

Multi-Agent Opinion Pooling by Voting for Bins: Simulations and Characterization

Lea Bauer¹ and Jonas Karge²

¹ TU Dresden, Germany

lea.bauer@mailbox.tu-dresden.de

² TU Dresden, Germany

jonas.karge@tu-dresden.de

Abstract. In the context of aggregating probabilistic opinions from multiple agents facing severe uncertainty, *imprecise probabilities* are commonly utilized to represent their beliefs. *Voting for Bins* (VfB) is a novel voting method enabling agents with imprecise probabilistic beliefs to vote for sets of probability intervals, or bins. Inspired by the *Condorcet Jury Theorem*, VfB allows for the derivation of probabilistic assurances regarding the likelihood of identifying the correct alternative among a set, assuming the independence of the electorate and given estimates of the agents’ average competence levels. VfB also facilitates direct computation of the maximal number of bins, thereby determining the precision permitted in the voting process. In this work, we compare VfB’s performance, assessed by assigning an *epistemic value* to each aggregate, against standard imprecise pooling methods through multi-agent voting simulations. To the best of our knowledge, this work provides the first empirical comparison of imprecise pooling methods utilizing parameterized imprecise beliefs generated through a randomized process. Furthermore, we formally integrate VfB into the probabilistic pooling framework by examining which desirable properties, identified in the pooling literature, are satisfied by VfB.

Keywords: Opinion Pooling · Agent-based Simulation · Jury Theorem

1 Introduction

When the beliefs of multiple agents are modeled using probabilistic representations and consensus needs to be reached among these beliefs, the process is termed *probabilistic opinion pooling*. In typical real-world scenarios involving the aggregation of probabilistic beliefs, large groups of experts such as around 50 climate scientists [12] or 140 epidemiologists [16] often hold heterogeneous opinions that need to be pooled. When events with substantial uncertainty are involved, it is often times assumed that the agents’ beliefs are best represented by *imprecise probabilities*, which are intervals of probability values assigned to an event.

The aggregation of probabilistic beliefs among a group of agents and the representation of beliefs using imprecise probabilities are central topics in the realm

of multi-agent systems. Recent studies on pooling explore the interplay between direct evidential updating and probability aggregation in multi-agent systems [14], develop a Bayesian approach to probability pooling, employing copulas to capture dependencies among agents [13], or delve into consensus formation for multi-agent systems where agents’ beliefs are both vague and uncertain [4]. Additionally, researchers have demonstrated the utility of imprecise probabilities in model-checking multi-agent systems [18, 19] and in defining a decidable multi-agent logic [6]. Furthermore, there is emerging interest in imprecise pooling itself within the context of multi-agent systems [14].

Pooling imprecise probabilistic beliefs represents a relatively novel approach, and determining the optimal pooling method remains a subject of contention [17]. Recently, a new method for aggregating imprecise beliefs, known as *Voting for Bins* (VfB), has emerged [9]. Inspired by the *Condorcet Jury Theorem* (CJT) [3], a fundamental principle in voting theory, VfB strives to provide probabilistic assurances for selecting the correct option from a set of alternatives under specific assumptions. Utilizing a generalization of the CJT, VfB operates under the assumption of independent voters with varying levels of competence [10]. What distinguishes VfB is its treatment of not only agents’ beliefs as imprecise probabilities, but that also the alternatives themselves form probability intervals. By assuming that the correct probability for an event lies within one of these intervals and with an estimate of the average competence levels of the agents, VfB enables the direct computation of the maximum precision permissible in the aggregation process. However, it’s important to note that VfB is introduced primarily as a voting method rather than a direct opinion pooling function.

Our Contribution. In this work, our objective is to establish VfB as an effective imprecise pooling method. We accomplish this goal by (i) conducting comparative analyses of its performance against established imprecise pooling functions through multi-agent simulations. To the best of our knowledge, our work presents the first empirical comparison of imprecise pooling methods based on parameterized imprecise beliefs generated through a randomized process. Additionally, (ii) we assess the extent to which VfB satisfies a set of desirable properties identified for pooling functions.

2 Preliminaries

In this section, we introduce the formal framework underlying (i) the representation of imprecise probabilistic beliefs, (ii) imprecise pooling functions in general, and (iii) imprecise pooling by Voting for Bins.

