
Diverse Randomized Agents Vote to Win

Albert Xin Jiang
Trinity University
xjiang@trinity.edu

Leandro Soriano Marcolino
USC
sorianom@usc.edu

Ariel D. Procaccia
CMU
arielpro@cs.cmu.edu

Tuomas Sandholm
CMU
sandholm@cs.cmu.edu

Nisarg Shah
CMU
nkshah@cs.cmu.edu

Milind Tambe
USC
tambe@usc.edu

Abstract

We investigate the power of voting among diverse, randomized software agents. With teams of computer Go agents in mind, we develop a novel theoretical model of two-stage noisy voting that builds on recent work in machine learning. This model allows us to reason about a collection of agents with different biases (determined by the first-stage noise models), which, furthermore, apply randomized algorithms to evaluate alternatives and produce votes (captured by the second-stage noise models). We analytically demonstrate that a uniform team, consisting of multiple instances of any single agent, must make a significant number of mistakes, whereas a diverse team converges to perfection as the number of agents grows. Our experiments, which pit teams of computer Go agents against strong agents, provide evidence for the effectiveness of voting when agents are diverse.

1 Introduction

Recent years have seen a surge of work at the intersection of social choice and machine learning. In particular, significant attention has been given to the learnability and applications of *noisy preference models* [16, 2, 1, 3, 24]. These models enhance our understanding of voters' behavior in elections, and provide a theoretical basis for reasoning about crowdsourcing systems that employ voting to aggregate opinions [24, 8]. In contrast, this paper presents an application of noisy preference models to the design of systems of *software agents*, emphasizing the importance of voting and diversity.

Our starting point is two very recent papers by Marcolino et al. [19, 20], which provide a new perspective on voting among multiple software agents. Their empirical results focus on Computer Go programs (see, e.g., [10]), which often use Monte Carlo tree search algorithms [7]. Taking the team formation point of view, Marcolino et al. establish that a team consisting of multiple (four to six) different computer Go programs that use plurality voting — each agent giving one point to a favorite alternative — 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 [12, 13, 6, 14], 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 to classical (deterministic) planning, not to mention that randomization is often used in many other AI applications.

Focusing on the computer Go application, we 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.

However, the model suggested by Marcolino et al. [19] is rather rudimentary: 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. Marcolino et al. [20] present a similar model, but study the effect of increasing the size of the action space (i.e., the board size in the Go domain). More importantly, Marcolino et al. [19, 20] — and other related work [6] — assume that each agent votes for a single alternative. In contrast, it is potentially possible to design agents that generate a ranking of multiple alternatives, calling for a principled way to harness this additional information.

1.1 Our Approach and Results

We introduce the following novel, abstract model of voting, and instantiate it using Computer Go. In each state, which corresponds to a board position in Go, there is a ground truth, which captures the true quality of different alternatives — feasible moves in Go. Heuristic agents have a noisy perception of the quality of alternatives. We model this using a *noise model* for each agent, which randomly maps the ground truth to a ranking of the alternatives, representing the agent’s *biased view* of their qualities. But if a single agent is presented with the same state twice, the agent may choose two different alternatives. This is because agents are assumed to be randomized. For example, as mentioned above, most computer Go programs, such as Fuego [10], rely on Monte Carlo Tree Search to randomly decide between different moves. We model this additional source of noise via a second noise model, which takes the biased ranking as input, and outputs the agent’s *vote* (another ranking of the alternatives). A *voting rule* is employed to select a single alternative (possibly randomly) by aggregating the agents’ votes. Our main theoretical result is the following theorem, which is, in a sense, an extension of the classic Condorcet Jury Theorem [9].

Theorem 2 (simplified and informal). (i) *Under extremely mild assumptions on the noise models and voting rule, a uniform team composed of copies of any single agent (even the “strongest” one with the most accurate noise models), for any number of agents and copies, is likely to vote for suboptimal alternatives in a significant fraction of states;* (ii) *Under mild assumptions on the noise models and voting rule, a diverse team composed of a large number of different agents is likely to vote for optimal alternatives in almost every state.*

We show that the assumptions in both parts of the theorem are indeed mild by proving that three well-known noise models — the Mallows- ϕ model [18], The Thurstone-Mosteller model [26, 21], and the Plackett-Luce model [17, 23] — satisfy the assumptions in both parts of the theorem. Moreover, the assumptions on the voting rule are satisfied by almost all prominent voting rules.

We also present experimental results in the Computer Go domain. As stated before, our key methodological contributions are a procedure for automatically generating diverse teams by using different parameterizations of a Go program, and a novel procedure for extracting rankings of moves from algorithms that are designed to output only a single good move. We show that the diverse team significantly outperforms the uniform team under the plurality rule. We also show that it is possible to achieve better performance by extracting rankings from agents using our novel methodology, and aggregating them via ranked voting rules.

2 Background

We use $[k]$ as shorthand for $\{1, \dots, k\}$. A vote is a total order (ranking) over the alternatives, usually denoted by σ . The set of rankings over a set of alternatives A is denoted by $\mathcal{L}(A)$. For a ranking σ , we use $\sigma(i)$ to denote the alternative in position i in σ , so, e.g., $\sigma(1)$ is the most preferred alternative in σ . We also use $\sigma([k])$ to denote $\{\sigma(1), \dots, \sigma(k)\}$. A collection of votes is called a *profile*, denoted by π . A *deterministic voting rule* outputs a winning alternative on each profile. For a *randomized voting rule* f (or simply a voting rule), the output $f(\pi)$ is a distribution over the alternatives. A

voting rule is *neutral* if relabeling the alternatives relabels the output accordingly; in other words, the output of the voting rule is independent of the labels of the alternatives. All prominent voting rules, when coupled with uniformly random tie breaking, are neutral.

Families of voting rules. Next, we define two families of voting rules. These families are quite wide, disjoint, and together they cover almost all prominent voting rules.

- *Condorcet consistency.* An alternative is called the Condorcet winner in a profile if it is preferred to every other alternative in a majority of the votes. Note that there can be at most one Condorcet winner. A voting rule is called *Condorcet consistent* if it outputs the Condorcet winner (with probability 1) whenever it exists. Many famous voting rules such as Kemeny’s rule, Copeland’s rule, Dodgson’s rule, the ranked pairs method, the maximin rule, and Schulze’s method are Condorcet consistent.
- *PD-c Rules [8].* This family is a generalization of positional scoring rules that include prominent voting rules such as plurality and Borda count. While the definition of Caragiannis et al. [8] outputs rankings, we naturally modify it to output winning alternatives. Let $T_\pi(k, a)$ denote the number of times alternative a appears among first k positions in profile π . Alternative a is said to *position-dominate* alternative b in π if $T_\pi(k, a) > T_\pi(k, b)$ for all $k \in [m - 1]$, where m is the number of alternatives in π . An alternative is called the *position-dominating winner* if it position-dominates every other alternative in a profile. It is easy to check that there can be at most one position-dominating winner. A voting rule is called *position-dominance consistent* (PD-c) if it outputs the position-dominating winner (with probability 1) whenever it exists. Caragiannis et al. [8] show that all positional scoring rules (including plurality and Borda count) and Bucklin’s rule are PD-c (as rules that output rankings). We show that this holds even when the rules output winning alternatives. This is presented as Proposition 1 in the online appendix (specifically, Appendix A).

Caragiannis et al. [8] showed that PD-c rules are disjoint from Condorcet consistent rules (actually, for rules that output rankings, they use a natural generalization of Condorcet consistent rules that they call PM-c rules). Their proof also establishes the disjointness of the two families for rules that output winning alternatives.

2.1 Noise Models

One view of computational social choice models the votes as noisy estimates of an unknown true order of the alternatives. These votes come from a distribution that is parametrized by some underlying ground truth. The ground truth can itself be the true order of alternatives, in which case we say that the noise model is of the *rank-to-rank* type. The ground truth can also be an objective true quality level for each alternative, which is more fine-grained than a true ranking of alternatives. In this case, we say that the noise model is of the *quality-to-rank* type. See [15] for examples of quality-to-rank models and how they are learned. Note that the output votes are rankings over alternatives in both cases. We denote the ground truth by θ . It defines a *true ranking* of the alternatives (even when the ground truth is a quality level for each alternative), which we denote by σ^* .

Formally, a noise model P is a set of distributions over rankings — the distribution corresponding to the ground truth θ is denoted by $P(\theta)$. The probability of sampling a ranking σ from $P(\theta)$ is denoted by $\Pr_P[\sigma; \theta]$.

Similarly to voting rules, a noise model is called *neutral* if relabeling the alternatives permutes the probabilities of various rankings accordingly. Formally, a noise model P is called neutral if $\Pr_P[\sigma; \theta] = \Pr_P[\tau\sigma; \tau\theta]$, for every permutation τ of the alternatives, every ranking σ , and every ground truth θ . Here, $\tau\sigma$ and $\tau\theta$ denote the result of applying τ on σ and θ , respectively.

Classic noise models. Below, we define three classical noise models:

- *The Mallows- ϕ model [18].* This is a rank-to-rank noise model, where the probability of a ranking decreases exponentially in its distance from the true ranking. Formally, the Mallows- ϕ model for m alternatives is defined as follows. For all rankings σ and σ^* ,

$$\Pr[\sigma; \sigma^*] = \frac{\phi^{d_{KT}(\sigma, \sigma^*)}}{Z_\phi^m}, \quad (1)$$

where d_{KT} is the Kendall-Tau distance that measures total pairwise disagreement between two rankings, and the normalization constant $Z_\phi^m = \prod_{k=1}^m \sum_{j=0}^{k-1} \phi^j$ is independent of σ^* .

- *The Thurstone-Mosteller (TM)* [26, 21] and the *Plackett-Luce (PL)* [17, 23] models. Both models are of the quality-to-rank type, and are special cases of a more general *random utility model* (see [2] for its use in social choice). In a random utility model, each alternative a has an associated true quality parameter θ_a and a distribution μ_a parametrized by θ_a . In each sample from the model, a noisy quality estimate $X_a \sim \mu_a(\theta_a)$ is obtained, and the ranking where the alternatives are sorted by their noisy qualities is returned.

