UvA-DARE (Digital Academic Repository)

Optimal Truth-Tracking Rules for the Aggregation of Incomplete Judgments

Terzopoulou, Z.; Endriss, U.

DOI

[10.1007/978-3-030-30473-7_20](https://doi.org/10.1007/978-3-030-30473-7_20)

Publication date

2019

Document Version

Submitted manuscript

Published in

Algorithmic Game Theory

[Link to publication](#)

Citation for published version (APA):

Terzopoulou, Z., & Endriss, U. (2019). Optimal Truth-Tracking Rules for the Aggregation of Incomplete Judgments. In D. Fotakis, & E. Markakis (Eds.), *Algorithmic Game Theory: 12th International Symposium, SAGT 2019, Athens, Greece, September 30-October 3, 2019 : proceedings* (pp. 298-311). (Lecture Notes in Computer Science; Vol. 11801). Springer. https://doi.org/10.1007/978-3-030-30473-7_20

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Optimal Truth-Tracking Rules for the Aggregation of Incomplete Judgments

Zoi Terzopoulou^(✉) and Ulle Endriss

Institute for Logic, Language and Computation
University of Amsterdam
{z.terzopoulou,u.endriss}@uva.nl

Abstract. Suppose you need to determine the correct answer to a complex question that depends on two logically independent premises. You can ask several agents to each evaluate either just one of those premises (which they can do with relatively high accuracy) or both premises (in which case their need to multitask will lower their individual accuracy). We first determine the optimal rule to aggregate the individual judgments reported by the agents and then analyse their strategic incentives, depending on whether they are motivated by (i) the group tracking the truth, by (ii) maximising their own reputation, or by (iii) maximising the agreement of the group’s findings with their own private judgments. We also study the problem of deciding how many agents to ask for two judgments and how many to ask for just a single judgment.

1 Introduction

Suppose a group of agents need to collectively determine the answer to a binary question that directly depends on the evaluation of several independent criteria. A correct *yes/no* answer—both on the different criteria and on the complex question—exists, but the agents are *a priori* unaware of it. Still, the agents can reflect on the possible answers and obtain a judgment which has a certain probability of being correct. But, most importantly, different agents may assess different parts of the question under consideration. We assume that, under time restrictions and cognitive constraints, the more criteria a given agent tries to assess, the less accurate her judgments are likely to be. This decrease in accuracy might be due to *time pressure* [3, 10, 16], *multitasking attempts* [1], or *speeded reasoning* [18]. How can the agents then, as a group, maximise the probability of discovering the correct answer to the complex question they are facing?

Example 1. An academic hiring committee needs to decide whether Alice should get the advertised research job. In order to do so, the committee members (professors 1, 2, and 3) have to review two of Alice’s papers—Alice will be hired if and only if both these papers are marked as “excellent”. Due to an urgent deadline the committee is given only one day to judge the quality of Alice’s papers. After the day passes, professor 1 has spent all her time on one of the two papers, while professors 2 and 3 have looked at both, and they express the following “yes” and “no” opinions, related to whether the relevant paper is excellent:

	Paper 1	Paper 2
Professor 1	Yes	—
Professor 2	No	Yes
Professor 3	No	Yes

Assuming that professor 1 has a higher probability to be correct than professor 2 about the first paper, but also taking into account that professor 3 agrees with professor 2, what is the best way to aggregate the given judgments if the committee wants to be as accurate as possible on Alice’s evaluation? Δ

Judgment aggregation [11, 12, 15] is a formal framework for group decision making concerned with the aggregation of individual judgments about several logically interconnected propositions into one collective judgment. Along the lines of Example 1, the propositions can be separated into the *premises* (e.g., excellency of Alice’s papers) and the *conclusion* (e.g., Alice’s hiring), where the conclusion is satisfied if and only if all premises are. Given two independent premises φ and ψ and a group of agents, each of which answers specific questions regarding the premises, in this paper we consider two cases of practical interest:

- (i) *Free assignment*: Each agent chooses with some probability whether to report an opinion only on the first premise, only on the second premise, or on both.
- (ii) *Fixed assignment*: Each agent is asked (and required) to report an opinion only on the first premise, only on the second premise, or on both.

Our main goal is to achieve an aggregate judgment on the conclusion that has—in expectation—high chances to reflect the truth. Notably, under the assumption that the agents are sincere about the judgments they obtain after contemplating their appointed premises, we find that the optimal aggregation rule is always a weighted majority rule assigning each agent a weight that depends on the size of her submitted judgment. We may think of this as a *scoring rule* [17].

But a further problem arises, namely that the agents may behave strategically, trying to manipulate the collective outcome to satisfy their own preferences. We examine the three most natural cases for the preferences of an agent in our context, i.e., preferences that prioritise outcomes that are close to (i) the truth, (ii) the agent’s reported judgment, or (iii) the agent’s sincere judgment. In addition, we study how an agent’s incentives to be insincere relate to the information the agent holds about the judgments reported by her peers.

Finally, knowing in which scenarios the agents are sincere, we ask (from a mechanism-design point of view): Which fixed assignment is the most efficient one, meaning that it achieves the highest probability of producing a correct collective judgment? Our answer here depends heavily on the number of agents in the group as well as on exactly how accurate the agents are individually.

