PROJECT SUMMARY

Overview:

Social choice theory is the field that studies the aggregation of individual preferences toward a collective choice, e.g., the design and theoretical evaluation of voting rules. The much younger field of computational social choice originates from a set of insightful observations by Bartholdi, Tovey, and Trick about the role of computational complexity in social choice; to date a significant part of the computational social choice literature directly extends these ideas. While the artificial intelligence (AI) community has so far played a dominant role in computational social choice research, the interaction between core AI paradigms and social choice theory has been surprisingly limited.

This proposal is ambitious in setting an agenda for computational social choice by taking a broad view of AI research, and calling attention to exciting synergies between AI and social choice theory. Indeed, the thesis of this proposal is that insights and techniques from social choice theory can be synthesized with AI research (broadly construed), directly leading to new and improved systems, and in time significantly influencing AI research while also extending the scope of social choice theory.

Specifically, the proposed research aims to explore the interaction between social choice and the following AI areas: (i) decision making under uncertainty, by synthesizing models studied in AI with social choice theory to create new ways to model, analyze, and make decisions in environments where preferences are dynamically changing; (ii) multiagent systems, by studying settings where agents randomly vote over multiple states, and investigating the connection between normative properties and system performance; and finally (iii) machine learning, by employing insights about strategic behavior under structured preferences, developed in the social choice literature, in order to design regression learning algorithms that discourage strategic manipulation, and applying these algorithms to real-world problems that arise, e.g., in business environments.

Intellectual Merit :

The proposed research tackles a variety of problems drawn from different subfields of AI. Each of these problems is important and merits specific attention; solutions can directly enable new technologies and improve existing systems.

Taking a broader, longer-term point of view, an overarching goal of this proposal is to demonstrate the potential of social choice theory to AI researchers, and ultimately to establish social choice theory as a standard paradigm in AI. The proposal serves as a proof of concept by spanning a diverse collection of areas, but may bring about a future where the deeper insights of social choice theory contribute to further areas such as robotics and planning.

Equally importantly, this proposal aims to greatly increase the scope of social choice theory. Indeed, traditional social choice theory is hampered by issues that are circumvented by the proposed research, and furthermore the proposed research gives rise to fundamentally new models of social choice.

Broader Impacts :

The proposal contains an elaborate education, dissemination, and outreach plan, which consists of a variety of interconnected components. Three of these components are key to the plan's impact. First, a new web-based voting system, which has the potential to serve and educate hundreds of thousands of users while gathering useful data that will be made available to the research community. Second, a new book on computational social choice, which will play a role in curriculum development and dissemination of knowledge within and outside of computer science. Third, chairmanship of the Fifth International Workshop on Computational Social Choice (COMSOC 2014), which will help set a new agenda for the field. Some other highlights of the proposed plan include an invited Scientific American article, summer schools and tutorials, outreach via a widely-read blog, curriculum development, and mentorship of graduate and undergraduate students with a focus on diversity.

1 Introduction

Social choice theory is the field that studies the aggregation of individual preferences toward a collective choice. For example, social choice theorists — who hail from different disciplines, including mathematics, economics, and political science — are interested in the design and theoretical evaluation of voting rules. This field has stimulated intellectual thought for centuries. Some trace its roots to Plato [135], and over the centuries it has fascinated many a great mind, from the Marquis de Condorcet and Pierre-Simon de Laplace, through Charles Dodgson (a.k.a. Lewis Carroll), to Nobel Laureates such as Kenneth Arrow.

Computational social choice, which by comparison is a very young field — only two decades old — originates from a set of insightful observations by Bartholdi, Tovey, and Trick: that computational complexity can serve as a barrier against manipulation of elections [20], but at the same time can prevent the use of complex voting rules in practice [21]. With several exceptions [19, 76], the field was largely dormant until it started gaining momentum a decade ago. To date a significant part of the computational social choice literature directly extends the ideas of Bartholdi, Tovey and Trick in dealing with the role of computational complexity in social choice [121, 106, 149, 49, 62, 63, 22, 80]; see the survey by Faliszewski and Procaccia [64]. More generally, most of the work on computational social choice has concentrated on applying *computational thinking* [143] to augment, analyze, or reevaluate social choice settings and ideas [65, 117, 29, 34, 95, 97, 48, 58, 83, 139]. Similar statements are true with respect to related fields, such as end-to-end verifiable voting [41, 42].

While the artificial intelligence (AI) community has so far played a dominant role in computational social choice research, the interaction between core AI paradigms and social choice theory has been surprisingly limited. Prominent examples of this interaction include papers that combine constraint handling and knowledge representation (KR) ideas — such as CP-nets [25] — with social choice [86, 145, 114]; see the monograph by Rossi et al. [130] for an overview. My own work with colleagues takes a learning-theoretic approach to the design of desirable voting rules [125, 124, 126], and further afield the machine learning literature on permutations [46, 79, 92] explores automated methods for producing good rankings. Other papers by AI researchers take a logic-based view of voting [127]. In addition, some AI artifacts use simple voting methods [75, 70, 134, 129], but there are relatively few applications of the deeper insights of social choice theory to the design of AI systems.

This proposal is ambitious in setting an agenda for computational social choice by taking a broad view of AI research, and calling attention to exciting synergies between AI and social choice theory. Indeed, my thesis is that

> *... insights and techniques from social choice theory can be synthesized with AI research (broadly construed), directly leading to new and improved systems, and in time significantly influencing AI research while also extending the scope of social choice theory.*

Specifically, the proposed research aims to explore the interaction between social choice and the following AI areas.

- *Decision making under uncertainty*. In some settings, such as collaborative computational science and public policy advocacy, sequential decision making is taking place based on dynamically changing preferences of individuals, where the preferences in the next stage can depend stochastically on the decision made in the current stage. Today decisions in these settings are made in an unprincipled way that does not take a long-term view. I propose to synthesize social choice theory with models of decision making under uncertainty studied in AI, with the goal of providing algorithms that link sequential decisions in an optimal way while keeping the decision making mechanism socially desirable. Ultimately I envision decision support systems for settings where currently there are none.

- *Multiagent systems*. Interactions between multiple software agents naturally occur in the context of electronic commerce [140], interdomain routing [89], and other Internet environments; or in teams of robots [111]. As a design principle, the multiagent paradigm, which has been applied in a variety of domains [88, 133], enjoys properties such as scalability and simplicity [40]. Voting offers a natural method for merging the opinions of teams of agents or reaching consensus among multiple agents more generally, especially in systems where the scores assigned by diverse agents to alternatives are hard to compare across agents; but which voting rule should be used? I start from exciting new settings that involve random voting over multiple states, with potential applications ranging from computer Go to disaster response. Then, taking a broader view of multiagent systems, I propose to investigate the relation between well-studied *normative* properties of voting rules, and the performance of multiagent systems that are governed by those rules. These research directions enrich social choice theory itself by inspiring new models of voting.

