

The More the Worst-Case-Merrier: A Generalized Condorcet Jury Theorem for Belief Fusion

Jonas Karge and Sebastian Rudolph

Computational Logic Group, TU Dresden

{jonas.karge, sebastian.rudolph}@tu-dresden.de

Abstract

In multi-agent belief fusion, there is increasing interest in results and methods from social choice theory. As a theoretical cornerstone, the *Condorcet Jury Theorem* (CJT) states that given a number of equally competent, independent agents where each is more likely to guess the true out of two alternatives, the chances of determining this objective truth by majority voting increase with the number of participating agents, approaching certainty. Past generalizations of the CJT have shown that some of its underlying assumptions can be weakened. Motivated by requirements from practical belief fusion scenarios, we provide a significant further generalization that subsumes several of the previous ones. Our considered setting simultaneously allows for heterogeneous competence levels across the agents (even tolerating entirely incompetent or even malicious voters), and voting for any number of alternatives from a finite set. We derive practical lower bounds for the numbers of agents needed to give probabilistic guarantees for determining the true state through approval voting. We also demonstrate that the non-asymptotic part of the CJT fails in our setting for arbitrarily high numbers of voters.

1 Introduction

When aggregating pieces of information or judgments from several sources, one can have one of two goals in mind: One goal is to ensure a *fair* aggregation procedure that outputs a collective judgment which best reconciles the preferences held by each individual source – in such a case, “truth” emerges from the process itself in form of the aggregated judgment. Another goal presumes the existence of a *true* state of the world, independent from the consulted sources’ inclinations, and strives for a procedure that has a high probability to determine this *objective truth*, even at the risk of “unfairly” discarding the information given by some sources. In this work, it is the second, *truth tracking* goal we have in mind.

The Condorcet Jury Theorem (CJT) is a venerable and prominent result from voting theory (Marquis de Condorcet 1785). In its simplest version, it states that a group of independent, equally competent, and reliable agents voting on two alternatives where one represents the true state of the world (but the agents do not directly know which), then majority voting best tracks this true state. The CJT’s original purpose was to provide a strong justification for majority voting as method of societal decision making, as its outcome in a democratic procedure binds the rest of society (Young 1988). The

CJT provides this justification by asserting that a large group of voters performs better at reaching the objectively correct decision – should one exist – than single individuals.

Lately, with waning focus on the human nature of voters, the CJT and its extensions have been the subject of great interest and found applications in various disciplines where a voter now is seen more generally as some “abstract agent” that could be instantiated by humans, but also by automated decision-making procedures drawing information, e.g., from sensors, pre-specified knowledge, and/or inferencing methods. Notable application fields include *information science*, e.g., for relating voting to link-based ranking of webpages (Masteron and Olsson 2016), and *machine learning*, where CJT-like results can be harnessed for evaluating the importance of diversity for the quality of consensus-clustering algorithms (Jain 2018) or for guaranteeing that sets of classifiers (typically referred to as an *ensembles*), whose individual decisions are aggregated by some sort of voting mechanism, outperform individual classifiers (Dietterich 2000; Lam 2000).

But also the *knowledge representation* area has seen significant recent influence from social choice theory. In particular, voting methods have been applied to multi-agent systems and information fusion, where pieces of information coming from potentially conflicting agents have to be merged into a collective knowledge base. In this setting, the aggregation process is realized by defining belief fusion operators that apply voting rules to propositional knowledge bases with the aim of tracking the correct piece of information (Pigozzi and Hartmann 2007). Subsequently, the CJT was formalized and generalized for the belief fusion framework, demonstrating how to best track the true state of the world when knowledge bases in an incomplete information setting are merged (Everaere, Konieczny, and Marquis 2010).

In the following, we will provide a high-level overview of the CJT and its generalizations.¹ In all settings, some number of agents votes on a finite set of alternatives. Thereby, it is always assumed that the true state of the world is among the alternatives that can be voted on.

Condorcet’s Theorem. Condorcet’s original setting is *dichotomic*, that is, the set of alternatives is restricted to two and each agent has to vote for exactly one of the two. Moreover,

¹For a rigid formalization of these results, see Section 4.

a particular agent is said to be *reliable* if the probability that she votes for the true state of the world is strictly greater than one half. In a *homogenous* setting, the probability of voting for a particular alternative is the same across all agents. Additionally, the agents choices are assumed to be *independent*. That is, intuitively, knowing the choices of some of the agents does not provide any information on the choices of the other agents. Upon voting, the *majority* rule stipulates that the alternative that receives more than half of the votes wins. Note that in a dichotomic setting with an odd number of voters, the majority always produces a unique winner. From the above assumptions, Condorcet derived the following result:

Theorem 1 (Marquis de Condorcet 1785). *For odd-numbered homogenous groups of independent and reliable agents in a dichotomic voting setting, the probability that majority voting identifies the correct alternative*

- (1) *increases monotonically with the number of agents and*
- (2) *converges to 1 as the number of agents goes to infinity.*

Note that only (2), commonly referred to as the *asymptotic part* of the theorem or *Crowd Infallibility* is featured in most generalizations of the CJT (Ben-Yashar and Paroush 2000), while (1), the *non-asymptotic part*, also called *Growing Reliability*, has been demonstrated to be violated for a small number of agents when the homogeneity assumption is weakened (Owen, Grofman, and Feld 1989).

Naturally, the underlying assumptions seem quite strong: Voters are typically not equally competent. Moreover, agents often times will be presented more than two alternatives. Also, an agent might wish to abstain or to vote for more than one particular alternative at a time. Guided by these observations, various generalizations of the CJT have been established, relaxing some of the underlying assumptions. In the following, we briefly discuss some existing generalizations of the CJT that are central to our contribution.

Choosing from More Options. The first generalization we discuss is due to List and Goodin (2001), who generalized the CJT to any finite number of alternatives. As before, their setting still presumes homogeneous groups of independent agents, and each agent is supposed to pick exactly one of the alternatives. However, each agent is now seen as *reliable* if the probability that she votes for the correct alternative is strictly greater than the probability that she votes for any other alternative. Moreover, instead of majority voting, *plurality voting* is applied. That is, an alternative wins if it receives strictly more votes than any other alternative.

Theorem 2 (List and Goodin 2001). *For a homogenous group of independent and reliable agents where each agent votes for exactly one alternative from a finite set of alternatives, the probability that plurality voting identifies the correct alternative*

- (1) *is higher than that for any other alternative winning and*
- (2) *converges to 1 as the number of agents goes to infinity.*

Still, for this result it is assumed that all voters are equally reliable and that agents can vote for one alternative only.

Choosing Several Options or None. More recently, Everaere et al. (2010) generalized the CJT to settings where agents

vote by choosing any subset from a finite set of alternatives. Again, they presume a homogeneous group of independent agents, but adjust the *reliability* criterion: It requires that the probability of the correct alternative being contained in the chosen set is strictly greater than that of any other.

The appropriate voting method is then *approval voting*, where an alternative wins if it is contained in more of the agent-picked-sets than any other.

Theorem 3 (Everaere, Konieczny, and Marquis 2010). *For a homogeneous group of independent and reliable agents, the probability that approval voting on finitely many alternatives identifies the correct alternative converges to 1 as the number of agents goes to infinity.*

While Everaere et al.'s result allows a more liberal voting scheme, it still relies on the – rather artificial – assumption that all agents are equally competent.

Heterogenous Competence. One has to acknowledge that, in most natural cases, there are agents more competent than others. Likewise, it would be practical if one could give (probabilistic) correctness guarantees even for cases where the group of agents contains a certain share of incompetent (or even malicious) agents. To this end, Owen et al. (1989) have shown a generalized CJT wherein, in a dichotomic setting, the homogeneity requirement is dropped and the reliability assumption is significantly weakened. They define *reliability* on a group-level rather than individually and merely require that the probability to vote for the correct alternative, *averaged across all agents*, is strictly greater than one half.