Prior work on judgment aggregation aiming at the tracking of the truth, which can be traced all the way back to the famous Condorcet Jury Theorem [8], has primarily focused on scenarios with two independent premises and one conclusion, like the one investigated in this paper. But such work has solely been concerned with the case of *complete* judgments, that is, the special case where

all agents report opinions on all propositions under consideration. Under this assumption, Bovens and Rabinowicz [4] and de Clippel and Eliaz [6] compare two famous aggregation rules (for uniform and varying individual accuracies, respectively): the *premise-based* rule (according to which the collective judgment on the conclusion follows from the majority's judgments on the premises) and the *conclusion-based* rule (which simply considers the opinion of the majority on the conclusion), concluding that most of the time the premise-based rule is superior. Strengthening this result, Hartmann and coauthors [13, 14] show that the premise-based rule is optimal across wider classes of aggregation rules too. Generalising the model further, Bozbay et al. [5] study scenarios with any number of premises and agents with incentives to manipulate the collective outcome, and design rules that are optimal truth-trackers, but again assuming complete reported judgments. Also focusing on strategic agents, Ahn and Oliveros [2] wonder “should two issues be decided jointly by a single committee or separately by different committees?” This question differs essentially from the one addressed in our work, since our model accounts for the lower accuracy of the agents who judge a greater number of issues.

This paper proceeds as follows. In Sect. 2 we present our basic model. In Sect. 3 we provide our central result about the optimal aggregation rule for truth-tracking with incomplete judgments and in Sect. 4 we conduct a game-theoretical analysis of our model. We then engage with finding the optimal fixed assignment for sincere agents in Sect. 5, and we conclude in Sect. 6.

2 The Model

Let φ and ψ be two logically independent premises and $c = (\varphi \wedge \psi)$ be the corresponding conclusion, and assume that all three propositions are associated with a correct answer “yes” or “no”, where a positive answer on the conclusion is equivalent to a positive answer on both premises. Each agent i in a group $N = \{1, \dots, n\}$ with $n \geq 2$ holds a *sincere* judgment $J_i^* \subseteq \{\varphi, \bar{\varphi}, \psi, \bar{\psi}\}$ that contains at most one formula from each pair of a premise and its negation: $\varphi \in J_i^*$ ($\bar{\varphi} \in J_i^*$) means that agent i judges φ as true (false). Clearly, agent i cannot judge the conclusion without having judged both premises, but her judgment on the conclusion would follow directly from her judgment on the two premises in case she had one. We denote by $\mathcal{J}$ the set of all admissible individual judgments. We say that two judgments J, J' *agree* on their evaluation of a proposition if they both contain either the non-negated or the negated version of that proposition.

An *aggregation rule* F is a function that maps every reported profile $\mathbf{J} = (J_1, \dots, J_n) \in \mathcal{J}^n$ of individual judgments to a set of collective judgments $F(\mathbf{J})$. F is *resolute* if $|F(\mathbf{J})| = 1$ for every profile $\mathbf{J}$. A collective outcome $J \in F(\mathbf{J})$ is a *logically consistent* set $J \subseteq \{\varphi, \bar{\varphi}, \psi, \bar{\psi}, c, \bar{c}\}$ that contains exactly one formula from each pair of a proposition and its negation (namely, it is *complete*). We write $J^\bullet \subseteq \{\varphi, \bar{\varphi}, \psi, \bar{\psi}, c, \bar{c}\}$ for the judgment that captures the correct evaluation on all three propositions.

We define N_1^φ (N_2^φ) to be the sets of agents who report a judgment on one (two) premise(s) and say “yes” (and analogously for “no”) on φ . We also define $n_1^\varphi = |N_1^\varphi|$ and $n_2^\varphi = |N_2^\varphi|$ to be the relevant cardinalities of these sets.

We denote by p the probability that agent i ’s judgment J_i^* is correct on a premise when i judges both premises and by q the relevant probability when i only judges a single premise (assuming that the probability of each agent i ’s judgment being correct on a premise φ is independent (*i*) of whether φ is true or false and (*ii*) of what i ’s judgment on premise ψ is). We assume that the probabilities p and q are the same for all agents, but the agents make their judgments independently of each other. We shall moreover suppose that all agents’ judgments are more accurate than a random guess, but not perfect, and that agents judging a single premise are strictly more accurate than those judging both premises, i.e., that $1/2 < p < q < 1$. Then, $P(\mathbf{J}^*)$ denotes the probability of the sincere profile $\mathbf{J}^*$ to arise and $P(\mathbf{J}_{-i}^* | J_i^*)$ the probability that the judgments of all agents besides i form the sincere (partial) profile $\mathbf{J}_{-i}^*$, given that i has the sincere judgment J_i^* . Formally, for a fixed assignment:

$$P(\mathbf{J}^*) = P(\varphi \text{ true}) \cdot P(\mathbf{J}^* | \varphi \text{ true}) + P(\varphi \text{ false}) \cdot P(\mathbf{J}^* | \varphi \text{ false})$$

where $P(\mathbf{J}^* | \varphi \text{ true}) = q^{n_1^\varphi} p^{n_2^\varphi} (1-q)^{n_1^\varphi} (1-p)^{n_2^\varphi}$, and similarly for $P(\mathbf{J}^* | \varphi \text{ false})$. The *accuracy* $P(F)$ of a resolute aggregation rule F is defined as:

$$P(F) = \sum_{\substack{\mathbf{J}^* \in \mathcal{J}^n \text{ s.t.} \\ F(\mathbf{J}^*) \text{ and } \mathbf{J}^* \text{ agree on } c}} P(\mathbf{J}^*)$$

3 Optimal Aggregation

We now define the (irresolute) aggregation rule F_{irr}^{opt} , such that for all profiles $\mathbf{J}$:

$$F_{irr}^{opt}(\mathbf{J}) = \underset{\substack{J \\ \text{complete} \\ \text{consistent}}}{\operatorname{argmax}} \sum_{i \in N} w_i \cdot |J \cap J_i|$$

where $w_i = \log \frac{q}{1-q}$ if $|J_i| = 1$, $w_i = \log \frac{p}{1-p}$ if $|J_i| = 2$, and $w_i = 0$ if $|J_i| = 0$. Observe that the base of the logarithm in the definition of w_i is irrelevant.

