

PEERBTS: Incentivizing Effort in Strategyproof Peer Selection

Harper Lyon
Tulane University
New Orleans, United States of
America
hlyon@tulane.edu

Omer Lev
Ben-Gurion University
Beersheba, Israel
omerlev@bgu.ac.il

Nicholas Mattei
Tulane University
New Orleans, United States of
America
nsmattei@tulane.edu

ABSTRACT

Background: Peer evaluation, the evaluation and selection of agents by their peers, is an important problem in the field of computational social choice and has a wealth of existing work focusing on strategyproof mechanisms. However, all published mechanism share one flaw in that they do not reward agents for any effort involved in the evaluation process. In cases where high quality evaluation has a cost - as in academic peer review - this missing incentive is a serious oversight.

Objectives and Research Questions: We address the following questions: (1) What are the theoretical requirements for a peer-selection mechanism that incentivizes effort? (2) How can we combine prediction scoring with existing peer-selection mechanisms in an incentive compatible way? (3) Which contexts allow for incentive compatible prediction scoring in peer-selection? (4) How does the addition of prediction scoring impact the performance of a peer-selection mechanism?

Methods: We first show that incentivizing effort in peer-selection requires information beyond a single report. In response we propose PeerBTS, combining a lottery over prediction scores from the Bayesian Truth Serum with existing peer-selection mechanisms to incentivize effort while remaining BNIC. We demonstrate the conditions under which the dual incentives of honesty and effort are preserved. Finally, we present a simulation based on Mallows models to validate the performance of PeerBTS.

Results: We demonstrate that an incentive-compatible peer-selection mechanism using agent predictions to incentive effort is possible, but with various trade-offs. Our prediction scoring function, the RBTS, requires a known shared prior over agent preferences to ensure honesty, and the combination of a prediction lottery with a traditional mechanism introduces in-exactness. In simulation we find that PeerBTS behaves as expected in regards to incentives with limited degradation of mechanism performance. Finally, we briefly discuss the results of a limited pilot experiment of including predictions in peer-review performed for AAMAS-GIAW 26.

Conclusions: We believe we are the first to suggest such a combination of both truth and effort incentives, and we hope it will result in further initiatives in this domain. Our mechanism is merely incentive compatible, and suffers from various downsides but as first

step towards combining incentives in peer-selection it represents a promising direction for both theoretical work in mechanism design and potentially practical improvements to real world peer selection systems.

KEYWORDS

Peer Selection, Peer Review, Prediction Scoring, Incentive Compatibility

ACM Reference Format:

Harper Lyon, Omer Lev, and Nicholas Mattei. 2026. PEERBTS: Incentivizing Effort in Strategyproof Peer Selection. In *Appears at the 8th Games, Agents, and Incentives Workshop (GAIW-26). Held as part of the Workshops at the 25th International Conference on Autonomous Agents and Multiagent Systems., Paphos, Cyprus, May 2026*, IFAAMAS, 11 pages.

1 INTRODUCTION

Peer evaluation – the process by which a group evaluates its own members and chooses the best of them – has been used by people for millennia. The various “mechanisms” people used ranged from hoping for divine interventions and lotteries¹, to attempts at forms of elections [25]. However, the salience of this problem became much more significant as technology enabled the groups conducting peer evaluation to grow dramatically, e.g., Massive Open Online Courses (MOOCs) like Coursera. Needing a way to grade the tasks done by huge classes (up to several million), practitioners have turned to students to rate each other [8, 21, 28]. Similarly, Computer Science conferences face an increasing demand for reviewers, and many have resorted to requiring those submitting papers to also review papers.

As the relevance of peer evaluation grew, the issues with existing systems became clearer. First and foremost, they are not strategyproof, and thus may incentivize agents to misreport their views, as it may improve their chances of getting chosen themselves [2, 14]. Consider, for example, the system used in many CS conferences, in which a submission’s “grade” is an average of its reviewers’ grades, and acceptance is for the top k , for some value of k . All reviewers will be better off giving 0 as their grade to anyone, as while they cannot affect the grade their submission receives, they can lower the threshold grade to be in the top- k , so their work might be selected. This is typically called “torpedo reviewing” in the literature and has been documented at conferences [15].

Several strategyproof mechanisms have been suggested for peer evaluation, in which no agent can improve their chance of selection by changing their review on others [1, 4, 18, 20] (more coverage is

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

Appears at the 8th Games, Agents, and Incentives Workshop (GAIW-26). Held as part of the Workshops at the 25th International Conference on Autonomous Agents and Multiagent Systems., Armstrong, Curry, Hosseini, Mattei, Tsang, Wqs (Chairs), May 2026, Paphos, Cyprus. © 2026 Copyright held by the owner/author(s). ... \$ACM ISBN 978-x-xxxx-xxxx-YY/MM

¹From the bible’s Numbers chapter 17 and Joshua chapter 7 to Delphic oracles studying goat entrails and the Athenian *Boule* [22] and the Venetian Doge selection mechanism [26], they all tried to improve their selection process with lotteries.

available in several overviews [19, 27]). The problem with strategyproof mechanisms is that, by definition, they force a separation between the evaluations an agent gives on others and the possibility of that same agent being selected. This separation is desirable in the mechanism design literature as the assumption is that removing this connection will cause agents to report truthfully [12].

However, when there is a cost – whether of time, effort, or monetary value (as with reviewing papers, grant proposals, or Coursera student work) – the disconnect between evaluations and agents’ personal benefit is detrimental to the quality of the peer evaluation since there is no *within mechanism incentive to exert effort*. Previous attempts at tying the evaluations to rewards for the agent providing reviews, e.g., Merrifield and Saari [25], Srinivasan and Morgenstern [35], use things like rewarding agents for giving a review close to the average of others agents’ reviews are inherently not strategyproof, as agents are being rewarded *not for giving their truthful views, but for being able to predict consensus opinion*. Consider a reviewer for a CS conference using such a reward scheme, who finds a fundamental problem with a paper’s proof. If the reviewer thinks other reviewers won’t pick up on the proof’s problem, they are better-off ignoring it and giving the proposal the same review as the others, as that will increase their own chances of selection.

Contribution. We propose PEERBTS, a solution using a Bayesian Truth Serum (BTS) mechanism [30], which we pair with PEERNOMINATION [20], a known strategyproof peer evaluation mechanism.² We show that existing strategyproof mechanisms do not incentivize effort and that, if the only data given by agents is their review of others, it is *impossible* to create a mechanism that incentivizes truth-telling while also incentivizing agents to put in effort to evaluate others. We give a general framework to design reward mechanisms as a combination of scoring and selection mechanisms, prove when and which properties can be retained; and instantiate and evaluate one such mechanism – PEERBTS – that is Bayes-Nash incentive compatible.