Theorem 4 (Owen, Grofman, and Feld 1989). *For a reliable group of independent agents in a dichotomic voting setting, the probability that majority voting identifies the correct alternative converges to 1 as the number of agents goes to infinity.*

Then again, while allowing for voter heterogeneity, this result is restricted to two alternatives where agents can vote for one alternative at a time only.

Our Contribution. As central contribution of this paper, we state (Section 4) and prove (Section 5) a generalization of the CJT that subsumes all the results discussed before. It establishes that the probability of approval voting identifying the correct alternative converges to 1 as the number of agents goes to infinity, in any setting where

- (i) the set of choices to pick from is finite and fixed, but of arbitrary size,
- (ii) upon voting, the agents can choose any number of alternatives from that set, and
- (iii) we require group reliability, i.e., the average probability to vote for the true alternative is (by a fixed margin $\Delta p > 0$) greater than the average probability to vote for any other alternative.

We note that the existence of a fixed positive probability margin Δp is an implicit requirement in all mentioned versions of the CJT. All these results fail if the difference between the probabilities is allowed to approach zero as the number of voters grows (see Berend and Paroush 1998).

Further contributions of this paper are the following:

- We establish a formally rigorous unifying probabilistic framework that allows for comparing the discussed results and makes all underlying assumptions explicit (Section 4).
- We derive bounds for the minimal number of agents required to guarantee that the approval voting is successful with a probability higher than a given value P_{\min} , depending on the number m of choices and the probability margin Δp . It turns out that for different values of these parameters, different formulae provide better estimates.
- We discuss why an analogon to part (1) of the CJT cannot be established in our setting and provide corresponding counterexamples for arbitrarily large numbers of agents.

2 Motivational Scenario

To further motivate the use of the CJT for information fusion, we briefly discuss a hypothetical scenario based on an emerging technology called *Smart Dust*.

Smart Dust is a micro-electro-mechanical system (MEMS) with wireless, dust-sized components, called motes, that can carry sensors in order to provide information on their domain of application (Sharma and Sultana 2020). Possible applications range from general engineering and health to environmental monitoring (e.g., to support forecasting of natural disasters) (Haenggi 2016).

For instance, when applied to monitoring geological activity, a Smart Dust system consisting of possibly thousands of motes (Nerkar and Kumar 2016), spread across a region of interest, can be used to detect patterns known to precede major geological events such as earthquakes or landslides. The goal then is to assemble the individual motes’ findings to derive predictions of such events occurring in the near future (Sharma and Sultana 2020). Key issues in collecting the data from the Smart Dust system include the aggregation of heterogeneous information and the management of uncertainty at various levels (Agogino, Granderson, and Qiu 2002).

In the CJT setting, we can see each sensor-carrying mote as an individual agent. In a hypothetical scenario, we can imagine a manufacturer of a Smart Dust system to give certain guarantees regarding the reliability of the provided motes. This may include a certain percentage of motes malfunctioning (through production errors or as a consequence of their deployment) as well as the probabilities of a functional mote to correctly identify patterns that precede earthquakes or landslides, and to distinguish these patterns from each other and from harmless geological activity in the area where it is applied. In this scenario, the probability of a mote correctly identifying such a pattern also depends on, for instance, its location in that area. The collection of motes/agents then has heterogeneously distributed levels of competence about which only statistical guarantees can be given.

Then, the data delivered by the plethora of motes has to be aggregated. To this end, it is conceivable to apply voting methods, where each potential prediction associated with a detected pattern constitutes a particular alternative to vote on. The generalization of the CJT proven in this paper justifies this strategy in principle and ensures that – given basic reliability guarantees – the probability of making the right prediction goes against 1 with a growing number of motes.

Even more, our results providing lower bounds on this success probability depending on the number of agents, the number of alternatives, and certain quantified statistical reliability guarantees, allow us to give practical concrete estimates for the number of motes (or “kilograms of smart dust”) that need to be deployed in order to successfully track the correct geodynamic situation.

3 Preliminaries

Probability Theory. While we assume the reader to be familiar with the foundations of probability, we briefly recall some basic notions and fix notations used. Following common practice we use X (usually with annotations) to denote random variables. For *Bernoulli random variables*, where X can only take values from $\{0, 1\}$, we let $p = \mathbb{P}(X = 1)$ denote the probability for X taking the value and obtain $\mathbb{P}(X = 0) = 1 - p$, where, as usual, $\mathbb{P}(\cdot)$ is used to denote the *probability* of X taking the indicated values. Given a random variable X , we let $\mathbb{E}(X)$ denote the *expected value* of X , defined as $\sum_x X \mathbb{P}(X = x)$ whenever X ranges over finitely many values. Likewise, we define the *variance* $\mathbb{V}\text{ar}(X)$ of X as $\mathbb{E}((X - \mathbb{E}(X))^2)$. For Bernoulli distributions with parameter p , we obtain $\mathbb{E}(X) = p$ and $\mathbb{V}\text{ar}(X) = p(1 - p)$.

Voting. We assume a finite set $\mathcal{W} = \{\omega_1, \dots, \omega_m\}$ of m items referred to as *choices*. We further assume a finite set $\mathcal{A} = \{a_1, \dots, a_n\}$ of n *agents*.

Given \mathcal{A} and \mathcal{W} , one *approval voting (instance)* is represented by the relation $V \subseteq \mathcal{A} \times \mathcal{W}$ where $(a_i, \omega_j) \in V$ means that agent a_i approves choice ω_j (meaning that she considers it more favourable or likely than any of the non-approved choices). Given an approval vote V , the *score* $\#_V \omega$ of some choice $\omega \in \mathcal{W}$ is defined as the overall number of votes that ω receives, i.e.,

$$\#_V \omega = |\{a_i \in \mathcal{A}_n \mid (a_i, \omega) \in V\}|.$$

We then say some choice $\omega \in \mathcal{W}$ *wins the approval vote* V if it receives strictly higher scores (that is: strictly more votes) than any other choice, that is,

$$\#_V \omega > \max_{\omega' \in \mathcal{W} \setminus \{\omega\}} \#_V \omega'.$$

Let us remark that other types of voting schemes can be obtained from approval voting. Whenever every agent is required to vote for exactly one choice (i.e., V is required to be a function), approval voting is identical to *plurality voting*. If one additionally requires $m = 2$, one obtains *majority voting*, since in the functional setting with two choices, a choice winning the approval vote means that it must obtain strictly more than half of all votes.

4 Truth Tracking through Voting

In the truth tracking setting, the choices represent possible *world states*, with the basic assumption that in each instantiation of the voting procedure, exactly one of these states $\omega_* \in \mathcal{W}$ is the true, actual world state. The agents are assumed not to have direct access to the information about the actual world state, but they might be able to obtain indications through some epistemic process (such as sensor readings).

Formal Probabilistic Model. As usual, the described scenario is modeled as a random process. This process generates ω_* as well as V and is governed by a joint probability distribution \mathbb{P} over the Bernoulli (i.e., $\{0, 1\}$ -valued) random variables

$$\begin{array}{ccc} X_*^{\omega_1}, \dots, X_*^{\omega_m}, \\ X_1^{\omega_1}, \dots, X_1^{\omega_m}, \\ \vdots & \ddots & \vdots \\ X_n^{\omega_1}, \dots, X_n^{\omega_m}. \end{array}$$

The values taken by these random variables represent the outcome of a voting event as follows: $X_*^{\omega_j}$ is 1 if ω_j is the actual world state (i.e., $\omega_j = \omega_*$), and 0 otherwise, whereas $X_i^{\omega_j}$ is 1 if the i th agent voted for the j th world state (i.e., $(a_i, \omega_j) \in V$) and 0 otherwise. For brevity, we let $[\omega_* = \omega_j]$ denote the expression $X_*^{\omega_j} = 1 \wedge \bigwedge_{\omega \in \mathcal{W} \setminus \{\omega_j\}} X_*^\omega = 0$.