F_{irr}^{opt} functions as a weighted-majority rule on each premise separately, assigning to the agents weights according to the size of their reported judgments, and subsequently picks that evaluation of the conclusion that is consistent with the collective evaluation of the premises [17]. Then, F^{opt} is a resolute version of F_{irr}^{opt} that, if the obtained collective judgments are more than one, randomly chooses one of them for the collective outcome.

For a resolute aggregation rule F , the probability $P(F)$ depends on the probabilities $P(F \text{ correct on } \varphi)$ and $P(F \text{ correct on } \psi)$, which, for simplicity, we call P_φ and P_ψ , respectively. For the remainder of this paper we will further assume that the prior probabilities of the two premises being true or false are equal (and independent of each other). That is, $P(\varphi \text{ true}) = P(\psi \text{ true}) = 1/2$. Then:

$$\begin{aligned}
P(F) &= \frac{1}{4} [(P_\varphi P_\psi) + (P_\varphi P_\psi + (1 - P_\varphi)(1 - P_\psi) + P_\varphi(1 - P_\psi)) + (P_\varphi P_\psi + \\
&\quad (1 - P_\varphi)(1 - P_\psi) + (1 - P_\varphi)P_\psi) + (P_\varphi P_\psi + P_\varphi(1 - P_\psi) + (1 - P_\varphi)P_\psi)] \\
&= \frac{1}{2} + \frac{P_\varphi P_\psi}{2}
\end{aligned} \tag{1}$$

Given a (fixed or free) assignment, let us denote by $\mathbf{J}_F^\varphi$ ($\mathbf{J}_F^{\bar{\varphi}}$) and $\mathbf{J}_F^\psi$ ($\mathbf{J}_F^{\bar{\psi}}$) the sets of all possible profiles of reported judgments that lead to a “yes” (“no”) collective answer on φ and ψ under the rule F , respectively. Then:

$$P_\varphi = \frac{1}{2} \sum_{\mathbf{J}^* \in \mathbf{J}_F^\varphi} P(\mathbf{J}^* \mid \varphi \text{ true}) + \frac{1}{2} \sum_{\mathbf{J}^* \in \mathbf{J}_F^{\bar{\varphi}}} P(\mathbf{J}^* \mid \varphi \text{ false}) \tag{2}$$

Now, for a fixed assignment and a profile $\mathbf{J}^*$, we have that:

$$\begin{aligned}
P(\mathbf{J}^* \mid \varphi \text{ true}) &> P(\mathbf{J}^* \mid \varphi \text{ false}) && \Leftrightarrow \\
q^{n_1^\varphi} p^{n_2^\varphi} (1 - q)^{n_1^{\bar{\varphi}}} (1 - p)^{n_2^{\bar{\varphi}}} &> (1 - q)^{n_1^\varphi} (1 - p)^{n_2^\varphi} q^{n_1^{\bar{\varphi}}} p^{n_2^{\bar{\varphi}}} && \Leftrightarrow \\
n_1^\varphi \log \frac{q}{1 - q} + n_2^\varphi \log \frac{p}{1 - p} &> n_1^{\bar{\varphi}} \log \frac{q}{1 - q} + n_2^{\bar{\varphi}} \log \frac{p}{1 - p}
\end{aligned} \tag{3}$$

Analogously, we consider a free assignment where agent i makes a sincere judgment on premise φ with probability p_i^φ , on premise ψ with probability p_i^ψ , and on both premises with probability $p_i^{\varphi, \psi}$. Given a sincere profile $\mathbf{J}^*$:

$$P(\mathbf{J}^* \mid \varphi \text{ true}) = \sum_{i \in N_1^\varphi \cup N(1, \bar{\varphi})} \sum_{j \in N_2^\varphi \cup N(2, \bar{\varphi})} p_i^\varphi p_j^{\varphi, \psi} q^{n_1^\varphi} p^{n_2^\varphi} (1 - q)^{n_1^{\bar{\varphi}}} (1 - p)^{n_2^{\bar{\varphi}}}$$

Defining $P(\mathbf{J}^* \mid \varphi \text{ false})$ similarly, we have as in as in Equation 3:

$$\begin{aligned}
P(\mathbf{J}^* \mid \varphi \text{ true}) &> P(\mathbf{J}^* \mid \varphi \text{ false}) && \Leftrightarrow \\
n_1^\varphi \log \frac{q}{1 - q} + n_2^\varphi \log \frac{p}{1 - p} &> n_1^{\bar{\varphi}} \log \frac{q}{1 - q} + n_2^{\bar{\varphi}} \log \frac{p}{1 - p}
\end{aligned} \tag{4}$$

Theorem 1 states the main results of this section. The proof technique we use is a standard method in research on *maximum likelihood estimators* [9].

Theorem 1. *For any fixed (or free) assignment and sincere judgments, $F^{opt} \in \arg\max_F P(F)$. For every other aggregation rule $F' \in \arg\max_F P(F)$, F' only differs from F^{opt} on the tie-breaking part.*

Proof. For a fixed assignment, it follows from Equations 2 and 3 that P_φ (and P_ψ) will be maximal if and only if F assigns to the agents weights as in F^{opt} . Equation 1 implies that $\max_F P(F) \leq \frac{1}{2} + \frac{\max_F P_\varphi \max_F P_\psi}{2}$, so $P(F)$ is maximal if and only if F^{opt} (or a rule that only differs from F^{opt} on the tie-breaking part) is used. The proof is analogous for a free assignment. $\square$

4 Strategic Behaviour

In this section we study the incentives of the agents to report insincere judgments when the most accurate rule F^{opt} is used. We examine in detail both fixed and free assignments. An agent’s incentives to be insincere will of course depend on the type of her preferences. We analyse three natural and disjoint cases regarding these preferences, assuming that all agents have the same preference type: First, the agents may want the group to reach a correct judgment—these preferences are called *truth-oriented*. Second, the agents may want to report an opinion that is close to the collective judgment, no matter what that judgment is—these preferences are called *reputation-oriented*. Third, the agents may want the group’s judgment to agree with their own sincere judgment—these preferences are called *self-oriented*.¹

To make things formal, we employ tools from Bayesian game theory. We wish to understand when sincerity by all agents is an equilibrium:

Given that agent i holds the sincere judgment J_i^ , and given that the rest of the agents are going to be sincere no matter what judgments they have, is sincerity (i.e., reporting J_i^*) a best response of agent i ?*

We examine the *interim* and the *ex-post* case. In both cases agent i already knows her own sincere judgment, but in the former case she is ignorant about the judgments of the rest of the group (only knowing that they have to be probabilistically compatible with her own judgment), while in the latter case she is in addition fully informed about them (this can happen, for example, after some communication action has taken place).