2 RELATED WORK

Peer selection has a rich history, but as an object of academic study has typically been modeled in the context of academic peer review. Peer review as done by independent panels has been an object of much study in the past [10, 24, 37], but in the past few decades there has been a move for a broad base of reviewers, leading many organizations to move beyond expert panels and adopt true peer selection by requiring applicants to also serve as reviewers. Moving towards community review and away from hand selected expert panels puts new stresses on the peer review process, including issues of incentive and dishonesty, which is where the tools of mechanism design have proved useful. Much of formal work was first inspired by Merrifield and Saari [25], who proposed a basic Borda mechanism to adjudicate the allocation of time on telescopes. A variety of more sophisticated mechanisms have been proposed, including CREDIBLESUBSET [18] and EXACTDOLLARPARTITION [4] (see Olckers and Walsh [27] for a recent survey). This research focuses primarily on strategyproofness, though there are further concerns, such as group fairness [5] or handling conflict of interest

²We will see our method can incorporate any peer evaluation mechanism, however for clarity we focus on PEERNOMINATION.

[15, 40]. Our work is similar, as we are concerned with an property other than strategyproofness, namely incentivization *within* the mechanism, which has not been studied.

Beyond the traditional research in peer selection we also pull from the literature on group predictions/forecasting. Mechanisms for high quality group predictions are not a new idea and have been explored for decades [11, 13]. There are a number of traditions in this space, such as the use of markets for predictions [3, 39], but we draw from a related tradition of survey scoring, where the goal is to solicit honest responses from participants. The leading work in this space is the Bayesian Truth Serum [30], which asks agents to to make a prediction about the population response rate. Further work on group prediction – most notably the Good Judgment Project [36] and related efforts – has produced a wealth of expansions on this idea, including the Robust Bayesian Truth Serum [31, 38] which we use as a core component of PEERBTS. Research in this area has proven that the techniques of mechanism design can be applied profitably to group decision making, and so it is a natural choice for our extension of traditional peer selection mechanisms.

A related, though distinct, line of research focuses on incentivizing effort/accuracy through more active means. In particular Chakraborty et al. [9] draw on work on Massive Open Online Courses (MOOCs) [29] to present a mechanism which incentives effort from self-interested participants. While our goals are similar our approaches are quite different. In particular, Chakraborty et al. use known-value probe papers to directly estimate agent effort and reward agents with bonuses exterior to the selection process itself. This approach is well designed and could be quite profitably implemented in cases where the center has some access to ground truth, making it a good fit for MOOCs. Our own mechanism is more setting agnostic, as we do not rely on having any access to the ground truth nor on exterior rewards, though we do accept several tradeoffs for this flexibility.

3 PRELIMINARIES AND NOTATION

Our running example is a conference where agents submit proposals, which they wish to see selected, and also serve as reviewers for the conference [34]. For ease of exposition, we suppose that agents are providing *reviews* on a set of *proposals* that are under review. The review is the act of submitting their belief about a proposal’s quality, and we will also ask agents for a prediction of the reviews of other agents as part of the mechanism. We will also refer to the *review board* as the set of agents reviewing a particular proposal, and the review assignment of a single agent as the set of proposals they are to review. Unlike previous peer evaluation research, we assume there is some cost (effort, time) for agents to generate their reviews, and an agent can choose to pay this cost and generate a quality review, or not and generate a low quality one. Additionally, taking a standard mechanism design approach, agents will only do what they are *incentivized* to do [7]. Note that we are not disparaging a typical conference reviewer who may have intrinsic or altruistic motives for submitting good reports, rather we are working from first principles in incentive and mechanism design.

Following the broader literature, we define a peer selection problem over a set of agents \mathcal{N} , with $|\mathcal{N}| = n$. Each agent $i \in \mathcal{N}$ has

a private evaluation function $\sigma_i : \mathcal{N} \rightarrow \mathbb{R}$, representing their perception (review) of the quality of all agents involved which can be viewed as a (strict) ordinal ranking function $\mathcal{N} \rightarrow \{0, 1, \dots, n\}$ by ordering the agents by quality (breaking ties in some manner). Taken together, we refer to the set of all evaluating functions of all agents as a *profile*, σ , and the set of all possible profiles as Σ .

In peer selection there is a specified number k of selections or winners, \mathcal{W} , where $\mathcal{W} \subseteq \mathcal{N}$ and, ideally, $|\mathcal{W}| = k$. Informally, we want to select the “best” k sized subset of proposals. To do so, we define a selection mechanism as a function from profiles to a selected set, formally $f : \Sigma \rightarrow 2^{\mathcal{N}}$. We call any such function f *exact* if $\forall \sigma \in \Sigma, |f(\sigma)| = k$.

In real life instances of peer selection we typically do not solicit a complete evaluation or ranking of all involved agents, as this would be far too onerous. Instead we define a review assignment function $A : \mathcal{N} \rightarrow 2^{\mathcal{N}}$ such that $\forall i \in \mathcal{N}, i \notin A(i)$. This assigns each agent to a review assignment $A(i)$, and we can also refer to the set of reviewers, review board, assigned to a specific proposal as $A^{-1}(i)$. We further restrict the review assignment functions to be m -regular, where $\forall i \in \mathcal{N}, |A(i)| = m$. In this way m parametrizes the number of proposals that any agent is expected to review, and m tends to be relatively small compared to n as a matter of practicality.

Unlike previous peer evaluation research, we assume there is some cost to agent i to obtain σ_i , akin to putting in time to read a paper for a conference review. The agent can chose not to put in this effort, resulting in low-quality or random data going into the mechanism. Hence our central question is *can we design mechanism incentives to encourage agents to put in the effort to generate high quality reviews without destroying the incentives to be truthful?*

3.1 Desirable Properties

Ideally, peer selection mechanism should satisfy **strategyproofness**, under the assumption that the only thing agents desire is to have their proposal selected. Informally, this means that no agent involved in the selection process has an incentive to misreport their preferences to increase their chance of selection; hence, we can ideally expect honest reports from every reviewer to a strategyproof mechanism. Formally, a mechanism f is strategyproof if

$$\begin{aligned} \forall i \in \mathcal{N}, \forall \sigma_i, \quad i \notin f((\sigma_1, \dots, \sigma_i^*, \dots, \sigma_n)) \\ \Rightarrow \quad i \notin f((\sigma_1, \dots, \sigma_i, \dots, \sigma_n)) \end{aligned}$$

where σ_i^* is agent i 's true private ranking function. In probabilistic mechanisms we require that i cannot increase their probability of being selected by misreporting.

A related, weaker, property is **incentive compatibility**, where truthfulness is a Nash equilibrium, i.e., when all are truthful it is never beneficial for any agent to deviate,

$$\begin{aligned} \forall i \in \mathcal{N}, \forall \sigma_i, \quad P(i \in f((\sigma_1^*, \dots, \sigma_i^*, \dots, \sigma_n^*))) \\ \geq P(i \in f((\sigma_1^*, \dots, \sigma_i, \dots, \sigma_n^*))) \end{aligned}$$

An additional desirable property is that no particular agent have more influence than others – **anonymity**.