- *Machine learning*. Recent applications [38] deploy machine learning techniques in business settings such as inventory management. However, existing applications are susceptible to strategic manipulations that decrease the efficiency of the system. I suggest that it is possible to apply social choice methodology for dealing with structured preferences in strategic settings in order to construct machine learning algorithms that discourage manipulation. This can directly enable the deployment of more efficient inventory management systems.

Intellectual merit. The proposed research tackles a variety of problems drawn from different subfields of AI, as discussed above (and in more detail below). Each of these problems is important and merits specific attention; solutions can directly enable new technologies and improve existing systems.

Taking a broader, longer-term point of view, I wish to promote the synthesis between AI and social choice as a new agenda for the field of computational social choice. An overarching goal of this proposal is to demonstrate the potential of social choice theory to AI researchers, and ultimately to establish social choice theory as a standard paradigm in AI. The proposal serves as a proof of concept by spanning a diverse collection of areas, but I envision a future where the deeper insights of social choice theory contribute to further areas such as robotics and planning.

It is equally exciting to note that this proposal has the potential to significantly increase the scope of social choice theory. Indeed, traditional social choice theory is hampered by the difficulty of changing voting systems in political settings, and by a plethora of impossibility results. I believe that the proposed research, as part of the emerging field of computational social choice, can help address these two shortcomings. Indeed, the first is a nonissue in the settings that I propose to study, because the designer of an AI system is free to employ any voting rule. The second obstacle is circumvented by the constructive approach that is prevalent in this proposal and calls for brand new models of social choice.

Broader impacts. This proposal contains an elaborate education, dissemination, and outreach plan, which consists of a variety of interconnected components. I view three of these components as key to the plan's impact. First, I propose to design, implement, and deploy a web-based voting system, which will enhance student education (in and out of class); aid the broader research community by gathering much-needed real-world ranking data; and serve as a collective decision making tool for small groups and larger organizations, with potentially hundreds of thousands of users. Second, I am editing (together with several prominent colleagues) a book on computational social choice that will be published by Cambridge University Press;

I explain how the proposed research can enhance the book, and how the book, in turn, will play a role disseminating the results to colleagues in CS and economics. Third, I am co-chair of the Fifth International Workshop on Computational Social Choice (COMSOC 2014), which will take place at Carnegie Mellon University in June 2014; this is an opportunity to set a new agenda for the computational social choice community and attract more US-based researchers to the field.

Some other highlights of the proposed plan include: (i) writing an invited article on computational social choice for the magazine *Scientific American*, which will present some of the proposed research, and leveraging this article for additional outreach; (ii) building on an NSF-sponsored summer school that I recently organized to hold summer schools and tutorials; (iii) enhancing dissemination and outreach through a widely-read blog that I regularly contribute to; (iv) developing a new course on computational social choice; (v) mentoring graduate and undergraduate students and encouraging diversity through organizations such as CMU's Women@SCS.

2 Background

Perhaps the most common social choice model includes a set of *agents* $N = \{1, \ldots, n\}$, and a set A of alternatives. The preferences of agent i over the alternatives are expressed as a linear order σ_i. In other words, each agent ranks the alternatives, from most preferred (position one) to least preferred (position $|A|$). I denote the set of linear orders over A by $\mathcal{L} = \mathcal{L}(A)$. A *preference profile* $\langle \sigma_1, \ldots, \sigma_n \rangle \in \mathcal{L}^n$ is a collection of the preferences of all agents.

A *voting rule* is a function $f : \mathcal{L}^n \to A$ that selects a single alternative given the preferences of the agents. For completeness I mention that a *rank aggregation rule* $h : \mathcal{L}^n \to \mathcal{L}$ also receives a preference profile as input, but returns a ranking over the alternatives.

Perhaps the best-known family of voting rules is *positional scoring rules*. Each positional scoring rule is defined by parameters $\alpha = \langle \alpha_1, \ldots, \alpha_{|A|} \rangle$. Each agent awards α_k points to the alternative ranked in position k in its preferences; the alternative with the largest total number of points is selected (this may require breaking ties). The most common voting rule in political elections, *plurality*, can be represented as a positional scoring rule with the vector $\langle 1, 0, \ldots, 0 \rangle$. The voting rule proposed by the Chevalier de Borda in the 18th century, known today as *Borda count*, is another famous positional scoring rule with parameters $\langle |A| - 1, |A| - 2, \ldots, 0 \rangle$. *Veto* is the positional scoring rule given by the parameters $\langle 1, \ldots, 1, 0 \rangle$.

Over the years many voting rules have been proposed. Social choice theorists usually compare different voting rules according to normative guidelines, formalized as *social choice properties* (or axioms). One such property is called *monotonicity*: pushing an alternative upwards in the preferences of the agents, while everything else stays fixed, should not harm the lucky alternative, in the sense that if it was a winner before the improvement it should remain a winner. *Condorcet consistency*, named after the Marquis de Condorcet (1743–1794), stipulates that if there is an alternative that is preferred by a majority of agents to every other alternative, that alternative should be selected by the voting rule. Two of the most fundamental properties are *anonymity*, which informally means that the voting rule is indifferent to the names of the agents, and *neutrality*, which means that the voting rule is indifferent to the names of the alternatives.[1]

A central topic in social choice theory, which lies at its intersection with game theory, is the prevention of manipulation in elections. Informally, a voting rule is *strategyproof* if an agent cannot benefit by misreporting its preferences, regardless of the preferences reported by other agents. Borda count, for example, is not strategyproof. To see this, let $n = 3$, $|A| = 4$, and assume alternative a is ranked first by agents 1 and 2

[1] Anonymity and neutrality are actually mutually exclusive for some values of n. Although voting rules like plurality seem to satisfy both, it is provably impossible to break ties in a neutral way.

and second by agent 3, and alternative b is ranked second by agents 1 and 2 and first by agent 3. Alternative a is the winner with 8 points, versus b with 7 points. However, by switching a with its bottom-ranked alternative, agent 3 can decrease the score of a to 6, thereby crowning b which is its most preferred alternative. More generally, the striking Gibbard-Satterthwaite Theorem [71, 131] informally states that when $|A| \geq 3$, a "reasonable" voting rule cannot be strategyproof.