Call $T \in \{\text{“truth”}, \text{“reputation”}, \text{“self”}\}$ the type of the agents’ preferences. Let us denote by $U_i^T((J_i, \mathbf{J}_{-i}^*), \mathbf{J}^*)$ the utility that agent i gets by reporting judgment J_i , when the sincere profile of the group is $\mathbf{J}^*$ and all other agents $j \neq i$ report their sincere judgments J_j^* . $EU_i^T((J_i, \mathbf{J}_{-i}^*), J_i^*)$ stands for the expected utility that agent i gets by reporting judgment J_i , when her sincere judgment is J_i^* and all other agents j report their sincere judgments for any possible such judgments. More precisely, we have that:

$$\begin{aligned} U_i^{truth}((J_i, \mathbf{J}_{-i}^*), \mathbf{J}^*) &= |F^{opt}(J_i, \mathbf{J}_{-i}^*) \cap J^\blacktriangle| \\ U_i^{reputation}((J_i, \mathbf{J}_{-i}^*), \mathbf{J}^*) &= |F^{opt}(J_i, \mathbf{J}_{-i}^*) \cap J_i| \\ U_i^{self}((J_i, \mathbf{J}_{-i}^*), \mathbf{J}^*) &= |F^{opt}(J_i, \mathbf{J}_{-i}^*) \cap J_i^*| \end{aligned}$$

Also, for any $T \in \{\text{“truth”}, \text{“reputation”}, \text{“self”}\}$:

¹For instance, doctors making judgments about their patients may simply care about the correctness of their collective judgment, participants of an experiment that are paid proportionally to their agreement with the group can be assumed to aim at being seen to agree with their peers, and people who like having their opinions confirmed might manipulate the group to agree with their own privately held judgment.

$$EU_i^T((J_i, \mathbf{J}_{-i}^*), J_i^*) = \sum_{\mathbf{J}_{-i}^*} U_i^T((J_i, \mathbf{J}_{-i}^*), \mathbf{J}^*) P(\mathbf{J}_{-i}^* | J_i^*)$$

We proceed with formally defining *strategyproofness* in our framework, namely the situation where all agents being sincere forms an equilibrium. Given a preference type $T \in \{\text{“truth”}, \text{“reputation”}, \text{“self”}\}$ and a (fixed or free) assignment, we say that *sincerity always gives rise to an interim equilibrium* if and only if $J_i^* \in \operatorname{argmax}_{J_i \in A_i} EU_i^T((J_i, \mathbf{J}_{-i}^*), J_i^*)$ for all agents i and sincere judgments J_i^* , where $A_i \subseteq \mathcal{J}$ is the set of all judgments that agent i is can potentially report under the given assignment. Similarly, *sincerity always gives rise to an ex-post equilibrium* if and only if the above holds, where EU_i^T is replaced by U_i^T .

Table 4 summarises our results, where “✓” stands for strategyproofness and “✗” designates the existence of a counterexample.

Assignment \ Preferences	Fixed		Free	
	<i>interim</i>	<i>ex-post</i>	<i>interim</i>	<i>ex-post</i>
truth-oriented	✓ _{Thm 4}	✓ _{Thm 4}	✓ _{Thm 4}	✓ _{Thm 4}
reputation-oriented	✓ _{Thm 5}	✗ _{Prp 6}	✓ _{Thm 7}	✗ _{Prp 6}
self-oriented	✓ _{Thm 8}	✓ _{Thm 8}	✗ _{Prp 9}	✗ _{Prp 9}

Table 1. Strategyproofness results.

Two fundamental lemmas are in order (the proofs are easy and thus omitted). First, we verify the basic intuition that when an agent holds more information about the reported judgments of the rest of the group, then her incentives to manipulate increase. Second, we stress that whenever we can find a counterexample of *ex-post* strategyproofness under fixed assignments, the same counterexample works for free assignments too.²

Lemma 2. *For any assignment and type of preferences, ex-post strategyproofness implies interim strategyproofness.*

Lemma 3. *For any type of preferences, ex-post strategyproofness under free assignments implies ex-post strategyproofness under fixed assignments.*

4.1 Truth-oriented Preferences

When all agents have truth-oriented preferences and when the rule F^{opt} is used to aggregate their reported judgments, it directly is in everyone’s best interest to be sincere—given that the rest of the group is sincere as well—irrespective of whether the assignment materialised is fixed or free and whether the agents know the judgments of their peers. Intuitively, the agents can trust that the rule F^{opt} will achieve a collective judgment that is as accurate as possible.

²The other direction does not hold. Importantly, a counterexample may go through under free but not fixed assignments because the agents have the option to manipulate by abstaining on some premise they have sincerely thought about.