3.2 PEERNOMINATION

While PEERBTS is a general mechanism that can be combined with any strategyproof (or even incentive compatible) mechanism, we instantiate and evaluate it with PEERNOMINATION [20]. Broadly, PEERNOMINATION is approval based – agents either approve others

or not, and do not rank them or give them more granular grades. Agents approve the top $\lfloor k/n \cdot m \rfloor$, i.e., the rounded down share that needs to be selected from the population in their review assignment. The agent in the next position, rank $\lfloor k/n \cdot m \rfloor + 1$, is considered partially approved – given $(k/n \cdot m - \lfloor k/n \cdot m \rfloor)$ points. The set of accepted agents are those who have been approved with $m/2$ points, i.e., approved by most of their reviewers.

Note that PEERNOMINATION is not an exact mechanism, as it does not guarantee that exactly k agents are chosen. The mechanism has further variants, allowing for a weighing of the reviews of different agents, but that is outside the scope of this paper. As PEERNOMINATION uses ranking information for its outcome, we shall assume from now on that $\sigma_i(j)$ returns the rank of agent j by agent i .

4 EXISTING MECHANISMS LACK INCENTIVES

Our first key result is that mechanisms in which agents are simply asked to review other agents (via grade, ranking, approval, etc.) can never be both strategyproof (or even the weaker incentive compatible) and incentivize agents to put in the effort that results in good reviews.

THEOREM 1. *A peer-evaluation mechanism cannot be anonymous and strategyproof (or incentive-compatible) and incentivize agents to produce a high quality (i.e., truthful) σ , if its input is only σ .*

PROOF. Since the mechanism is anonymous, there is a profile $\bar{\sigma} \in \Sigma$ for which agent i 's review will influence the outcome. For any such profile, because the mechanism is strategyproof (incentive-compatible), and agent i 's evaluation is private and does not depend on others, then agent i 's probability of being selected is not influenced by their review. So whether i reports some randomized σ_i' or an exact σ_i^* , its chances of being selected remain the same. Since the only thing the agent cares about is being selected, and we assume there is a cost to finding σ_i^* , there is no incentive to report σ_i^* , and the mechanism has no way of incentivizing agents. \square

COROLLARY 1. *A peer-evaluation mechanism that is anonymous, strategyproof (or incentive-compatible) and incentivizes agents to produce a high quality σ , requires agents to provide at least one more piece of information besides σ_i .*

Throughout the paper we assume that providing a high quality review $\sigma \in \Sigma$, requires some sort of effort, e.g., reading a paper to give a review on it, and that if this effort is made, it will result in a truthful and high quality σ^* . This holds even if that effort was motivated by a different incentive. For example, if a mechanism gives a bonus to an agent able to estimate what the average reviewer will say of a proposal (as is the case in the NSF pilot study [25]), the agent will need to read the proposal to give an answer, but by reading it, they are now able to write a high quality review, even though they have not directly been rewarded for their review. Therefore any incentive towards effort will have a side effect of improving review quality.

5 PEERBTS: PEER SELECTION WITH BAYESIAN TRUTH SERUM

As Corollary 1 points out, we must add another piece of information to the mechanism in order to have one that incentivizes both truth-telling and high-quality reviews. To do so, we combine a Bayesian Truth Serum mechanism with an existing strategyproof peer-evaluation algorithm. We used PEERNOMINATION here, in order to present an explicit implementation, but one could use others.

5.1 Mechanism Overview

PEERBTS works at a high level by taking any existing strategyproof mechanism $f(\sigma)$ and a reward mechanism $h(\sigma, b)$ which relies on σ and another piece of information b . We build PEERBTS based on the Bayesian Truth Serum (BTS) family of mechanisms which ask agents for their opinions and their estimates of the average of others' opinions, rewarding agents based the accuracy of their predictions.

PEERBTS first constructs S_f , the set of selected agents using $f(\sigma)$, with the number of selected being $k - d$, where $d < k$ is some integer parameter. Then, using $h(\sigma, b)$, it creates a set S_h of d agents rewarded by h (see Figure 1). The set of winners is $S_f \cup S_h$, notice that even if f was an exact mechanism, the combined set may be smaller than k , as sometimes $S_f \cap S_h \neq \emptyset$.

5.2 Reward Mechanism: Predict and Evaluate

In PEERBTS we use approval functions and predictions of approval rates as the additional information required by Corollary 1. We first define the necessary part for PEERNOMINATION: an approval function, indicating which agents rank proposal i at rank q or above: $\alpha_i : A(i) \rightarrow \{0, 1\}$. Typically this is defined over a quota $q \in (0, m)$ such that

$$\alpha_i(j) = \begin{cases} 1 & \text{if } \sigma_j(i) \leq q \\ 0 & \text{otherwise} \end{cases}$$

with variations allowing for probabilistic approval when q is a non-integer. We denote the set of all such approval functions as the approval profile α .

We also collect reviewers' predictions of approval rates. We define the belief function $b_i : A(i) \rightarrow [0, 1]$ such that $b_i(j)$ represents the proportion of reviewers j believes will approve proposal i . Formally, $b_i(j) = \mathbb{E}_j \left[\frac{1}{m} \sum_{s \in A^{-1}(i)} \alpha_s(j) \right] \forall j \in A(i)$. We call the collection of belief functions a belief profile $b = \{b_1, b_2, \dots, b_n\}$.

5.3 Reward Mechanism: Prediction Scoring

To reward agents for their predictions we define a score function, which assigns a numerical score to each agent according to the reviews and predictions. Formally: $h : \Sigma \times \beta \rightarrow \mathbb{R}^{|N|}$.

We use the Robust Bayesian Truth Serum (RBTS) [38], a variant of the well known Bayesian Truth Serum [30], to score agent reviews and predictions. RBTS is preferable in this context as it is incentive compatible even when dealing with small populations. RBTS takes as input the set of approval functions α and belief functions b and assigns each reviewer a score for each response and prediction within their review assignment. Considering the reviews of a single agent i , each agent j in the review board $A^{-1}(i)$ is assigned a reference agent $k = j + 1 \bmod |A^{-1}(i)|$ and peer agent

Algorithm 1: Robust Bayesian Truth Serum

Input: Agent $i \in N$, m -regular review assignment A , approval profile α , belief profile b , proper scoring rule R
Output: Score vector $\langle s_j \in [0, 2] \mid \forall j \in A^{-1}(i) \rangle$

```

1 Scores :=  $\langle 0 \mid \forall j \in A^{-1}(i) \rangle$ 
2 // Score agents in review board
3 for  $j \in A^{-1}(i)$  do
4    $k := j + 1 \bmod m$ 
5    $l := j + 2 \bmod m$ 
6    $\delta := \min(b_k(i), 1 - b_k(i))$ 
7   if  $\alpha_j(j) = 1$  then
8      $w_j := b_i(k) + \delta$ 
9   else if  $\alpha_i(k) = 0$  then
10     $w_j := b_i(k) - \delta$ 
11   Scores( $j$ ) =  $R(w_j, \alpha_i(l)) + R(b_i(j), \alpha_i(l))$ 
12 end
13 return Scores
```

$l = j + 2 \bmod |A^{-1}(i)|$, both their review and prediction are scored against j and k . Algorithm 1 gives our formal implementation.

