

Empathetic Social Choice on Social Networks

Amirali Salehi-Abari and Craig Boutilier

Abstract

Social and economic networks play a fundamental role in facilitating interactions and behaviors between individuals, businesses, and organizations. It is widely recognized that such networks can correlate behaviors (and arguably preferences) among connected agents. We introduce a model for social choice—specifically, consensus decision making—on such networks that reflects certain interdependencies among agent utilities. Specifically, we define an *empathetic social choice framework* in which agents derive utility based on both their own intrinsic preferences and the satisfaction of their neighbors. We show how this problem translates into a weighted form of classical preference aggregation (e.g., social welfare maximization or certain forms of voting), and develop effective algorithms for consensus decision making that we believe should scale to large-scale (online) social or economic networks. Preliminary experiments validate the effectiveness of our proposed algorithms.

1 Introduction

Social networks play a central role in individual interactions and decision making. Indeed, it is widely acknowledged that the behaviors [7], and to a lesser extent the preferences, of individuals connected in a social network are correlated in ways that can be explained, in part, by network structure [10, 13]. Because of this, and the increasing availability of data that allows one to infer such relationships, the study of social choice problems on social networks is one of tremendous practical import. In fact, arguably most group decision problems, whether social, corporate, or policy-oriented, involve people at least some of whom are linked via myriad social ties. However, social choice in the context of social networks is something that has received, until recently, relatively little attention. Recent work has examined, for example, the formation of (hedonic) coalitions on social networks [6, 5], and stable matching on social networks [3, 16], in which the network captures one’s affinity for potential partners. The influence of social networks on voting behavior has received considerable attention in the social sciences (e.g., [1, 14, 15]), and the emergence of online social network has even spawned computational research on the mechanisms to support delegation of votes in an online network [4].

In this paper, we consider the problem of *consensus decision making on social networks*, for example, in the form of voting over some option space. Specifically, we consider the problem of selecting a single option from a set of alternatives, for some group connected by a social network—e.g., a local constituency electing a political representative, or colleagues selecting a venue for a corporate retreat. While individuals have, as usual, personal *intrinsic utility* over the option space, we also incorporate a novel form of *empathetic utility* on social networks: in our model, the utility (or satisfaction) of an individual with a winning alternative a is a function of both her intrinsic utility for a and her *empathetic utility* for the “happiness” of her neighbors in the network. This use of empathetic utility can be seen as reflecting recent findings that suggest a person’s happiness is influenced by the happiness of others with whom they are connected [11].

We consider two varieties of empathetic preference. In the first, the *local empathetic model*, the utility of individual i for alternative a combines her intrinsic preference for a with the *intrinsic* preference of i ’s neighbors for a , where the weight given to the preference of any neighbor j depends on the strength of the relationship between i and j . For instance,

in selecting a restaurant, i may be willing to sacrifice some of her own intrinsic preference for the chosen restaurant if her colleagues are happier with the cuisine, and she defers more strongly to her closest friends. In the second, *global empathetic model*, i 's utility for a depends on her intrinsic preference and the *total utility* that her neighbors have for a (not just their intrinsic preference). In other words, she doesn't just want her neighbors to be satisfied with a , she wants them to have high utility, which depends on the utility of *their* neighbors, and so on. For example, in voting for a political candidate, i may have a mild preference for a over b , but if b is strongly preferred by not only her closest neighbors, but also by their neighbors and many others in the community, she might prefer to see b elected so she won't have to interact with grumpy neighbors for the next five years.

Our main contributions in this paper are to develop a model for preference aggregation (e.g., certain forms of voting) that select consensus alternatives in a way that is sensitive to both intrinsic and empathetic preferences. Of course, we don't expect voters to actually compute such combined preferences; indeed, they may not have direct knowledge of the preferences of their neighbors. Instead voters specify their preferences for options and for the satisfaction of their neighbors (the latter could be inferred or estimated directly from the social network in some settings). We then propose methods for computing optimal alternatives under both the local and the global models. The former, unsurprisingly, corresponds to a simple form of weighted preference aggregation or weighted voting in which each voter implicitly "delegates" a portion of her vote to her neighbors. The latter, because individual utilities are co-dependent—indeed, utility spreads throughout the network much like PageRank values—requires the solution of a linear system to determine the optimal (fixed-point) option for the group. We describe (mild) conditions under which a fixed point is guaranteed to exist, and show that it too results in a form of weighted voting, where the weights assigned to each voter's intrinsic preference is readily derived from the solution to this linear system. Experiments explore various properties of our model and algorithms.

2 Social Empathetic Model

We begin by outlining our basic social choice model, motivating two notions of empathetic preference on social networks, and then defining socially optimal outcomes within this model. We also briefly discuss related work.

2.1 The Social Choice Setting

Apart from empathetic preferences on a social network, which we specify below, the choice framework we adopt is standard. We assume a set of alternatives $\mathcal{A} = \{a_1, \dots, a_m\}$ and a set of agents $\mathcal{N} = \{1, \dots, n\}$. Each agent j has *intrinsic preferences* over \mathcal{A} in the form of either a (strict) preference ranking \succ_j^I or a utility function u_j^I . For ease of presentation, we describe preferences in terms of utility functions, but discuss below on how to interpret voting procedures within our model. For example, in our experiments we use simple utility functions based on rankings of alternatives and score-based voting rules (specifically, Borda and plurality) to define "utility" for alternatives.