Theorem 4. *For any fixed (or free) assignment and truth-oriented preferences:*

- (i) *sincerity always gives rise to an interim equilibrium*
- (ii) *sincerity always gives rise to an ex-post equilibrium*

Proof. By Lemma 2, we only need to prove case (ii). For an arbitrary sincere profile $\mathbf{J}^* = (J_i^*, \mathbf{J}_{-i}^*)$, we have that $U_i^{truth}((J_i, \mathbf{J}_{-i}^*), \mathbf{J}^*) = |F^{opt}(J_i, \mathbf{J}_{-i}^*) \cap \mathbf{J}^\blacktriangle|$, where $\mathbf{J}^\blacktriangle$ captures the true evaluation of the propositions. Now suppose, aiming for a contradiction, that there is an insincere judgment J_i of agent i such that $|F^{opt}(J_i, \mathbf{J}_{-i}^*) \cap \mathbf{J}^\blacktriangle| > |F^{opt}(J_i^*, \mathbf{J}_{-i}^*) \cap \mathbf{J}^\blacktriangle|$. This means that $F^{opt}(J_i, \mathbf{J}_{-i}^*) \neq F^{opt}(J_i^*, \mathbf{J}_{-i}^*)$. Then, for a suitable aggregation rule $F' \neq F^{opt}$, we can write $F^{opt}(J_i, \mathbf{J}_{-i}^*) = F'(J_i, \mathbf{J}_{-i}^*)$ and derive that $|F(J_i^*, \mathbf{J}_{-i}^*) \cap \mathbf{J}^\blacktriangle| > |F^{opt}(J_i^*, \mathbf{J}_{-i}^*) \cap \mathbf{J}^\blacktriangle|$. But this is impossible, because according to Theorem 1 F^{opt} has to maximise agreement with $\mathbf{J}^\blacktriangle$. Hence, it holds that $|F^{opt}(J_i, \mathbf{J}_{-i}^*) \cap \mathbf{J}^\blacktriangle| \leq |F^{opt}(J_i^*, \mathbf{J}_{-i}^*) \cap \mathbf{J}^\blacktriangle|$ for all J_i , which implies that $U_i^{truth}((J_i, \mathbf{J}_{-i}^*), \mathbf{J}^*) \leq U_i^{truth}((J_i^*, \mathbf{J}_{-i}^*), \mathbf{J}^*)$ for all J_i and concludes the proof. $\square$

4.2 Reputation-oriented Preferences

When the agents care about the positive reputation they obtain by agreeing with the collective judgment of the group, their incentives to behave insincerely heavily depend on whether they already know the judgments of their peers. Of course: if an agent knows precisely what the collective judgment of the group will be, she can simply change her reported judgment to fully match that collective judgment. On the other hand, we will see that if an agent does not know exactly what the sincere judgments of her peers are, it is more attractive for her to remain sincere (as—knowing that her sincere judgment is more accurate than random—she can reasonably expect the group to agree with her).

Theorem 5. *For any fixed assignment and reputation-oriented preferences, sincerity always gives rise to an interim equilibrium.*

Proof. Given a fixed assignment, an agent i , and a sincere judgment J_i^* , let us call P_{dis} the probability that agent i will disagree with the group on the evaluation of premise φ . Let us assume that agent i 's judgment J_i^* concerns both premises φ and ψ (the proof is analogous when J_i^* concerns only premise φ). Now, let us denote by P_g the probability that the group is collectively correct on their evaluation of φ . Recalling that $p > 1/2$ is the probability that agent i is correct on her evaluation of φ , it holds that:

$$P_{dis} \leq p(1 - P_g) + (1 - p)P_g = p + P_g(1 - 2p)$$

Now, we have that $p + P_g(1 - 2p) \leq 1/2$ if and only if $P_g \geq \frac{p-1/2}{2p-1} = 1/2$, which holds since all members of the group are more accurate than random. So, $P_{dis} \leq 1/2$, which means that it is more probable for the group's judgment to agree with J_i^* on premise φ than to disagree with it, and the same

holds for premise ψ as well. Therefore, agent i has no better option than to report her sincere judgment on the premises that are assigned to her. Formally, $EU_i^{\text{reputation}}((J_i, \mathbf{J}_{-i}^*), J_i^*)$ is maximised when $J_i = J_i^*$. $\square$

Proposition 6. *For reputation-oriented preferences, there exists a fixed (and thus a free, from Lemma 3) assignment where sincerity does not always give rise to an ex-post equilibrium.*

Proof. Consider a fixed assignment where all agents in the group are asked about both premises φ and ψ , agent i has the sincere judgment $J_i^* = \{\varphi, \psi\}$, and all other agents j have sincere judgments $J_j = \{\varphi, \bar{\psi}\}$. Agent i would increase her utility by reporting the insincere judgment $J_i = \{\varphi, \bar{\psi}\}$. $\square$

Theorem 7. *For any free assignment and reputation-oriented preferences, sincerity always gives rise to an interim equilibrium.*

Proof. Consider an arbitrary free assignment and an agent i with sincere judgment J_i^* . Since the given assignment is uncertain, agent i can potentially report any judgment set she wants, that is, $A_i = \mathcal{J}$. Suppose that her sincere judgment J_i^* has size $|J_i^*| = k \in \{1, 2\}$. First, following the same argument as that in the proof of Theorem 5, we see that agent i cannot increase her expected utility by reporting a judgment $J_i \neq J_i^*$ with $|J_i| = k$. We omit the formal details, but the intuition is clear: the group has higher probability to agree with the sincere evaluation of agent i on each premise than to disagree with it.