For our purposes what is most important is that RBTS is Bayes-Nash incentive compatible for any population of 3 or greater [38, Theorem 9], so as long as every agent $a \in N$ has a review board $|A(a)| \geq 3$, we can safely use RBTS as a scoring mechanism.

5.4 Combining Reward & Quality Selections

Combining the output of the peer selection mechanism and the Robust-BTS score in a way that maintains incentive compatibility is not trivial: There are subtle issues in many of the obvious ways in which one might combine these two functions. The most straightforward approach is to use $h(\sigma, b)$ to select the top d agents from $N \setminus f(\sigma)$, to avoid double selection and guarantee exactness. This, however, is not strategyproof nor incentive compatible. As a counter example, consider a 2-selection setting with four agents (1, 2, 3, 4), strategyproof selection mechanism f , and scoring function h . On truthful reviews from all agents we get profile σ and similarly on effortful predictions we get b . Suppose that these two inputs give us the rankings: $f(\sigma) = [1, 2, 3, 4]$ and $h(\sigma, b) = [2, 3, 1, 4]$. Following our simple combination we take the top agent from $1 \in f(\sigma)$ and then the top agent of $h(\sigma, b) \setminus \{1\}$. On the truthful rankings this would give us the set 1, 2, the two top agents of their respective rankings.

Now say that by misreporting, agent 3 can produce the profile σ' and ranking $f(\sigma') = [2, 1, 3, 4]$. This is not a violation of f 's strategyproofness since in this ranking 3 is indifferent between $f(\sigma)$ and $f(\sigma')$. However, this results in a different outcome as the rankings are combined: under $f(\sigma') = [2, 1, 3, 4]$ and $h(\sigma', b) = [2, 3, 1, 4]$ we would first select $2 \in f(\sigma)$, and then select $3 \in h(\sigma, b) \setminus \{2\} = [3, 1]$. By promoting an agent in one ranking, they can be removed from consideration in the other, breaking strategyproofness. Similar counter examples can be constructed in a probabilistic context as agents are incentivized to shift probability mass to higher ranked agents to improve their odds of selection.

To achieve Bayes-Nash incentive compatible (BNIC) combination we accept a measure of inexactness in the final selected set by allowing agents to be selected twice. This can be done deterministically, but a randomized approach is superior as it ensures any increase in prediction score increases selection probability. We treat the prediction scores as stakes in a joint lottery, inspired by




Figure 1: Overview of information flow in the PeerBTS mechanism.

Algorithm 2: Lottery Formalization

Input: Agents \mathcal{N} , m -regular review assignment A , approval function α , prediction function b , target selection size d , weighting exponent ϵ

Output: Selected set $S \subset \mathcal{N} \cup \{\emptyset\}$

```

1  $\mathcal{B} := []$  // Initialize list to hold winners for each sublottery
2 for  $i \in \mathcal{N}$  do
3   // Collect reviews of  $i$ 
4   Actual :=  $\langle \alpha_j(j) | j \in A^{-1}(i) \rangle$ 
5   // Collect predictions
6   Predicted :=  $\langle b_j(j) | j \in A^{-1}(i) \rangle$ 
7   // Score using Algorithm 1
8   Scores $_i :=$  RBTS( $i$ , Actual, Predicted)
9   for  $j \in A^{-1}(i)$  do
10     $L_i(j) := \frac{\text{Scores}(j)^\epsilon}{2^{\epsilon \cdot m}}$ 
11  end
12   $L_i(\emptyset) := 1 - \sum_{j \in A^{-1}(i)} L_i(j)$ 
13   $s_i \sim L_i$ 
14   $\mathcal{B} \leftarrow s_i$ 
15 end
16 // Draw  $d$  agents from  $\mathcal{B}$  without replacement to create winning set
17  $W \sim \binom{\mathcal{B}}{d} \setminus \{\emptyset\}$ 
18 return  $W$ 

```

Aziz et al. [4], and draw d participants. This gives us a complete mechanism in Algorithm 2.

5.5 Conditions for Incentive Compatibility

Due to space constraints we defer to Appendix A.1 the proofs of the RBTS mechanism requirements for an admissible prior.

THEOREM 2. *If agent preferences are drawn from a distribution in T and $m \geq 3$ then the RBTS lottery (Algorithm 2) is Bayes-Nash incentive compatible.*

These conditions are met both where there is a common knowledge of how beliefs are formed by all agents, which is a very strong condition. But they are also met when all agents derive their preferences from the same source/process, including that agents have the same initial chance of developing any particular belief, and that any given belief has at least some chance of being developed [32].

6 THEORETICAL PROPERTIES OF PEERBTS

In this section we give an intuition for PEERBTS guarantees.

6.1 Agent Incentives

From Theorem 2 agents have an incentive to maximize the score returned by Algorithm 1. We know that there is an equilibrium where all agents review honestly, but it is useful to examine this more directly. Specifically, we are interested in any incentive that pushes an agent to report their reviews honestly. A detailed model for agent utility or reliability is beyond the scope of this paper, but

even a simple model of agent incentives is useful for understanding the functionality of PEERBTS.

6.1.1 Incentive Example. To build intuition we present a simple example calculation of RBTS scores to demonstrate the incentives. We assume for the following that the honest rate of approval for agents is .2, representing an uncommon but not highly unlikely review. To simplify the calculations, we also assume that all agents other than one – agent j – predict randomly, drawing their prediction uniformly from $[0, 1]$. This gives us consistent expected values for most of the score calculation, $\mathbb{E}(b_i(k)), \mathbb{E}(\delta) = .5$ which gives us $\mathbb{E}(w_j) = \alpha_i(j)$ by the definition of w_j .

We first compute agent j 's expected information score $\mathbb{E}(R_q(w_j, \alpha_i(k)))$ if they report randomly (for simplicity reporting 0 or 1 with equal probability). Splitting the expected score into these two cases we get $.5(\mathbb{E}(R_q(0, \alpha_i(k))) + \mathbb{E}(R_q(1, \alpha_i(k))))$. We also split the expected value for $\alpha_k(i)$, giving us

$$.5(.2(R_q(0, 1) + R_q(1, 0)) + .8(R_q(0, 0) + R_q(1, 0)))$$

with a final score of $.5(.2(0+1) + .8(1+0)) = .5$. If, instead, j reports honestly according to p we would instead see a final expected score of

$$.2^2(R_q(0, 1) + R_q(1, 0)) + .8^2(R_q(0, 0) + R_q(1, 0)) = .68$$

meaning that in this case j has an incentive to report honestly.