For the Thurstone-Mosteller model, $\mu_a(\theta_a)$ is taken to be the normal distribution $\mathcal{N}(\theta_a, \nu^2)$ with mean θ_a , and variance ν^2 . Its PDF is

$$f(x) = \frac{1}{\sqrt{2\pi\nu^2}} e^{-\frac{(x-\theta_a)^2}{2\nu^2}}.$$

For the Plackett-Luce model, $\mu_a(\theta_a)$ is taken to be the Gumbel distribution $\mathcal{G}(\theta_a)$. Its PDF follows $f(x) = e^{-(x-\theta_a)-e^{-(x-\theta_a)}}$. The CDF of the Gumbel distribution $\mathcal{G}(\theta_a)$ is given by $F(x) = e^{-e^{-(x-\theta_a)}}$. Note that we do not include a variance parameter because this subset of Gumbel distributions is sufficient for our purposes.

The Plackett-Luce model has an alternative, more intuitive, formulation. Taking $\lambda_a = e^{\theta_a}$, the probability of obtaining a ranking is the probability of sequentially choosing its alternatives from the pool of remaining alternatives. Each time, an alternative is chosen among a pool proportional to its λ value. Hence, $\Pr[\sigma; \{\lambda_a\}] = \prod_{i=1}^m \frac{\lambda_{\sigma(i)}}{\sum_{j=i}^m \lambda_{\sigma(j)}}$, where m is the number of alternatives.

3 Theoretical Results

In this section, we present our theoretical results. But, first, we develop a novel model that will provide the backdrop for these results. Let $N = \{1, \dots, n\}$ be a set of agents. Let S be the set of states of the world, and let $|S| = t$. These states represent different scenarios in which the agents need to make decisions; in Go, these are board positions. Let μ denote a probability distribution over states in S , which represents how likely it is to encounter each state. Each state $s \in S$ has a set of alternatives A_s , which is the set of possible actions the agents can choose in state s . Let $|A_s| = m_s$ for each $s \in S$. We assume that the set of alternatives is fixed in each state. We will later see how our model and results can be adjusted for varying sets of alternatives. The ground truth in state $s \in S$ is denoted by θ_s , and the true ranking in state s is denoted by σ_s^* .

Votes of agents. The agents are presented with states sampled from μ . Their goal is to choose the true best alternative, $\sigma_s^*(1)$, in each state $s \in S$ (although we discuss why our results also hold when the goal is to maximize expected quality). The inability of the agents to do so arises from two different sources: the suboptimal heuristics encoded within the agents, and their inability to fully optimize according to their own heuristics — these are respectively modeled by two noise models P_i^1 and P_i^2 associated with each agent i .

The agents inevitably employ heuristics (in domains like Go) and therefore can only obtain a noisy evaluation of the quality of different alternatives, which is modeled by the noise model P_i^1 of agent i . The biased view of agent i for the true order of the alternatives in A_s , denoted σ_{is} , is modeled as a sample from the distribution $P_i^1(\sigma_s^*)$. Moreover, we assume that the agents' decision making is randomized. For example, top computer Go programs use Monte Carlo tree search algorithms [7]. We therefore assume that each agent i has another associated noise model P_i^2 such that the final ranking that the agent returns is a sample from $P_i^2(\sigma_{is})$. To summarize, agent i 's vote is obtained by first sampling its biased truth from P_i^1 , and then sampling its vote from P_i^2 . It is clear that the composition $P_i^2 \circ P_i^1$ plays a crucial role in this process.

Agent teams. Since the agents make errors in estimating the best alternative, it is natural to form a team of agents and aggregate their votes. We consider two team formation methods: a *uniform* team comprising of multiple copies of a single agent that share the same biased truths but have different final votes due to randomness; and a *diverse* team comprising of a single copy of each agent with different biased truths and different votes. We show that the diverse team outperforms the uniform team irrespective of the choice of the agent that is copied in the uniform team.

3.1 Restrictions on Noise Models

No team can perform well if the noise models P_i^1 and P_i^2 lose all useful information. Hence, we impose intuitive restrictions on the noise models; our restrictions are mild, as we demonstrate (Theorem 1) that the three classical noise models presented in Section 2.1 satisfy all our assumptions.

PM- α Noise Model For $\alpha > 0$, a neutral noise model P is called *pairwise majority preserving with strength α* (or PM- α) if for every ground truth θ (and the corresponding true ranking σ^*) and every $i < j$, we have

$$\Pr_{\sigma \sim P(\theta)}[\sigma^*(i) \succ_{\sigma} \sigma^*(j)] \geq \Pr_{\sigma \sim P(\theta)}[\sigma^*(j) \succ_{\sigma} \sigma^*(i)] + \alpha, \quad (2)$$

where \succ_{σ} is the preference relation of a ranking σ sampled from $P(\theta)$. Note that this definition applies to both quality-to-rank and rank-to-rank noise models. In other words, in PM- α noise models every pairwise comparison in the true ranking is preserved in a sample with probability at least α more than the probability of it not being preserved.

PD- α Noise Model For $\alpha > 0$, a neutral noise model is called *position-dominance preserving with strength α* (or PD- α) if for every ground truth θ (and the corresponding true ranking σ^*), every $i < j$, and every $k \in [m - 1]$ (where m is the number of alternatives),

$$\Pr_{\sigma \sim P(\theta)}[\sigma^*(i) \in \sigma([k])] \geq \Pr_{\sigma \sim P(\theta)}[\sigma^*(j) \in \sigma([k])] + \alpha. \quad (3)$$

That is, for every $k \in [m - 1]$, an alternative higher in the true ranking has probability higher by at least α of appearing among the first k positions in a vote than an alternative at a lower position in the true ranking.

Compositions of noise models with restrictions. As mentioned above, compositions of noise models play an important role in our work. The next lemma shows that our restrictions on noise models are preserved, in a sense, under composition; its proof appears in Appendix B.

Lemma 1. *For $\alpha_1, \alpha_2 > 0$, the composition of a PD- α_1 noise model with a PD- α_2 noise model is a PD- $(\alpha_1 \cdot \alpha_2)$ noise model.*

Unfortunately, a similar result does not hold for PM- α noise models; the composition of a PM- α_1 noise model and a PM- α_2 noise model may yield a noise model that is not PM- α for any $\alpha > 0$. In Appendix C, we give such an example. While this is slightly disappointing, we show that a stronger assumption on the first noise model in the composition suffices.

PPM- α Noise Model For $\alpha > 0$, a neutral noise model P is called *positional pairwise majority preserving* (or PPM- α) if for every ground truth θ (and the corresponding true ranking σ^*) and every $i < j$, the quantity

$$\Pr_{\sigma \sim P(\theta)}[\sigma(i') = \sigma^*(i) \wedge \sigma(j') = \sigma^*(j)] - \Pr_{\sigma \sim P(\theta)}[\sigma(j') = \sigma^*(i) \wedge \sigma(i') = \sigma^*(j)] \quad (4)$$

is non-negative for every $i' < j'$, and at least α for some $i' < j'$. That is, for $i' < j'$, the probability that $\sigma^*(i)$ and $\sigma^*(j)$ go to positions i' and j' respectively in a vote should be at least as high as the probability of them going to positions j' and i' respectively (and at least α greater for some i' and j'). Summing Equation (4) over all $i' < j'$ shows that every PPM- α noise model is also PM- α .

Lemma 2. *For $\alpha_1, \alpha_2 > 0$, if noise models P^1 and P^2 are PPM- α_1 and PM- α_2 , respectively, then their composition $P^2 \circ P^1$ is PM- $(\alpha_1 \cdot \alpha_2)$.*

The lemma's proof is relegated to Appendix D.

3.2 Team Formation and the Main Theoretical Result

Let us explain the process of generating votes for the uniform team and for the diverse team. Consider a state $s \in S$. For the uniform team consisting of k copies of agent i , the biased truth σ_{is} is drawn from $P_i^1(\theta_s)$, and is common to all the copies. Each copy j then individually draws a vote π_{is}^j from $P_i^2(\sigma_{is})$; we denote the collection of these votes by $\pi_{is}^k = (\pi_{is}^1, \dots, \pi_{is}^k)$. Under a voting rule f , let $X_{is}^k = \mathbb{I}[f(\pi_{is}^k) = \sigma_s^*(1)]$ be the indicator random variable denoting whether the uniform

team selects the best alternative, namely $\sigma_s^*(1)$. Finally, agent i is chosen to maximize the overall accuracy $\mathbb{E}[X_{is}^k]$, where the expectation is over the state s and the draws from P_i^1 and P_i^2 .

The diverse team consists of one copy of each agent $i \in N$. Importantly, although we can take multiple copies of each agent and a total of k copies, we show that taking even a single copy of each agent outperforms the uniform team. Each agent i has its own biased truth σ_{is} drawn from $P_i^1(\theta_s)$, and it draws its vote ψ_{is} from $P_i^2(\sigma_{is})$. This results in the profile $\psi_s^n = (\psi_{1s}, \dots, \psi_{ns})$. Let $Y_s^n = \mathbb{I}[f(\psi_s^n) = \sigma_s^*(1)]$ be the indicator random variable denoting whether the diverse team selects the best alternative, namely $\sigma_s^*(1)$.

Below we put forward a number of assumptions on noise models; different subsets of assumptions are required for different results. We remark that each agent $i \in N$ has two noise models *for each possible number of alternatives* m . However, for the sake of notational convenience, we refer to these noise models as P_i^1 and P_i^2 irrespective of m . This is natural, as the classic noise models defined in Section 2.1 describe a noise model for each m .

A1 For each agent $i \in N$, the associated noise models P_i^1 and P_i^2 are neutral.

A2 There exists a universal constant $\eta > 0$ such that for each agent $i \in N$, every possible ground truth θ (and the corresponding true ranking σ^*), and every $k \in [m]$ (where m is the number of alternatives), $\Pr_{\sigma \sim P_i^1(\theta)}[\sigma^*(1) = \sigma(k)] \leq 1 - \eta$.

In words, assumption A2 requires that the true best alternative appear in any particular position with probability at most a constant which is less than 1. This ensures that the noise model indeed introduces a non-zero constant amount of noise in the position of the true best alternative.

A3 There exists a universal constant $\alpha > 0$ such that for each agent $i \in N$, the noise models P_i^1 and P_i^2 are PD- α .

A4 There exists a universal constant $\alpha > 0$ such that for each agent $i \in N$, the noise models P_i^1 and P_i^2 are PPM- α and PM- α , respectively.