Our goal is to select a single consensus alternative $a^* \in \mathcal{A}$ that implements some social choice function f relative to the preferences of \mathcal{N} . For example, if agents' utilities were dictated solely by intrinsic preference and f were (utilitarian) social welfare, we would select $a^* = \arg \max \sum_j u_j^I(a)$. If preferences were given by intrinsic preference rankings, f would typically be represented by some voting rule (e.g., plurality or Borda).¹

¹Our model below applies directly to more general social choice problems, such as assignment/segmentation problems with network externalities (where individuals may be assigned different al-

2.2 Empathetic Preference on Social Networks

We depart now from the typical social choice framework by considering *empathetic preferences*, in which the preferences of one agent are dependent on those of others. We consider the specific case in which these influences are induced by connections in a social network (though the notion of empathetic preference need not be confined to networks). We focus on agent utility functions rather than preference rankings, since these allow the straightforward expression of quantitative tradeoffs between intrinsic and empathetic preference.²

Before discussing additional motivation, we introduce our model and notation. We assume a directed weighted graph $G = (\mathcal{N}, E)$ over agents, with an edge jk indicating that j 's utility is dependent (in a way to be specified below) on its neighbor k 's preference, the strength of this dependence given by edge weight w_{jk} . Naturally j 's utility will usually depend on its own intrinsic preferences, so loops jj will usually be present. We assume that $w_{jk} \geq 0$ for any edge jk , and that $\sum_k w_{jk} = 1$ for any j (though allowing variable weightings to reflect, say, weighted voting schemes is also possible). For convenience, we treat missing edges as if they had weight zero (and vice versa). Thus, we represent the graph with a weight matrix $\mathbf{W} = [w_{ij}]$. We generally think of these edges as corresponding directly to some relationship in a social network, or possibly induced from such relationships. See Fig. 1(a) for an illustration.

We take j 's utility for a to be a linear combination of its own intrinsic preference for a and the empathetic preference derived from each of its neighbors—recall that we consider pure consensus/single-winner voting scenarios in which a single option a is selected for all $j \in \mathcal{N}$ —where network weights determine the relative importance of each.³ Letting $e_{jk}(a)$ denote the *empathetic utility derived by j from k* , we define j 's utility $u_j(a)$ to be

$$u_j(a) = w_{jj}u_j^I(a) + \sum_{k \neq j} w_{jk}e_{jk}(a).$$

The ratio of w_{jj} to $\sum_{k \neq j} w_{jk}$ captures the relative importance of intrinsic and empathetic utility to j .

We consider two ways in which to define empathetic preferences e_{jk} . In the *local empathetic model*, we simply define $e_{jk}(a) = u_k^I(a)$; in other words, j 's utility for a is simply a linear combination of intrinsic utilities of j 's neighbor (including its own):

$$u_j(a) = \sum_k w_{jk}u_k^I(a). \quad (1)$$

This model captures the fact that an agent j is concerned about the “direct” preference of a neighbor k for alternative a ; but the fact that k 's utility may depend on k 's *own* neighbors does not impact j . For instance, consider a family or a group of friends deciding on a movie (or restaurant or outing): the preferences of certain family members (e.g., parents) for a specific film may depend on the preferences of others (e.g., children, whom they want to be entertained by the choice of film).

In the *global empathetic model*, we define $e_{jk}(a) = u_k(a)$, so that k 's complete utility for a —which *may depend on k 's own neighbors*—influences j 's utility for a , giving rise to

$$u_j(a) = w_{jj}u_j^I(a) + \sum_{k \neq j} w_{jk}u_k(a). \quad (2)$$

ternatives), matching problems, and so on, without difficulty. Our algorithms are, however, specific to the “single-choice” assumption.

²Suitable *qualitative* expression of such tradeoffs is an important ongoing research direction.

³More general non-linear models are possible as well.

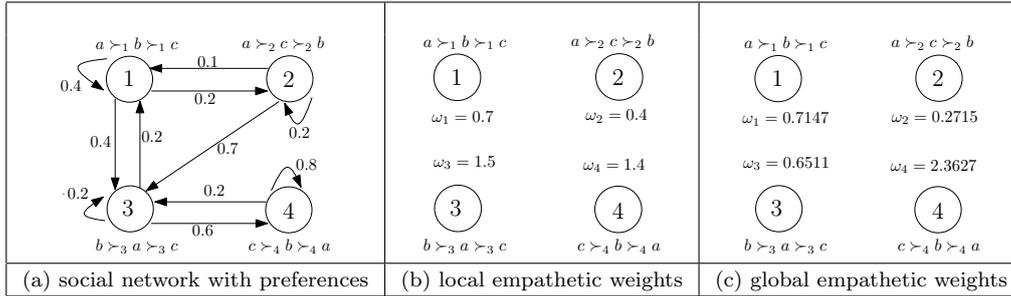


Figure 1: Social network and ranked preferences, with weights under the local and global empathetic model. Using Borda or plurality-based utility, the consensus winner is different in each model: a under intrinsic; b under local empathetic; c under global empathetic.

In this model, j 's utility for a depends on the utility (not just intrinsic preference) of its neighbors for a . For example, a voter may care about the overall level of satisfaction of her neighbors when voting for a political representative, but recognize that there is a larger societal effect at work, where their satisfaction also depends on their neighbors, etc. More concretely, companies linked in complex supply chain may well care about the overall success of their suppliers and customers, and consider adopting industry-specific or economic policies in that light. In the global model, the circular dependence of utilities requires a fixed point solution to the linear system defined by Eq. 2 (see below).