However, we also need to show that agent i cannot increase her expected utility by reporting a judgment $J_i \neq J_i^*$ with $|J_i| \neq k$. The case where $|J_i| > k$ is straightforward: if agent i has no information about one of the premises, the best she could do is reporting a random judgment on that premise, but this would not increase her expected utility. Thus, we need to consider the case where $|J_i| < k$, and more specifically the only interesting scenario with $|J_i^*| = 2$ and $|J_i| = 1$. Say, without loss of generality, that $J_i = \{\varphi\}$. Let us call $P_{ag,2}^\varphi$ ($P_{ag,2}^\psi$) the probability that the group will agree with agent i on her evaluation of φ (ψ) given that agent i reports her sincere judgment on both premises, and let us call $P_{ag,1}^\varphi$ the probability that the group will agree with agent i on her evaluation of φ given that she reports her a judgment only on premise φ . We have that:

$$\begin{aligned} EU_i^{\text{reputation}}((J_i^*, \mathbf{J}_{-i}^*), J_i^*) &= P_{ag,2}^\varphi + P_{ag,2}^\psi & \text{and} \\ EU_i^{\text{reputation}}((J_i, \mathbf{J}_{-i}^*), J_i^*) &= P_{ag,1}^\varphi \end{aligned}$$

Analogously to the proof of Theorem 5, it holds that $P_{ag,2}^\varphi \geq 1/2$ and $P_{ag,2}^\psi \geq 1/2$, so $P_{ag,2}^\varphi + P_{ag,2}^\psi \geq 1$. This means that $P_{ag,1}^\varphi \leq P_{ag,2}^\varphi + P_{ag,2}^\psi$ (because $P_{ag,1}^\varphi \leq 1$ is a probability value), so, it is the case that $EU_i^{\text{reputation}}((J_i, \mathbf{J}_{-i}^*), J_i^*) \leq EU_i^{\text{reputation}}((J_i^*, \mathbf{J}_{-i}^*), J_i^*)$. We conclude that agent i cannot increase her expected utility by reporting J_i instead of J_i^* . $\square$

4.3 Self-oriented Preferences

Suppose the agents would like the collective outcome to agree with their own sincere judgment. Now, having a fixed or a free assignment radically changes their strategic considerations: under fixed assignments the agents can never increase their utility by lying, while there are free assignments where insincere behaviour is profitable. The critical difference is that when the agents are free to submit judgments of variable size, they can increase the weight that the optimal aggregation rule F^{opt} will assign to their judgment on one of the two premises by avoiding to report a judgment on the other premise, thus having more opportunities to manipulate the outcome in favour of their private judgment.

Theorem 8. *For any fixed assignment and self-oriented preferences:*

- (i) *sincerity always gives rise to an interim equilibrium*
- (ii) *sincerity always gives rise to an ex-post equilibrium*

Proof. By Lemma 2, we only need to show case (ii). Given a fixed assignment, an agent can only report an insincere opinion by flipping her sincere judgment on some of the premises she is asked about. But if she did so, F^{opt} could only favour a judgment different from her own, not increasing her utility. Thus, every agent always maximises her utility by being sincere. $\square$

Proposition 9. *For self-oriented preferences, there exists a free assignment such that sincerity does not always give rise to an interim (and thus also not to an ex-post, from Lemma 2) equilibrium.*

Proof. Consider a group of three agents and a free assignment as follows: Agent 1 reports an opinion on both premises φ, ψ with probability $1/2$, and only on premise φ or only on premise ψ with probability $1/4$ and $1/4$, respectively. Agent 2 reports a judgment on both premises φ, ψ with probability 1, and agent 3 reports a judgment only on premise φ with probability 1. Suppose that agent 1's truthful judgment is $J_1^* = \{\varphi, \psi\}$. Suppose additionally that $q > \frac{p^2}{p^2 + (1-p)^2}$. In such a case, if agent 1 decides to report her sincere judgment on both premises, she will always be unable to affect the collective outcome on φ according to the rule F^{opt} , and she will obtain an outcome that agrees with her sincere judgment on ψ with probability $p(p + \frac{1-p}{2}) + (1-p)(1 - p + \frac{p}{2}) = p^2 + p + 1 < 1$. However, if agent 1 reports the insincere judgment $J_1 = \{\psi\}$ instead, she will always obtain a collective outcome on ψ that is identical to her own sincere judgment, corresponding to a higher expected utility of value 1. Thus, we can conclude that $J_1^* \notin \arg\max_{J_1 \in A_1} EU_1^{\text{self}}((J_1, J_2^*, J_3^*), J_1^*)$. $\square$

5 Optimal Fixed Assignment

Having a group of n agents, different choices for assigning agents to questions concerning the premises induce different fixed assignments, which in turn yield

the correct answer on the conclusion with different probability. In this section we are interested in finding the optimal (viz., the most accurate) such assignment.

Let us denote by $n_1 \leq \lfloor \frac{n}{2} \rfloor$ the number of agents that will be asked to report a judgment only on premise φ . For symmetry reasons, we assume that the same number of agents will be asked to report a judgment only on premise ψ , and the remaining $n - 2n_1$ agents will be asked to report a judgment on both premises. Given n_1 , $P_{\varphi, n_1}(F^{opt})$ is the probability of the aggregation rule F^{opt} producing a correct evaluation of premise φ , and since we assume that the same number of agents that will judge φ will also judge ψ , it will be the case that $P_{\varphi, n_1}(F^{opt}) = P_{\psi, n_1}(F^{opt})$. As in Sect. 3, the accuracy of F^{opt} regarding the conclusion is $P_{n_1}(F^{opt}) = \frac{1}{2} + \frac{P_{\varphi, n_1}(F^{opt})P_{\psi, n_1}(F^{opt})}{2} = \frac{1}{2} + \frac{P_{\varphi, n_1}(F^{opt})^2}{2}$. Thus, we will maximise $P_{n_1}(F^{opt})$ if and only if we maximise $P_{\varphi, n_1}(F^{opt})$:

$$\operatorname{argmax}_{0 \leq n_1 \leq \lfloor \frac{n}{2} \rfloor} P_{n_1}(F^{opt}) = \operatorname{argmax}_{0 \leq n_1 \leq \lfloor \frac{n}{2} \rfloor} P_{\varphi, n_1}(F^{opt})$$

The optimal assignment depends on the specific number of agents n , but also on the values p and q of the individual accuracy. For small groups of at most four agents, we calculate exactly what the optimal assignment is for any p and q ; for larger groups, we provide results for several indicative values of p and q .