We next specify basic assumptions we make about the above joint distribution. Obviously, considering the described meaning of the joint distribution, we always require that $[\omega_j = \omega_*]$ holds for exactly one $\omega_j \in \mathcal{W}$.

One central and crucial assumption that we impose is that the agents do not influence each other in their decision whether or not to approve some choice.

Definition 1. A joint distribution satisfies agent approval independence if, conditioned on the actual world state, the decision to approve any given ω_j is made independently across all agents, i.e., for any $\omega, \omega_j \in \mathcal{W}$ and any sequence v_1, \dots, v_n of values from $\{0, 1\}$ the following holds:

$$\mathbb{P}\left(\bigwedge_{i=1}^n X_i^{\omega_j} = v_i \mid [\omega_* = \omega]\right) = \prod_{i=1}^n \mathbb{P}\left(X_i^{\omega_j} = v_i \mid [\omega_* = \omega]\right).$$

In other words, in a setting satisfying agent approval independence, the joint probability of any given approval voting outcome referring to a given world state can be obtained as the product of the corresponding marginal probabilities, when controlling for the actual world state.

Next, we introduce a formal notion that will allow us to quantify how well the group of agents as a whole is able to estimate the actual true world state.

Definition 2. A joint probability distribution satisfies Δp -group reliability for some $\Delta p > 0$, if the probability to approve the true world state, averaged across all agents, is at least by Δp higher than the averaged probability for approving any other state, i.e., for every $\omega, \omega' \in \mathcal{W}$ with $\omega \neq \omega'$ the following holds:

$$\frac{1}{n} \sum_{i=1}^n \mathbb{P}\left(X_i^\omega = 1 \mid [\omega_* = \omega]\right) \geq \Delta p + \frac{1}{n} \sum_{i=1}^n \mathbb{P}\left(X_i^{\omega'} = 1 \mid [\omega_* = \omega]\right).$$

The previous two assumptions are required to obtain our result (and any of the prior results discussed in this paper), so we refer to this setting as *I&R* (for independent and reliable).

The assumptions introduced next are optional; they allow to constrain the setting further in order to formalize existing results in a uniform framework. The first such constraint requires the conditional approval probabilities to be uniform across all agents.

Definition 3. A joint distribution satisfies homogeneity if, upon fixing the actual world state, the marginal probability to approve any given ω_j is the same for every agent, i.e., for any $\omega, \omega' \in \mathcal{W}$ and all $i, k \in \{1, \dots, n\}$ the following holds:

$$\mathbb{P}(X_i^\omega = 1 \mid [\omega_* = \omega']) = \mathbb{P}(X_k^\omega = 1 \mid [\omega_* = \omega']).$$

Another constraint realizes the plurality voting scheme by imposing that every agent must approve exactly one of the choices.

Definition 4. A joint distribution satisfies (vote) completeness if for every $i \in \{1, \dots, n\}$ the following holds:

$$\sum_{j=1}^m X_i^{\omega_j} = 1.$$

We conclude our framework by formally defining an appropriate notion characterizing the chance that the true world state is identified via an approval vote.

Definition 5. Given a family \mathcal{P} of joint probability distributions for n agents and a set \mathcal{W} of m choices, the approval vote worst-case success probability $P_{m,n}^{\text{wcs}}$ is defined by

$$\min_{\substack{\mathbb{P} \in \mathcal{P} \\ \omega \in \mathcal{W}}} \mathbb{P}\left(\bigwedge_{\omega' \in \mathcal{W} \setminus \{\omega\}} \sum_{k=1}^n X_k^{\omega'} > \sum_{k=1}^n X_k^\omega \mid [\omega_* = \omega]\right).$$

In words, for m choices and n agents, given a family \mathcal{P} of distributions adhering to a specific set of assumptions about the voting setting, $P_{m,n}^{\text{wcs}}$ expresses the guaranteed minimal probability that approval voting identifies the true world state.

Formulating CJT Versions. With the above notions in place, we can represent the previously shown versions of the CJT in a uniform way as follows, where the number $m \geq 2$ of choices and the group reliability parameter $\Delta p > 0$ are arbitrary but fixed throughout the limit process:

- Marquis de Condorcet: In any complete, homogeneous I&R setting holds $P_{2,n}^{\text{wcs}} \xrightarrow{n \rightarrow \infty} 1$.
- List and Goodin: In any complete, homogeneous I&R setting holds $P_{m,n}^{\text{wcs}} \xrightarrow{n \rightarrow \infty} 1$.
- Everaere, Konieczny, and Marquis: In any homogeneous I&R setting holds $P_{m,n}^{\text{wcs}} \xrightarrow{n \rightarrow \infty} 1$.
- Owen, Grofman, and Feld: In any complete I&R setting holds $P_{2,n}^{\text{wcs}} \xrightarrow{n \rightarrow \infty} 1$.

In the same vein, we can express our main theorem, which generalizes all of the above:

Theorem 5 (Generalized CJT). In any I&R setting with fixed $m \geq 2$ and $\Delta p > 0$ holds $P_{m,n}^{\text{wcs}} \xrightarrow{n \rightarrow \infty} 1$.

5 Proof of the Asymptotic Result

For the proof of our generalization of the (asymptotic part of) CJT, it is convenient to consider ω_* fixed (yet still unknown to the agents) and the remaining elements of \mathcal{W} appropriately renamed, i.e., we let $\mathcal{W} = \{\omega_*, \omega_1, \dots, \omega_{m-1}\}$. This is without loss of generality and allows for a more succinct formulation without conditional probabilities.

For every $\omega \in \mathcal{W}$, we let $p_1^\omega, \dots, p_n^\omega$ denote the Bernoulli parameters of the random variables $X_1^\omega, \dots, X_n^\omega$, respectively, which are independently distributed by assumption.

Then $\mathbb{E}(X_k^\omega) = p_k^\omega$ for all $k \in \{1, \dots, n\}$. We will make use of the average of the p_k^ω given by

$$\bar{p}^\omega = \frac{1}{n} \sum_{k=1}^n p_k^\omega.$$

We define for all ω , n , and k

$$\tilde{X}_k^\omega = X_k^\omega - \mathbb{E}(X_k^\omega) = X_k^\omega - p_k^\omega$$

the centralized versions of all considered distributions. Observe that $\mathbb{E}(\tilde{X}_k^\omega) = 0$ and

$$\text{Var}(\tilde{X}_k^\omega) = \text{Var}(X_k^\omega) = p_k^\omega(1 - p_k^\omega) \leq \frac{1}{4}. \quad (1)$$

Moreover, for every $\omega \in \mathcal{W}$, we define the (original and centralized) distribution for the score received by ω by

$$X^\omega = \sum_{k=1}^n X_k^\omega \quad \text{and} \quad \tilde{X}^\omega = \sum_{k=1}^n \tilde{X}_k^\omega,$$

respectively, and note that

$$\begin{aligned} X^\omega &= \sum_{k=1}^n X_k^\omega = \sum_{k=1}^n \mathbb{E}(X_k^\omega) + \sum_{k=1}^n (X_k^\omega - \mathbb{E}(X_k^\omega)) \\ &= \sum_{k=1}^n p_k^\omega + \sum_{k=1}^n (X_k^\omega - p_k^\omega) \\ &= n\bar{p}^\omega + \sum_{k=1}^n \tilde{X}_k^\omega = n\bar{p}^\omega + \tilde{X}^\omega. \end{aligned} \quad (2)$$