When considering the prediction score the incentive is even easier to see. If j predicts randomly with $\mathbb{E}(b_i(j)) = .5$ then the expected prediction score $\mathbb{E}(R_q(b_i(j), \alpha_i(k)))$ does not even need to be broken up based on $\mathbb{E}(\alpha_i(k))$, as $R_q(.5, 0) = R_q(.5, 1) = .75$. In all cases a random prediction has an expected value of .75, so any score higher than this is preferable for agent j . If j predicts $b_i(j) = .2$, we have

$$\mathbb{E}(R_q(b_i(j), \alpha_i(k))) = .2R_q(.2, 1) + .8R_q(.2, 0) = .84$$

so in this case we also see a clear incentive for agent j to improve their predictions.

6.1.2 Formal Analysis. As shown in Theorem 2 under ideal conditions (in the honest equilibrium) the RBTS lottery incentivizes honest reporting and accurate predictions. Our example calculations show that these incentives can exist in at least one case when other agents are not playing the equilibrium strategy but the question remains: what are the limits of the incentives for honesty and accuracy when using the RBTS lottery? Assuming the agent population reports according to some distribution p^* (even if this is a composite distribution) each agent has the following general incentives:

Reports. Agents are incentivized to report according to p^* . This can be shown by careful application of a result from the original description of RBTS [38, Lemma 8], but this is not necessary. Instead note that if all other agents are reporting according to p^* , then their behavior is equivalent to that of honest agents in world where the




Figure 2: Mean lottery share for agent populations with Deciles treatment in simulation with $m = 3, e = 8.0$, and $\varphi = 0.5$.

true report distribution $p = p^*$. We know that honest agents form an equilibrium around truthfully reporting according to p , the same is true for dishonest agents all reporting p^* , so every agent has an incentive to report according to p^* .

Predictions. Agents are incentivized to accurately predict p^* , a well known property of quadratic scoring rules [33] and a result proved directly by the creators of the RBTS [38, Lemma 7]. Hence, the incentives in the RBTS (and therefore our lottery) are fundamentally incentives to follow the pack. In the case of honest agents reporting from a shared prior over observations, this incentivizes honesty. In other cases it incentivizes understanding and following the behavior of the broader agent population.

While obviously an incentive towards honesty is ideal, the pack-following nature of the RBTS can still be useful. Recall that the primary goal of PEERBTS is to provide agents an incentive to expend some effort in the reviewing process. While we do not model in detail the application of effort – as any such model would be highly ad-hoc – it is clear that an agent j acting with no effort whatsoever will need to either report a constant approval value (0 or 1) or consistently randomize their report. Since all agents have an incentive to follow the reporting of other agents as long as $p_j \neq p^*$ agent j has an incentive to expend effort to move closer to p^* – so even in cases where PEERBTS fails to be BNIC it still incentivizes effort. This is equivalent to a Keynesian guessing game [17], and should exhibit similar dynamics.

7 EMPIRICAL EVALUATION OF PEERBTS

Following the experimental methods of previous work on strategyproof peer selection [4] we evaluate PEERBTS in a series of simulations.³

7.1 Experimental Setup

We extend the framework of Lev et al. [20] and only briefly describe the pre-existing methods here; a more complete specification of certain details can be found in that paper. Complete parameters used can be found in Appendix A.2.

³All experimental code and data – including PEERBTS implementation – will be published on GitHub in the event of publication.

7.1.1 Noisy Profile Generation. Following the Condorcet perspective on social choice we assume that there exists an underlying true ranking of agent quality. For simplicity we assume that the true ranking of an agent is identical to their number, so that agent 1 is the highest quality agent, agent 2 the second, and so on. To model potential inaccuracies in agent reviews we construct our profiles from noisy observations of this ranking according to a Mallows model [23]. We sample permutations of a base ranking ϕ based on a dispersion parameter $\varphi \in [0, 1]$, creating a distribution where the probability of an alternative ranking ϕ' is $\pi_{R,\varphi} \propto \varphi^{KT(\phi,\phi')}$, where $KT(R, R')$ is the Kendall-Tau distance between ϕ to ϕ' [16].

Also note that as a continuous, shared distribution the Mallows model satisfies our requirements for an admissible prior as shown in Lemma 1 and Lemma 2. For this reason we do not simulate strategic behavior among agents, though this would be an interesting direction for future research.

7.1.2 Prediction Generation. PEERBTS solicits predictions from the entire agent population, and as such the generation of agent predictions is a key part of our simulation. Following our profile sampling methodology, we treat agent predictions as (potentially) noisy observations of a true profile σ which represents the preferences of all agents. Again mirroring our profile sampling methodology we use a Mallows model to sample agent observations, meaning that in our simulation an individual agent i 's prediction takes the form of a predicted profile σ_i^* , from which we can easily compute the belief profile β_i . Note that this is not required by our formal statement of belief profiles in PEERBTS, but it is a useful simplification for empirical simulation. Under this model we define four simple agent population models. **Clairvoyant** agents with perfect predictive accuracy; **Random** agents who predict random values; **Divided** where the top 50% of agents are Clairvoyant and the bottom 50% Random; and **Deciles** where agents are divided in deciles with decreasing predictive accuracy. Additional details regarding these models can be found in Appendix A.3.

7.2 Evaluation Metrics

Previous work on peer selection has mostly evaluated mechanism performance via the proportion agents selected from an assumed ground truth ranking, i.e., recall at k [6]. This is a useful metric, but is insufficient to interpret the impact of different prediction models. To fill this gap we define the mean lottery share for agent i as the average probability share i was awarded across all predictions in their review assignment, formally $\bar{l}_i = 1/m \sum_{j \in A(i)} L_j(i)$. This allows us to quantify an agents overall performance in the RBTS lottery and check that the mechanism is incentivizing effortful predictions by rewarding skilled predictors.

7.3 Results

Mean Lottery Share. Examining the measure of mean lottery share, i.e., amount of reward for predicting well, across all experimental values shows a small but positive impact when agents are better predictors. In all cases the mean lottery share consistently discriminates between agents based on prediction quality, however the absolute differences are small. An example of this dynamic in the Deciles population model can be seen in Figure 2 with additional results in Appendix A.4. Selecting for parameter values




Figure 3: Mean recall and anti-recall of ground truth for combined PEER NOMINATION and PEERBTS with Decile prediction model.

which produce close to the maximal difference the gap in lottery share between highly accurate predictors and random or low skill predictors is ≈ 0.2 , meaning about a 20% increase in reward. While this is not as strong as we might like it does demonstrate the proven theoretical incentive for agents to improve their predictions.