Correlations of behavior and/or preferences among agents connected in social network is widely accepted, and can be explained by a variety of mechanisms [10, 13]. Among these are: *technology/information diffusion*, in which agents become aware of opportunities or innovations from connections to their neighbors; *network externalities*, in which the benefits of adopting some behavior increase with the number of neighbors doing the same; or *homophily*, in which people with similar characteristics (say, preferences) more readily form social ties. Our empathetic model is somewhat different in that a person's intrinsic preferences over options \mathcal{A} are not presumed to be correlated with their neighbors, but their revealed preferences for \mathcal{A} might be: their choices (or stated utilities) will generally reflect some consideration, however estimated, of their neighbors' preferences as well.

2.3 Social Welfare as Weighted Intrinsic Utilities

In realistic social choice situations, agents with empathetic preferences must often perform sophisticated reasoning not only about their intrinsic preferences for alternatives, but also about those of their neighbors. Thus, even in the local empathetic setting, expressing preferences (e.g., voting) is difficult since agents usually have incomplete (and in some cases, no) information about the preferences of their friends, neighbors, or colleagues. The global empathetic setting is even more complex, since an agent is further required to reason about her neighbors' network connections as well as their intrinsic/empathetic tradeoffs.

In our models, preference aggregation and optimization can be performed by simply having agents specify their *intrinsic preferences*, as is standard in social choice, and the *weights* they assign to neighbors in their local network. In social scenarios, this can remove a considerable informational and cognitive burden from agents who might otherwise be required to determine their total utility for alternatives. In other situations, agents might not wish to reveal their preferences to their neighbors, but might still want their neighbors to obtain a favorable result (consider, for example, a collection of companies, voting over some economic policy alternatives, that are linked together in complex supply chain relationships which correlates their stability or profitability). It turns out that, given a known

network G , the problem of consensus decision making with empathetic preferences can be recast as a *weighted preference aggregation problem over intrinsic preferences alone*. Not only does this ease the burden on agents, it also allows one to recast the problem as one of simple weighted voting, or of weighted (utilitarian) social welfare maximization, rendering the decision making process itself fully transparent. Here we focus on social welfare maximization.

For the local model, determining the weights associated with each agents' intrinsic preference is straightforward. Assume network weights \mathbf{W} . Let $\mathbf{u}(a)$ be the n -vector of agent utilities to be computed as a function of the corresponding vector $\mathbf{u}^I(a)$ of intrinsic utilities for some fixed alternative a . By Eq. 1, we have $\mathbf{u}(a) = \mathbf{W}\mathbf{u}^I(a)$. Then letting $\boldsymbol{\omega} = \mathbf{e}^\top \mathbf{W}$ (where \mathbf{e} is a vector of ones), the social welfare of any alternative a under the local empathetic model is given by

$$sw_l(a, \mathbf{u}^I) = \boldsymbol{\omega}^\top \mathbf{u}^I(a). \quad (3)$$

Thus social welfare maximization under local empathetic utility is simply weighted maximization of intrinsic preference, where the weight of j 's intrinsic utility ω_j is simply the sum of the weights of its incoming edges.

Fig. 1(b) illustrates the local model in action. The derived weights for each agent are shown. We assume preference rankings, and suppose utilities are derived from these using either Borda or plurality scores. We see that the decision can be different under the local model than using voting based on intrinsic preferences along (a wins in the intrinsic model, while b wins in the local model). Indeed, using score-based voting rules, we can readily interpret this model as a form of *empathetic voting*, where the weight one assigns to a neighbor can be interpreted as the extent to which one would sacrifice one's own preferences to improve a neighbor's intrinsic satisfaction with the winning alternative.

Things are slightly more subtle in the global empathetic model. Computing the utility vector $\mathbf{u}(a)$ for alternative a requires solving a linear system to compute the fixed point of Eq. 2. Unfortunately, a unique solution is not guaranteed to exist.⁴ However, in addition to our assumptions above of *non-negativity* (i.e., $\mathbf{W} \geq \mathbf{0}$) and *normalization* (i.e., $\sum_k w_{jk} = 1$ for all j), a third mild condition on the social network (weight matrix \mathbf{W}) is sufficient to ensure a unique fixed point solution, namely, *positive self-loop*: $w_{jj} > 0$ for all j . Let \mathbf{D} be the $n \times n$ diagonal matrix with $d_{jj} = w_{jj}$. We can write Eq. 2 as

$$\mathbf{u}(a) = (\mathbf{W} - \mathbf{D})\mathbf{u}(a) + \mathbf{D}\mathbf{u}^I(a). \quad (4)$$

As a consequence,

Theorem 2.1 (Fixed-point Utility) *Assuming nonnegativity, normalization, and positive self-loop, Eq. 4 has a unique fixed-point solution $\mathbf{u}(a) = (\mathbf{I} - \mathbf{W} + \mathbf{D})^{-1}\mathbf{D}\mathbf{u}^I(a)$.*

(Proofs of all results are included in the longer version of this paper.) As in the local model, social welfare maximization in the global model can be interpreted as weighted maximization of intrinsic preference (though with a less straightforward interpretation):

Corollary 2.1 *In the global empathetic model, social welfare of alternative a is given by $sw(a, \mathbf{u}^I) = \boldsymbol{\omega}^\top \mathbf{u}^I$ where $\boldsymbol{\omega}^\top = \mathbf{e}^\top (\mathbf{I} - \mathbf{W} + \mathbf{D})^{-1}\mathbf{D}$.*

Once again, in (score-rule based) voting contexts, one can interpret the global empathetic model as trading off one's own satisfaction with a winning alternative with the "overall" (not just intrinsic) satisfaction of one's neighbors: see Fig. 1(c) for an illustration. We discuss weight computation in Sec. 3.