Proposition 10. *For $n = 2$, $\operatorname{argmax}_{n_1} P_{\varphi, n_1}(F^{opt}) = 1$. Thus, when there are just two agents, it is optimal to ask each of them to evaluate one of the two premises ($n_1 = 1$) rather than asking both to evaluate both premises ($n_1 = 0$).*

Proof. For $n = 2$ we have two options: $n_1 = 0$ or $n_1 = 1$. We consider them separately. It is the case that $P_{\varphi, 0}(F^{opt}) = p^2 + \frac{1}{2}2p(1-p) = p$, while $P_{\varphi, 1}(F^{opt}) = q > p$. Thus, $\operatorname{argmax}_{n_1} P_{\varphi, n_1}(F^{opt}) = 1$. $\square$

Proposition 11. *For $n = 3$, $\operatorname{argmax}_{n_1} P_{\varphi, n_1}(F^{opt}) = 1$ if and only if $q \geq p^2(3 - 2p)$.*

Proof. For $n = 3$ we have two options: $n_1 = 0$ or $n_1 = 1$. We consider them separately. It is the case that $P_{\varphi, 0}(F^{opt}) = p^3 + \binom{3}{2}p^2(1-p) = p^2(3 - 2p)$, while $P_{\varphi, 1}(F^{opt}) = q$ (because the judgment of the agent who reports only on premise φ will always prevail over the judgment of the agent who reports on both premises). Thus, $\operatorname{argmax}_{n_1} P_{\varphi, n_1}(F^{opt}) = 1$ if and only if $q \geq p^2(3 - 2p)$. $\square$

Thus, if agents who evaluate both premises are correct 60% of the time, then in case there are three agents, you should ask two of them to focus on a single premise each if and only if their accuracy for doing so is at least 64.8%.

Proposition 12. *For $n = 4$, $\operatorname{argmax}_{n_1} P_{\varphi, n_1}(F^{opt}) = 1$ if $q < \frac{p^2}{(1-p)^2 + p^2}$ and $\operatorname{argmax}_{n_1} P_{\varphi, n_1}(F^{opt}) = 2$ otherwise.*

Proof. We provide a sketch. We have three options: $n_1 = 0$, $n_1 = 1$, or $n_1 = 2$.

$$\begin{aligned}
P_{\varphi,0}(F^{opt}) &= p^4 + \binom{4}{3} p^3(1-p) + \frac{1}{2} \binom{4}{2} p^2(1-p)^2 = p^2(3-2p) \\
P_{\varphi,1}(F^{opt}) &= \begin{cases} q(p^2 + 2p(1-p) + \frac{(1-p)^2}{2}) + (1-q)\frac{p^2}{2} & \text{if } q = \frac{p^2}{(1-p)^2+p^2} \\ q & \text{if } q > \frac{p^2}{(1-p)^2+p^2} \\ q(p^2 + 2p(1-p)) + (1-q)p^2 & \text{if } q < \frac{p^2}{(1-p)^2+p^2} \end{cases} \\
P_{\varphi,2}(F^{opt}) &= q^2 + \frac{1}{2}q(1-q) = q
\end{aligned}$$

The claim now follows after some simple algebraic manipulations, by distinguishing cases regarding the relation of q to $\frac{p^2}{(1-p)^2+p^2}$. $\square$

Now, for an arbitrary number of agents n and a number of agents n_1 who judge only premise φ (and the same for ψ), we have the following:

$$P_{n_1,\varphi}(F^{opt}) = \sum_{k=0}^{n-2n_1} \left(\sum_{\substack{\ell=0 \\ \text{s.t. } \ell \in W}}^{n_1} P(k, \ell, n, n_1, p, q) + \frac{1}{2} \sum_{\substack{\ell=0 \\ \text{s.t. } \ell \in T}}^{n_1} P(k, \ell, n, n_1, p, q) \right)$$

- k counts how many of the agents that judge both premises are right on φ .
- ℓ counts how many of the agents that judge only φ are right on φ .
- $W = \{\ell \mid \ell \log \frac{q}{1-q} + k \log \frac{p}{1-p} > (n_1 - \ell) \log \frac{q}{1-q} + (n - 2n_1 - k) \log \frac{p}{1-p}\}$.
- $T = \{\ell \mid \ell \log \frac{q}{1-q} + k \log \frac{p}{1-p} = (n_1 - \ell) \log \frac{q}{1-q} + (n - 2n_1 - k) \log \frac{p}{1-p}\}$.
- $P(k, \ell, n, n_1, p, q) = \binom{n-2n_1}{k} \binom{n_1}{\ell} p^k (1-p)^{n-2n_1-k} q^\ell (1-q)^{n_1-\ell}$.

For large groups with $n \geq 5$ it is too complex to calculate the optimal assignment analytically in all cases. We instead look at some representative values of p and q . For that purpose, we define a parameter α that intuitively captures the agents' *multitasking ability*, as follows: $\alpha = \frac{p-0.5}{q-0.5}$. Clearly, $0 < \alpha < 1$, and the smaller α is, the worse multitaskers the agents can be assumed to be.

We consider three types for the agents' multitasking ability: *good*, *average*, and *bad*, corresponding to values for α of 0.8, 0.5, and 0.2, respectively. In addition, we consider four types for the agents' accuracy on a single question: *very high*, *high*, *medium*, and *low*, corresponding to values for q of 0.9, 0.8, 0.7, and 0.6, respectively. Table 2 demonstrates our findings regarding the optimal assignment in terms of the number n_1 for these characteristic cases, for groups with at most 15 members.^{3,4} In general, we can observe that the better multitaskers the agents are, the lower the number n_1 that corresponds to the best assignment is. This verifies an elementary intuition suggesting that if the agents are good at multitasking, then it is profitable to ask many of them about both premises, while if the agents are bad at multitasking, it is more beneficial to ask them about a single premise each.