We show that the preceding assumptions are indeed very mild in that classical noise models introduced in Section 2.1 satisfy all four assumptions. The proof of the following result appears in Appendix E.

Theorem 1. *With a fixed set of alternatives (such that the true qualities of every two alternatives are distinct in the case where the ground truth is the set of true qualities), the Mallows- ϕ model with $\phi \in [\rho, 1 - \rho]$, the Thurstone-Mosteller model with variance parameter $\sigma^2 \in [L, U]$, and the Plackett-Luce model all satisfy assumption A1, A2, A3, and A4, given that $\rho \in (0, 1/2)$, $L > 0$, and $U > L$ are constants.*

We are now ready to present our main result; its proof appears in Appendix F.

Theorem 2. *Let μ be a distribution over the state space S . Let the set of alternatives in all states $\{A_s\}_{s \in S}$ be fixed.*

1. *Under the assumptions A1 and A2, and for any neutral voting rule f , there exists a universal constant $c > 0$ such that for every k and every $N = \{1, \dots, n\}$, it holds that $\max_{i \in N} \mathbb{E}[X_{is}^k] \leq 1 - c$, where the expectation is over the state $s \sim \mu$, the ground truths $\sigma_{is} \sim P_i^1(\theta_s)$ for all $s \in S$, and the votes $\pi_{is}^j \sim P_i^2(\sigma_{is})$ for all $j \in [k]$.*
2. *Under each of the following two conditions, for a voting rule f , it holds that $\lim_{n \rightarrow \infty} \mathbb{E}[Y_s^n] = 1$, where the expectation is over the state $s \sim \mu$, the biased truths $\sigma_{is} \sim P_i^1(\theta_s)$ for all $i \in N$ and $s \in S$, and the votes $\psi_{is} \sim P_i^2(\sigma_{is})$ for all $i \in N$ and $s \in S$: (i) assumptions A1 and A3 hold, and f is PD- c ; (ii) assumptions A1 and A4 hold, and f is Condorcet consistent.*

4 Experimental Results

We now present our experimental results in the Computer Go domain. We use a novel methodology for generating large teams, which we view as one of our main contributions. It is fundamentally

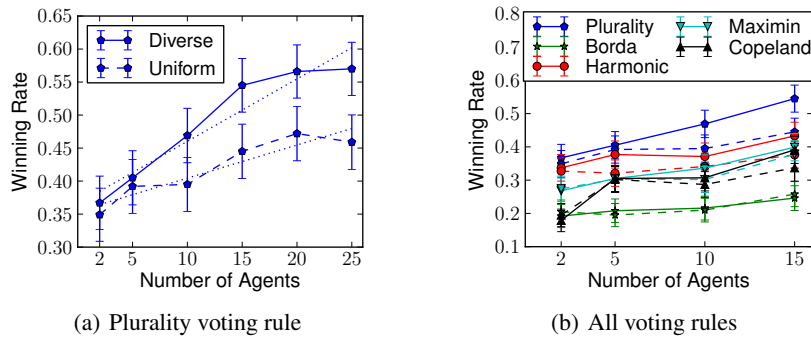


Figure 1: Winning rates for Diverse (continuous line) and Uniform (dashed line), for a variety of team sizes and voting rules.

different from that of Marcolino et al. [19, 20], who created a diverse team by combining four different, independently developed Go programs. Here we automatically create arbitrarily many diverse agents by parameterizing one Go program. Specifically, we use different parametrizations of Fuego 1.1 [10]. Fuego is a state-of-the-art, open source, publicly available Go program; it won first place in 19×19 Go in the Fourth Computer Go UEC Cup, 2010, and also won first place in 9×9 Go in the 14th Computer Olympiad, 2009. We sample random values for a set of parameters for each generated agent, in order to change its behavior. In Appendix G we list the sampled parameters, and the range of sampled values. The original Fuego is the strongest agent, as we show in Appendix H.

All results were obtained by simulating 1000 9×9 Go games, in an HP dl165 with dual dodeca core, 2.33GHz processors and 48GB of RAM. We compare the winning rates of games played against a fixed opponent. In all games the system under evaluation plays as white, against the original Fuego playing as black. We evaluate two types of teams: *Diverse* is composed of different agents, and *Uniform* is composed of copies of a specific agent (with different random seeds). In order to study the performance of the uniform team, for each sample (which is an entire Go game) we construct a team consisting of copies of a randomly chosen agent from the diverse team. Hence, the results presented for Uniform are approximately the mean behavior of all possible uniform teams, given the set of agents in the diverse team. In all graphs, the error bars show 99% confidence intervals.

Fuego (and, in general, all programs using Monte Carlo tree search algorithms) is not originally designed to output a ranking over all possible moves (alternatives), but rather to output a single move — the best one according to its search tree (of course, there is no guarantee that the selected move is in fact the best one). In this paper, however, we wish to compare plurality (which only requires each agent’s top choice) with voting rules that require an entire ranking from each agent. Hence, we modified Fuego to make it output a ranking over moves, by using the data available in its search tree (we rank by the number of simulations per alternative). We ran games under 5 different voting rules: plurality, Borda count, the harmonic rule, maximin, and Copeland. Plurality, Borda count (which we limit to the top 6 positions in the rankings), and the harmonic rule (see Appendix A) are PD-c rules, while maximin and Copeland are Condorcet-consistent rules (see, e.g., [24]).

We first discuss Figure 1(a), which shows the winning rates of Diverse and Uniform for a varying number of agents using the plurality voting rule. The winning rates of both teams increase as the number of agents increases. Diverse and Uniform start with similar winning rates, around 35% with 2 agents and 40% with 5 agents, but with 25 agents Diverse reaches 57%, while Uniform only reaches 45.9%. The improvement of Diverse over Uniform is not statistically significant with 5 agents ($p = 0.5836$), but is highly statistically significant with 25 agents ($p = 8.592 \times 10^{-7}$). We perform linear regression on the winning rates of the two teams to compare their rates of improvement in performance as the number of agents increases. Linear regression (shown as the dotted lines in Figure 1(a)) gives the function $y = 0.0094x + 0.3656$ for Diverse ($R^2 = 0.9206$, $p = 0.0024$) and $y = 0.0050x + 0.3542$ for Uniform ($R^2 = 0.8712$, $p = 0.0065$). In particular, the linear approximation for the winning rate of Diverse increases roughly twice as fast as the one for Uniform as the number of agents increases.

Despite the strong performance of Diverse (it beats the original Fuego more than 50% of the time), it seems surprising that its winning rate converges to a constant that is significantly smaller than 1, in light of Theorem 2. There are (at least) two reasons for this apparent discrepancy. First, Theorem 2 deals with the probability of making good moves in individual board positions (states), whereas the figure shows winning rates. Even if the former probability is very high, a bad decision in a single state of a game can cost Diverse the entire game. Second, our diverse team is formed by randomly sampling different parametrizations of Fuego. Hence, there might still exist a subset of world states where all agents would play badly, regardless of the parametrization. In other words, the parametrization procedure may not be generating the idealized diverse team (see Appendix H).

Figure 1(b) compares the results across different voting rules. As mentioned above, to generate ranked votes, we use the internal data in the search tree of an agent’s run (in particular, we rank using the number of simulations per alternative). We can see that increasing the number of agents has a positive impact for all voting rules under consideration. Moving from 5 to 15 agents for Diverse, plurality has a 14% increase in the winning rate, whereas other voting rules have a mean increase of only 6.85% ($std = 2.25\%$), close to half the improvement of plurality. For Uniform, the impact of increasing the number of agents is much smaller: Moving from 5 to 15 agents, the increase for plurality is 5.3%, while the mean increase for other voting rules is 5.70% ($std = 1.45\%$). Plurality surprisingly seems to be the best voting rule in these experiments, even though it uses less information from the submitted rankings. This suggests that the ranking method used does not typically place good alternatives in high positions other than the very top.

Hence, we introduce a novel procedure to generate rankings, which we view as another major methodological contribution. To generate a ranked vote from an agent on a given board state, we run the agent on the board state 10 times (each run is independent of other runs), and rank the moves by the number of times they are played by the agent. We use these votes to compare plurality with the four other voting rules, for Diverse with 5 agents. Figure 2 shows the results. All voting rules outperform plurality; Borda and maximin are statistically significantly better ($p < 0.007$ and $p = 0.06$, respectively). All ranked voting rules are also statistically significantly better than the non-sampled (single run) version of plurality.




Figure 2: All voting rules, for Diverse with 5 agents, using the new ranking methodology.

5 Discussion

While we have focused on computer Go for motivation, we have argued in Section 1 that our theoretical model is more widely applicable. At the very least, it is relevant to modeling game-playing agents in the context of other games. For example, random sampling techniques play a key role in the design of computer poker programs [25]. A complication in some poker games is that the space of possible moves, in some stages of the game, is infinite, but this issue can likely be circumvented via an appropriate discretization.

Our theoretical model does have (at least) one major shortcoming when applied to multistage games like Go or poker: it assumes that the state space is “flat”. So, for example, making an excellent move in one state is useless if the agent makes a horrible move in a subsequent state. Moreover, rather than having a fixed probability distribution μ over states, the agents’ strategies actually determine which states are more likely to be reached. To the best of our knowledge, existing models of voting do not capture sequential decision making — possibly with a few exceptions that are not relevant to our setting, such as the work of Parkes and Procaccia [22]. From a theoretical and conceptual viewpoint, the main open challenge is to extend our model to explicitly deal with sequentiality.

Acknowledgments: Procaccia and Shah were partially supported by the NSF under grants IIS-1350598 and CCF-1215883, and Marcolino by MURI grant W911NF-11-1-0332.

