second, we were interested in the average utility of each collective at convergence. Finally, we were interested in the characteristics of individuals who switched votes during the simulation. In particular, we were interested in which collectives they moved from and to.

Results: Simulation parameters:

- *weights*: $\{0.05, 0.46, 0.49\}$
- N: 10,000
- C: 100

In this run, the leftmost candidate begins with a small portion of the votes, while the rightmost candidate has

RO	Size	Left	Center	Right
CLR	3601	-0.12 (0.030)	4.13 (0.36)	-2.86 (1.94)
RLC	2901	-0.23 (0.006)	-2.30 (1.40)	4.51 (0.12)
RCL	2093	-0.13 (0.015)	1.50 (0.94)	4.51 (0.09)
CRL	1091	-0.08 (0.030)	4.17 (0.24)	1.46 (1.02)
LCR	313	0.32 (0.120)	-2.22 (1.70)	-3.54 (1.35)
LRC	1	5.38 (0.000)	-1.11 (0.00)	-0.91 (0.00)

Table 2: Average expected utility for each candidate, for each rank order (RO).

the highest proportion of the votes by a margin that is smaller than the leftmost candidate's votes. We repeated the simulation with these parameters several times with results similar to the following. The following table shows the final vote count and average expected utility (from eq. **??**) of each candidate at convergence.

	Left	Center	Right
Final vote count	314	4692	**4994**
Expected utility	-0.13	**1.52**	1.27

The most significant discovery in this simulation is that the candidate with the highest average expected utility is *not* the candidate with the highest vote count, even though each individual is voting according to her highest expected utility. (**?**) discusses the range of values that result in inconsistent behavior. Interestingly, the most inconsistent results occur when the *weights* ratio is Right=Left+Center.

We now show the results of our collective choice function. Each collective is defined by their rank order, shown in the left column in Table **??**. The second column shows the collective's size and the remaining columns show the collective belief, computed from eq. **??**, followed by the variance in parentheses.

We note that the average expected utilities for each candidate over the whole population P can be derived from these results using the formula:

$$U(o) = \sum_{j=1}^{m} U_j(o) * \frac{|C_j|}{|P|} \tag{5}$$

Where m is the number of collectives. While the average expected utilities over P incorrectly predicts the outcome, we observe that the average expected utility of each collective's preferred candidate results in collective beliefs that accurately reflect the outcome of the election (*Right* is highest *and Right* wins), shown in the following table:

Left	.0313*.32+.0001*5.38 = **0.011**
Center	.3601*4.13+.1091*4.17 = **1.94**
Right	.2901*4.51+.2093*4.51 = **2.25**

We now demonstrate that the rank order of the individuals that switch votes is a useful predictor of voter behavior. In the run described, 180 individuals switched votes during the simulation because the candidate with the highest expected utility changed. We observe that all of the individuals that switched moved from the collective represented by the rank order LCR, to the collective represented by the rank order CLR. The fact that all individuals who switched were from the same collective supports our theory that our collective choice function is a better predictor of voting behavior than the single consensus social choice function. This social choice function does not distinguish these potential "vote switchers" from the rest of the population. As we might expect, the individuals that switched collectives and votes had expected utilities that indicated that they were nearly indifferent to their first and second preferred options. As in a real election, "swing votes" prove to be particularly important.

Finally, our results imply the Nash equilibrium solution for this election simulation. This game theoretic solution occurs when all players are aware of the strategy of their opponents, and are maintaining a strategy that maximizes their utility given their opponents' strategies (**?**). In our simulation, we consider the strategy of a collective (or individual) to be the candidate that the collective has selected. In the simulation, each individual makes a selection with the awareness of the *vote count*, which represents the strategy of the rest of the population. During the simulation, individuals may change their strategy to maximize their expected utility given the vote count. However, once individuals have stopped switching votes, they have settled on a strategy that maximizes their utility given the strategies of the rest of the population. Each collective in Table **??** reflects the strategy of its members through its expected utility.

5 Future Work and Conclusions

Our results indicate that interesting behavior occurs at the borders of collectives, in other words in situations in which individuals are nearly indifferent to multiple options. In our continued research, we will investigate these situations more closely as they may reveal additional constraints on collectives worth considering. A natural next step from extracting symbolic preference orders from quantitative results would be to represent the preference orders quantitatively. For instance, perhaps each individual or collective could have a distribution over the possible preference orders. Complementing the research described in this paper, we have developed an algorithm to partition a population into collectives and form collective belief models.

By couching the social choice function in a competitive environment, we are able to overcome many of the theoretical limitations of the single consensus model. Our collective choice function discovers a set of Pareto optimal solutions in situations for which there is no single solution, and in some strategic situations it is able to infer the Nash equilibrium solution. Our election simulation extended irrationality results by identifying a situation in which the traditional social choice function fails to predict the outcome of an election. The simulations also imply that our collective choice function produces solutions that more accurately predict the behavior of large groups of decision-makers.

References

Arrow, K. J. 1950. A difficulty in the concept of social welfare. *Journal of Political Economy* 58.

Chinchuluun, A.; Pardalos, P.; Migdalas, A.; and Pitsoulis, L., eds. 2008. *Pareto Optimality, Game Theory and Equilibria*. Springer.

Feldman, A. M., and Serrano, R. 2007. Arrows̋ impossibility theorem: Two simple single-profile versions. Working Papers 2007-07, Instituto Madrileo de Estudios Avanzados (IMDEA) Ciencias Sociales.

Greene, K., and Luger, G. 2009. Aggreeing to disagree: Leveraging consensus and divergence in bayesian belief aggregation. In *Papers from the AAAI Spring Symposium, Technosocial Predictive Analytics Workshop*, 49–54.

Greene, K.; Kniss, J.; Luger, G.; and Stern, C. 2009. Satisficing the masses: Applying game theory to large-scale, democratic decision problems. In *Proceedings of the IEEE Conference on Social Computing (SocialCom-2009)*.

Greene, K. A.; Kniss, J. M.; and Luger, G. F. 2010. Enabling rational democratic decision-making with collective belief models and game theoretic analysis. Technical report, UNM Department of Computer Science.

Hylland, A., and Zeckhauser, R. J. 1979. The impossibility of bayesian group decision making with separate aggregation of beliefs and values. *Econometrica* 47(6):1321–36.

Murphy, K. 1998. A brief introduction to graphical models and bayesian networks.

Pennock, D. M., and Wellman, M. P. 1999. Graphical representations of consensus belief. In *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence*.

Seidenfeld, T.; Kadane, J.; and Schervish, M. 1989. On the shared preferences of two bayesian decision makers. *The Journal of Philosophy* 86(5):225–244.

Shachter, R. D. 1986. Evaluating influence diagrams. *Operations Research* 34(6):871–882.

Watkins, T. 2009. Arrow's impossibility theorem for aggregating individual preferences into social preferences.