Fortunately, it is possible to circumvent the Gibbard-Satterthwaite theorem in settings where the preferences of agents can be assumed to possess a structure that is more specific than arbitrary rankings. The best-known example is *single-peaked preferences* [23], which means that the alternatives can be ordered so that each agent i has a *peak* $p_i \in A$, and if $p_i \preceq x \prec y$ (according to the underlying order on the alternatives) then agent i prefers x to y, and similarly if $p_i \succeq x \succ y$ then agent i prefers x to y. When the preferences of agents are single-peaked there are socially desirable strategyproof voting rules, and in fact the family of all strategyproof voting rules in this setting has been characterized [109].[2]

3 Research Plan

One of the overarching long-term goals of this proposal is to demonstrate the potential of a *broad* synthesis of AI and social choice. To do this, I propose research challenges that span three subfields of AI, corresponding to three subsections. The challenges in each subsection are related and together can bring about significant innovations in their respective subfield. Crucially, though, here the whole is greater than the sum of its parts: Only the complete picture suggests an impact on AI research as a whole through the integration of social choice theory.

3.1 Social Choice and Decision Making Under Uncertainty

Social choice theory, whose basics were formally presented in Section 2, is inherently static in nature, in that a one-shot collective decision is being made based on a fixed set of agents, a fixed set of alternatives, and fixed preferences. Despite the theory's wide scope, it falls short when considering *dynamic* settings, in particular settings where decisions must be made in the context of a population with constantly changing preferences, and the evolution of future preferences depends (stochastically) on past preferences and past decisions.

Consider, for example, decision making in collaborative projects. Examples include collaborative computational science involving several groups of researchers, where the alternatives represent the different experiments that can be run on shared computational resources, and open source software projects such as Debian [1]. In collaborative science projects, it is natural that preferences over which experiment to run next change in a way that depends, likely probabilistically, on previous experiments. Debian already employs an elaborate voting rule in order to facilitate decisions on issues ranging from the choice of a logo to which features to implement. Debian, as an example of open source software projects, naturally supports dynamic social choice, where the goal is to sequentially implement features in a way that improves the system as a whole, while keeping the many developers satisfied with the decision making process.

As another example, consider online public policy advocacy groups, which are quickly gaining popularity and influence on the web and in social networks such as Facebook (via the application Causes). Some of these groups[3] boast millions of members but employ only a handful of staffers. Ideally the causes or issues

[2] Arguably the most natural and desirable strategyproof voting rule under single-peaked preferences orders the agents' peaks according to the underlying order on alternatives, and returns the median peak.

[3] I do not mention specific groups to avoid the appearance of expressing political preferences.

that are advocated by a group directly stem from the collective preferences of the members. A salient feature of public policy advocacy groups is that the time frame between deciding on a cause and acting on it is very short. Crucially, when a cause is chosen and advocated, the preferences of the members will usually change, and this should have an impact on the next cause to be chosen. So, we are faced with a situation where both the current cause and the preferences of the members are constantly shifting. This calls for a consistent and socially desirable mechanism that sequentially selects the current cause given the current preferences of the group's members.

In very recent work with David Parkes of Harvard University [113] we introduced dynamic preferences into the traditional, static social choice setting by representing the preferences of each agent as a random variable whose current value is the agent's current ranking of the alternatives. Each agent is associated with a Markovian transition model that stochastically determines how the agent's preferences are updated based on the agent's current preferences and the currently selected alternative (c.f. [39] for dynamic incentive mechanisms). This model can be expressed using the formalism of *Markov decision processes (MDPs)* — specifically, *factored* MDPs [27]. Note that the Markov assumption is essentially without loss of generality because histories of bounded length can be handled at the cost of enlarging the state space.

Parkes and I introduced *social choice MDPs*, where a state — which represents the current value of each of the agents' random variables — corresponds to a preference profile. A deterministic (stationary) policy in an MDP maps each state to the action taken in that state. Crucially, a deterministic policy in a social choice MDP corresponds to a voting rule: it is a function from preference profiles to alternatives.

The final component of a social choice MDP is a reward function, which associates a given action taken in a given state with a reward; the goal is to optimize the infinite sum of discounted rewards. The existence of such an explicit objective is novel from a social choice perspective. Interesting reward functions depend (in full or in part) on which preferences agents hold. In the context of public policy advocacy, for example, a staffer may wish to steer society toward states where there is social consensus (because consensus empowers advocacy), and is rewarded when consensus (e.g., in the Condorcet sense) is achieved.

Unfortunately, a policy that optimizes the discounted rewards, when viewed as a voting rule, may not satisfy basic social choice properties, and hence may be undesirable as a social decision making mechanism. The algorithmic goal is therefore to *tractably compute an optimal deterministic policy subject to given social choice constraints*. Policies in this setting are simply traditional voting rules and hence it is possible to employ the social choice properties introduced in Section 2, thereby leveraging the traditional theory of social choice to this novel, dynamic setting. The traditional static properties are indeed still meaningful because, unlike the center, the agents do not necessarily take a long-term view and look for social fairness (e.g., selecting a consensus alternative when one exists) on a per-decision (i.e., per-stage) basis.

Our existing algorithmic results [113], which build on techniques developed in AI for exploiting symmetries in MDPs [148, 128], establish that when the number of alternatives is constant, optimal policies that satisfy a variety of given social choice constraints can be computed in polynomial time in the number of agents. These preliminary results are meant to establish that the problem is provably tractable in the worst case, and while I believe them to be of significant theoretical interest, the existing algorithms were not designed to scale well when the number of alternatives is not very small. I therefore propose to:

Research Challenge 1. *Devise heuristic algorithms and approximation algorithms that compute (almost) optimal policies under social choice constraints, and scale well with the number of agents and alternatives.*

One approach for designing exact heuristic algorithms builds on techniques introduced by Dolgov and Durfee [56]. Indeed, an optimal randomized policy in an MDP can be obtained as the solution to a linear program (LP). The variables in the dual LP, known as *occupation measures*, naturally allow for introducing constraints. Dolgov and Durfee observe that it is possible to search only deterministic policies by adding binary variables and some additional constraints.[4] Some combinations of axioms, e.g., anonymity and monotonicity, can then be encoded as constraints using the occupation measures. I will validate the scalability of this approach via simulations.

The construction of approximation algorithms would start from the work of Guestrin et al. [74]. They present algorithms that exploit the structure of factored MDPs, and use an approximate value function that can be represented as a linear combination of basis functions, where each basis function is defined over a small number of variables. While their techniques are not directly applicable to Research Challenge 1, I am optimistic that it will be possible to develop new techniques for computing socially desirable policies under the assumption of an approximate value function representation. My goal is to design computationally-efficient algorithms that provide worst-case approximation guarantees[5] as well as good empirical performance in terms of solution quality.