References

- [1] H. Azari Soufiani, W. Z. Chen, D. C. Parkes, and L. Xia. Generalized method-of-moments for rank aggregation. In *Proc. of 27th NIPS*, pages 2706–2714, 2013.
- [2] H. Azari Soufiani, D. C. Parkes, and L. Xia. Random utility theory for social choice. In *Proc. of 26th NIPS*, pages 126–134, 2012.
- [3] H. Azari Soufiani, D. C. Parkes, and L. Xia. Computing parametric ranking models via rank-breaking. In *Proc. of 31st ICML*, 2014. Forthcoming.
- [4] P. Baudiš and J. I. Gailly. PACHI: State of the art open source go program. In *Proc. of 13th ACG*, pages 24–38, 2011.
- [5] C. Boutilier, I. Caragiannis, S. Haber, T. Lu, A. D. Procaccia, and O. Sheffet. Optimal social choice functions: A utilitarian view. In *Proc. of 13th EC*, pages 197–214, 2012.
- [6] Y. Braouezec. Committee, expert advice, and the weighted majority algorithm: An application to the pricing decision of a monopolist. *Computational Economics*, 35(3):245–267, 2010.
- [7] C. Browne, E. J. Powley, D. Whitehouse, S. M. Lucas, P. I. Cowling, P. Rohlfshagen, S. Tavener, D. Perez, S. Samothrakis, and S. Colton. A survey of Monte Carlo tree search methods. *IEEE Transactions on Computational Intelligence and AI in Games*, 4(1):1–43, 2012.
- [8] I. Caragiannis, A. D. Procaccia, and N. Shah. When do noisy votes reveal the truth? In *Proc. of 14th EC*, pages 143–160, 2013.
- [9] M. de Condorcet. *Essai sur l’application de l’analyse à la probabilité de décisions rendues à la pluralité de voix*. Imprimerie Royal, 1785. Facsimile published in 1972 by Chelsea Publishing Company, New York.
- [10] M. Enzenberger, M. Müller, B. Arneson, and R. Segal. Fuego — an open-source framework for board games and Go engine based on Monte Carlo tree search. *IEEE Transactions on Computational Intelligence and AI in Games*, 2(4):259–270, 2010.
- [11] E. Hellinger. Neue begründung der theorie quadratischer formen von unendlichvielen veränderlichen. *Journal für die reine und angewandte Mathematik*, 136:210–271, 1909. In German.
- [12] L. Hong and S. E. Page. Groups of diverse problem solvers can outperform groups of high ability problem solvers. *Proceedings of the National Academy of Sciences of the United States of America*, 101(46):16385–16389, 2004.
- [13] L. Hong and S. E. Page. Some microfoundations of collective wisdom. In H. Landemore and J. Elster, editors, *Collective Wisdom*, pages 56–71. Cambridge University Press, 2009.
- [14] M. LiCalzi and O. Surucu. The power of diversity over large solution spaces. *Management Science*, 58(7):1408–1421, 2012.
- [15] T.-Y. Liu. *Learning to Rank for Information Retrieval*. Springer, 2011.
- [16] T. Lu and C. Boutilier. Learning Mallows models with pairwise preferences. In *Proc. of 28th ICML*, pages 145–152, 2011.
- [17] R. D. Luce. *Individual choice behavior: A theoretical analysis*. Wiley, 1959.
- [18] C. L. Mallows. Non-null ranking models. *Biometrika*, 44:114–130, 1957.
- [19] L. S. Marcolino, A. X. Jiang, and M. Tambe. Multi-agent team formation — diversity beats strength? In *Proc. of 23rd IJCAI*, pages 279–285, 2013.
- [20] L. S. Marcolino, H. Xu, A. X. Jiang, M. Tambe, and E. Bowring. Give a hard problem to a diverse team: Exploring large action spaces. In *Proc. of 28th AAAI*, 2014.
- [21] F. Mosteller. Remarks on the method of paired comparisons: I. the least squares solution assuming equal standard deviations and equal correlations. *Psychometrika*, 16(1):3–9, 1951.
- [22] D. C. Parkes and A. D. Procaccia. Dynamic social choice with evolving preferences. In *Proc. of 27th AAAI*, pages 767–773, 2013.
- [23] R. Plackett. The analysis of permutations. *Applied Statistics*, 24:193–202, 1975.
- [24] A. D. Procaccia, S. J. Reddi, and N. Shah. A maximum likelihood approach for selecting sets of alternatives. In *Proc. of 28th UAI*, pages 695–704, 2012.
- [25] T. Sandholm. The state of solving large incomplete-information games, and application to Poker. *AI Magazine*, 31(4):13–32, 2010.
- [26] L. L. Thurstone. A law of comparative judgement. *Psychological Review*, 34:273–286, 1927.

A Voting Rules

We first define positional scoring rules and Bucklin's rule, and then prove Proposition 1.

Positional Scoring Rules A positional scoring rule is given by a scoring vector $(\alpha_1, \dots, \alpha_m)$ where $\alpha_i \geq \alpha_{i+1}$ for all $i \in \{1, \dots, m\}$ and $\alpha_1 > \alpha_m$. Under this rule for each vote σ and $i \in [m]$, α_i points are awarded to the alternative $\sigma(i)$. The alternative with the most points overall is selected as the winner. The proof of Proposition 1 holds irrespective of the tie-breaking rule used. Special positional scoring rules include *plurality* with scoring vector $(1, 0, 0, \dots, 0)$, *Borda count* with scoring vector $(m, m-1, \dots, 1)$, the *veto* rule with scoring vector $(1, 1, \dots, 1, 0)$, and the *harmonic* rule [5] with scoring vector $(1, 1/2, \dots, 1/m)$.

Bucklin's rule The Bucklin score of an alternative a is the minimum k such that a is among the first k positions in the majority of input votes. Bucklin's rule outputs the alternative with the lowest Bucklin score, and breaks ties among alternatives with the same Bucklin score ℓ according to the number of rankings that have the alternative in the first ℓ positions.

Proposition 1. *All positional scoring rules (including plurality and Borda count) and Bucklin's rule are PD-c rules.*

Proof. Consider a profile π with n rankings and a position-dominating winner a . We show that any positional scoring rule as well as Bucklin's rule outputs a on π . For any $j \in \{1, \dots, m-1\}$, let $T_\pi(c, j)$ denote the number of votes where alternative c is among the first j positions in π .

For Bucklin's rule, consider arbitrary alternative $a' \neq a$. Let k denote the Bucklin score of a and k' denote the Bucklin score of a' . If $k > k'$, then $T_\pi(a', k') > n/2$ and $T_\pi(a, k') \leq T_\pi(a, k-1) < n/2$, which is impossible since the a is the position-dominating winner in π . If $k < k'$, then Bucklin's rule would select a over a' , as required.

If $k = k' \neq m$, then we have $T_\pi(a, k) > T_\pi(a', k)$ because a is the position-dominating winner. Hence, the tie is broken in favor of a . Lastly, we note that $k = k' = m$ is not possible because it would imply that the total number of appearances of a and a' in the last position is $n - T_\pi(a, m-1) + n - T_\pi(a', m-1) > 2 \cdot n - 2 \cdot T_\pi(a, m-1) > n$. Thus, Bucklin's rule would choose a over every other alternative a' , i.e., it would output a as the winner, as required.

Consider a positional scoring rule with scoring vector $(\alpha_1, \dots, \alpha_m)$. As shown in the proof of Theorem 3.10 in [8], the score of an alternative a' in π is equivalently given by $\sum_{k=1}^{m-1} \beta_k \cdot T_\pi(a', k)$, where $\beta_i = \alpha_i - \alpha_{i-1} \geq 0$. It is now easy to see that the position-dominating winner a would have strictly higher score than every other alternative because $\beta_i > 0$ for some i . Hence, every positional scoring rule would also output a , as required. \square

B Proof of Lemma 1

Let m be the number of alternatives. The result is trivial for $m = 1$. For $m \geq 2$, let P^1 and P^2 be PD- α_1 and PD- α_2 noise models, respectively. For $i, j \in [m]$, let $T^1(i, j)$ and $T^2(i, j)$ be the probabilities that the i^{th} alternative in the true ranking is placed in position j in a sample from P^1 and P^2 , respectively.¹ For $i, j \in [m]$, let $F^1(i, j) = \sum_{l=1}^j T^1(i, l)$ and $F^2(i, j) = \sum_{l=1}^j T^2(i, l)$. Now, fix $1 \leq p < q \leq m$. The difference between the probabilities of the p^{th} and q^{th} alternatives in the true ranking appearing among the first k positions in a sample from $P^2 \circ P^1$ is

$$\begin{aligned} & \sum_{i=1}^k \sum_{j=1}^m T^1(p, j) T^2(j, i) - \sum_{i=1}^k \sum_{j=1}^m T^1(q, j) T^2(j, i) \\ &= \sum_{i=1}^k \sum_{j=1}^m T^2(j, i) (T^1(p, j) - T^1(q, j)) = \sum_{j=1}^m F^2(j, k) (T^1(p, j) - T^1(q, j)) \\ &= \sum_{j=1}^m (F^2(j, k) - F^2(j+1, k)) (F^1(p, j) - F^1(q, j)) \geq \alpha_1 \cdot \alpha_2, \end{aligned}$$

¹Note that due to neutrality of P^1 and P^2 , these probabilities are independent of the true ranking.

where the second transition follows by interchanging the order of summation, the third transition follows by simple algebra (we let $F^2(m+1, k) = 0$ because there are only m alternatives), and the last transition holds because for $j = 1$, the two terms in the summation are at least α_2 and α_1 , respectively, as P^1 and P^2 are PD- α_1 and PD- α_2 noise models, respectively, and for $j > 1$, both terms are non-negative. \square

C Example: Composition of PM Noise Models

Consider a neutral noise model P over 3 alternatives that is of the rank-to-rank type. We describe the probabilities of various rankings when the ground truth is $a \succ b \succ c$ (the probabilities of various permutations of the true ranking are independent of the true ranking): $\Pr_P[a \succ b \succ c; a \succ b \succ c] = 0.51$, $\Pr_P[b \succ a \succ c; a \succ b \succ c] = 0.09$, $\Pr_P[c \succ b \succ a; a \succ b \succ c] = 0.4$. That is, the true ranking stays unchanged with probability 0.51, the top two alternative are swapped with probability 0.09, and the first and the last alternatives are swapped with probability 0.4.

Clearly, P is PM-0.02 because every pairwise comparison is preserved with probability at least 0.51. Let us consider the composition $P \circ P$. We evaluate the probability that the top alternative in the true ranking stays above the second alternative in the true ranking in a vote sampled from $P \circ P$. This probability is precisely

$$0.51 \cdot 0.51 + 0.09 \cdot 0.49 + 0.4 \cdot 0.4 = 0.4642 < 0.5.$$

Hence, in $P \circ P$ the pair of top two alternatives in the true ranking is flipped with probability more than 0.5. It follows that $P \circ P$ is not PM- α for any $\alpha > 0$.