Imprecise Probabilistic Beliefs. The standard approach to representing the probabilistic belief of an agent is to utilize a single probability function, \mathbb{P} , satisfying the Kolmogorov axioms which maps events to real numbers between 0 and 1, reflecting the agent’s confidence in the truth of those propositions: $\mathbb{P} : \mathcal{A} \rightarrow \mathbb{R}$. Here, \mathcal{A} represents an algebra of events over the event space Ω , defined as a

set of subsets of Ω that is closed under complement and finite unions [17]. The value assigned to an event signifies the agent’s degree of belief in that proposition [12]. However, for certain applications, it may be implausible to expect agents to maintain precise degrees of belief. Consider, for instance, the event predicting that the global sea level will rise by at least 1.5 meters by the year 2100, relative to the 2000 level. To address such events characterized by substantial uncertainty, the standard representation has been extended to encompass so-called *Imprecise Probabilities*:

Definition 1 (Imprecise Probabilities). *Imprecise probabilities are sets of probability functions [2].*

We denote a specific set of probability functions by \mathcal{P} . To represent the set of values \mathcal{P} assigns to a specific proposition more compactly, we define the *imprecise degree of belief* in an event as follows:

Definition 2 (Imprecise Degree of Belief). *An agent’s imprecise degree of belief in a proposition A is represented by a function, $\mathcal{P}(A)$, where $\mathcal{P}(A) = \{\mathbb{P}(A) : \mathbb{P} \in \mathcal{P}\}$ [1].*

Throughout this work, and as is common in the imprecise pooling literature, we assume the imprecise degree of belief to be *convex*. In other words, if the set of probability functions assigns different values to an event, all values between those are also within the imprecise degree of belief. Thus, the belief in a proposition is represented by an interval of values of the form $[a, b]$, where $a, b \in [0, 1]$, with a referred to as the *lower probability* of the interval and b as the *upper probability*.

Example 1. Suppose, an agent’s belief is represented by three probability functions that assign some event A values from the set $\{0.4, 0.6, 0.8\}$. By convexity, we may represent the agent’s imprecise degree of belief with $\mathcal{P}(A) = [0.4, 0.8]$. Thus, our agent is 40 – 80% confident that event A will occur.

Imprecise Pooling. An imprecise pooling function, denoted as \mathcal{F} , operates by taking as input a *profile* of sets of probability functions, one for each agent, represented as $(\mathcal{P}_1, \dots, \mathcal{P}_n)$, forming a set of sets of probability functions. The pooling function then maps this profile to a single set of probability functions, known as the aggregate or pool of \mathcal{F} . Typically, profiles are pooled with respect to a given proposition A . In this case, the input to \mathcal{F} consists of sets of imprecise degrees of belief, intervals of probability values, with the output being a single such interval. Numerous pooling functions have been defined in the literature. In this study, we compare the recently proposed Voting for Bins framework to the three standard pooling functions for imprecise probabilities: convex pooling, linear pooling, and logarithmic pooling. We formally introduce these three methods before subsequently presenting the Voting for Bins framework. For each pooling method, we define a finite set $\mathcal{A} = \{a_1, \dots, a_n\}$ comprising n agents.

Let H denote the *convex hull* of a set, representing the smallest set that contains the original set and all elements in between. Convex pooling is defined as follows:

Definition 3 (Convex Pooling). $\mathcal{F}(\mathcal{P}_1, \dots, \mathcal{P}_n)(A) = H(\bigcup_{i=1}^n \mathcal{P}_i(A))/[15]$.

Intuitively, for a given proposition A , convex pooling takes as input n imprecise degrees of belief and returns a single interval. The lower probability of the interval is the lowest endpoint from the input profile, while the upper probability is the greatest endpoint from the profile.

For linear pooling, the input profile is represented by specifying the lower and upper probabilities of each imprecise degree of belief. Let λ_i denote a weight assigned to the belief of each agent i [11].

Definition 4 (Linear Pooling). $\mathcal{F}([a_1, b_1], \dots, [a_n, b_n])(A) = [\sum_i \lambda_i a_i, \sum_i \lambda_i b_i]$.

In other words, linear pooling provides a weighted average of the lower and upper probabilities of the input profile.

Finally, logarithmic pooling can be understood as the weighted geometric mean of the lower and upper probabilities of the input profile [12]:

Definition 5 (Logarithmic Pooling).

$$\mathcal{F}([a_1, b_1], \dots, [a_n, b_n])(A) = \left[\frac{\prod_i a_i^{\lambda_i}}{\prod_i a_i^{\lambda_i} + \prod_i (1-a_i)^{\lambda_i}}, \frac{\prod_i b_i^{\lambda_i}}{\prod_i b_i^{\lambda_i} + \prod_i (1-b_i)^{\lambda_i}} \right].$$