⁴Consider two individuals j and k , with $w_{jj} = w_{kk} = 0$, $w_{jk} = w_{kj} = 1$, $u_j^I(a) = 0.1$, and $u_k^I(a) = 1$. The induced system does not have a unique fixed-point solution.

2.4 Related Work

We are unaware of other formal models which consider the dependency between agent utilities in a social network using the type of empathetic utility we introduce above. However, empathetic utilities might be viewed as a form of network externality in an agent’s utility function, though unlike typical models of externalities, an agent’s utility depends on the (latent) utility of its neighbors for the chosen alternative rather than the behavior of, or allocation made to (at least directly), her neighbors (or others).

Decision making in the presence of network externalities has recently attracted attention. Bodine-Baron *et al.* [3] study stable matchings (e.g., of students to residences) with *peer effects*: these local network externalities reflect the fact that students prefer to be assigned to the same residence as their friends in a social network. Brânzei and Larson address coalition formation on social networks in two different settings: (a) agent utility for a coalition depends on its affinity weights with others in the coalition [5]; and (b) agent utility depends on her distance to others on the induced social network [6]. The problem of auction design in social networks with positive network externalities is studied in [12].

Boldi *et al.* [4] consider voting on social networks, describing a form of *delegative democracy* in which an individual can either express her preferences directly, or to delegate her vote to a proxy from among her neighbors. In our model, individuals are not asked to delegate their votes or preferences: we simply consider the dependency of their preferences on those of others, though this can be viewed loosely as *implicit, partial* delegation of preferences.

3 Computing Winners in the Empathetic Models

We now consider the question of computing the social welfare maximizing alternative in both the local and global empathetic models. In Sec. 2.3, we observed that—for both the local and global empathetic models—social welfare can be expressed as $sw(a, \mathbf{u}^I) = \boldsymbol{\omega}^\top \mathbf{u}^I(a)$ for an appropriate weight vector $\boldsymbol{\omega}$. Given the vectors $\mathbf{u}^I(a)$ for any $a \in \mathcal{A}$, we can readily compute the optimal alternative $a^* = \arg \max_{a \in \mathcal{A}} \boldsymbol{\omega}^\top \mathbf{u}^I(a)$, requiring $O(nm)$ time. Of course, this presupposes access to $\boldsymbol{\omega}$, which has different meanings in each model, and hence requires different approaches for its computation. In the global model, this suggests a different method for computing a^* as well, without (necessarily) requiring the full computation of $\boldsymbol{\omega}$.

We first consider the local model, where $\boldsymbol{\omega}^\top$ can be calculated easily with a single vector-matrix multiplication, $\boldsymbol{\omega}^\top = \mathbf{e}^\top \mathbf{W}$, in time $O(n^2)$. However, social networks are generally extremely sparse, with the number of outgoing edges associated with any node j in the graph bounded by some small constant c which is independent of the network size (generally, social networks, while potentially locally dense, are sparse in a global sense). In sparse networks, $\boldsymbol{\omega}$ can be computed much more efficiently: ω_j is simply the sum of j ’s outgoing edges weights. If the outgoing neighbors of any node are bounded by a constant, $\boldsymbol{\omega}$ can be computed in $O(n)$ time and a^* can be determined in the straightforward fashion mentioned above in $O(nm)$ time. Thus the complexity of computing optimal alternatives in the local empathetic model is no different than that of straightforward social welfare maximization of straightforward (e.g., scoring rule-based) voting.

In the global model, $\boldsymbol{\omega}^\top$ has a more complicated expression, $\boldsymbol{\omega}^\top = \mathbf{e}^\top \mathbf{A}^{-1} \mathbf{D}$ where $\mathbf{A} = \mathbf{I} - \mathbf{W} + \mathbf{D}$ (see Cor. 2.1). The difficulty lies largely in matrix inversion: \mathbf{A}^{-1} can be computed via Gauss-Jordan elimination, which has complexity $O(n^3)$. This implies that straightforward computation of a^* requires $O(n^3 + nm)$ time. In general, matrix inversion is no harder than matrix multiplication (see, e.g., [9, Thm. 28.2]). Although efficient matrix multiplication is the topic of ongoing research (e.g., [8]), its complexity cannot be less than $O(n^2)$ since all n^2 entries must be computed. Therefore, straightforward computation of a^* in the global model cannot have complexity less than $O(n^2 + nm)$.

We expect n to be extremely large in at least some social choice problems on social networks, e.g., in the tens of thousands (number of people in a small town), the millions (large big cities), or hundreds of millions (large country, number of Facebook or Twitter users). This makes algorithms that scale more than linearly in n problematic, both from the perspective of time and memory. Of course, many iterative methods have been proposed for matrix inversion and solving linear systems (e.g., Jacobi, Gauss-Siedel) which have $O(n)$ complexity (in non-sparse systems) per iteration and tend to converge very quickly in practice. We now briefly describe the use of a standard Jacobi method for computing a^* in the global model. We first show how to compute the utility vector $\mathbf{u}(a)$ for each alternative a , and then propose an algorithm called *iterated candidate elimination (ICE)* that will compute the optimal a^* without (necessarily) computing each $\mathbf{u}(a)$ fully.