D Proof of Lemma 2

The result is trivial for $m = 1$. Let $m \geq 1$. Let P^1 and P^2 be PPM- α_1 and PM- α_2 noise models respectively. Let m denote the number of alternatives. Fix $1 \leq i < j \leq m$. For $1 \leq i' < j' \leq m$, define $T^1(i, j, i', j')$ to be the probability that the alternatives in positions i and j in the true ranking appear in positions i' and j' , respectively, in a vote sampled from P^1 . Since P^1 is PPM- α_1 , we have $T^1(i, j, i', j') \geq T^1(i, j, j', i')$ for all $1 \leq i' < j' \leq m$ and $T^1(i, j, i', j') \geq T^1(i, j, j', i') + \alpha_1$ for some $1 \leq i' < j' \leq m$.

Let $T^2(i, j)$ denote the probability that the alternative in position i in the true ranking is preferred to the alternative in position j in the true ranking in a vote sampled from P^2 . Since P^2 is PM- α_2 , we know that $T^2(i, j) \geq T^2(j, i) + \alpha_2$. Now, the difference between the probability of the alternative in position i in the true ranking being preferred to the alternative in position j in the true ranking in a vote from $P^2 \circ P^1$ and the probability of its converse is

$$\begin{aligned} & \sum_{i', j' \in [m]} T^1(i, j, i', j') \cdot T^2(i', j') - \sum_{i', j' \in [m]} T^1(i, j, i', j') \cdot T^2(j', i') \\ &= \sum_{i', j' \in [m]} T^1(i, j, i', j') \cdot (T^2(i', j') - T^2(j', i')) \\ &= \sum_{1 \leq i' < j' \leq m} (T^1(i, j, i', j') \cdot (T^2(i', j') - T^2(j', i')) + T^1(i, j, j', i') \cdot (T^2(j', i') - T^2(i', j'))) \\ &= \sum_{1 \leq i' < j' \leq m} (T^1(i, j, i', j') - T^1(i, j, j', i')) \cdot (T^2(i', j') - T^2(j', i')) \geq \alpha_1 \cdot \alpha_2, \end{aligned}$$

where the last transition holds because our assumptions on P^1 and P^2 imply that there exist $1 \leq i' < j' \leq m$ for which $T^1(i, j, i', j') \geq T^1(i, j, j', i') + \alpha_1$, and for those values of i' and j' , we have $T^2(i', j') \geq T^2(j', i') + \alpha_2$. Thus, $P^2 \circ P^1$ is PM- $(\alpha_1 \cdot \alpha_2)$. \square

E Proof of Theorem 1

We show that three classical noise models—the Mallows- ϕ model, the Thurstone-Mosteller model, and the Plackett-Luce model—satisfy our four assumptions. For assumption A4, we only show that the noise models are PPM- α ; this implies that they are also PM- α .

Proof of Theorem 1. Let m denote the number of alternatives. Let $\mathcal{L}(A)$ be the set of all rankings over m alternatives. In all the proofs below, P will denote the noise model under consideration, and $p_{i,j}$ will denote the probability that the alternative in position i in the true ranking appears in position j in a vote sampled from P . We begin with the proof for the Mallows- ϕ model.

Part I: The Mallows- ϕ model.

We prove that the Mallows- ϕ model with $\phi \in [\rho, 1 - \rho]$ and constant $\rho \in (0, 1/2)$ satisfies assumptions A1, A2, A3, and A4.

Assumption A1: It is well-known and easy to check that neutrality of the Kendall-Tau distance implies neutrality of the Mallows- ϕ model for all $\phi \in (0, 1)$. Hence, the Mallows- ϕ model satisfies assumption A1.

Assumption A2: We need to show that there exists a constant $\eta > 0$ such that $p_{1,k} \leq 1 - \eta$ for all $k \in [m]$. Lemma 3.8 in [8] shows that when σ is sampled from the Mallows- ϕ model with true ranking σ^* and m alternatives,

$$p_{i,1} = \frac{\phi^{i-1}}{\sum_{j=0}^{m-1} \phi^j}.$$

The proof explicitly evaluates the probability by summing the probabilities of all rankings where $\sigma(1) = \sigma^*(i)$. Using an almost identical proof technique, we evaluate the similar probability $p_{1,i}$ used in assumption A2. First, we show that for any $i \in [m - 1]$, we have $p_{1,i+1} = \phi \cdot p_{1,i}$. To see this,

$$\begin{aligned} p_{1,i} - p_{1,i+1} &= \frac{\sum_{\sigma \in \mathcal{L}(A) | \sigma(i) = \sigma^*(1)} \phi^{d_{KT}(\sigma, \sigma^*)} - \sum_{\sigma \in \mathcal{L}(A) | \sigma(i+1) = \sigma^*(1)} \phi^{d_{KT}(\sigma, \sigma^*)}}{Z_\phi^m} \\ &= \frac{\sum_{\sigma \in \mathcal{L}(A) | \sigma(i) = \sigma^*(1)} (\phi^{d_{KT}(\sigma, \sigma^*)} - \phi^{d_{KT}(\sigma_{\sigma(i) \leftrightarrow \sigma(i+1)}, \sigma^*)})}{Z_\phi^m} \\ &= \sum_{\sigma \in \mathcal{L}(A) | \sigma(i) = \sigma^*(1)} \frac{\phi^{d_{KT}(\sigma, \sigma^*)} \cdot (1 - \phi)}{Z_\phi^m} = (1 - \phi) \cdot p_{1,i}, \end{aligned}$$

where the second transition holds because $\sigma \leftrightarrow \sigma_{\sigma(i) \leftrightarrow \sigma(i+1)}$ is a bijection between the set of rankings where $\sigma^*(1) = \sigma(i)$ and the set of rankings where $\sigma^*(1) = \sigma(i+1)$. The third transition holds because swapping $\sigma(i) = \sigma^*(1)$ with the alternative $\sigma(i+1)$ does not change any pairwise comparisons between σ and σ^* , except that of $\sigma^*(1)$ and $\sigma(i+1)$. The latter is mismatched with σ^* after the exchange. Hence, the Kendall Tau distance to σ^* , which is equal to the number of pairwise mismatches with σ^* , increases by exactly 1 after the exchange.

Hence, $p_{1,i} - p_{1,i+1} = (1 - \phi) \cdot p_{1,i}$, which implies that $p_{1,i+1} = \phi \cdot p_{1,i}$. Applying this repeatedly, we have that $p_{1,i} = p_{1,1} \cdot \phi^{i-1}$, for every $i \in [m]$. Summing over $i \in [m]$ and observing that $\sum_{i=1}^m p_{1,i} = 1$, we get that

$$p_{1,i} = \frac{\phi^{i-1}}{\sum_{j=0}^{m-1} \phi^j} \leq \frac{\phi^{i-1}}{\sum_{j=0}^{\infty} \phi^j} = \phi^{i-1} \cdot (1 - \phi).$$

Hence, for all $i \in [m]$, $p_{1,i} \leq p_{1,1} \leq 1 - \phi \leq 1 - \rho$. Hence, the Mallows- ϕ model satisfies assumption A2 with $\eta = \rho$.

Assumption A3: We need to show that for all $i, j \in [m]$ with $i < j$ and $k \in [m - 1]$,

$$\Pr_P[\sigma^*(i) \in \sigma([k])] > \Pr_P[\sigma^*(j) \in \sigma([k])]. \quad (5)$$

We take the difference of the two terms, and remove the set of common rankings where both $\sigma^*(i)$ and $\sigma^*(j)$ are in $\sigma([k])$. Thus, we get

$$\begin{aligned}
\Pr_P[\sigma^*(i) \in \sigma([k])] - \Pr_P[\sigma^*(j) \in \sigma([k])] &= \sum_{\substack{\sigma \in \mathcal{L}(A) \\ \sigma^*(i) \in \sigma([k]) \\ \sigma^*(j) \notin \sigma([k])}} \Pr_P[\sigma; \sigma^*] - \sum_{\substack{\sigma \in \mathcal{L}(A) \\ \sigma^*(j) \in \sigma([k]) \\ \sigma^*(i) \notin \sigma([k])}} \Pr_P[\sigma; \sigma^*] \\
&= \sum_{\substack{\sigma \in \mathcal{L}(A) \\ \sigma^*(i) \in \sigma([k]) \\ \sigma^*(j) \notin \sigma([k])}} (\Pr_P[\sigma; \sigma^*] - \Pr_P[\sigma_{\sigma^*(i) \leftrightarrow \sigma^*(j)}; \sigma^*]) \quad (6) \\
&\geq (1 - \phi) \cdot \sum_{\substack{\sigma \in \mathcal{L}(A) \\ \sigma^*(i) \in \sigma([k]) \\ \sigma^*(j) \notin \sigma([k])}} \Pr_P[\sigma; \sigma^*] \\
&\geq (1 - \phi) \cdot \frac{\phi^{-m^2}}{Z_\phi^m} \geq \rho^{1-m-m^2},
\end{aligned}$$

where the second transition holds because $\sigma \leftrightarrow \sigma_{\sigma^*(i) \leftrightarrow \sigma^*(j)}$ is a bijection between the set of rankings where $\sigma^*(i) \in \sigma([k])$ and $\sigma^*(j) \notin \sigma([k])$, and the set of rankings where $\sigma^*(i) \notin \sigma([k])$ and $\sigma^*(j) \in \sigma([k])$. The third transition holds because σ matches with σ^* in the pairwise comparison of $\sigma^*(i)$ and $\sigma^*(j)$ (because $\sigma^*(i) \in \sigma([k])$ and $\sigma^*(j) \notin \sigma([k])$), thus swapping them increases its distance from σ^* by at least 1 due to the swap-increasing property of the Kendall-Tau distance (Lemma 3.5 in [8]). Hence, the probability drops at least by a factor of ϕ from Equation (1). The fourth transition holds because there is at least one ranking where $\sigma^*(i) \in \sigma([k])$ and $\sigma^*(j) \notin \sigma([k])$, and the ranking has probability at least ϕ^{-m^2}/Z_ϕ^m . Finally, the last transition holds because $1/Z_\phi^m \geq (1 - \phi)^m$ (which is easy to show), and $\phi \in [\rho, 1 - \rho]$.