The model described above takes the agents' preference transition models as a given. That is, I assume above that the algorithm has access to a stochastic model that provides the probability of an agents' preferences transitioning between every pair of rankings, conditioned on the action that is currently taken. In the public policy advocacy example, I imagine categorizing agents according to carefully selected features, and eliciting the vector of features from each agent. Each vector of features can then be associated with a transition model in an approximate way. Finding good features and constructing the associated transition models is a nontrivial problem.

Research Challenge 2. *Develop transition models that can be automatically configured, given vectors of features, to predict how agents' preferences evolve.*

There are a variety of techniques that may be suitable, ranging from automated methods such as machine learning [51] and hidden Markov models, to marketing and psychological approaches. I wish to emphasize though that the purpose of this research is to provide the first principled approach to decision making in environments where currently one is not available. Even a coarse partition of the agents into just several transition models, and subsequent application of the algorithms obtained via Research Challenge 1, would be an improvement over the status quo. Over time this initial partition can be gradually refined to yield better and better approximations of reality.

Above I have implicitly assumed that the sets of agents and alternatives remain fixed throughout the dynamic process. These assumptions can be justified (I omit the details due to space constraints), but the assumption about the fixed set of alternatives does have some bite, and therefore a theoretical model that allows for dynamic alternatives would certainly be superior.

Research Challenge 3. *Extend the preliminary dynamic social choice model to allow for a dynamic set of alternatives.*

I conclude with two comments. First, the dynamic social choice model is special in that it marries the optimization approach popular in computer science with the normative approach that is common in social choice theory, by using normative properties as constraints for the optimization problem. Second, I believe that the scope of this research challenge is even wider. For example, in the political context, one can imagine providing a decision support system to a party leader in the US Senate or House of Representatives, although this is admittedly ambitious.

[4] Although there is some literature on randomized voting rules (see, e.g., [72, 52, 117]), deterministic voting rules are nevertheless considered significantly more desirable; the property of being deterministic can be viewed as an additional, implicit, social choice constraint. In contrast, in Section 3.2 I discuss randomized voting where the (software) agents themselves are randomized.

[5] A number of existing papers, by us [33, 34] and others [14, 83], design approximation algorithms (with worst-case multiplicative guarantees) for hard-to-compute voting rules.

3.2 Social Choice and Multiagent Systems

Some believe that one of the goals of computational social choice should be leveraging centuries of research on social choice to facilitate preference aggregation in multiagent systems. However, our understanding of the advantages of this approach is still relatively limited. For example, researchers have designed teams of software agents that vote over movie recommendations [70]; weighted majority voting (as one of several *ensemble methods* [55]) has been shown to be competitive for tasks such as gene function prediction [129]; and voting has been used for *data fusion* in sensor networks [112]. My own previous work with Ioannis Caragiannis of the University of Patras [35] shows that simple voting rules can be employed to reduce communication in multiagent systems. These papers and others establish the potential of teams of voting agents, and voting in multiagent systems more generally. I wish to take the next step in realizing this potential by employing *modern* social choice results, and developing new models of social choice.

I propose to start from a very recent paper by Marcolino et al. [100], which inspires a novel way of thinking about voting among multiple software agents. Their empirical results focus on Computer Go programs (see, e.g., [60]), which often use Monte Carlo tree search algorithms [32]. Taking the team formation [69, 91, 13, 101] point of view, Marcolino et al. establish that a team consisting of multiple (four to six) different computer Go programs that use plurality voting (see Section 2) to decide on each move outperforms a team consisting of multiple copies of the strongest program (which is better than a single copy because the copies are initialized with different random seeds). The insight is that even strong agents are likely to make poor choices in some states, which is why diversity beats strength. And while the benefits of diversity in problem solving are well studied [77, 78, 141, 31, 90], the setting of Marcolino et al. combines several ingredients. First, performance is measured across multiple states; as they point out, this is also relevant when making economic decisions (such as stock purchases) across multiple scenarios, or selecting item recommendations for multiple users. Second, agents' votes are based on randomized algorithms; this is also a widely applicable assumption, and in fact even Monte Carlo tree search specifically is used for problems ranging from traveling salesman [115] to classical (deterministic) planning [146], not to mention that randomization is often used in many other AI applications such as motion planning [87]. Therefore, the setting is relevant to futuristic multiagent systems that rely on planning and search, such as the ones designed for disaster response [133].

Focusing on the computer Go application for concreteness, I find it exciting because it provides an ideal example of voting among teams of software agents: It is difficult to compare quality scores assigned by heterogeneous agents to different moves, so optimization approaches that rely on *cardinal utilities* fall short while voting provides a natural aggregation method. More generally the setting's new ingredients call for a novel model of social choice, which should be rich enough to explain the empirical finding that diversity beats strength. But Marcolino et al. [100] suggest a theoretical model that is rather restricted: They prove that a diverse team would outperform copies of the strongest agent *only if* one of the weaker agents outperforms the strongest agent in at least one state; their model cannot quantify the advantage of diversity. Moreover, the work of Marcolino et al. [100] — and other related work [31] — assumes that each agent votes for a single alternative. In contrast, strong open-source Go programs like Fuego [60] can potentially generate a ranking of multiple moves (at most 361) on a full 19×19 Go board, calling for a principled way to harness this additional information.

Research Challenge 4. *In the above setting, construct and evaluate a model that quantifies and harnesses the benefits of voting among diverse (randomized) software agents.*

I propose the following model as a starting point. The strength of an agent is a parameter $\phi_i \in [\alpha, 1-\alpha]$ for a constant $\alpha > 0$. Assume that in each state $s \in S$ there is an underlying ranking σ^*_s of the alternatives according to their true strength. The bias of agent i in state s is captured by a noisy ranking that is drawn from a standard noise model known as the *Mallows model* [98, 96]. Specifically, the probability of agent i being biased toward a ranking σ in state s is proportional to $(\phi_i)^{d_{KT}(\sigma_s^*,\sigma)}$, where d_{KT} is the *Kendall tau distance*, which counts the number of pairs of alternatives on which the two rankings disagree; note that the noise level is lower (i.e., rankings that are similar to the true one are more likely) the closer the strength parameter ϕ_i is to zero. Now, when agent i actually votes in state s, it again draws a sampled ranking from the Mallows model, but this time with the *biased* ranking as the underlying ranking. Hence, even strong agents may be biased in some states, and multiple copies of the same agent cast different votes.