Mechanism Recall. The recall at k of PEERBTS is not a primary concern of our analysis. However, it is still useful to understand how PEERBTS performs, especially as the share of selections assigned to the RBTS lottery (reward) increases. This analysis is heavily dependent on the assumption that higher quality agents are superior predictors. The impact on recall caused by selecting from the RBTS lottery is entirely determined by the relationship between agent quality and predictive effort/skill. For this reason we only examine recall under the Deciles prediction model, as this model most directly connects agent rank and effort.

As can be seen in Figure 3, the overall mechanism performance declines as additional selections are allocated to the RBTS (reward) lottery. In particular, the prediction lottery selects more agents outside of the true $\{1, \dots, k\}$, leading to lower recall as d increases. This is consistent with the earlier measurements of lottery share under the Decile model (Figure 2), as the relatively low differences in selection probability between agents means that PEERBTS has a limited ability to select agents in the recall set if the link between effort and quality is not sufficiently strong.

One advantage of PEERBTS is that its performance does not greatly degrade as the overall noise Φ increases. In contrast, PEER-NOMINATION (and by extension all other traditional peer selection mechanisms) struggles as agent reviews diverge more dramatically from the ground truth. Despite this advantage, our results suggest that in most cases the prediction lottery share d should be kept relatively small relative to k . That is, a few slots should be reserved for the reward, but most selections should be made based on quality.

7.4 Results Discussion

Our simulations reveal distinct practical limitations in PEERBTS, especially in regards to consistently rewarding/selecting skilled predictors. Even large differences in agent predictive accuracy, such as

those in the Divided and Decile models, had very small differences in the mean lottery share given to the best and worst predictors at .1041 and .3226 respectively. While this has no bearing on the theoretical properties of PEERBTS, it suggests that any actual implementation of the hybrid mechanisms we propose should be done cautiously and with the knowledge that any agents selected through the RBTS lottery are not necessarily of particularly high quality in the ground truth. The most obvious step to take to mitigate this concern is to limit the number of selections made using the prediction lottery, as only one selection is technically required to create our desired incentive for high quality predictions.

8 CONCLUSIONS AND FUTURE WORK

We have presented a novel peer selection mechanism – PEERBTS – that both incentivizes agents to report their true beliefs, as well as incentivizing them to put in effort to “discover” their predictions and beliefs. We show that to do so, the mechanism designer must require agents to report more than just agent beliefs on their peers, and we suggest a prediction algorithm based on RBTS [38] for which the additional information is an estimate of others’ beliefs, side-by-side by one’s own review. This prediction algorithm can be combined with *any* other peer-selection mechanism, resulting in a reward structure, such that the combined mechanism keeps the truth-incentivizing properties of its component mechanisms, while rewarding agents which have put in an effort.

We believe we are the first to suggest such a combination of both truth and effort incentives, and we hope it will result in further initiatives in this domain. Our mechanism is merely incentive compatible, but we hope that a strategyproof mechanism is out there as well. As our empirical evaluations show, there is plenty more to do: first and foremost, trying to find an exact variant (or show it is impossible), and secondly, constructing reward mechanisms that more sharply differentiate between better and worse reviewers. Additionally, it is fascinating to see what can be done with further information requests beyond beliefs, and beliefs about others’ beliefs.

ACKNOWLEDGMENTS

Mattei was supported in part by NSF Awards, IIS-RI-2134857, IIS-RI-2339880 and CNS-SCC-2427237 as well as the Harold L. and Heather E. Jurist Center of Excellence for Artificial Intelligence at Tulane University and the Tulane University Center of Excellence for Community-Engaged Artificial Intelligence (CEAI). Portions of this research were conducted with high performance computational resources provided by the Louisiana Optical Network Infrastructure (LONI) (<http://www.loni.org>).