Hence, there exists a constant $\alpha = \rho^{1-m-m^2} > 0$ such that $\Pr[\sigma^*(i) \in \sigma([k])] \geq \Pr[\sigma^*(j) \in \sigma([k])] + \alpha$ for all $i, j \in [m]$ with $i < j$ and $k \in [m - 1]$, as desired.

Assumption A4: We want to show that there exists a constant $\alpha > 0$ such that for all $i, j \in [m]$ with $i < j$, the quantity

$$\Pr_P[\sigma^*(i) = \sigma(i') \wedge \sigma^*(j) = \sigma(j')] - \Pr_P[\sigma^*(i) = \sigma(j') \wedge \sigma^*(j) = \sigma(i')] \quad (7)$$

is non-negative for all $i', j' \in [m]$ with $i' < j'$ and at least α for some $i', j' \in [m]$ with $i' < j'$.

Fix $i, j, i', j' \in [m]$ such that $i < j$ and $i' < j'$. Similarly to Equation (6), we note that $\sigma \leftrightarrow \sigma_{\sigma^*(i) \leftrightarrow \sigma^*(j)}$ is also a bijection between the set of rankings where $\sigma^*(i) = \sigma(i') \wedge \sigma^*(j) = \sigma(j')$, and the set of rankings where $\sigma^*(i) = \sigma(j') \wedge \sigma^*(j) = \sigma(i')$. Hence, following the same steps, we can derive

$$\begin{aligned}
&\Pr_P[\sigma^*(i) = \sigma(i') \wedge \sigma^*(j) = \sigma(j')] - \Pr_P[\sigma^*(i) = \sigma(j') \wedge \sigma^*(j) = \sigma(i')] \\
&\geq (1 - \phi) \cdot \Pr_P[\sigma^*(i) = \sigma(i') \wedge \sigma^*(j) = \sigma(j')].
\end{aligned}$$

Thus, the difference is always non-negative. Further, note that there always exists a ranking σ where $\sigma^*(i) = \sigma(i') \wedge \sigma^*(j) = \sigma(j')$. Thus, using the same bound as in the case of assumption A3, we get that there exists a constant $\alpha > 0$ depending only on ρ and m such that the quantity in Equation (7) is at least α for all $1 \leq i' < j' \leq m$.

PART II: The Thurstone-Mosteller and the Plackett-Luce models.

We give a common proof by viewing both noise models as special cases of a random utility model. Let θ_a denote the true quality of alternative a . Let $\mu_a(\theta_a)$ denote the distribution from which noisy estimate of the quality of alternative a is sampled. Let f_{θ_a} and F_{θ_a} denote the PDF and CDF, respectively, of $\mu_a(\theta_a)$, i.e., the noisy quality estimate comes from a distribution that only depends on the true quality. We assume the following two properties on f_θ .

(P1) f_θ shifts with θ , i.e., $f_\theta(x) = f_{\theta'}(x + \theta' - \theta)$ for all $x, \theta, \theta' \in \mathbb{R}$.

(P2) It is more likely that a higher quality estimate emerged from higher true quality and lower quality estimate emerged from lower true quality, than vice-versa. Formally, for all $\theta > \theta'$ and $x > x'$, $f_\theta(x) \cdot f_{\theta'}(x') - f_\theta(x') \cdot f_{\theta'}(x) > 0$. Further, for $x, x', \theta, \theta', x - x'$, and $\theta - \theta'$ bounded from above and below by constants the difference is at least a constant.

(P3) A random variable X distributed according to f_θ satisfies that $\Pr[|X - \theta| \leq c]$, $\Pr[X > \theta + c]$, and $\Pr[X < \theta - c]$ are all positive constants less than 1 if c is a positive constant. Further, the set of values $\{f_\theta(\theta + x)\}_{x \in [-c, c]}$ is bounded from above and below by positive constants if c is a constant.

Lemma 3. *The normal distribution (with variance parameter bounded from above and below by positive constants) and the Gumbel distribution satisfy properties P1, P2, and P3.*

Proof. Property P1 follows directly from the definition of the two distributions. For property P2, first consider the normal distribution with fixed variance ν .

$$\begin{aligned} & f_\theta(x) \cdot f_{\theta'}(x') - f_\theta(x') \cdot f_{\theta'}(x) \\ &= \frac{1}{2\pi\nu^2} \left(e^{-\frac{(x-\theta)^2}{2\nu^2}} \cdot e^{-\frac{(x'-\theta')^2}{2\nu^2}} - e^{-\frac{(x'-\theta)^2}{2\nu^2}} \cdot e^{-\frac{(x-\theta')^2}{2\nu^2}} \right) \\ &= \frac{1}{2\pi\nu^2} \left(e^{-\frac{(x-\theta)^2 + (x'-\theta')^2}{2\nu^2}} - e^{-\frac{(x'-\theta)^2 + (x-\theta')^2}{2\nu^2}} \right) > 0, \end{aligned}$$

where the last transition holds because

$$\begin{aligned} & (x' - \theta)^2 + (x - \theta')^2 - (x - \theta)^2 - (x' - \theta')^2 \\ &= 2(x\theta + x'\theta' - x\theta' - x'\theta) \\ &= 2(x - x')(\theta - \theta') > 0. \end{aligned}$$

Similarly, for the Gumbel distribution, we have

$$\begin{aligned} & f_\theta(x) \cdot f_{\theta'}(x') - f_\theta(x') \cdot f_{\theta'}(x) \\ &= e^{-(x-\theta)-e^{-(x-\theta)}} \cdot e^{-(x'-\theta')-e^{-(x'-\theta')}} - e^{-(x'-\theta)-e^{-(x'-\theta)}} \cdot e^{-(x-\theta')-e^{-(x-\theta')}} \\ &= e^{-x+\theta-x'+\theta'} \cdot \left(e^{-e^{-(x-\theta)}-e^{-(x'-\theta')}} - e^{-e^{-(x'-\theta)}-e^{-(x-\theta')}} \right). \end{aligned}$$

Finally, we have that

$$\begin{aligned} & e^{-(x-\theta)} + e^{-(x'-\theta')} - e^{-(x'-\theta)} - e^{-(x-\theta')} \\ &= \left(e^{-x'} - e^{-x} \right) \left(e^{\theta'} - e^\theta \right) < 0. \end{aligned}$$

In both of these cases, it can easily be checked that the difference is at least a positive constant if $x, x', \theta, \theta', x - x'$, and $\theta - \theta'$ are bounded from both sides by constants.

For property P3, this is a well-known fact for the normal distribution when the standard deviation σ itself is bounded from above and below by constants. For the Gumbel distribution, this can be checked using its explicit PDF $f_\theta(x) = e^{-(x-\theta)-e^{-(x-\theta)}}$ and its explicit CDF $F_\theta(x) = e^{-e^{-(x-\theta)}}$.

Hence, both distributions satisfy all three properties. \square

Next, we show that any random utility model where the PDF satisfies these three properties satisfies our four assumptions. Recall that the set of alternatives and therefore their true qualities are fixed. We use a slightly different notation for the following proofs. For $i \in [m]$, let the true quality of alternative $\sigma^*(i)$ be θ_i , and let its noisy quality estimate be the random variable X_i whose value is drawn from $\mu_a(\theta_i)$, where $a = \sigma^*(i)$. Let f_{X_i} and F_{X_i} denote the PDF and CDF of X_i respectively. Let X_{-i} denote the set of random variables $\{X_1, \dots, X_{i-1}, X_{i+1}, \dots, X_m\}$.

For a set S of t random variables and $k \in [t]$, let $\text{top}_k(S)$ be the random variable denoting the k^{th} highest value among the random variables in S , i.e., it is the $t - k + 1^{\text{th}}$ order statistic of S . Correspondingly, let $f_{\text{top}_k(S)}$ and $F_{\text{top}_k(S)}$ denote its CDF and PDF respectively.

Assumption A1: Neutrality is evident because the noisy quality estimates (and therefore the ranking of the alternatives in a sample vote) depend on the true qualities of the alternatives, but are independent of the identities of the alternatives.

Assumption A2: We want to show that there exists a constant $\eta > 0$ such that $p_{1,i} \leq 1 - \eta$ for all $i \in [m]$. Note that

$$p_{1,1} \geq \Pr \left[X_1 > \frac{\theta_1 + \theta_2}{2} \right] \cdot \prod_{i=2}^m \Pr \left[X_i < \frac{\theta_1 + \theta_2}{2} \right].$$

On the right hand side, we use the fact that X_i are independent of each other. Each term represents the probability of an X_i bounded within a constant amount on one side. Hence, due to property (P3) and the fact that m is constant, we have that $p_{1,1} \geq \eta_1$ for some constant $\eta_1 > 0$. This immediately implies that $p_{1,i} \leq 1 - \eta_1$ for all $i \in \{2, \dots, m\}$. We now prove that $p_{1,1} \leq 1 - \eta_2$ for some constant $\eta_2 > 0$. This is sufficient because it then follows that assumption A2 is satisfied with $\eta = \min(\eta_1, \eta_2) > 0$.

Note that

$$p_{1,1} \leq 1 - \Pr \left[X_1 < \frac{\theta_1 + \theta_2}{2} \right] \cdot \Pr \left[X_2 > \frac{\theta_1 + \theta_2}{2} \right].$$

Once again, both probability terms on the right hand side are positive constants due to property P3. Hence, $p_{1,1}$ is a constant less than 1. Hence, we have that assumption A2 is satisfied.

Assumption A3: We want to show that there exists a constant $\alpha > 0$ such that for all $i, j \in [m]$ with $i < j$ and $k \in [m - 1]$, we have $\Pr_P[\sigma^*(i) \in \sigma([k])] \geq \Pr_P[\sigma^*(j) \in \sigma([k])] + \alpha$.