Furthermore, note that (due to the independence of all $X_1^\omega, \dots, X_n^\omega$)

$$\text{Var}(\tilde{X}^\omega) = \text{Var}(X^\omega) = \sum_{k=1}^n \text{Var}(X_k^\omega) \leq \frac{n}{4}. \quad (3)$$

According to the group reliability assumption, there exists some (uniform) $\Delta p > 0$ such that for every n and $\omega_\dagger \in \mathcal{W} \setminus \{\omega_*\}$ holds

$$\bar{p}^{\omega_*} \geq \Delta p + \bar{p}^{\omega_\dagger}.$$

In the following, let $\omega_\dagger \in \mathcal{W} \setminus \{\omega_*\}$ denote an arbitrary but fixed ‘‘competitor’’ of ω_* in the approval vote. The strategy used toward showing the desired result starts by establishing a lower bound for the probability of ω_* winning against ω_\dagger . Intuitively, this probability increases with growing n , since the distributions for $\frac{X^{\omega_*}}{n}$ and $\frac{X^{\omega_\dagger}}{n}$ will get concentrated more and more narrowly around \bar{p}^{ω_*} and \bar{p}^{ω_\dagger} , respectively.²

To separate the scores of ω_* and ω_\dagger , we make use of the threshold values

$$\theta^* = n(\bar{p}^{\omega_*} - \frac{\Delta p}{2}) \quad \text{as well as} \quad \theta^\dagger = n(\bar{p}^{\omega_\dagger} + \frac{\Delta p}{2})$$

and observe that $\theta^* \geq \theta^\dagger$ holds by construction, as the difference between \bar{p}^{ω_*} and \bar{p}^{ω_\dagger} is at least Δp . For the following, we will employ a popular inequality by Chebyshev.

Lemma 1 (Chebyshev 1867). *For any distribution X with finite $\mathbb{E}(X)$ and $\text{Var}(X)$ and any real number r holds*

$$\mathbb{P}(|X - \mathbb{E}(X)| > r) \leq \frac{\text{Var}(X)}{r^2}.$$

²This general observation regarding growing sums of independently distributed random variables also shows up in the popular Central Limit Theorem, but, as we will argue in Section 7, that theorem itself is not applicable in our case.

Then we can obtain a lower bound for the probability that in the course of an approval vote the correct choice ω_* receives more votes than some fixed competitor ω_\dagger as follows:

$$\mathbb{P}(X^{\omega_*} > X^{\omega_\dagger})$$

introduce fixed separating thresholds

$$> \mathbb{P}(X^{\omega_*} > \theta^* \geq \theta^\dagger > X^{\omega_\dagger})$$

worst case, even if $X^{\omega_*}, X^{\omega_\dagger}$ maximally dependent

$$\geq 1 - \mathbb{P}(X^{\omega_*} \leq \theta^*) - \mathbb{P}(\theta^\dagger \leq X^{\omega_\dagger})$$

Equation (2) and Definitions of θ^* and θ^\dagger

$$= 1 - \mathbb{P}(n\bar{p}^{\omega_*} + \tilde{X}^{\omega_*} \leq n(\bar{p}^{\omega_*} - \frac{\Delta p}{2}))$$

$$- \mathbb{P}(n(\bar{p}^{\omega_\dagger} + \frac{\Delta p}{2}) \leq n\bar{p}^{\omega_\dagger} + \tilde{X}^{\omega_\dagger})$$

$$= 1 - \mathbb{P}(\tilde{X}^{\omega_*} \leq -n\frac{\Delta p}{2}) - \mathbb{P}(n\frac{\Delta p}{2} \leq \tilde{X}^{\omega_\dagger}) \quad (4)$$

$$\geq 1 - \mathbb{P}(|\tilde{X}^{\omega_*}| \geq n\frac{\Delta p}{2}) - \mathbb{P}(n\frac{\Delta p}{2} \leq |\tilde{X}^{\omega_\dagger}|)$$

Chebyshev’s inequality (2 \times)

$$\geq 1 - \frac{\text{Var}(\tilde{X}^{\omega_*})}{(n\frac{\Delta p}{2})^2} - \frac{\text{Var}(\tilde{X}^{\omega_\dagger})}{(n\frac{\Delta p}{2})^2}$$

Equation (3) (2 \times)

$$\geq 1 - \frac{\frac{n}{4}}{(n\frac{\Delta p}{2})^2} - \frac{\frac{n}{4}}{(n\frac{\Delta p}{2})^2} = 1 - \frac{2}{n\Delta p^2} \quad (5)$$

Now, in order to obtain the probability for ω_* winning the approval vote against *all* competing $\omega_\dagger \in \mathcal{W} \setminus \{\omega_*\}$ simultaneously, we state

$$\mathbb{P}(\bigwedge_{\omega_\dagger \in \mathcal{W} \setminus \{\omega_*\}} X^{\omega_*} > X^{\omega_\dagger})$$

worst case, even if $X^{\omega_1}, X^{\omega_2}, \dots$ maximally dependent

$$\geq 1 - \sum_{i=1}^{m-1} \mathbb{P}(X^{\omega_*} \leq X^{\omega_i})$$

$$= 1 - \sum_{i=1}^{m-1} (1 - \mathbb{P}(X^{\omega_*} > X^{\omega_i})) \quad (6)$$

Equation (5)

$$> 1 - \sum_{i=1}^{m-1} (1 - (1 - \frac{2}{n\Delta p^2})) = 1 - \frac{2(m-1)}{n\Delta p^2} \quad (7)$$

As this argument holds for every choice of ω_* and is valid for every I&R joint probability distribution, this allows us to conclude our desired result

$$P_{m,n}^{\text{wcs}} \xrightarrow{n \rightarrow \infty} 1, \quad (8)$$

which finishes the proof of the main theorem.

6 Estimates for Required Number of Agents

For the given proof, not much care was given to arrive at particularly tight bounds. Still, beyond the convergence behavior in the infinite, the obtained result allows to derive concrete guarantees in the finite: Provided the average probability of the true state being approved is known to be by Δp better than that of any competitor, the number n of agents needed to guarantee a success probability of at least P_{\min} when choosing among m options is

$$n \geq \frac{2(m-1)}{\Delta p^2(1 - P_{\min})}. \quad (9)$$

That is, the number of agents needed grows linearly with the number of wrong alternatives offered, is inversely proportional to the admissible failure probability $1 - P_{\min}$, and depends inversely quadratically on the guaranteed margin between the average probabilities of right and wrong choice.

Better Bounds for High Values of P_{\min} and/or m . The guarantees established via Equation (9) are still unsatisfactory, in particular when the required success probability and/or the number of choices are high. Assuming a generous margin of $\Delta p = 0.5$, the number of required voters for $m = 11$ and $P_{\min} = 0.9$ is 800 and grows to 80,000 when requiring $P_{\min} = 0.99$ and letting $m = 101$. This raises the question of whether lower, more practical bounds can be obtained. Fortunately this is the case.

To this end, a crucial observation is that for Bernoulli random variables, their potential values are confined to the interval $[0, 1]$, so the probability of a value outside that interval being taken is 0. For the sum of (not necessarily identically distributed) independent random variables with this property, Hoeffding's inequality provides a tail estimate.