REFERENCES

- [1] Noga Alon, Felix A. Fischer, Ariel D. Procaccia, and Moshe Tennenholtz. 2011. Sum of us: strategyproof selection from the selectors. In *Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-2011), Groningen, The Netherlands, July 12-14, 2011*, Krzysztof R. Apt (Ed.). ACM, 101–110. <https://doi.org/10.1145/2000378.2000390>
- [2] Parinaz Naghizadeh Ardabili and Mingyan Liu. 2013. Incentives, Quality, and Risks: A Look Into the NSF Proposal Review Pilot. *CoRR* abs/1307.6528 (2013). arXiv:1307.6528 <http://arxiv.org/abs/1307.6528>
- [3] Kenneth J Arrow, Robert Forsythe, Michael Gorham, Robert Hahn, Robin Hanson, John O Ledyard, Saul Levmore, Robert Litan, Paul Milgrom, Forrest D Nelson, et al. 2008. The promise of prediction markets. , 877–878 pages.
- [4] Haris Aziz, Omer Lev, Nicholas Mattei, Jeffrey S. Rosenschein, and Toby Walsh. 2019. Strategyproof peer selection using randomization, partitioning, and apportionment. *Artificial Intelligence* 275 (2019), 295–309. <https://doi.org/10.1016/j.artint.2019.06.004>
- [5] Haris Aziz, Evi Micha, and Nisarg Shah. 2023. Group Fairness in Peer Review. In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine (Eds.). http://papers.nips.cc/paper_files/paper/2023/hash/ccba10dd4e80e7276054222bb95d467c-Abstract-Conference.html
- [6] Niclas Boehmer, Piotr Faliszewski, Lukasz Janeczko, Andrzej Kaczmarczyk, Grzegorz Lisowski, Grzegorz Pierczyński, Simon Rey, Dariusz Stolicki, Stanisław Szufa, and Tomasz Was. 2024. Guide to numerical experiments on elections in computational social choice. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*. 7962–7970.
- [7] Tilman Börgers. 2015. *An introduction to the theory of mechanism design*. Oxford university press.
- [8] Ioannis Caragiannis, George A Krimpas, and Alexandros A Voudouris. 2015. Aggregating Partial Rankings with Applications to Peer Grading in Massive Online Open Courses. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*. 675–683.
- [9] Anujit Chakraborty, Jatin Jindal, and Swaprava Nath. 2024. Removing Bias and Incentivizing Precision in Peer-grading. *J. Artif. Intell. Res.* 79 (2024), 1001–1046. <https://doi.org/10.1613/JAIR.1.15329>
- [10] S Cole, JR Cole, and GA Simon. 1981. Chance and consensus in peer review. *Science* 214, 4523 (1981), 881–886. <https://doi.org/10.1126/science.7302566>
- [11] Morris H DeGroot. 1974. Reaching a consensus. *Journal of the American Statistical Association* 69, 345 (1974), 118–121.
- [12] Piotr Faliszewski and Ariel D Procaccia. 2010. AI’s war on manipulation: Are we winning? *AI Magazine* 31, 4 (2010), 53–64.
- [13] Christian Genest and James V Zidek. 1986. Combining probability distributions: A critique and an annotated bibliography. *Statist. Sci.* 1, 1 (1986), 114–135.
- [14] Iryna Gurevych, Anna Rogers, Nihar B Shah, and Jingyan Wang. 2024. Reviewer No. 2: Old and New Problems in Peer Review (Dagstuhl Seminar 24052). *Dagstuhl Reports* 14, 1 (2024), 130–161.
- [15] Steven Jecmen, Hanrui Zhang, Ryan Liu, Nihar B. Shah, Vincent Conitzer, and Fei Fang. 2020. Mitigating Manipulation in Peer Review via Randomized Reviewer Assignments. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (Eds.). <https://proceedings.neurips.cc/paper/2020/hash/93fb39474c51b8a82a68413e2a5ae17a-Abstract.html>
- [16] Maurice G Kendall. 1938. A new measure of rank correlation. *Biometrika* 30, 1/2 (1938), 81–93.
- [17] John Maynard Keynes. 1936. *The State of Long-Term Expectation*. Macmillan, London, Chapter 12, 147–164.
- [18] David Kurokawa, Omer Lev, Jamie Morgenstern, and Ariel D. Procaccia. 2015. Impartial Peer Review. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, Qiang Yang and Michael J. Wooldridge (Eds.). AAAI Press, 582–588. <http://ijcai.org/Abstract/15/088>
- [19] Omer Lev, Harper Lyon, and Nicholas Mattei. 2024. Impartial Peer Selection: An Annotated Reading List. *ACM SIGecom Exchanges* 22, 1 (2024), 113–117.
- [20] Omer Lev, Nicholas Mattei, Paolo Turrini, and Stanislav Zhydkov. 2023. Peer-Nomination: A novel peer selection algorithm to handle strategic and noisy assessments. *Artificial Intelligence* 316 (2023), 103843.
- [21] Heng Luo, Anthony C Robinson, and Jae-Young Park. 2014. Peer grading in a MOOC: Reliability, validity, and perceived effects. *Journal of Asynchronous Learning Networks* 18, 2 (2014), n2.
- [22] Irad Malkin and Josine Blok. 2024. *Drawing Lots: From Egalitarianism to Democracy in Ancient Greece*. Oxford University Press.
- [23] Colin L Mallows. 1957. Non-null ranking models. I. *Biometrika* 44, 1/2 (1957), 114–130.
- [24] Robert A McNutt, Arthur T Evans, Robert H Fletcher, and Suzanne W Fletcher. 1990. The Effects of Blinding on the Quality of Peer Review: A Randomized Trial. *The Journal of the American Medical Association* 263, 10 (1990), 1371–1376.
- [25] Michael R Merrifield and Donald G Saari. 2009. Telescope time without tears: a distributed approach to peer review. *Astronomy & Geophysics* 50, 4 (2009), 4–16.
- [26] Miranda Mowbray and Dieter Gollmann. 2007. Electing the Doge of Venice: Analysis of a 13th century protocol. In *Proceedings of the IEEE Symposium on Computer Security Foundations*. 295–310.
- [27] Matthew Olckers and Toby Walsh. 2022. Manipulation and Peer Mechanisms: A Survey. *CoRR* abs/2210.01984 (2022). <https://doi.org/10.48550/ARXIV.2210.01984> arXiv:2210.01984
- [28] Chris Piech, Jonathan Huang, Zhenghao Chen, Chuong Do, Andrew Ng, and Daphne Koller. 2013. Tuned models of peer assessment in MOOCs. *arXiv preprint arXiv:1307.2579* (2013).
- [29] Chris Piech, Jonathan Huang, Zhenghao Chen, Chuong B. Do, Andrew Y. Ng, and Daphne Koller. 2013. Tuned Models of Peer Assessment in MOOCs. *CoRR* abs/1307.2579 (2013). arXiv:1307.2579 <http://arxiv.org/abs/1307.2579>
- [30] Dražen Prelec. 2004. A Bayesian Truth Serum for Subjective Data. *Science* 306, 5695 (2004), 462–466. <https://doi.org/10.1126/science.1102081> arXiv:https://www.science.org/doi/pdf/10.1126/science.1102081
- [31] Goran Radanovic and Boi Faltings. 2013. A Robust Bayesian Truth Serum for Non-Binary Signals. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 27, 833–839.
- [32] Michel Regenwetter. 2006. *Behavioral social choice: probabilistic models, statistical inference, and applications*. Cambridge University Press.
- [33] Reinhard Selten. 1998. Axiomatic characterization of the quadratic scoring rule. *Experimental Economics* 1, 1 (1998), 43–61.
- [34] Nihar B. Shah. 2022. Challenges, experiments, and computational solutions in peer review. *Commun. ACM* 65, 6 (2022), 76–87. <https://doi.org/10.1145/3528086>
- [35] Siddharth Srinivasan and Jamie Morgenstern. 2021. Auctions and peer prediction for academic peer review. *arXiv preprint arXiv:2109.00923* (2021).
- [36] Philip E Tetlock and Dan Gardner. 2016. *Superforecasting: The art and science of prediction*. Random House.
- [37] Christine Wenneras and Agnes Wold. 1997. Nepotism and sexism in peer-review. *Nature* 387 (1997), 341–343. <https://doi.org/10.1038/387341a0>
- [38] Jens Witkowski and David C. Parkes. 2012. A Robust Bayesian Truth Serum for Small Populations. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence, July 22-26, 2012, Toronto, Ontario, Canada*, Jörg Hoffmann and Bart Selman (Eds.). AAAI Press, 1492–1498. <https://doi.org/10.1609/AAAI.V26I1.8261>
- [39] Justin Wolfers and Eric Zitzewitz. 2004. Prediction markets. *Journal of economic perspectives* 18, 2 (2004), 107–126.
- [40] Yichong Xu, Han Zhao, Xiaofei Shi, and Nihar B. Shah. 2019. On strategyproof conference peer review. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (Macao, China) (IJCAI’19)*. AAAI Press, 616–622.

A APPENDIX

A.1 Proofs for Conditions for Incentive Compatability

For the lottery mechanism described in Algorithm 2 to inherit the strict Bayes-Nash incentive compatibility of the RBTS, RBTS' conditions need to be met: All agents involved need to have a shared admissible prior and there are at least 3 agents involved in any given round of the lottery, which is equivalent to requiring $m \geq 3$. Of these, the admissible prior is the only difficulty, as we have not previously defined the beliefs of agents involved in the peer selection in probabilistic terms.

The requirements for an admissible prior are detailed in Witkowski and Parkes [38], Definition 1, but for our purposes we need to define a set of states T with $|T| \geq 2$ (T will be, for us, the state of nature – any specific instantiation of agents' beliefs) and common priors $p_{j,t} = P(\alpha_j(i) = 1|t)$ such that:

- (1) $\forall j, k \in \mathcal{N}, t \in T : p_{j,t} = p_{k,t}$.
- (2) $\forall j \in \mathcal{N}, t \in T : 0 < p_{j,t} < 1$.
- (3) $\forall j \in \mathcal{N} \exists t_1, t_2 \in T : p_{j,t_1} \neq p_{j,t_2}$

We will show that assuming agent beliefs are sampled from the same distribution fulfill these requirements. Such a distribution can be a noisy variant of a ground truth, such as a Mallows' distribution [23] (the common theoretical assumption made by peer-selection papers), or any other relevant distribution to the particular situation.