Consider first a simple iterative method for computing $\mathbf{u}(a)$. Let $\mathbf{u}^{(t)}(a)$ be the estimated utilities for alternative a after t iterations.

Theorem 3.1 *Consider the following iteration:*

$$\mathbf{u}^{(t+1)}(a) = (\mathbf{W} - \mathbf{D})\mathbf{u}^{(t)}(a) + \mathbf{D}\mathbf{u}^I(a).$$

Assuming nonnegativity, normalization, and positive self-loop, this method converges to $\mathbf{u}(a)$, the solution to Eq. 4.

For each $j \in \mathcal{N}$, the method computes:

$$u_j^{(t+1)}(a) = w_{jj}u_j^I(a) + \sum_{k \neq j} w_{jk}u_k^{(t)}(a). \quad (5)$$

We can interpret $u_j^{(t)}(a)$ as agent j 's estimated utility for alternative a after t iterations. This updating scheme has a natural interpretation in terms of agent behavior: suppose that each individual is able to repeatedly observe her friends' revealed utilities, and updates her own utility for various alternatives in response. This process will eventually converge (this is true even if the updates are "asynchronous"). One can readily bound the error in the estimated utilities at the t^{th} iteration:

Theorem 3.2 *In the iterative scheme above,*

$$\left\| \mathbf{u}(a) - \mathbf{u}^{(t)}(a) \right\|_{\infty} \leq (1 - \tilde{w})^t \left\| \mathbf{u}(a) - \mathbf{u}^{(0)}(a) \right\|_{\infty},$$

where $\tilde{w} = \min_{1 \leq i \leq n} w_{ii}$.

Hence, societies in which individuals have self-loops with relatively larger weight (i.e., less empathy) converge to fixed-point utilities faster societies with greater empathy (our empirical results below support this).

This error bound allows one to bound the error in estimated social welfare if the utilities of all alternatives are estimated in this fashion. Let $sw^{(t)}(a) = \sum_j u_j^{(t)}(a)$.

Theorem 3.3 *Assume $u_j^I(a) \in [c, d]$ and $u_j^{(0)}(a) \in [c, d]$, for all $j \in \mathcal{N}$. Under the conditions above, for any t : $|sw(a) - sw^{(t)}(a)| \leq n(d - c)(1 - \tilde{w})^t$, where $\tilde{w} = \min_{1 \leq i \leq n} w_{ii}$.*

As a result, we know that (under the same assumptions):

Proposition 3.4 *If $sw^{(t)}(b) - sw^{(t)}(a) \geq 2n(d - c)(1 - \tilde{w})^t$ then $sw(b) > sw(a)$.*

We can exploit Prop. 3.4 in a simple iterative algorithm for computing a^* we call *iterated candidate elimination (ICE)*. The intuition behind ICE is to iteratively update the estimated utilities of the subset $C \subset \mathcal{A}$ of candidates that are non-dominated, and gradually prune away any candidate that is dominated by another until only one, a^* , remains. Roughly, ICE first initializes $C = \mathcal{A}$ and $u_j^{(0)}(a) = c$ for all $j \in \mathcal{N}, a \in \mathcal{A}$. An iteration of ICE consists of: (1) updating estimated utilities using Eq. 5 for all j and $a \in C$; (2) computing estimated social welfare of each $a \in C$; (3) determining the maximum estimated social welfare $\hat{sw}^{(t)}$; (4) testing each $a \in C$ for domination, i.e., $\hat{sw}^{(t)} - sw^{(t)}(a) \geq 2n(d - c)(1 - \bar{w})^t$; and (5) eliminating all dominated candidates from C . The algorithm terminates when only one candidate (i.e., a^*) remains in C (the pseudo-code for the algorithm is provided in the longer version of this paper). The running time of ICE is at most $O(tmn^2)$ where t is the number of iterations required. More precisely, ICE runs in $O(tm|E|)$ time; and if the number of outgoing edges is bounded, $O(tmn)$. Our hope is that in practice, the methods converges in relatively few iterations, a fact indeed borne out in our preliminary experiments below. ICE also provides a natural means of approximation in large problems.

4 Empirical Results

We now describe some preliminary experiments on randomly generated networks and intrinsic preferences designed to test the differences in the decisions that result under standard non-empathetic, local empathetic and global empathetic models, the impact of these decisions on different agents, and the performance of the ICE algorithm.

Experimental Setup. Our test scenarios require generation of intrinsic preferences and a social network. We assume that individual intrinsic utilities arise from an underlying preference ordering over \mathcal{A} . In all experiments, we assume $m = 5$ or $m = 10$ alternatives, and draw a random preference ordering for each agent j under the impartial culture assumption (all permutations are equally likely). For simplicity, and to draw connections to voting on social networks, we assume j 's utility is given by the Borda or plurality score of the alternative in its ranking. If treating these strictly as utility, they embody very different, extreme assumptions: Borda treats utility differences as smooth and linear, whereas plurality views utility in a more "all or nothing fashion."

We generate random social networks using a *preferential attachment* model for scale-free networks [2] (this is only one of many models that can be adopted). The model works in the following iterative fashion: start with n_0 initial nodes; we repeatedly add nodes (until we have a graph with n nodes), where each new node added is connected to $k \leq n_0$ existing nodes, and an existing node i is selected as a neighbor with probability $P_i = \frac{\deg(i)}{\sum_j \deg(j)}$. We set $n_0 = 2$ and $k = 1$ or $k = 2$ in all our experiments. We then convert the resulting undirected graph to an directed graph by replacing each undirected edge with the two corresponding directed edges; add a self-loop to each node with weight α ; then add normalized weights to all other edges (all outgoing edges from j excluding the self loop have equal weights that sum to $1 - \alpha$). The parameter $\alpha \in (0, 1]$ represents the degree of self-interest, and $1 - \alpha$ the degree of empathy in the society.