In this model, I can prove that the expected number of states where a diverse team of agents selects the optimal alternative via plurality voting approaches $|S|$ as the number of agents n grows; this can be viewed as an extension of the Condorcet Jury Theorem [53]. Crucially, though, this is not the case for multiple copies of the same agent: If $|S| = \Omega(\log n)$, i.e., the number of states is at least logarithmic in the number of agents, almost surely[6] the expected number of states where n copies of any single agent select the optimal alternative via plurality is upper-bounded by $\beta|S|$ for a constant $\beta < 1$. Intuitively, one can expect each agent $i \in N$ to be weak — in the sense of having a highly biased ranking — in a constant fraction of states $s \in S$, so a team consisting of copies of any individual agent is likely to vote suboptimally in many states.

This preliminary model, which was concocted during the preparation of this proposal, is meant to establish the feasibility and potential of Research Challenge 4; it is limited in that it does not take into account the sequential nature of the domain (cf. [61]). More importantly, although agents submit rankings, this model is still restricted to plurality voting and its predictions hold only in the limit. What can social choice theory tell us about which voting rule to choose when the number of agents is realistic?

As before, think of the votes as noisy estimates of an underlying ground truth, drawn from a noise model such as the Mallows model. A promising approach, which was advocated by the Marquis de Condorcet and modernized by Young [147], views voting rules as *estimators*. From this viewpoint, one voting rule is better than another if it is more likely to output the true underlying ranking; the best voting rule is a maximum likelihood estimator (MLE) of the true ranking. Building on work in AI [48, 47, 59, 18], I have shown in work with colleagues (based on theory, simulations, and real data from Amazon Mechanical Turk) that this approach can inform the choice of voting rules in human computation systems [120, 99].

But what is the right noise model? The Mallows model implicitly assumes that an agent is as likely to mistakenly swap a pair of good alternatives (positioned near the top of the true ranking) as a pair of bad alternatives (positioned near the bottom). In contrast, agents that rely on techniques like Monte Carlo tree search are more likely to correctly compare good alternatives and less likely to distinguish between bad alternatives. This calls for noise models where the probability of a vote can depend in nuanced ways on its distance from the true ranking. In recent work with Ioannis Caragiannis and my student Nisarg Shah [36] we explored general noise models and characterized the connection between distance functions between rankings (which determine the probability of a vote) and good voting rules. I plan to leverage our results to address Challenge 4 on a theoretical level. For the purposes of empirical evaluation, note that the computer Go domain provides an excellent testbed because several strong Go programs are open source [100].

As promising as the MLE point of view is, the common way of thinking about voting rules may turn out to be even more exciting. Indeed, let us reconsider the above discussion on a higher level of generality. Social choice theorists usually take a *normative* approach to the comparison of voting rules: They propose and study intuitive, desirable properties that voting rules should satisfy, as described in Section 2. Designing for such normative properties is very natural when one has political elections or voting in committees in mind. However, in multiagent systems one usually seeks to optimize a measure of system performance.

[6] With probability that approaches 1 as n grows.

These measures can be context-specific, e.g., the number of conflicts in a scheduling system [75] or the winning rate of a computer Go system, or context-independent, like the system's robustness to random failures (which for example plays an important role in sensor networks [88]). In any case, intuitive notions of social justice may seem meaningless in multiagent systems. Nevertheless, I believe that social choice properties can be rigorously related to system performance.

Research Challenge 5. *Relate well-studied normative social choice properties, such as monotonicity and Condorcet-consistency, to prominent operational measures of system performance.*

Progress on this research challenge has the potential to significantly impact the interaction between social choice and AI, and would require conceptual breakthroughs. Fortunately some existing papers provide possible springboards. In particular, a paper by Dwork et al. [57] connects a generalized version of Condorcet-consistency to aggregation of web search results in a way that is resistant to spam, and although this connection is exciting, here system performance is measured in a rather restricted way.

More generally, but in a somewhat different setting, a body of work in combinatorics and the theory of computer science considers Boolean functions $f : \Omega_n \to \{-1, 1\}$, where $\Omega_n = \{-1, 1\}^n$ is the Hamming cube, and studies their *noise sensitivity*, i.e., the probability that a random flip of several bits in the input of f would flip its output. It is known [108] that *monotone* Boolean functions — where if $x_i \leq y_i$ for all i then $f(\mathbf{x}) \leq f(\mathbf{y})$ — provide some guarantees with respect to noise sensitivity. Boolean functions are also known to be closely related to voting rules [81]. Moreover, in work with colleagues, we have previously studied the degree of robustness to failures exhibited by common voting rules [122], in settings where communication of preferences is noisy. By building on this chain of connections I hope to take a first step in tackling Research Challenge 5 by formally relating the monotonicity of a voting rule to the robustness to failures of a multiagent system governed by that rule.

3.3 Social Choice and Machine Learning

A regression learning setting includes an input space $\mathcal{X}$ (e.g., Euclidean space), and a class $\mathcal{F}$ of functions $f : \mathcal{X} \to \mathbb{R}$ (e.g., linear functions). A regression learning algorithm is given a dataset of examples $\langle \mathbf{x}, y \rangle$, where $\mathbf{x} \in \mathcal{X}$ is an input point and $y \in \mathbb{R}$ is a *label*, drawn according to an unknown distribution μ on $\mathcal{X} \times \mathbb{R}$. The goal is to return a function $f \in \mathcal{F}$ such that when a pair $\langle \mathbf{x}, y \rangle$ is drawn according to μ, $f(\mathbf{x})$ is close to y on average. Put another way, a regression learning algorithm attempts to fit the observed data in a way that generalizes well.

Consider for a moment the distribution process of a large retail chain; to be concrete I focus on the Spanish fashion company Zara. Store managers typically report their predicted demand to the central warehouse, where global shipments of inventory are optimized. In recent years Zara has reengineered its distribution process [37, 38]. Crucially, regression learning is now employed to predict the upcoming weekly demand. The prediction is based on past sales data, but also on requests by store managers. This introduces incentives for the store managers, whose salaries depend largely on the sales in their own stores. Caro et al. [38] write that "this caused store managers to frequently request quantities exceeding their true needs, particularly when they suspected that the warehouse might not hold enough inventory of a top-selling article to satisfy all stores. ... Zara might in time consider introducing formal incentives for store managers to provide accurate forecasts, adding to its more traditional sales-related incentives" [38, p. 74]. Similar issues in the context of automotive industry supply chains have been studied by Altintas, Montgomery, and Trick [16, 17].