Lemma 2 (Hoeffding 1963). *Let X_1, \dots, X_n be independent random variables satisfying $\mathbb{P}(l_i \leq X_i \leq u_i) = 1$ for reals l_i, u_i . Consider the sum of these random variables, $X = \sum_{i=1}^n X_i$. Then for every real number $t > 0$ holds*

$$\mathbb{P}(X - \mathbb{E}(X) \geq t) \leq e^{-\frac{2}{\sum_{i=1}^n (u_i - l_i)^2} t^2}.$$

This allows for improving the established lower success probability bound for winning against one competitor:

$$\mathbb{P}(X^{\omega_*} > X^{\omega_{\dagger}})$$

Equation (4)

$$\geq 1 - \mathbb{P}(\tilde{X}^{\omega_*} \leq -n\frac{\Delta p}{2}) - \mathbb{P}(n\frac{\Delta p}{2} \leq \tilde{X}^{\omega_{\dagger}})$$

Hoeffding's theorem (2 \times) noting that $u_i - l_i = 1$ for all i

$$\geq 1 - e^{-\frac{2}{n}(n\frac{\Delta p}{2})^2} - e^{-\frac{2}{n}(n\frac{\Delta p}{2})^2} = 1 - 2e^{-\frac{1}{2}n\Delta p^2} \quad (10)$$

Then we obtain for the winning against all competitors:

$$\mathbb{P}(\bigwedge_{\omega_{\dagger} \in \mathcal{W} \setminus \{\omega_*\}} X^{\omega_*} > X^{\omega_{\dagger}})$$

Equation (6)

$$\geq 1 - \sum_{i=1}^{m-1} (1 - \mathbb{P}(X^{\omega_*} > X^{\omega_i}))$$

Equation (10)

$$\begin{aligned} &= 1 - \sum_{i=1}^{m-1} (1 - (1 - 2e^{-\frac{1}{2}n\Delta p^2})) \\ &= 1 - (m-1)2e^{-\frac{1}{2}n\Delta p^2} \end{aligned} \quad (11)$$

And as estimate for the number of independent agents needed to surpass a success probability of P_{\min} :

$$n \geq \frac{2}{\Delta p^2} \ln \frac{2(m-1)}{1 - P_{\min}} \quad (12)$$

Note that this significantly lowers the bounds in the cases considered above: for $\Delta p = 0.5$, $m = 11$ and $P_{\min} = 0.9$ the number of required voters is 42. For $P_{\min} = 0.99$ and $m = 101$, it grows only very moderately to 80. It is also easy to see that Equation (12) yields better estimates than Equation (9) for all possible values of Δp , m , and P_{\min} due to the fact that $\ln(x) < x$ always holds.

Better Bounds for Large Δp . The improved bound still yields unintuitively high values under certain conditions. As an informative cornercase, consider $\Delta p = 1$. This means

that the average probability to approve the true world state is by 1 higher than approving any other state. This condition can only be satisfied if every voter behaves perfectly, i.e., she always approves ω_* and never approves any competing ω_{\dagger} . Under such fortunate circumstances, it is obvious that consulting just one agent (i.e., $n = 1$) must be sufficient, no matter how high P_{\min} and m . However, upon inspection of Equation (12), even when fixing $\Delta p = 1$, we can obtain arbitrarily large values for n upon choosing high P_{\min} and/or m . Note that very large values of Δp are not just of academic interest since there might be belief fusion scenarios wherein close-to-perfect sensors are deployed.

To appropriately handle such situations, we resort to a more careful analysis and refined estimates of tail probabilities. We state Jensen's inequality, recalling that a function $f : [l, h] \rightarrow \mathbb{R}$ from some real interval $[l, h]$ to the real numbers is *concave* if for any $x_1, x_2 \in [l, h]$ and every $t \in [0, 1]$ holds

$$f(tx_1 + (1-t)x_2) \geq tf(x_1) + (1-t)f(x_2).$$

Lemma 3 (Jensen 1906). *For any concave function $f : [l, u] \rightarrow \mathbb{R}$, real values $x_1, \dots, x_n \in [l, u]$ and positive weights $w_1, \dots, w_n > 0$ holds*

$$\frac{\sum_{i=1}^n w_i f(x_i)}{\sum w_i} \leq f\left(\frac{\sum_{i=1}^n w_i x_i}{\sum w_i}\right).$$

In particular, choosing $w_i = \frac{1}{n}$ for every $i \in \{1, \dots, n\}$ yields

$$\frac{1}{n} \sum_{i=1}^n f(x_i) \leq f\left(\frac{1}{n} \sum_{i=1}^n x_i\right).$$

That is, for concave functions, applying the function to the average of a collection of numbers yields a larger value than taking the average over the corresponding collection of function values. In particular, given values p_1, \dots, p_n and letting $\bar{p} = \frac{1}{n} \sum_{i=1}^n p_i$, we can use the observation that the function $f : x \mapsto x(1-x)$ is concave to obtain

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n p_i(1-p_i) &= \frac{1}{n} \sum_{i=1}^n f(p_i) \leq f\left(\frac{1}{n} \sum_{i=1}^n p_i\right) \\ &= f(\bar{p}) = \bar{p}(1-\bar{p}). \end{aligned} \quad (13)$$

This insight allows us to obtain a better bound on the variance of the sum of the approvals for a given $\omega \in \mathcal{W}$:

$$\begin{aligned} \text{Var}(\tilde{X}^{\omega}) &= \text{Var}(X^{\omega}) \\ &= \sum_{k=1}^n \text{Var}(X_k^{\omega}) = \sum_{k=1}^n p_k^{\omega}(1-p_k^{\omega}) \\ \text{Equation (13)} \\ &\leq np^{\omega}(1-p^{\omega}). \end{aligned} \quad (14)$$

As another ingredient, we make use of the *Chebyshev-Cantelli inequality*, a less commonly known variant of Chebyshev's inequality, which allows to infer good estimates for one-sided tail bounds of probability distributions.

Lemma 4 (Chebyshev 1867, Cantelli 1928). *For any distribution X with finite $\mathbb{E}(X)$ and $\text{Var}(X)$ and any real number $\lambda > 0$ holds*

$$\mathbb{P}(X - \mathbb{E}(X) \geq \lambda) \leq \frac{\text{Var}(X)}{\text{Var}(X) + \lambda^2}.$$

Applying the above to $-X$ yields

$$\mathbb{P}(X - \mathbb{E}(X) \leq -\lambda) \leq \frac{\text{Var}(X)}{\text{Var}(X) + \lambda^2}.$$




Figure 1: Left: Lower bound for n , given Δp and P_{\min} for fixed $m = 2$. Plot cut off at $n = 10^6$, note that $n \xrightarrow{\Delta p \rightarrow 0} \infty$ and $n \xrightarrow{P_{\min} \rightarrow 1} \infty$. Right: Contour lines for $n = 10^4, 10^3, 100, 50, 20, 10, 5$, and 2 (left to right), given $m = 2$ (top) and $m = 3$ (bottom). The dotted lines separate the area for large Δp , where Equation (12) provides the better estimate.