To prove Theorem 2, we first need several lemmata.

LEMMA 1. *If all agent beliefs (i.e., their preferences over others) are drawn from shared distribution $D \in \Delta(\mathcal{N}^{\mathcal{N}})$ (i.e., $\forall j \in \mathcal{N} : \sigma_j \sim D$), then there exists a shared prior $\forall i, j \in \mathcal{N}, d \in D(\mathcal{N}^{\mathcal{N}}) : P(\alpha_j(i) = 1|D)$.*

PROOF. Recall that $\alpha_i(j) = \begin{cases} 1 & \text{if } \sigma_j(i) \leq q \\ 0 & \text{otherwise.} \end{cases}$. Thus, it is fully determined by $\sigma_j - \alpha_i(j)$ inherits the distribution from D and we have a shared prior over D . \square

As this result suggests, we will define states T as a set of distributions over beliefs. To satisfy the requirement that the prior be fully mixed we take $T \subseteq \Delta^+(\mathcal{N}^{\mathcal{N}})$ ($\Delta^+(X)$ being items with full support over the simplex of elements of X).

LEMMA 2. *If $T \subseteq \Delta^+(\mathcal{N}^{\mathcal{N}})$ and $t_1, t_2 \in T$ exist such that $\forall j \in \mathcal{N} : p_{j,t_1} \neq p_{j,t_2}$, then any distribution over T that supports t_1, t_2 forms an admissible prior.*

PROOF. Lemma 1 shows that any shared distribution over preferences forms a shared prior for agent approval, satisfying condition 1. We can see that condition 2 since elements in T have a positive probability for any element in $\mathcal{N}^{\mathcal{N}}$ (yet only k elements are selected from \mathcal{N}), each element $i \in \mathcal{N}$ has a positive probability of not being selected and of being selected, i.e., $0 < p_{j,t} < 1$, satisfying condition 2. Finally, condition 3 is satisfied as it is explicitly assumed by the lemma. \square

We can now finish proving Theorem 2:

PROOF OF THEOREM 2. Lemmas 1 and 2 show that our use of the RBTS algorithm, maintains its Bayes-Nash incentive compatible property. So what is left is to show that our combination of the RBTS with the regular mechanism (Algorithm 2) does so as well. The probability of selecting agent j in the final lottery over \mathcal{B} is proportional to the number of j entries in \mathcal{B} . Therefore, j maximizes their chance of being selected by maximizing the number of entries they receive in \mathcal{B} , i.e., by maximizing the number of expected wins of lotteries L_i where $j \in A^{-1}(i)$. As each of these lotteries are independent, we can consider j 's goal as maximizing $L_i(j)$ for some arbitrary i . Since $L_i(j) = \frac{\text{Scores}_i(j)^\epsilon}{2^{\epsilon \cdot m}}$, and both ϵ and m are set in advance, this means agent j seeks to maximize their entry in Scores_i , i.e., their RBTS score.

We know that for agent populations ≥ 3 and under admissible priors that the RBTS scoring mechanism is Bayes-Nash incentive compatible [38, Theorem 9], and therefore that our RBTS lottery mechanism is as well. \square

COROLLARY 2. *PEERBTS incentivizes agents to report σ_i^* .*

PROOF OF COROLLARY 2. Agent i 's chance of being selected in S_f is not influenced by its σ_i , thanks to f 's strategyproofness. Agent i 's score in H is also not improved by changing σ_i^* . So no agent would intentionally report $\sigma_i' \neq \sigma_i^*$, when σ_i^* is known

Reporting a good b_i is strictly beneficial for the score in H . Thus, producing a good estimate of others, requires putting in the effort required to produce a good σ_i^* (as we noted in the comment following Corollary 1). \square

A.2 Simulation Parameters

Simulations were run with constant $n = 120, k = 20$ and varying $m \in \{3, 4, 6, 9\}, d \in \{2, 5, 10\}, e \in \{1, 2, 4, 8\}, \varphi \in \{0.0, 0.2, 0.5\}$ (review board size, RBTS lottery selection size, lottery exponentiation term, and dispersion parameter for agent profiles respectively).

A.3 Agent Prediction Model Details

In general, we can define an agent's predicted profile with the single dispersion parameter φ_i^* as

$$\sigma_i^*(j) = \begin{cases} \sigma_i & j = i \\ \sigma_i^* \sim M_{\varphi_i^*}(\sigma_j) & j \neq i \end{cases} \forall j \in \mathcal{N}$$

Where $M_{\varphi_i^*}(\sigma_j)$ is the Mallows distribution on permutations of σ_j under the dispersion parameter $\varphi_i^* \in [0, 1]$ as discussed previously. This is not a complete model of all ways agents could make predictions, but it is a flexible framework for studying the impact of prediction quality on agent outcomes under PEERBTS. While any set of \mathcal{N} φ^* values are sufficient to define the prediction skills of an agent population, we restricted our simulations to four simplified population models,

- Clairvoyant: $\forall i \in \mathcal{N} : \varphi_i^* = 0.0$, representing a population of agents with perfect predictive accuracy.
- Random: $\forall i \in \mathcal{N} : \varphi_i^* = 1.0$, representing a population of agents who make predictions uniformly at random.
- Divided: $\varphi_i^* = \begin{cases} 0.0 & i < \frac{|\mathcal{N}|}{2} \\ 1.0 & i \geq \frac{|\mathcal{N}|}{2} \end{cases}$, representing an agent population where the top half of agents $1, \dots, \frac{|\mathcal{N}|}{2}$ are perfect

predictors and the bottom half of agents are random predictors.

- Deciles: $\varphi_i^* = D(i) \in \{0.1, 0.2, \dots, 1.0\}$, representing a division of the agent population into discrete deciles where agent ranking determines predictive accuracy.

The two uniform population models (Clairvoyant and Random) are proof of concept population concepts intended to create scenarios where there is no connection between agent quality and prediction ability, while the two mixed population models reflect an assumption that high quality agents will be high quality predictors. We do not claim either that these are the only potentially

interesting population models, or that they reflect a realistic population of agents, but they allow us to perform a basic performance evaluation and ensure that PEERBTS functions well in stylized context where underlying agent quality is directly correlated with their predictive ability.

A.4 Expanded Lottery Share Figure

This is an extended version of Figure 2 which includes all tested agent prediction models.




Figure 4: Lottery share for all population models at $m = 3$, $e = 8.0$, and $\varphi = 0.5$.