In work with colleagues [54], we proposed a novel model of incentives in regression learning. Under this model, strategic agents are interested in different outcomes for the learning process. The labels are obtained from the agents; an agent can misreport its label in order to steer the learning process toward its desired outcome. The goal is to design regression learning algorithms that are *strategyproof*, i.e., immune to manipulation. My work with a different group introduces and studies a closely related model of incentives in classification [102, 103, 104, 105], but here I focus on the regression model for ease of exposition.

The reader may be wondering why this setting is at all related to the themes of this proposal. As it turns out, the literature on single-peaked preferences (see Section 2) is crucial for the design of strategyproof regression learning algorithms. To gain some intuition, consider a very special case where $\mathcal{X} = \mathbb{R}^d$ and $\mathcal{F}$ is the class of constant functions over $\mathcal{X}$. For simplicity, assume that each agent i controls a set of examples $\langle \mathbf{x}_{i1}, y_{i1}\rangle, \ldots, \langle \mathbf{x}_{im}, y_{im}\rangle$ and (unlike our general model) there is no sampling or uncertainty. Given a constant function $f(\mathbf{x}) \equiv c$ that is the outcome of the machine learning process, the *cost* (or disutility) of an agent is induced by the *absolute loss* function: $\sum_{j=1}^{m} |y_{ij} - f(\mathbf{x}_{ij})| = \sum_{j=1}^{m} |y_{ij} - c|$. It is easy to see that an agent's preferences over $\mathcal{F}$ are single peaked; the peak of agent i is the constant function $f^*(\mathbf{x}) = c^*$, where c^* is the median of the labels $y_{i1}, \ldots, y_{im}$. This insight, which also extends to other function classes (e.g., homogeneous linear functions over $\mathbb{R}$), has inspired our design of Project-and-Fit, a regression learning algorithm that in some settings is strategyproof and provides guarantees with respect to the quality of the solution[7] [54].

The strategyproofness guarantees provided by Project-and-Fit only hold for specific function classes, which induce single-peaked preferences. But other prominent function classes (e.g., nonhomogeneous linear functions over $\mathbb{R}^d$ for $d > 1$) induce highly structured preferences that are not single peaked.

Research Challenge 6. *Design strategyproof regression learning algorithms for prominent function classes.*

Specifically, I am interested in strategyproof algorithms that perform well in terms of loss. I will explore both worst-case approximation guarantees — strategyproof algorithms that are almost optimal in that sense (i.e., have an approximation ratio that is very close to 1) would need to be randomized [54], unlike Project-and-Fit — and empirical performance.[8]

A first approach draws on deep results from the social choice literature. For example, Border and Jordan [24] give possibility results for star-shaped preferences[9] over $\mathbb{R}^d$ that can help address Challenge 6.

Another approach, which I view as timely and even more promising, is informed by recent, seemingly unrelated results on locating a facility on the real line. To see the connection on an intuitive level, assume again that $\mathcal{F}$ is the class of constant functions. Since only the labels of examples matter, we can project each example onto the y axis (so an example corresponds to its label), and think of a constant function $f(\mathbf{x}) \equiv c$ as locating the facility at the point $y = c$ on the y axis. In my work with Moshe Tennenholtz [123] we asked whether one can design strategyproof mechanisms for various facility location problems in a way that provides worst-case approximation guarantees, e.g., with respect to the sum of distances between agent locations and the facility location. Since an early version of that paper became publicly available in November 2008, a considerable amount of work has been devoted to improving and extending our facility location results; some published papers include [15, 94, 93, 110, 66, 136, 137, 44, 67, 68, 142]. Many of these papers explore facility location variations where preferences are not single peaked; some of them (e.g., [15, 137, 67]) draw on insights from social choice theory (e.g., [109, 132]). I am optimistic that some of these cutting-edge techniques can be mapped back to the regression learning setting.

While the basic model introduced in our previous work [54] can inform the design of real-world inventory management systems and supply chains, it is not sufficiently powerful to fully capture the intricacies of such applications. Once the model is fully understood, I plan to make it richer.

[7] Specifically, the loss is at most three times higher than the optimal solution in the worst case.

[8] The design of strategyproof algorithms that perform well empirically is sometimes called *heuristic mechanism design* [50, 85].

[9] Intuitively, preferences over $\mathbb{R}^d$ are star-shaped when there is an ideal point $\mathbf{x}$ such that for every $\mathbf{y}$ and every point $\mathbf{z}$ on the line between $\mathbf{x}$ and $\mathbf{y}$, $\mathbf{z}$ is preferred to $\mathbf{y}$.

Research Challenge 7. *Extend the model of incentives in machine learning to accurately capture real-world inventory management processes, and ultimately other learning-based business environments such as supply chains.*

This challenge is, well, challenging, because in real-world applications it is nontrivial to identify the function class that is being learned. In particular, under Zara's process, optimization and learning are separate. Rather than directly analyzing existing processes, I hope to redesign them, and ultimately build new systems to validate my approach.

To put Challenge 7 in context, there are several other approaches that are relevant. For example, *peer prediction* methods [116, 107, 144] aim to truthfully elicit predictions, but the agents do not typically care about how the information they provide affects a global decision making process.[10] *Mechanism design* also offers relevant techniques, and in fact we have shown [54] that an extension of the Vickrey-Clarke-Groves Mechanism [138, 45, 73] provides guarantees for regression learning. However, this approach relies on monetary transfers and makes strong assumptions with respect to the relative value that agents attach to outcomes and money. While I believe that the social choice approach to Challenge 7 is especially exciting, I will draw on all of the pertinent literature.

3.4 Research Timeline

I plan to focus my initial efforts (years 1–2) on Research Challenges 1, 4, and 6, which are especially timely, very likely to admit feasible solutions, and can serve as springboards for addressing other challenges. Years 3–5 will be devoted to Research Challenges 2, 3, 5, and 7, which are longer-term in nature.

4 Education and Dissemination

Broadly speaking, one of the main goals of this proposal is to convey my enthusiasm about the proposed research, and, more generally, computational social choice and AI, to my students, to colleagues in academia, and to the public. With this goal in mind, I present below a plan of education, dissemination, and outreach, which contains multiple synergistic components and objectives.