Performance Metrics. To measure whether the different models result in difference decisions, we assume the agents *actual utility model* is one of intrinsic (non-empathetic), local or global. We then consider making decisions using any of these models as an *assumed utility model*, and measure the effect on actual utility (e.g., global empathetic utility) of making a decision using the assumed model (e.g., intrinsic). Since decisions might be different in each case, we measure the loss in social welfare due to making a decision using the incorrect model. Let $sw^{ac}(\cdot)$ and $sw^{as}(\cdot)$ be social welfare under the actual and assumed models, respectively, and a_w and a_s be the corresponding optimal alternatives (or winners). We

Actual Utility	Assumed Utility			WSWL
	intrinsic	local	global	
intrinsic	0%(0%)	1.45%(9.95%)	1.10%(8.00%)	5.59%(14.63%)
local	2.95%(19.28%)	0%(0%)	0.09%(3.21%)	11.22%(25.10%)
global	1.78%(12.73%)	0.07%(2.73%)	0%(0%)	9.01%(20.97%)

Table 1: Average (maximum) RSWL and WSWL: 2500 runs, Borda scoring, $m = 5$, $n = 1000$, $k = 1$, $\alpha = 0.25$.

Actual Utility	Assumed Utility		
	intrinsic	local	global
intrinsic	0.0%(0.0%)	28.4%(100.0%)	22.6%(100.0%)
local	28.5%(100.0%)	0.0%(0.0%)	1.2%(86.9%)
global	22.3%(100.0%)	1.1%(97.0%)	0.0%(0.0%)

Table 2: Average (maximum) NSWL: 2500 runs, Borda, $m = 5$, $n = 1000$, $k = 1$, $\alpha = 0.25$.

define *relative social welfare loss (RSWL)* to be $L(as, ac) = [sw^{ac}(a_w) - sw^{ac}(a_s)] / sw^{ac}(a_w)$ (we sometimes report it as a percentage). RSWL has a lower bound that is independent of the assumed model: let the alternative a^- have *minimum social welfare* under the actual model (so it is no better than the decision under the assumed model). *Worst-case social welfare loss (WSWL)* is defined as $W(ac) = [sw^{ac}(a_w) - sw^{ac}(a^-)] / sw^{ac}(a_w)$. Finally, it usually makes sense to normalize RSWL by considering the range of possible social welfare values actually attainable: *normalized social welfare loss (NSWL)* is simply $N(as, ac) = [sw^{ac}(a_w) - sw^{ac}(a_s)] / [sw^{ac}(a_w) - sw^{ac}(a^-)]$. This offers a more realistic picture of loss due to using an incorrect assumed utility model (by comparing it to the loss associated with making the *worst possible decision* under the actual model).

Social Welfare Loss. We first consider RSWL, WSWL and NSWL for all nine combinations of assumed and actual utility models. We fix $\alpha = 0.25$, $n = 1000$, $m = 5$, $k = 1$, and the scoring rule to Borda. We generate 50 random networks, and for each generate 50 intrinsic utility profiles (2500 problem instances), and compute RSWL and WSWL. Average (with maximum in parentheses) RSWL for various combinations of actual and assumed models is reported in Table 1 as are average (maximum) WSWL. Maximum RSWL is more than 19% and 12% when intrinsic utility is assumed but actual utility is local or global, respectively. Moreover, we can see that global vs. local and local vs. global are quite close. Notice that average differences are quite slight: this is because the impartial culture model, in essence, renders alternatives quite close in terms of Borda or plurality score. By normalizing for the fact that most decisions are reasonably good, we get a more accurate picture of the loss incurred by using non-empathetic voting. NSWL is reported in Table 2, which shows that making the wrong assumptions can be quite damaging; e.g., the intrinsic model loses 22.3–28.5% of empathetic social welfare on average.

Since impartial culture is generally viewed as an unrealistic model of real-world preferences, we also tested our methods using preferences drawn from 2002 Irish electoral data from the Dublin West constituency, with 9 candidates and 29,989 ballots of top- t form, of which 3800 are complete rankings.⁵ Generating 1000-node networks as above, we randomly assign full rankings to nodes from this set of 3800 complete rankings. Results on RSWL and WSWL for plurality scoring are shown in Table 3. As above, average RSWL is slight; but the maximum values show significant social welfare loss in certain instances, especially when using the intrinsic model to make decisions for empathetic preferences.

Utility and Societal Weights. We now examine how individual utility—and its intrinsic and empathetic components—and computed weights depend on their degree of nodes in the social networks in global empathetic model. Using data from the previous experiment,

⁵See www.dublincountyrturningofficer.com.

Actual Utility	Assumed Utility			WSWL
	intrinsic	local	global	
intrinsic	0%(0%)	1.82%(33.62%)	1.30%(19.55%)	97.22%(99.62%)
local	2.64%(39.25%)	0%(0%)	0.01%(6.80%)	97.28%(99.85%)
global	1.53%(31.26%)	0.10%(8.4%)	0%(0%)	97.24%(99.77%)

Table 3: Average (maximum) RSWL and WSWL for West Dublin data set: 2500 runs, plurality scoring, $m = 9$, $n = 1000$, $k = 1$, $\alpha = 0.25$.