These tools allow us to arrive at an improved lower bound for the probability of ω_* winning against one competitor:

$$\mathbb{P}(X^{\omega_*} > X^{\omega_{\dagger}})$$

Equation (4)

$$= 1 - \mathbb{P}(\tilde{X}^{\omega_*} \leq -n \frac{\Delta p}{2}) - \mathbb{P}(n \frac{\Delta p}{2} \leq \tilde{X}^{\omega_{\dagger}})$$

Chebyshev-Cantelli inequality (2x)

$$\begin{aligned} &\geq 1 - \frac{\text{Var}(\tilde{X}^{\omega_*})}{\text{Var}(\tilde{X}^{\omega_*}) + (n \frac{\Delta p}{2})^2} - \frac{\text{Var}(\tilde{X}^{\omega_{\dagger}})}{\text{Var}(\tilde{X}^{\omega_{\dagger}}) + (n \frac{\Delta p}{2})^2} \\ &= \frac{n^2 \Delta p^2}{4 \text{Var}(\tilde{X}^{\omega_*}) + n^2 \Delta p^2} + \frac{n^2 \Delta p^2}{4 \text{Var}(\tilde{X}^{\omega_{\dagger}}) + n^2 \Delta p^2} - 1 \end{aligned}$$

Equation (14)

$$\geq \frac{n^2 \Delta p^2}{4np^{\omega_*}(1-p^{\omega_*}) + n^2 \Delta p^2} + \frac{n^2 \Delta p^2}{4np^{\omega_{\dagger}}(1-p^{\omega_{\dagger}}) + n^2 \Delta p^2} - 1$$

value minimal for $p^{\omega_*} = \frac{1+\Delta p}{2}$ and $p^{\omega_{\dagger}} = \frac{1-\Delta p}{2}$, cf. appendix

$$\begin{aligned} &\geq \frac{n \Delta p^2}{(1-\Delta p^2) + n \Delta p^2} + \frac{n \Delta p^2}{(1-\Delta p^2) + n \Delta p^2} - 1 \\ &= \frac{2n \Delta p^2}{1+(n-1)\Delta p^2} - 1 = \frac{(n+1)\Delta p^2 - 1}{1+(n-1)\Delta p^2} \\ &= 1 - 2 \frac{1-\Delta p^2}{1+(n-1)\Delta p^2} \end{aligned} \quad (15)$$

This allows us to obtain a lower bound for the probability of ω_* winning the approval vote:

$$\mathbb{P}(\bigwedge_{\omega_{\dagger} \in \mathcal{W} \setminus \{\omega_*\}} X^{\omega_*} > X^{\omega_{\dagger}})$$

Equation (6)

$$\geq 1 - \sum_{i=1}^{m-1} (1 - \mathbb{P}(X^{\omega_*} > X^{\omega_i}))$$

Equation (15)

$$\begin{aligned} &\geq 1 - \sum_{i=1}^{m-1} (1 - (1 - 2 \frac{1-\Delta p^2}{1+(n-1)\Delta p^2})) \\ &= 1 - \frac{2(m-1)(1-\Delta p^2)}{1+(n-1)\Delta p^2} \end{aligned} \quad (16)$$

Consequently, we obtain the following estimate for the number of independent agents needed to surpass a given success probability of P_{\min} :

$$n \geq 1 + 2 \left(\frac{1}{\Delta p^2} - 1 \right) \left(\frac{m-1}{1-P_{\min}} \right). \quad (17)$$

Note that now we obtain $n \geq 1$ whenever $\Delta p = 1$ as desired. It can also be checked that this estimate is superior to Equation (9) for all values of Δp , m , and P_{\min} . However, between the estimates Equation (12) and Equation (17) none dominates the other for all values; therefore it is advisable to determine the minimum of the two in every case. Thus we arrive at our final bound:

Theorem 6. *In a Δp -group reliable setting with m choices, the worst case approval vote success probability is at least P_{\min} whenever the number of agents is equal or higher than*

$$\min \left(\frac{2}{\Delta p^2} \ln Q, 1 + \left(\frac{1}{\Delta p^2} - 1 \right) Q \right), \quad (18)$$

where $Q = 2 \frac{m-1}{1-P_{\min}}$ is the twofold ratio between the number of incorrect alternatives and the admissible error probability.

Figure 1 visualizes the thus established lower bounds of n depending on the parameters. While, in conformance with intuition, n goes against infinity as m grows, as P_{\min} approaches 1, or as Δp approaches 0, the asymptotic behaviour is qualitatively different: ensuring very high success probabilities or managing large numbers of choices has a much more moderate influence on the number of agents to be employed than coping with small group reliability margins.

7 Failure of the Non-Asymptotic Statement

We recall the non-asymptotic part of the original CJT stating that for any two odd natural numbers $i, j \in \mathbb{N}$ with $i < j$, the probability that ω_* wins the majority vote in a setting with i agents is strictly smaller than the probability for a setting with j agents.

It has been observed before that this proposition fails when the homogeneity assumption is dropped, even when upholding completeness and confining the choices to two (Owen, Grofman, and Feld 1989). We will now formally show that this failure is not just an artifact for small agent numbers, but that it occurs for arbitrarily large n .

Theorem 7. *Let $k \in \mathbb{N}$. Given $\mathcal{W} = \{\omega_*, \omega_\dagger\}$ as well as average approval probabilities \bar{p}^{ω_*} and $\bar{p}^{\omega_\dagger} = 1 - \bar{p}^{\omega_*}$ with $\bar{p}^{\omega_*} - \bar{p}^{\omega_\dagger} = \Delta p$ satisfying $1 > \Delta p > 0$, then there exist $i, j \in \mathbb{N}$ with $j > i > k$ and corresponding joint distributions \mathbb{P} and \mathbb{P}' for complete approval voting settings over \mathcal{W} with j and i agents, respectively, such that*

$$\mathbb{P}(X^{\omega_*} > X^{\omega_\dagger}) < \mathbb{P}'(X^{\omega_*} > X^{\omega_\dagger}).$$

In words, no matter how large the number of agents, if the concrete success probability is unknown, one can never be sure it will improve upon picking a setting with more agents.

We will show this by providing, for any given k and Δp concrete joint distributions where the success probability decreases despite an increasing number of agents. These distributions will be obtained, on the one hand, from the homogeneous setting underlying the original CJT and, on the other hand, from a scenario that we call the *experts & deniers setting*. In the following, let \bar{p}^{ω_*} and \bar{p}^{ω_\dagger} be fixed with $\bar{p}^{\omega_*} - \bar{p}^{\omega_\dagger} = \Delta p$ satisfying $1 > \Delta p > 0$. Due to the complete setting, we also know that $\bar{p}^{\omega_\dagger} = 1 - \bar{p}^{\omega_*}$ and thus $\bar{p}^{\omega_*} = \frac{1+\Delta p}{2}$ and $\bar{p}^{\omega_\dagger} = \frac{1-\Delta p}{2}$.

Definition 6. *For any $j \in \mathbb{N}$, let \mathbb{P}_j^- denote the unique joint probability distribution for the j -agent homogeneous complete setting with choices $\mathcal{W} = \{\omega_*, \omega_\dagger\}$.*

This is the probability distribution underlying Condorcet's original theorem. We obtain the following coarse estimate of an upper bound for the probability of the approval vote leading to a correct outcome:

$$\begin{aligned} \mathbb{P}_j^-(X^{\omega_*} > X^{\omega_\dagger}) &< 1 - \mathbb{P}_j^-(X^{\omega_\dagger} = j) \\ &= 1 - (\bar{p}^{\omega_\dagger})^j \\ &= 1 - \left(\frac{1-\Delta p}{2}\right)^j. \end{aligned}$$

Hence, given $\Delta p < 1$, we note that $\mathbb{P}_j^-(X^{\omega_*} > X^{\omega_\dagger})$ is strictly smaller than 1 for every j .

This homogeneous setting will now be contrasted with the *experts & deniers setting*, which in a way represents the ‘‘most inhomogeneous’’ way of realizing a joint probability distribution with the given average probabilities. In such a setting, every but possibly one agent falls into one of two categories: she can be an *expert*, who always infallibly picks the right choice, or a *denier*, who persistently chooses the incorrect option. Notably, in such a setting, the outcome of the majority vote is almost deterministic: it can take at most two different values, namely the floor and the ceiling of the expected value.