4.1 Design, Implementation, and Potential Impact of a New Voting Tool

Over the centuries, social choice theorists have proposed a plethora of intuitively appealing and theoretically grounded voting rules. Each of the prominent rules has its own advantages and disadvantages, which can be formalized. Although plurality is typically employed in political elections, sophisticated organizations often implement their own voting tools that leverage their favorite voting rule. For example, as mentioned in Section 3.1, thousands of developers participating in the Debian project (an open source software project responsible for the Debian Linux distribution) make binding general decisions on everything from project leaders to project logos using the elaborate Schulze method; the International Foundation for Autonomous Agents and Multiagent Systems (IFAAMAS) elects its board of directors using a voting rule called approval; and, somewhat ironically, the National Science Foundation SSS program is trying a new peer review method (for proposals submitted in October 2013) based on the Borda count rule [2].

A central component of this proposal is the design and implementation of a web-based voting tool that will be able to support the abovementioned small elections and many others. This system will allow

[10] More distantly related work on *prediction markets* has recently taken a first step toward dealing with "outside incentives" [43].

users to create elections by specifying alternatives and voters. Crucially, the system will also offer a list of voting rules, provide explanations for each one, and allow the election administrator to choose which rule to employ. Votes will be collected for a prespecified time frame, and ultimately a winner will be selected by aggregating the submitted votes via the chosen voting rule.

This system will serve three interconnected purposes. First, the system will enhance student education. In lectures on computational social choice, the tool will be used to perform more engaging demonstrations of different voting rules. Nowadays I collect votes only from a few students for demonstration purposes, whereas the system will allow the whole class to participate. The system will also enable demonstrations of the (sometimes surprising) consequences of increasing the number of voters or alternatives in the election. Outside of class, students are typically required to carry out small research projects, and I expect the system to provide a wonderful platform with which students can experiment.

Second, the system will be a resource for the wider scientific community. Currently there are few real-world datasets that contain rankings of alternatives. I will endeavor to make summary statistics of the data collected on the system freely available for academic research, subject to IRB approval, and given the express permission of users; permission will be sought when a new election is created (see also the data management plan). In addition, the system will be designed to painlessly integrate new voting rules; researchers will be able implement new voting rules and request to add them to the system. I expect that this feature will inform experimental research in social choice, leading to the development of new methods that work well in practice.

Third, the system will serve as a collective decision making tool for small groups and larger organizations. In addition to organization officers and admissions, as discussed above, applications can range from selecting award recipients and invited speakers to party menus and book club choices. I believe that a polished and user-friendly system can conceivably attract hundreds of thousands of users. The potential benefit is huge, in terms of exposing the public to the ideas of social choice, as well as collecting unprecedented amounts of real-world voting data.

A well-known existing system resembles that which I am proposing: the Condorcet Internet Voting System (CIVS) [3], maintained by Andrew Myers of Cornell's Computer Science Department. Since 2003, CIVS has elicited more than a hundred thousand votes, demonstrating the potential appeal of web-based voting systems. However, to realize the full potential, I believe that a significantly more streamlined user experience must be achieved. From the education and research viewpoints, CIVS employs a few similar voting rules rather than allowing users to select from a wide range of fundamentally different rules (a shortcoming that, in my view, limits its impact as a research and education tool); and the data collected via CIVS has not been made available to the research community. Smaller-scale web-based voting systems include OpenSTV, VoteFair, Simply Voting, and Whale.

4.2 Editing a Book on Computational Social Choice

I am currently working on the first book devoted to computational social choice, tentatively entitled "Handbook of Computational Social Choice", to be published by Cambridge University Press. This 600 page book was conceived as a collection of contributed chapters. I am editing the book together with four prominent colleagues: Felix Brandt (Technische Universität München, Germany), Vincent Conitzer (Duke University), Ulle Endriss (University of Amsterdam, The Netherlands), and Jérôme Lang (Université Paris-Dauphine, France). The invited chapter authors are in the process of preparing the first drafts of their chapters.

I expect that my work on the book will be closely intertwined with the other activities and goals of this proposal. The benefits in terms of education (see also Section 4.5) and dissemination are clear. In particular, in addition to serving as an essential tool for computer scientists who want to carry out research in computational social choice, one of our primary goals for the book is to draw in researchers in social choice theory and economics more broadly.

An observation that is perhaps less obvious is that the book will likely have a special role in promoting the research themes proposed above. Indeed, the current structure of the book reflects the research focus to date on bringing computational thinking to social choice, rather than exploring applications of social choice. In fact, one of the anonymous reviewers of our book proposal remarked that " ... the additional topic that would be good to promote is applications of computational social choice ... Of course there is the early work related to rank aggregation on the web [57] ... but I wonder whether the editors have in mind more recent papers that showcase some of the ways in which social choice theory can help to guide applications and real-world design?" Although the current plan is to discuss the relatively sparse work on applications in the book's introduction, I am hopeful that by the time the book is ready for publication the proposed research, together with related research thrusts by myself and others on applying computational social choice to human computation [120, 99] and small-scale elections [26], will have matured sufficiently to warrant a separate chapter devoted to applications of (computational) social choice.

4.3 Organizing a Workshop on Computational Social Choice

COMSOC is the International Workshop on Computational Social Choice, held biannually since 2006 (in Amsterdam, Liverpool, Düsseldorf, and Kraków). The workshop typically attracts close to a hundred attendees, and aims to bring together computer scientists, economists, and political scientists. The workshop's upcoming edition — COMSOC 2014 [4] — will take place on June 23-25, 2014, at Carnegie Mellon University. I am serving as the workshop's co-chair (together with Toby Walsh of NICTA and the University of New South Wales, Australia) and organizer.

Chairing COMSOC 2014 is a wonderful opportunity to set a new agenda for the computational social choice community, for example, through invited talks and tutorials. And while the 40–50 accepted papers are selected for presentation via a thorough peer-review process, I will be able to emphasize certain directions. I will prioritize novel ideas and paradigms over more traditional directions, thereby (gently) nudging the community toward forward-looking research with potential for novel applications. In particular, I hope to be in a position to accept a number of high-quality papers related to the interaction between AI and social choice, building enough critical mass to inspire attendees to contribute to this interaction.

Additionally, the center of gravity of the computational social choice community has so far been in Europe (as demonstrated by the venues of the previous four COMSOCs). By locating COMSOC 2014 at Carnegie Mellon University I plan to attract a new cohort of US-based researchers to the field.

4.4 Mentoring Students — and Diversity

The award will support a CMU Ph.D. student, who will be involved in all aspects of the proposed research. Section 3.1 in particular provides an exceptional starting point for a thesis, because the research challenges therein possess two highly desirable properties: building on firm existing foundations, and at the same time providing ample opportunity for ambitious research. Carnegie Mellon has strong CS undergraduates interested in doing research (for credit), and I plan to involve them in the proposed research as well. Once enough progress is made on their theoretical side, Research Challenges 2 and 7 will provide excellent undergraduate research opportunities. I also imagine that undergraduates will help with the implementation of the web-based voting system proposed in Section 4.1.