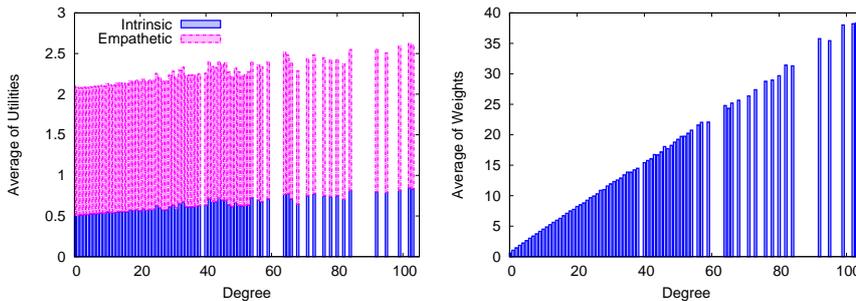


Figure 2: Average (a) intrinsic and empathetic utilities and (b) individual weights as function of node degree (global model, 2500 runs, $n = 1000$, $m = 5$, $\alpha = 0.25$, Borda scoring).

we show average utility and average weight in Fig. 2. From Fig. 2(a) we see that as node degree increases (each node has identical in/out-degree), overall utility tends to increase; moreover most of this increase is due in large part to an increase in intrinsic utility. Fig. 2(b) also illustrates a strong correlation between degree and agent weight. Nodes with higher degree are more powerful and “influential” in the choice of the consensus alternative. This correlation might be an artifact of the specific social networks we generate. However, the relationship between Figs. 2(a) and (b)—which is independent of the specifics of our experiments—shows that individuals with higher weight tend to prefer the consensus winner more than individuals with lower weight.

The effect of m , k , and scoring rule. We now explore the impact on RSWL of changing the numbers of agents m , the number of initial nodes k when generated the network, and difference between Borda and plurality scoring. We set $\alpha = 0.25$, $n = 1000$, and run 2500 instances for each parameter setting (as above).

Fig. 4 shows average (and maximum, minimum) RSWL for three actual, assumed model combinations for various combinations of rule, m and k , denoted by rule(m, k) (e.g., Plura(5, 1) represents $m = 5$, $k = 1$, and plurality). Comparing Borda(5, 1) and Plura(5, 1), and Borda(10, 1) and Plura(10, 1), we see plurality is more susceptible to social welfare loss than Borda. Increasing m has negligible effect on RSWL when Borda is used, but this is not true of plurality. Surprisingly, increasing k from 1 to 2 decreases RSWL (see Borda(5, \cdot)): this occurs because, when $k = 2$, the resulting network is denser since each node has at least two neighbors. This connectivity, causes the number of “very influential” agents to increase; but since weights are normalized (the sum of all weights sums to n), their overall influence decreases as they “share their influence,” and weight variance over \mathcal{N} decreases.

Self-loop weight α . When we vary the self-loop weight α , it has a significant effect on RSWL when the actual utility model is global but the intrinsic utility model is assumed. We fix $n = 1000$, $m = 5$, $k = 1$ and vary α over $\{0.05, 0.1, 0.25, 0.5, 0.75\}$ (2500 instances for each setting). Table 4 shows that, for both Borda and plurality, increasing α (i.e., decreasing overall degree of empathy) decreases RSWL.

Number of Iterations of ICE. Finally we examine how the self-loop weight α and Borda/plurality utilities affect the expected number of iterations by the iterated candidate

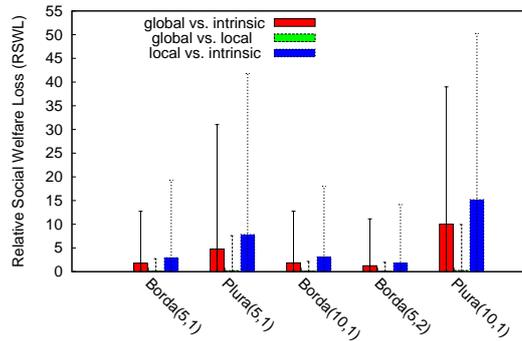


Figure 3: The average (maximum, minimum) RSWL (2500 runs).

	$\alpha = 0.05$	$\alpha = 0.1$	$\alpha = 0.25$	$\alpha = 0.5$	$\alpha = 0.75$	$\alpha = 1$
Borda	15.84%	14.82%	12.73%	9.79%	5.42%	0%
Plurality	39.12%	36.29%	31.02%	22.46%	13.09%	0%

Table 4: Maximum values of RSWL, global vs. intrinsic models.

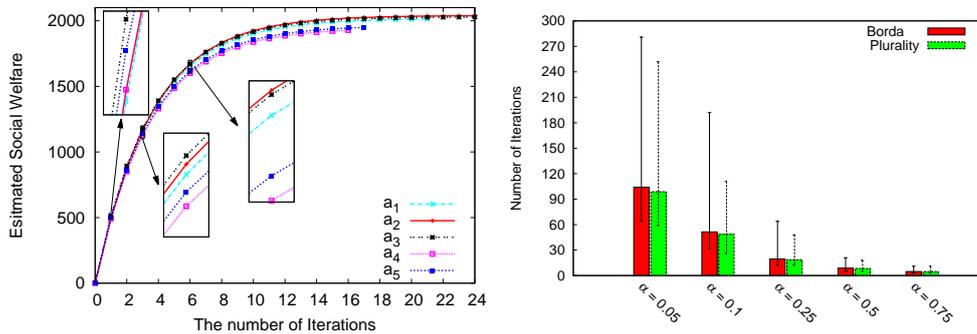
elimination algorithm. We fix $n = 1000$ and $m = 5$, and vary α over $\{0.05, 0.1, 0.25, 0.5, 0.75\}$ (2500 instances). Fig. 4(a) illustrates estimated social welfare for each alternative in one representative instance, with $\alpha = 0.25$ and Borda scoring (this instance of ICE converges in under 2 milliseconds). It converges completely in 24 iterations (n.b. $n = 1000$). Alternatives a_4 and a_5 are eliminated at iterations 16 and 17, respectively; a_1 after 20 iterations; and a_2 after 24 iterations; hence a_3 is optimal. We note that the relative ordering of the alternatives is fixed after 6 iterations (in this instance), which might suggest new methods for early termination.