Note that $\sigma^*(i) \in \sigma([k])$ is the probability that the quality estimate of $\sigma^*(i)$ is among the k highest quality estimates. This is equivalent to the quality estimate X_i of $\sigma^*(i)$ being higher than the k^{th} highest quality estimate among X_{-i} . Hence,

$$\Pr_P[\sigma^*(i) \in \sigma([k])] = \int_{t=-\infty}^{\infty} f_{X_i}(t) F_{\text{top}_k(X_{-i})}(t) dt. \quad (8)$$

Let $\Delta = \theta_i - \theta_j$. Now, similarly to Equation (8) we also have

$$\begin{aligned} \Pr_P[\sigma^*(j) \in \sigma([k])] &= \int_{t=-\infty}^{\infty} f_{X_j}(t) F_{\text{top}_k(X_{-j})}(t) dt \\ &= \int_{t=-\infty}^{\infty} f_{X_j}(t - \Delta) F_{\text{top}_k(X_{-j})}(t - \Delta) dt. \end{aligned} \quad (9)$$

Due to our assumption P1, we have $f_{X_i}(t) = f_{X_j}(t - \Delta)$ for all $t \in \mathbb{R}$. Now, we also have that for all $t \in \mathbb{R}$,

$$F_{\text{top}_k(X_{-j})}(t) \leq F_{\text{top}_k(X_{-i})}(t). \quad (10)$$

To see this, note that X_{-j} and X_{-i} have identical components except the former has X_i and the latter has X_j in its place. Note that X_i first-order stochastically dominates X_j because of our assumption P1. Hence, it follows that every order statistic of X_{-j} first-order stochastically dominates the corresponding order statistic of X_{-i} . In particular, $\text{top}_k(X_{-j})$ first-order stochastically dominates $\text{top}_k(X_{-i})$, which is exactly Equation (10).

Substituting Equation (10) in the difference of Equations (8) and (9), we get

$$\begin{aligned} &\Pr_P[\sigma^*(i) \in \sigma([k])] - \Pr_P[\sigma^*(j) \in \sigma([k])] \\ &\geq \int_{\theta_j}^{\theta_i} f_{X_i}(t) \cdot (F_{\text{top}_k(X_{-j})}(t) - F_{\text{top}_k(X_{-j})}(t - \Delta)) dt \\ &= \int_{\theta_j}^{\theta_i} f_{X_i}(t) \cdot \Pr[\text{top}_k(X_{-j}) \in (t - \Delta, t]] dt. \end{aligned}$$

Now, due to property P3 and because a continuous function achieves its minimum on a closed interval, both $f_{X_i}(t)$ and $\Pr[\text{top}_k(X_{-j}) \in (t - \Delta, t]]$ are bounded from below by positive constants

for $t \in [\theta_j, \theta_i]$. Hence, the integral of their product is bounded from below by a positive constant, which implies that assumption A3 is satisfied.

Assumption A4: We show a stronger condition that there exists a positive constant $\alpha > 0$ such that for all $i, j, i', j' \in [m]$ with $i < j$ and $i' < j'$,

$$\Omega = \Pr_P[\sigma^*(i) = \sigma(i') \wedge \sigma^*(j) = \sigma(j')] - \Pr_P[\sigma^*(i) = \sigma(j') \wedge \sigma^*(j) = \sigma(i')] \geq \alpha. \quad (11)$$

Fix $i, j \in [m]$ with $i < j$ and $i', j' \in [m]$ with $i' < j'$. We explicitly evaluate the probability difference Ω in Equation (11) by conditioning on the values of $X_S = \{X_k\}_{k \in S}$, where $S = [m] \setminus \{i, j\}$. Consider two cases.

Case 1: $j' > i' + 1$ In this case, we want that $X_i \in (\text{top}_{i'}(X_S), \text{top}_{i'-1}(X_S)]$ and $X_j \in (\text{top}_{j'-1}(X_S), \text{top}_{j'-2}(X_S)]$. (The latter interval has shifted indices because insertion of X_i would shift the rank of X_j by one.) Thus, we can evaluate Δ in Equation (11) as follows.

$$\begin{aligned} \Omega &= \int_{t_S} \int_{t_i=\text{top}_{i'}(X_S)}^{\text{top}_{i'-1}(X_S)} \int_{t_j=\text{top}_{j'-1}(X_S)}^{\text{top}_{j'-2}(X_S)} f_{X_i}(t_i) f_{X_j}(t_j) f_{X_S}(t_S) dt_j dt_i dt_S \\ &\quad - \int_{t_S} \int_{t_j=\text{top}_{i'}(X_S)}^{\text{top}_{i'-1}(X_S)} \int_{t_i=\text{top}_{j'-1}(X_S)}^{\text{top}_{j'-2}(X_S)} f_{X_i}(t_i) f_{X_j}(t_j) f_{X_S}(t_S) dt_j dt_i dt_S \\ &= \int_{t_S} \int_{q=\text{top}_{i'}(X_S)}^{\text{top}_{i'-1}(X_S)} \int_{q'=\text{top}_{j'-1}(X_S)}^{\text{top}_{j'-2}(X_S)} (f_{X_i}(q) f_{X_j}(q') - f_{X_i}(q') f_{X_j}(q)) f_{X_S}(t_S) dq' dq dt_S. \end{aligned}$$

Now, define $\Delta = (1/3) \cdot \min_{k, l \in [m], k \neq l} |\theta_k - \theta_l|$. Then, the intervals $\{[\theta_k - \Delta, \theta_k + \Delta]\}_{k \in [m]}$ do not intersect. Further, there is a constant probability that X_k is sampled from the interval $[\theta_k - \Delta, \theta_k + \Delta]$ is a positive constant for all $k \in S$ due to property (P3). Let us denote by R the region in which this happens for all $k \in S$. Hence, $\Pr[t_S \in R]$ is also a positive constant. We lower bound Ω by restricting the integration over t_S to R . Further, due to property (P2), we have that over that region,

$$\beta = \min_{\substack{t_S \in R \\ q \in (\text{top}_{i'}(t_S), \text{top}_{i'-1}(t_S)] \\ q' \in (\text{top}_{j'-1}(t_S), \text{top}_{j'-2}(t_S)]}} f_{X_i}(q) f_{X_j}(q') - f_{X_i}(q') f_{X_j}(q)$$

is a positive constant. Hence, we get that Ω is lower bounded by a positive constant, as required.

Case 2: $j' = i' + 1$. This case is similar to the previous case, except that the conditions on X_i and X_j change slightly. In this case, we need to have $X_i, X_j \in (\text{top}_{i'}(X_S), \text{top}_{i'-1}(X_S)]$ along with $X_i > X_j$. Hence, the evaluation of Ω in Equation (11) changes to

$$\begin{aligned} \Omega &= \int_{t_S} \int_{q=\text{top}_{i'}(X_S)}^{\text{top}_{i'-1}(X_S)} \int_{q'=\text{top}_{i'}(X_S)}^q (f_{X_i}(q) f_{X_j}(q') - f_{X_i}(q') f_{X_j}(q)) f_{X_S}(t_S) dq' dq dt_S \\ &\geq \int_{t_S} \int_{q=\frac{(1/3)\text{top}_{i'}(X_S) + (2/3)\text{top}_{i'-1}(X_S)}{(2/3)\text{top}_{i'-1}(X_S)}}^{\text{top}_{i'-1}(X_S)} \int_{q'=\text{top}_{i'}(X_S)}^{\frac{(2/3)\text{top}_{i'}(X_S) + (1/3)\text{top}_{i'-1}(X_S)}{(1/3)\text{top}_{i'-1}(X_S)}} (f_{X_i}(q) f_{X_j}(q') - f_{X_i}(q') f_{X_j}(q)) f_{X_S}(t_S) dq' dq dt_S \\ &\geq \beta > 0, \end{aligned}$$

where in the third transition,

$$\beta = \min_{\substack{t_S \in R \\ q \in \left[\frac{(1/3)\text{top}_{i'}(t_S) + (2/3)\text{top}_{i'-1}(t_S)}{(2/3)\text{top}_{i'-1}(t_S)}, \text{top}_{i'-1}(t_S) \right] \\ q' \in \left[\text{top}_{i'}(t_S), \frac{(2/3)\text{top}_{i'}(t_S) + (1/3)\text{top}_{i'-1}(t_S)}{(1/3)\text{top}_{i'-1}(t_S)} \right]}} f_{X_i}(q) f_{X_j}(q') - f_{X_i}(q') f_{X_j}(q),$$

where R is the same region as defined in Case 1. Once again, property (P2) implies that β is a positive constant. Hence, Ω is lower bounded by a constant.

Hence, both the Thurstone-Mosteller and the Plackett-Luce models are PPM- α for a constant $\alpha > 0$. This completes the proof that the three classical noise models, under suitable assumptions, satisfy our four assumptions. \square \square

F Proof of Theorem 2

We first prove an upper bound on the accuracy of the uniform team. Fix an agent $i \in N$. Consider a state $s \in S$. For an alternative $a \in A_s$, $\Pr[f(\boldsymbol{\pi}_{is}^k) = a \mid \sigma_{is}]$ is the probability that the uniform team chooses a when the biased truth is σ_{is} . The winner is chosen by applying f to k sampled votes from $P_i^2(\sigma_{is})$. Therefore, the neutrality of P_i^2 (assumption A1) and of f (assumption in the theorem) imply that for every permutation τ of the alternatives,

$$\Pr[f(\boldsymbol{\pi}_{is}^k) = \tau a \mid \tau \sigma_{is}] = \Pr[f(\boldsymbol{\pi}_{is}^k) = a \mid \sigma_{is}]. \quad (12)$$

Consider the set of rankings $Good_{is} \subseteq \mathcal{L}(A_s)$ such that the true best alternative $\sigma_s^*(1)$ has the highest winning probability among all alternatives in A_s if and only if $\sigma_{is} \in Good_{is}$. From Equation (12), we can see that permuting the alternatives in σ_{is} permutes the winning probabilities of the alternatives accordingly. Hence, the rankings in $Good_{is}$ are obtained by taking one ranking in $Good_{is}$ and applying all possible permutations that fix (i.e., do not relabel) $\sigma_s^*(1)$, and thus do not change its position in the ranking. Thus, there exists a $k \in [m_s]$ such that $Good_{is}$ is the set of rankings where $\sigma_s^*(1)$ is in position k .