Voting for Bins. The presentation of *Voting for Bins* (VfB) follows and simplifies the one of Karge (2023) [9]. VfB, being a voting method, assumes a finite set $\mathcal{W} = \{\omega_1, \dots, \omega_m\}$ of m items referred to as *alternatives*. Among these alternatives, there exists one designated as the correct alternative, denoted by ω_* . The underlying voting method employed is referred to as approval voting, wherein each agent is permitted to vote for any number of alternatives from \mathcal{W} . The alternative that garners strictly more votes than any other wins the approval vote. Since the correct alternative is unknown to the agents, each agent i possesses a certain probability $p_i^{\omega_*}$ of voting for the correct alternative, along with probabilities of voting for every incorrect alternative $p_i^{\omega_{\dagger}}$. In order to apply the Condorcet Jury Theorem (CJT), it is necessary to assume that, on average, agents are more inclined to vote for the correct alternative than for any incorrect one: Let \bar{p}^{ω_k} denote the average probability across agents to approve a given alternative k and let $\Delta p > 0$, then for every n and $\omega_{\dagger} \in \mathcal{W} \setminus \{\omega_*\}$ it must hold that

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

Furthermore, it is necessary to assume the independence of the electorate, meaning that agents neither influence one another in the voting process nor are influenced by external factors. In VfB, each alternative is interpreted as a subinterval of the unit interval of equal length, where each subinterval is referred to as a *bin*:

Definition 6 (Bin). Each $\omega_k \in \mathcal{W} = \omega_1, \dots, \omega_m$ represents a subinterval (bin) of the form $[a_k, a_{k+1})$, obtained by partitioning the unit interval ensuring that each ω_k has equal Lebesgue measure. The final subinterval is of the form $[a_m, 1]$.

Note that the Lebesgue measure is the standard method for measuring the length of an interval: For any closed $[a, b]$, open (a, b) , or half-open $(a, b]$ or $[a, b)$ interval, its Lebesgue measure is defined as the length $l = b - a$. Moreover, it's important to note that the number of bins in the voting process depends on the desired precision. As the winner of the approval vote is a single bin, the smaller its Lebesgue measure - indicating more bins - the more precise the outcome of the election. Next, we explore how agents vote for bins. Inspired by *modified supervaluationism* [7, 8], a philosophical theory to address vagueness, intuitively, each agent votes for the set of bins they are predominantly confident in:

Definition 7 (Predominant Confidence - Bins). *Let A be a proposition, and $\mathcal{P}(A) = [a, b]$ represent an agent's imprecise degree of belief in A . Given a set of bins, \mathbb{B} , we say that an agent is predominantly confident in B_j if the intersection of $\mathcal{P}(A)$ and B_j has a greater Lebesgue measure than the intersection of $\mathcal{P}(A)$ and any other bin, B_k , denoted as $l(\mathcal{P}(A) \cap B_j) \geq \max_{B_k \in \mathbb{B} \setminus B_j} l(\mathcal{P}(A) \cap B_k)$ for all $B_j, B_k \in \mathbb{B}$.*

From this, it is straightforward to define how agents vote in VfB:

Definition 8 (Voting for Bins). *We say that an agent a_i votes for an alternative ω_j if she is predominantly confident in that alternative.*

Example 2. Let there be two bins for proposition A with $B_1 = [0, 0.5]$ and $B_2 = [0.5, 1]$ and two agents with $\mathcal{P}_1(A) = [0.3, 0.9]$ and $\mathcal{P}_2(A) = [0, 1]$. We have $\mathcal{P}_1(A) \cap B_1 = [0.3, 0.5]$ and $\mathcal{P}_1(A) \cap B_2 = [0.5, 0.9]$ as well as $\mathcal{P}_2(A) \cap B_1 = [0, 0.5]$ and $\mathcal{P}_2(A) \cap B_2 = [0.5, 1]$. This yields $l(\mathcal{P}_1(A) \cap B_1) = 0.2$, $l(\mathcal{P}_1(A) \cap B_2) = 0.4$ as well as $l(\mathcal{P}_2(A) \cap B_1) = l(\mathcal{P}_2(A) \cap B_2) = 0.5$. Thus, agent 1 votes for B_2 only, whereas agent 2 votes for both bins. Hence, B_2 wins the approval vote with 2 votes.

Note that when there is a tie among alternatives with no bin having strictly more votes than any other, $\text{VfB}(\mathcal{P}_1, \dots, \mathcal{P}_n) = \emptyset$. With the direct correspondence between the alternatives in an election and subintervals of the unit interval (the bins), we can establish a bound on the maximal number of bins permitted for a given set of input parameters, as derived in Karge (2023) [9].

Theorem 1. *For n independent agents where $\Delta p \in (0, 1)$, the worst case approval vote success probability is at least P_{\min} whenever the number of alternatives is equal or lower than*


$$\max\left(\frac{(1-P_{\min})}{(2e^{-\frac{1}{2}n\Delta p^2})} + 1, \frac{(1-P_{\min})(1+(n-1)\Delta p^2)}{2(1-\Delta p^2)} + 1\right). \quad (1)$$

Finally, this translates straightforwardly into the maximal allowed precision in percent, denoted by C where we define C to be the proportion of the unit interval covered by a subinterval given by $C = \frac{100}{m}$.