Over the last two years I have mentored a female CMU Ph.D. student, XXXX ; our joint work has led to two publications [28, 29]. This summer I am also working with a female undergraduate, YYYY . However, I am keen on stepping up my efforts to encourage diversity in the coming years. At CMU's School of Computer Science we have excellent mechanisms for doing so. First, Women@SCS [5] —established in 1999 — is a professional organization supporting (among other things) academic opportunities for women and minorities in computer science along with a rich program of outreach activities. In particular, Women@SCS is sponsoring a workshop for undergraduate women in computer science (OurCS), which will next take place (at CMU) in October 2013. The workshop brings together undergraduate women from across the nation and allows participants to explore the CS research experience by working hands-on with researchers in a team. In coordination with ZZZZ , the director of Women@SCS, I plan to contribute to similar workshops in the future as a faculty panelist and research team leader. I have also accepted an invitation to participate as a session leader in the Women@SCS TechNights [6] program for middle school girls in 2013–2014. Second, CMU is a partner in the recently launched, NSF-funded Institute for African-American Mentoring in Computing Sciences (iAAMCS), and this summer the Computer Science Department is hosting several African-American students. I plan to leverage this existing infrastructure, as well as my connections in historically black universities (such as Howard University), to recruit excellent African-American undergraduates through the research experience for undergraduates (REU) program.

4.5 Curriculum Development

In spring 2013 I have designed and taught a new course called "Algorithms, Games, and Networks" [7], jointly with SSSSS . The course included a very significant component — five lectures — devoted to computational social choice. I have also taught Carnegie Mellon University's undergraduate and graduate AI courses, which likewise included a prominent computational social choice component [8].

I believe that these computational social choice lectures can easily be developed into a full-blown course, which would be unique in North America. I envision a mix between the course given by UUUU at the University of Amsterdam [9], which puts an emphasis on computational aspects of social choice and the AI perspective, and TTTTT's course "Social Choice and Networks" at UC Berkeley [10] which, among other things, showcases some of the key results of the (noncomputational) social choice literature. As the course textbook I expect to be able to use the new Handbook of Computational Social Choice (see Section 4.2), or at least a draft.

4.6 Broad Dissemination in Academia

One of the objectives of this proposal is to reach and significantly influence not only AI researchers, but also the mathematicians, economists, and political scientists who comprise the social choice theory community. To achieve this goal, I plan to disseminate the results of my research using a four-pronged approach. First, via the new book, as discussed above (see Section 4.2).

Second, via summer schools and tutorials. I recently organized (together with RRRRR of Stanford) an NSF-sponsored summer school [11], which took place at CMU in August 2012. Roughly 160 graduate students applied to the summer school; we selected 64, and invited nine prominent speakers. The summer school included a mini-course on computational social choice. Building on the lessons I have learned from leading this initiative, I plan to be involved in additional summer schools in the coming years, as an organizer and speaker. In addition, several years ago I presented (together with VVVVV of Duke) a half-day tutorial about computational social choice at large in AAMAS 2010 and EC 2010. The proposed research will lead to the creation and presentation of more specific tutorials and mini-courses that concentrate on the interplay between social choice and AI.

Third, I am one of the regular contributors to our blog *Turing's Invisible Hand* [12], which focuses on computation and economics. New posts are typically read by thousands, and in particular (as far as I can tell) most of the active researchers in the community follow the blog. The blog provides a wonderful forum for discussing exciting research (see, e.g., [2] for a recent, relevant example).

Fourth, perhaps least surprisingly, I plan to publish the results of my research in major AI conferences and journals, as well as interdisciplinary conferences and workshops. I have been deeply involved with several interdisciplinary venues, including the ACM Conference on Electronic Commerce (EC), the International Workshop on Computational Social Choice (COMSOC, see Section 4.3), and the Conference on Theoretical Aspects of Rationality and Knowledge (TARK).

4.7 Outreach

The main thrust of my outreach activities is the publication of stories in popular science magazines. I have written articles for *Communications of the ACM* [119], *AI Magazine* [64], and *ACM XRDS* [118], but in the near future I expect to reach an even larger audience through a story in *Scientific American (SciAm)* about computational social choice. Following detailed discussions with — and an official invitation from — SciAm's technology editor, Michael Moyer, a first draft of the story was accepted for publication. Roughly one quarter of the current version of the story is devoted to my existing work on applying simple voting rules to reduce communication in multiagent systems [35], but recently I was actually asked to revise the story in a way that increases the focus on multiagent systems. SciAm's editorial process is typically extremely long (several years), so as the proposed research begins to bear fruit, I hope to leverage the SciAm story to channel my excitement about this line of work to the public through the media.

Our blog, *Turing's Invisible Hand* (see Section 4.6), also serves as an effective vehicle for outreach. Indeed, posts of general interest are occasionally picked up by websites such as Slashdot, and are viewed by tens of thousands. For example, in February 2012 (which was admittedly an especially busy month) we had 79,951 views. Finally, the voting system proposed above has huge outreach potential (see Section 4.1).

Results of Prior NSF Support

I am a PI on two NSF awards. The first is CCF-1215883, *ICES: Small: Computational Fair Division: From Cake Cutting to Cloud Computing*. (Total award amount: $390,000; start date: July 1, 2012; duration: 36 months.) This award supports my research on computational fair division, focusing on topics such as the allocation of divisible goods, and fair resource allocation in data centers. These topics are disjoint from the topics proposed above. The ICES award partially supports my graduate advisees, AAAA and BBBBB . Since the award's start date a year ago we have made substantial progress directly on the proposed challenges; this work has already led to the publication of several research papers in top-tier venues [84, 30, 82], as well as an accessible survey paper in the Communications of the ACM [119], and several working papers. Moreover, by partially funding my own summer effort and my graduate students the award has helped support many of my 2013 publications, which include 4 AAAI papers, 3 IJCAI papers, 4 EC papers, 2 papers in Games and Economic Behavior (the top game theory journal), as well as other conference and journal papers.

The second award is IIS-1212499, *Summer School on Algorithmic Economics* (Total award amount: $50,000; start date: May 1, 2012; duration: 6 months.) This award, together with an award of $30,000 from ARO, supported the Summer School on Algorithmic Economics [11], which was held at CMU on August 6–10, 2012 (see also Section 4.6). This award did not support me, my students, or my research.