Fig. 4(b) shows the average (and max, min) number of iterations of ICE for various α , for both Borda and plurality. In all cases, the number of iterations is small compared to the size of the network. ICE is relatively insensitive to the scoring rule, and convergence time decreases dramatically with increasing α , as is typical for iterative algorithms (e.g., for Markov chains). (i.e., for a specific α , the average required iterations is almost the same for Borda and plurality).

5 Concluding Remarks and Future Work

We have presented a new model for social choice situations in which an individual’s intrinsic preference for alternatives is combined with their *empathetic* preferences, reflecting their desire to see others satisfied with the selected alternative. Treating a social network as one possible measure of strength of empathetic preference, we developed models and algorithms, for both local and global empathetic settings, that allow one to compute social welfare maximizing outcomes efficiently by weighting the contribution of each agent. Our models have a natural interpretation as empathetic voting models when scoring rules are used. Critically, we require only that individuals specify their intrinsic preferences (and network weights): they need not reason about their neighbor’s preferences.

This model, while novel, is merely a starting point for a broader investigation into the role of empathetic preferences in social choice. We are currently exploring more realistic processes for simultaneous generation of networks and preferences that are even better suited to empathetic voting than preferential attachment networks. While our focus has been the choice of a single alternative/winner, our model can also be applied to matching, assignment, and other group decision problems; each will require its own analysis and algorithmic de-



(a) Estimated social welfare over iterations (b) Number of iterations for different α

Figure 4: (a) Estimated social welfare over iterations of ICE for 1 run and (b) average (with maximum and minimum) number of iterations of ICE.

velopments. More importantly is the question of the prevalence and strength of empathetic preferences, the extent to which social network structure is indicative of such preferences, and how one can best discover these preferences in practical settings without an excessive burden on users. Two other important directions are: voting schemes in which agents can specify their tradeoffs between intrinsic and empathetic preference in a more qualitative fashion; and considering the possibility that agents are not truthful and fully aware of their utility functions. These questions require both social scientific and computational insight.

References

- [1] S. L. Althaus. Information effects in collective preferences. *Amer. Political Sci. Rev.*, 92(3):545–558, 1998.
- [2] A. Barabási and R. Albert. Emergence of Scaling in Random Networks. *Science*, 286:509–512, 1999.
- [3] E. Bodine-Baron, C. Lee, A. Chong, B. Hassibi, and A. Wierman. Peer effects and stability in matching markets. *SAGT-10*, pp.117–129, 2011.
- [4] P. Boldi, F. Bonchi, C. Castillo, and S. Vigna. Voting in Social Networks. *CIKM’09*, pp.777–786, 2009.
- [5] S. Brânzei and K. Larson. Coalitional affinity games and the stability gap. *IJCAI’09*, pp.79–84, 2009.
- [6] S. Brânzei and K. Larson. Social distance games. *IJCAI’11*, pp.91–96, 2011.
- [7] N. Christakis and J. Fowler. *Connected: The Surprising Power of Our Social Networks and How They Shape Our Lives*. Back Bay Books, 2011.
- [8] D. Coppersmith and S. Winograd. Matrix multiplication via arithmetic progressions. *Journal of Symbolic Computation*, 9(3):251–280, 1990.
- [9] T. H. Cormen, C. E. Leiserson, and R. L. Rivest. *Introduction to Algorithms*. 1990.
- [10] D. A. Easley and J. M. Kleinberg. *Networks, Crowds, and Markets*. Cambridge, 2010.
- [11] J. H. Fowler and N. A. Christakis. Dynamic spread of happiness in a large social network: longitudinal analysis over 20 years in the Framingham Heart Study. *Brit. Med. J.*, 337:a2338, 2008.
- [12] N. Haghpanah, N. Immorlica, V. Mirrokni, K. Munagala. Optimal auctions with positive network externalities. *EC’11*, pp.11–20, 2011.
- [13] M. O. Jackson. *Social and Economic Networks*. Princeton, Princeton, NJ, 2008.
- [14] D. C. Mutz. Cross-cutting social networks: Testing democratic theory in practice. *American Political Science Review*, 96(1):111–126, 2002.
- [15] A. E. Sokhey and S. D. McClurg. Social networks and correct voting. Preprint, 2008.
- [16] A. Wilson, M. Baccara, A. Imrohroglu, and L. Yariv. A field study on matching with network externalities. *Behavioral and Quantitative Game Theory (BQGT’10)*, 2010.

Amirali Salehi-Abari
 Department of Computer Science
 University of Toronto
 Email: abari@cs.toronto.edu

Craig Boutilier
 Department of Computer Science
 University of Toronto
 Email: cebly@cs.toronto.edu