Definition 7. *For any $i \in \mathbb{N}$, we let \mathbb{P}_i^+ denote the i -agent joint probability distribution with agent-wise Bernoulli parameters $p_1^{\omega_*}, p_1^{\omega_\dagger}, \dots, p_i^{\omega_*}, p_i^{\omega_\dagger}$ defined by:*

$$p_\ell^{\omega_*} = \begin{cases} 1 & \ell \leq \lfloor i\bar{p}^{\omega_*} \rfloor \\ i\bar{p}^{\omega_*} - \lfloor i\bar{p}^{\omega_*} \rfloor & \lfloor i\bar{p}^{\omega_*} \rfloor < \ell \leq \lceil i\bar{p}^{\omega_*} \rceil \\ 0 & \lceil i\bar{p}^{\omega_*} \rceil < \ell \end{cases}$$

while $p_\ell^{\omega_\dagger} = 1 - p_\ell^{\omega_*}$. We refer to this type of distribution as the *experts & deniers setting*.

Inspecting \mathbb{P}_i^+ , we see that \bar{p}^{ω_*} and \bar{p}^{ω_\dagger} are indeed the average approval probabilities of ω_* and ω_\dagger , respectively. Also note that by construction, due to the near-deterministic nature of this setting, we have

$$\begin{aligned} \mathbb{P}_i^+(\lfloor i\bar{p}^{\omega_*} \rfloor \leq X^{\omega_*} \leq \lceil i\bar{p}^{\omega_*} \rceil) &= 1 \text{ as well as} \\ \mathbb{P}_i^+(\lfloor i\bar{p}^{\omega_\dagger} \rfloor \leq X^{\omega_\dagger} \leq \lceil i\bar{p}^{\omega_\dagger} \rceil) &= 1. \end{aligned}$$

Then, whenever we pick $i \geq \frac{2}{\Delta p}$, we obtain

$$\lceil i\bar{p}^{\omega_\dagger} \rceil < i\bar{p}^{\omega_\dagger} + 1 \leq i\bar{p}^{\omega_\dagger} + i\Delta p - 1 = i\bar{p}^{\omega_*} - 1 < \lfloor i\bar{p}^{\omega_*} \rfloor,$$

and therefore with the above

$$\mathbb{P}_i^+(X^{\omega_*} > X^{\omega_\dagger}) = 1.$$

In words, whenever the number of agents is large enough, the experts & deniers setting ensures that the true world state will be identified with a probability of 1. From these observations, the proof of Theorem 7 follows immediately by picking $j > i > \max(k, \frac{2}{\Delta p})$ and letting $\mathbb{P}' = \mathbb{P}_i^+$ as well as $\mathbb{P} = \mathbb{P}_j^-$.

As an aside, the experts & deniers setting discussed here also serves as a good case in point to demonstrate that in the non-homogeneous setting, the Central Limit Theorem (CLT, which, in fact, comes handy to show homogeneous versions of the CJT) is not applicable. This holds even when resorting to the CLT formulations of Lyapunov (1900; 1902) or Lindeberg (1922) which can cope with sums of not identically distributed random variables. Both versions rely on the sum of the variances to grow indefinitely for $n \rightarrow \infty$. This is not the case for sequences of experts & deniers settings where most of the underlying Bernoulli parameters p_i^ω are 0 or 1, leading to a variance of 0. Indeed, such sequence do not converge against the normal distribution. This is the reason why we employed tail estimates in our proofs which are more robust against extremal variances.

8 Conclusion and Future Work

Motivated by multi-agent belief fusion scenarios, this paper established a generalized Condorcet Jury Theorem that holds under remarkably weak assumptions. That is, in the generalization presented here, agents are allowed to have heterogeneously distributed competence levels as long as the group of agents is on average more likely to vote for the correct alternative than any other alternative. It was shown that this holds for any finite number of alternatives under approval voting, where every agent is allowed to vote for any number of alternatives.

Beyond establishing this result, we were able to derive estimates for the number of independent agents necessary to guarantee a prescribed minimal probability, P_{\min} , of the aggregation process being successful, given the number m of alternatives, and the minimum distance Δp between the average approval probabilities for correct and wrong alternative.

Finally, we proved an impossibility result: When weakening the homogeneity assumption by allowing for average competence levels, part (1) of Condorcet's original jury theorem fails for any arbitrarily large number of agents. That is, the probability that the voting procedure yields the correct alternative cannot be guaranteed to increase when choosing a larger team of agents.

In future work, we intend to perform experiments to estimate whether there still is a significant margin between the theoretically established guarantees we provide and the statistical behavior observed when belief fusion through voting aggregation is simulated. If a significant difference were found, this would clearly motivate the search for even better, tighter theoretical bounds.

We also aim to integrate our generalized Condorcet Jury Theorem into the formal belief fusion framework: following the argument by Everaere, Konieczny, and Marquis (2010), our generalization gives rise to a new postulate for belief merging operators, based on group-reliability. Consequently, it is worthwhile to investigate which of the well-known operators used in the belief fusion literature satisfy this new postulate. The operators found to satisfy this new postulate are capable of tracking the true state of the world under the weak assumptions underlying our generalization of the CJT.

Moreover, it would be important to investigate to what extent our results can be maintained when allowing for a certain degree of mutual influence among the voters, thereby weakening the assumption of *independence*. Indeed, there exist results that extend the CJT in that direction. Ladha (1992) generalized the CJT in the dichotomic setting with majority voting by only restricting the average coefficients of correlations. On another note, Pivato (2017) generalized the CJT to a class of voting rules for more than two alternatives such as plurality voting and approval voting allowing for significant correlations among voters. Note that both these works weaken the requirement toward agent independence but impose that all participating agents are individually reliable, i.e., they cannot handle scenarios where unreliable or even malicious agents are involved. For future work, it is thus interesting to see to what extent the result shown here can be generalized to accommodate a moderate amount of correlation among voters, along the lines of the approaches mentioned above.

As a further generalization, it seems promising to look into scenarios where agents are allowed to give more fine-grained feedback than the simple approval or non-approval of alternatives. For example, Morreau (2021) recently proved a variant of the CJT that considers the aggregation of grades into the median grade assigned to some object in order to track the correct grade: In a peer review scenario, the reviewers assign scores to a given proposal, each of which represents a grade, i.e. a gradual evaluation of the object in question. This approach being restricted to finitely many

grades, a natural generalization to consider would be to allow for real-valued evaluations of objects such as fuzzy values or probabilities. If such a generalization can be shown, one could, for instance, aggregate probabilistic statements with the aim of tracking the correct probability assessment of an event directly. Clearly, approaches along those lines will need to be compared to well-established methods of quantitative belief fusion for heterogeneous sources such as Bayesian frameworks, Dempster-Shafer theory as well as fuzzy and possibilistic approaches (Bloch et al. 2001).

Appendix

We briefly present the determination of the extremal value determination used in Equation (12). Our goal is to determine the minimal value of the following expression over all $p = \bar{p}^{\omega^*}$:

$$\begin{aligned} & \frac{n^2 \Delta p^2}{4np^{\omega^*}(1-p^{\omega^*})+n^2\Delta p^2} + \frac{n^2 \Delta p^2}{4np^{\omega^\dagger}(1-p^{\omega^\dagger})+n^2\Delta p^2} - 1 \\ &= \frac{1}{\frac{4np^{\omega^*}(1-p^{\omega^*})}{n^2\Delta p^2} + 1} + \frac{1}{\frac{4np^{\omega^\dagger}(1-p^{\omega^\dagger})}{n^2\Delta p^2} + 1} - 1 \\ & \text{substitute } u = \frac{4np^{\omega^*}(1-p^{\omega^*})}{n^2\Delta p^2} \text{ and } v = \frac{4np^{\omega^\dagger}(1-p^{\omega^\dagger})}{n^2\Delta p^2} \\ &= \frac{1}{u+1} + \frac{1}{v+1} - 1 \end{aligned}$$