Let $Bad_{is} = \mathcal{L}(A_s) \setminus Good_{is}$. By assumption A2, there exists a constant $\eta > 0$ such that

$$\Pr[\sigma_{is} \in Bad_{is}] \geq \eta, \forall s \in S. \quad (13)$$

Further, when $\sigma_{is} \in Bad_{is}$, we know that there exists an alternative $a \in A_s$ with winning probability at least as high as that of $\sigma_s^*(1)$. Hence, the winning probability of $\sigma_s^*(1)$ is at most $1/2$. That is,

$$\mathbb{E}[X_{is}^k \mid \sigma_{is} \in Bad_{is}] \leq 1/2. \quad (14)$$

Putting everything together, we have that

$$\begin{aligned} \mathbb{E}[X_{is}^k] &= \sum_{s \in S} \mu(s) \cdot \left[\Pr[\sigma_{is} \in Bad_{is}] \cdot \mathbb{E}[X_{is}^k \mid \sigma_{is} \in Bad_{is}] \right. \\ &\quad \left. + \Pr[\sigma_{is} \in Good_{is}] \cdot \mathbb{E}[X_{is}^k \mid \sigma_{is} \in Good_{is}] \right] \\ &\leq \sum_{s \in S} \mu(s) \cdot \left[\Pr[\sigma_{is} \in Bad_{is}] \cdot \frac{1}{2} + \Pr[\sigma_{is} \in Good_{is}] \cdot 1 \right] \leq \sum_{s \in S} \mu(s) \cdot \left[\eta \cdot \frac{1}{2} + (1 - \eta) \cdot 1 \right] \\ &= 1 - \frac{\eta}{2}, \end{aligned}$$

where the second transition holds due to Equation (14), and the third transition holds due to Equation (13). Taking $c = \eta/2$ proves the first part of the theorem.

For the results regarding the diverse team, recall that in a state $s \in S$, every agent $i \in N$ first draws its biased truth $\sigma_{is} \sim P_i^1(\boldsymbol{\theta}_s)$, and then draws a sample $\psi_{is} \sim P_i^2(\sigma_{is})$. Equivalently, we can say that each agent $i \in N$ draws its vote $\psi_{is} \sim (P_i^2 \circ P_i^1)(\boldsymbol{\theta}_s)$. Let $P_i = P_i^2 \circ P_i^1$. Thus, $\boldsymbol{\psi}_s^n$ is a profile consisting of one sample from $P_i(\boldsymbol{\theta}_s)$ for each $i \in N$. We want to show that $\lim_{n \rightarrow \infty} \Pr[f(\boldsymbol{\psi}_s^n) = \sigma_s^*(1)] = 1$. We establish this under each of the two conditions mentioned in the second part of the theorem.

First, let assumptions A1 and A3 hold, and let f be a PD-c voting rule. Fix a state $s \in S$. Since P_i^1 and P_i^2 are PD- α (assumption A3), P_i is PD- α^2 due to Lemma 1. Thus, by definition we have that for every agent $i \in N$, every alternative $a \in A_s \setminus \{\sigma_s^*(1)\}$, and every $j \in [m_s - 1]$,

$$\Pr[\sigma_s^*(1) \in \psi_{is}([j])] \geq \Pr[a \in \psi_{is}([j])] + \alpha^2.$$

Recall that $T_{\boldsymbol{\psi}_s^n}(a, j)$ denotes the number of times alternative a appears among first j positions in $\boldsymbol{\psi}_s^n$. Due to Kolmogorov's strong law, we have that for all $a \in A_s \setminus \{\sigma_s^*(1)\}$ and for all $j \in [m_s - 1]$,

$$\lim_{n \rightarrow \infty} \Pr[T_{\boldsymbol{\psi}_s^n}(\sigma_s^*(1), j) \leq T_{\boldsymbol{\psi}_s^n}(a, j)] = 0.$$

Taking the union over $a \in A_s \setminus \{\sigma_s^*(1)\}$ and $j \in [m_s - 1]$, and applying the union bound, we get that the probability that $\sigma_s^*(1)$ is *not* the position-dominating winner in ψ_s^n is zero in the limit when $n \rightarrow \infty$. Since this holds for each state $s \in S$, it holds in expectation over the state space S . Moreover, the PD-c rule f always outputs the position-dominating winner (with probability 1) whenever it exists. Hence, we have $\lim_{n \rightarrow \infty} \Pr[f(\psi_s^n) = \sigma_s^*(1)] = 1$, as required.

The proof for Condorcet consistent rules is almost identical to the proof for PD-c rules. Let assumptions A1 and A4 hold, and let f be a Condorcet consistent voting rule. Fix a state $s \in S$. Since P_i^1 and P_i^2 are PPM- α and PM- α respectively (assumption A4), P_i is PM- α^2 due to Lemma 2. Thus, by definition we have that for every agent $i \in N$ and alternative $a \in A_s \setminus \{\sigma_s^*(1)\}$,

$$\Pr[\sigma_s^*(1) \succ_{\psi_{i_s}} a] \geq \Pr[a \succ_{\psi_{i_s}} \sigma_s^*(1)] + \alpha^2.$$

Due to Kolmogorov’s strong law, a majority of votes in ψ_s^n would rank $\sigma_s^*(1)$ above a with probability 1 as $n \rightarrow \infty$, for every $a \in A_s \setminus \{\sigma_s^*(1)\}$. Hence, in the limit, $\sigma_s^*(1)$ becomes the Condorcet winner with probability 1. Since this holds for each state $s \in S$, it holds in expectation over the state space S . Moreover, the Condorcet consistent voting rule f must output the Condorcet winner (with probability 1) whenever it exists. Hence, once again we have $\lim_{n \rightarrow \infty} \Pr[f(\psi_s^n) = \sigma_s^*(1)] = 1$, as required. \square (Proof of Theorem 2) \square

While Theorem 2 is already quite general, it is possible to generalize it even further. We preferred to present the simpler version for ease of exposition; but let us informally say that it is also possible to handle the case where the ground truth is an objective true quality for each alternative (and then one would rather know the expected quality of the chosen alternative), and the case where the number of alternatives in each state is not fixed.

G Parametrized Agents

In Table 1 we present the parameters that were sampled to generate parametrized versions of Fuego. For each random draw, we used a uniform random distribution, defined in the interval shown in the column “Range”. Also, depending on the domain of each parameter, we sample integers or floating point numbers. A detailed description of these parameters is available in the Fuego documentation, at <http://fuego.sourceforge.net/fuego-doc-1.1/>.

H Additional Experimental Results

In this section we present further analysis of our experimental results. First we show that the original Fuego is, indeed, stronger than the parametrized agents. Like in our experiments, we ran 1000 9×9 Go games, with the system under evaluation playing as white, against the original Fuego playing as black. In Figure 3 we can see the winning rate of Fuego and of each one of the parametrized agents. The original Fuego is the strongest agent (with $p < 0.01$ for all but 3 agents), having a winning rate close to 50%. The parametrized agents, on average, have a winning rate of 32.3% (std: 10.4%).

We also evaluate the diversity of a team of parametrized agents, by analyzing a sample of 10 parametrized agents. We use the metric proposed in [19], where diversity is defined as the average Hellinger Distance [11] between the probability density functions (PDFs) of all possible combinations of pairs of agents across different world states. We show three different results: *Control*





Figure 3: Winning rate of Fuego and of the parametrized agents.

Parameter	Domain	Range
uct_param_globalsearch_mercy_rule	Integer	[0,1]
uct_param_globalsearch_territory_statistics	Integer	[0,1]
uct_param_globalsearch_length_modification	Float	[0, 0.5]
uct_param_globalsearch_score_modification	Float	[0,0.5]
uct_param_player_forced_opening_moves	Integer	[0,1]
uct_param_player_reuse_subtree	Integer	[0,1]
uct_param_player_use_root_filter	Integer	[0,1]
uct_param_policy_nakade_heuristic	Integer	[0,1]
uct_param_policy_fillboard_tries	Integer	[0, 5]
uct_param_rootfilter_check_ladders	Integer	[0,1]
uct_param_search_check_float_precision	Integer	[0,1]
uct_param_search_prune_full_tree	Integer	[0,1]
uct_param_search_rave	Integer	[0,1]
uct_param_search_virtual_loss	Integer	[0,1]
uct_param_search_weight_rave_updates	Integer	[0,1]
uct_param_search_bias_term_constant	Float	[0, 1.0]
uct_param_search_expand_threshold	Integer	[1,4]
uct_param_search_first_play_urgency	Integer	[1,10000]
uct_param_search_knowledge_threshold	Integer	[0,10000]
uct_param_search_number_playouts	Integer	[1,3]
uct_param_search_prune_min_count	Integer	[1,128]
uct_param_search_randomize_rave_frequency	Integer	[0,200]
uct_param_search_rave_weight_final	Integer	[1000,10000]
uct_param_search_rave_weight_initial	Integer	[0,999]

Table 1: Parameters sampled to generate different versions of Fuego.

compares each agent with a second sample of itself, in order to measure the noise in our evaluation; *Parametrized Agents* compares all possible pairs of parametrized agents, in order to estimate the diversity of our team; and *Independent Agents* compares each parametrized agent with Pachi [4], an independently developed Computer Go program. In order to perform the analysis, we estimate the PDFs of Pachi and 10 agents from the diverse team, using 100 different board states. For each board state we sample 100 moves for each agent. The results are shown in Figure 4(a). These results indicate that the level of diversity is especially high when the parametrized agents are compared with Pachi, suggesting that the current parametrization methodology falls short of creating an idealized diverse team. That said, the methodology does lead to some diversity, as indicated by the statistically significant difference between the Control bar and Parametrized Agents bar.

We also evaluate the level of diversity by testing whether there is a set of board states where all parametrized agents have a low probability of playing the best action. Again, we evaluate a sample of 10 agents from the diverse team. We first estimate the best move for each of 100 board states. To this end, we use Fuego to evaluate the given board state, but with a time limit 50x higher than the default one. Then, based on the previous estimated PDFs of the parametrized agents, we can obtain the probability of each agent playing the optimal action. Finally, we calculate the proportion of board states in which all parametrized agents play the best action with probability below a certain threshold. The results are shown in Figure 4(b). It turns out that all parametrized agents play the optimal action with probability smaller than $1/2$ in 40% of the board states. Moreover, in 10% of the board states, the probability of playing the optimal action is lower than 10%. Hence, there is still a large set of board states in which all agents play badly, regardless of the parametrization.



(a) Diversity of the parametrized agents, compared with a second sample and with the diversity between agents have probability of playing the best action independently developed agents
 (b) Percentage of world states where all parametrized agents have probability of playing the best action below the given threshold

Figure 4: Evaluation of the diversity of the parametrized agents, and the fraction of states in which all of them have a low probability of playing the optimal action. The error bars show 99% confidence intervals.