³For *any* assignment, collective accuracy converges to 1 as the size of the group grows larger. Thus, our analysis is most interesting for groups that are not very large.

⁴The calculations were performed using a computer program in R.

$n \backslash \begin{matrix} q \\ \alpha \end{matrix}$	Low			Medium			High			Very high		
	bad	avg.	good	bad	avg.	good	bad	avg.	good	bad	avg.	good
5	2	2	2	2	2	2	2	2	2	2	2	2
6	3	3	3	3	3	3	3	3	3	3	3	3
7	3	3	2	3	3	2	3	3	3	3	3	3
8	4	4	3	4	3	3	4	4	3	4	4	3
9	4	4	4	4	4	4	4	4	4	4	4	4
10	5	5	5	5	5	5	5	5	5	5	5	5
11	5	5	4	5	5	4	5	5	5	5	5	5
12	6	6	5	6	5	5	6	6	5	6	6	5
13	6	6	6	6	6	6	6	6	6	6	6	6
14	7	7	7	7	7	7	7	7	7	7	7	7
15	7	7	6	7	7	6	7	7	7	7	7	7

Table 2. Optimal assignment in terms of the number of agents who should be asked about premise φ only, for different group size, individual accuracy, multitasking ability.

6 Conclusion

We have contributed to the literature on the truth-tracking of aggregation rules by considering scenarios where the agents may not all judge the same number of issues that need to be decided by the group. Assuming that multitasking is detrimental to the agents’ accuracy, we have found what the optimal method to aggregate the judgments of the agents is, and we have analysed the incentives for strategic behaviour that the agents may exhibit in this new context.

For this first study on the topic a few simplifying assumptions have been made. First, we have assumed that all agents have the same accuracy and that there are only two premises. Our analysis can be naturally extended beyond this case, considering different accuracies and more than two premises: the optimal aggregation rule would still be a weighted majority rule, agents with truth-oriented preferences would still always be sincere, while agents who care about their reputation or their individual opinion would still find reasons to lie—but careful further work is essential here in order to clarify all relevant details. Second, we have assumed that the exact values of the agents’ accuracies are known, but this is often not true in practice. Thus, to complement our theoretical work, our results could be combined with existing experimental research that measures the accuracy of individual agents on specific application domains, ranging from human-computer interaction [1] to crowdsourcing [7].

References

1. Adler, R.F., Benbunan-Fich, R.: Juggling on a high wire: Multitasking effects on performance. *International Journal of Human-Computer Studies* **70**(2), 156–168 (2012)

2. Ahn, D.S., Oliveros, S.: The Condorcet jur(ies) theorem. *Journal of Economic Theory* **150**, 841–851 (2014)
3. Ariely, D., Zakay, D.: A timely account of the role of duration in decision making. *Acta Psychologica* **108**(2), 187–207 (2001)
4. Bovens, L., Rabinowicz, W.: Democratic answers to complex questions—an epistemic perspective. *Synthese* **150**(1), 131–153 (2006)
5. Bozbay, I., Dietrich, F., Peters, H.: Judgment aggregation in search for the truth. *Games and Economic Behavior* **87**, 571–590 (2014)
6. de Clippel, G., Eliaz, K.: Premise-based versus outcome-based information aggregation. *Games and Economic Behavior* **89**, 34–42 (2015)
7. Cohensius, G., Porat, O.B., Meir, R., Amir, O.: Efficient crowdsourcing via proxy voting. In: *Proceedings of the 7th International Workshop on Computational Social Choice (COMSOC-2018)* (2018)
8. de Condorcet, M.: *Essai sur l’application de l’analyse à la probabilité des décisions rendues à la pluralité des voix*. Imprimerie Royale (1785)
9. Dawid, A.P., Skene, A.M.: Maximum likelihood estimation of observer error-rates using the EM algorithm. *Applied Statistics* **28**(1), 20–28 (1979)
10. Edland, A., Svenson, O.: Judgment and decision making under time pressure. In: *Time Pressure and Stress in Human Judgment and Decision Making*, pp. 27–40. Springer (1993)
11. Endriss, U.: Judgment aggregation. In: Brandt, F., Conitzer, V., Endriss, U., Lang, J., Procaccia, A.D. (eds.) *Handbook of Computational Social Choice*, pp. 399–426. Cambridge University Press (2016)
12. Grossi, D., Pigozzi, G.: *Judgment Aggregation: A Primer*, Synthesis Lectures on Artificial Intelligence and Machine Learning, vol. 8. Morgan & Claypool Publishers (2014)
13. Hartmann, S., Pigozzi, G., Sprenger, J.: Reliable methods of judgement aggregation. *Journal of Logic and Computation* **20**(2), 603–617 (2010)
14. Hartmann, S., Sprenger, J.: Judgment aggregation and the problem of tracking the truth. *Synthese* **187**(1), 209–221 (2012)
15. List, C., Pettit, P.: Aggregating sets of judgments: An impossibility result. *Economics and Philosophy* **18**(1), 89–110 (2002)
16. Payne, J.W., Bettman, J.R., Johnson, E.J.: Adaptive strategy selection in decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition* **14**(3), 534 (1988)
17. Terzopoulou, Z., Endriss, U., de Haan, R.: Aggregating incomplete judgments: Axiomatisations for scoring rules. In: *Proceedings of the 7th International Workshop on Computational Social Choice (COMSOC-2018)* (2018)
18. Wilhelm, O., Schulze, R.: The relation of speeded and unspeeded reasoning with mental speed. *Intelligence* **30**(6), 537–554 (2002)