We obtain the value of the $p = \bar{p}^{\omega^*}$ producing the minimum of this expression as the zeros of its first derivative. Consequently, letting $(\cdot)'$ denote $\frac{d}{dp}(\cdot)$, we find

$$\begin{aligned} 0 &= \left(\frac{1}{u+1} + \frac{1}{v+1} - 1\right)' = \left(\frac{1}{u+1}\right)' + \left(\frac{1}{v+1}\right)' \\ &= \frac{-u'}{(u+1)^2} + \frac{-v'}{(v+1)^2}. \end{aligned}$$

Multiplying with $(u+1)^2(v+1)^2$, we obtain:

$$0 = -v'(u+1)^2 - u'(v+1)^2,$$

and further

$$\begin{aligned} 0 &= (2p-1-2\Delta p)(p-p^2+(n\frac{\Delta p}{2})^2)^2 \\ &+ (2p-1)((n\frac{\Delta p}{2})^2+p-p^2-\Delta p+2p\Delta p-\Delta p^2)^2 \\ &= (2p-\Delta p-1) \cdot \\ & (2p^4-4p^3\Delta p-4p^3+4p^2\Delta p^2+6p^2\Delta p \\ & -4p^2(n\frac{\Delta p}{2})^2+2p^2-2p\Delta p^3-4p\Delta p^2 \\ & +4p\Delta p(n\frac{\Delta p}{2})^2-2p\Delta p+4p(n\frac{\Delta p}{2})^2 \\ & +\Delta p^3+\Delta p^2-2\Delta p(n\frac{\Delta p}{2})^2+2(n\frac{\Delta p}{2})^4). \end{aligned}$$

As the second factor does not produce zeros under the given assumptions, we proceed with

$$\begin{aligned} 0 &= 2p - \Delta p - 1 \\ p &= \frac{1+\Delta p}{2}, \end{aligned}$$

which yields

$$\begin{aligned} p^{\omega^*} &= p = \frac{1+\Delta p}{2} \text{ as well as} \\ p^{\omega^\dagger} &= p^{\omega^*} - \Delta p = \frac{1-\Delta p}{2}. \end{aligned}$$

Acknowledgments

Sebastian Rudolph has received funding from the European Research Council (Grant Agreement no. 771779, DeciGUT). Jonas Karge was supported by the Bundesministerium für Bildung und Forschung (BMBF, Federal Ministry of Education and Research) in the Center for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI). We are grateful for comments by Lucía Gómez Álvarez as well as the three anonymous reviewers.

References

- Agogino, A. M.; Granderson, J.; and Qiu, S. 2002. Sensor validation and fusion with distributed ‘smart dust’ motes for monitoring and enabling efficient energy use. In *Proceedings of the AAAI 2002 Spring Symposium, 2002*, 51–57.
- Ben-Yashar, R., and Paroush, J. 2000. A nonasymptotic Condorcet jury theorem. *Social Choice and Welfare* 17:189–199.
- Berend, D., and Paroush, J. 1998. When is Condorcet’s Jury Theorem valid? *Social Choice and Welfare* 15:481–488.
- Bloch, I.; Hunter, A.; Appriou, A.; Ayoun, A.; Benferhat, S.; Besnard, P.; Cholvy, L.; Cooke, R.; Cuppens, F.; Dubois, D.; Fargier, H.; Grabisch, M.; Kruse, R.; Lang, J.; Moral, S.; Prade, H.; Saffiotti, A.; Smets, P.; and Sossai, C. 2001. Fusion: General concepts and characteristics. *International Journal of Intelligent Systems* 16(10):1107–1134.
- Cantelli, F. P. 1928. Sui confini della probabilita. *Atti del Congresso Internazionale dei Matematici* 6:47 – 59.
- Chebyshev, P. L. 1867. Des valeurs moyennes. *Liouville’s Journal de Mathématiques Pures et Appliquées* 12:177–184.
- Marquis de Condorcet, M. J. A. N. C. 1785. *Essai sur l’application de l’analyse à la probabilité des décisions rendues à la pluralité des voix*. Imprimerie Royale, Paris.
- Dietterich, T. G. 2000. Ensemble methods in machine learning. In *Multiple Classifier Systems*, 1–15. Berlin, Heidelberg: Springer Berlin Heidelberg.
- Everaere, P.; Konieczny, S.; and Marquis, P. 2010. The epistemic view of belief merging: Can we track the truth? In Coelho, H.; Studer, R.; and Wooldridge, M. J., eds., *Proceedings of the 19th European Conference on Artificial Intelligence (ECAI 2010)*, volume 215 of *FAIA*, 621–626. IOS Press.
- Haenggi, M. 2016. *Smart Dust: Sensor Network Applications, Architecture, and Design*. Boca Raton: Taylor & Francis. Chapter 1.
- Hoeffding, W. 1963. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association* 58(301):13–30.
- Jain, B. J. 2018. Condorcet’s jury theorem for consensus clustering. In Trollmann, F., and Turhan, A., eds., *Proceedings of the 41st German Conference on AI (KI 2018)*, volume 11117 of *LNCS*, 155–168. Springer.
- Jensen, J. L. W. V. 1906. Sur les fonctions convexes et les inégalités entre les valeurs moyennes. *Acta Mathematica* 30:175 – 193.
- Ladha, K. K. 1992. The Condorcet jury theorem, free speech, and correlated votes. *American Journal of Political Science* 36(3):617–634.
- Lam, L. 2000. Classifier combinations: Implementations and theoretical issues. In Kittler, J., and Roli, F., eds., *Proceedings of the First International Workshop on Multiple Classifier Systems (MCS 2000)*, volume 1857 of *LNCS*, 77–86. Springer.
- Lindeberg, J. W. 1922. Eine neue Herleitung des Exponentialgesetzes in der Wahrscheinlichkeitsrechnung. *Mathematische Zeitschrift* 15:211–225.
- List, C., and Goodin, R. E. 2001. Epistemic democracy: Generalizing the Condorcet jury theorem. *Journal of Political Philosophy* 9(3):277–306.
- Lyapunov, A. 1900. Sur une proposition de la théorie des probabilités. *Bulletin de l’Académie impériale des sciences de St. Pétersbourg* 13:359–386.
- Lyapunov, A. 1902. Nouvelle forme du théorème sur la limite de probabilité. *Mémoires de l’Académie impériale des sciences de St. Pétersbourg* 12:1–24.
- Masterton, G., and Olsson, E. J. 2016. Linking as voting: how the Condorcet jury theorem in political science is relevant to webometrics. *Scientometrics* 106:945–966.
- Morreau, M. 2021. Democracy without enlightenment: A jury theorem for evaluative voting. *Journal of Political Philosophy* 29(2):188–210.
- Nerkar, M. H., and Kumar, N. 2016. Smart Dust. *International Journal of Advanced Research in Computer and Communication Engineering* 5:685–691.
- Owen, G.; Grofman, B.; and Feld, S. L. 1989. Proving a distribution-free generalization of the Condorcet Jury Theorem. *Mathematical Social Sciences* 17(1):1–16.
- Pigozzi, G., and Hartmann, S. 2007. Aggregation in multiagent systems and the problem of truth-tracking. In Duffee, E. H.; Yokoo, M.; Huhns, M. N.; and Shehory, O., eds., *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2007)*, 219–221. IFAAMAS.
- Pivato, M. 2017. Epistemic democracy with correlated voters. *Journal of Mathematical Economics* 72(C):51–69.
- Sharma, N., and Sultana, P. 2020. Earthquake monitoring and prediction system using smart dust – a novel approach. *International Journal of Scientific and Technology Research* 9(4):3103 – 3109.
- Young, H. P. 1988. Condorcet’s theory of voting. *The American Political Science Review* 82(4):1231–1244.