Example 3. Suppose $\Delta p = 0.3$, $P_{\min} = 0.9$, and $n = 150$. Then, we may allow for 43 bins of equal Lebesgue measure. This translates to a precision of 2.23%.

3 Simulations


This section presents a comparative analysis of VfB as a pooling method against three standard pooling methods through multi-agent simulations. The performance of a pooling method is assessed by assigning an *epistemic value* to its outcome, defined as $\mathcal{V}(K, w) = \alpha \mathbb{E}(K, w) + (1 - \alpha) \mathbb{T}(K, w)$ [11]. Here, $\mathbb{E}(K, w)$ represents the truth value of belief K at state of world w , set to 1 if the pooled belief contains the correct probability for the event. The informativeness, $\mathbb{T}(K, w)$, is defined as $1 - l(\mathcal{F}(\mathcal{P}_1, \dots, \mathcal{P}_n))$, where $l(\mathcal{F}(\mathcal{P}_1, \dots, \mathcal{P}_n))$ denotes the Lebesgue measure of the aggregated belief function, reflecting the principle that smaller pooled beliefs are more informative. We balance \mathbb{E} and \mathbb{T} through α .




(a) $n = 100, \Delta p = 0.4, t = 10\%$




(b) $n = 100, \Delta p = 0.4, t = 100\%$




(c) $n = 750, \Delta p = 0.15, t = 10\%$



(d) $n = 25, \Delta p = 0.7, t = 50\%$



(e) $n = 10, \Delta p = 0.9, t = 100\%$



(f) $n = 50, \Delta p = 0.5, t = 100\%$

For our simulations, we generate imprecise probabilistic beliefs for a single proposition using input parameters $\Delta p, P_{min}, l, \alpha, n$, and t , along with a randomly generated true probability $p^* \in [0, 1]$ for the event. In the simulations, Δp represents the margin by which the agents are more likely to have their belief centered around p^* . P_{min} denotes the minimal success probability for VfB to identify p^* . n indicates the number of agents, while l denotes the number of rounds for a single parameter input, during which new imprecise beliefs are




Fig. 3: Plots (a)-(g) show the average epistemic values of the four pooling functions for $\alpha = 0.5$ and a high number of voting simulations, plot (h) considers different α -values, and (i) shows epistemic values for VfB for single voting rounds.

generated and aggregated. α represents the weight assigned to either truth or informativeness in computing the epistemic value, and t represents the percentage of the maximal number of bins, given by Theorem 1, to be constructed.

We begin by computing the number of bins, representing the desired precision, for a given set of input parameters, ensuring that p^* falls within one of these bins. Subsequently, we construct imprecise beliefs that not only correlate with Δp but also allow for heterogeneity, meaning that some beliefs may be unlikely to include p^* as long as Δp remains respected. The construction of these imprecise beliefs in our algorithm is rather complex; due to space constraints, we refer readers to the provided code repository³ for a detailed pseudocode description and the Rust implementation. Once the imprecise beliefs and bins are established, we conduct a single voting round of VfB and compute the aggregate for the same beliefs using the three standard pooling functions. Subsequently, we compute the epistemic value for each aggregate. This process is iterated 1,000 times for each parameter setting, generating new sets of imprecise beliefs in each round while maintaining p^* fixed. The average epistemic value of each pooling function over these rounds is then determined. For linear and logarithmic pooling, we employed equal weights for all beliefs across all simulations to ensure a fair comparison, although VfB could potentially be extended to weighted approval voting. The computation of optimal weights based on input beliefs for each method is left for future investigation. Additionally, to provide a visual representation of the underlying beliefs and computed pools, we present the average imprecise degree of belief across all agents and voting rounds alongside the average aggregate for all pooling functions. These averages are computed by averaging the lower and upper probabilities of all imprecise beliefs and aggregates, respectively.

Plots (a) through (g) are generated with $P_{min} = 0.9$ and $\alpha = 0.5$, while varying n , Δp , and t , aiming for large and small n as well as high and low Δp .

³ For a detailed implementation and pseudocode of our algorithm, please refer to the GitHub repository at <https://github.com/lea-bauer/multi-agent-opinion-pooling-by-voting-for-bins-simulations/tree/main>.

The plots consistently demonstrate VfB outperforming the other three pooling methods, with linear pooling achieving relatively high epistemic values for large n . In contrast, logarithmic and convex pooling yield epistemic values centered around 0.5 due to the large aggregates produced by both methods when n is large, but perform quite well for small number of agents. Plot (h) depicts the variation in epistemic values across 20 different α values spanning the unit interval, with $P_{min} = 0.8$. Plot (i) focuses solely on VfB in 10 voting rounds with $P_{min} = 0.1$, showcasing the epistemic value and average aggregate per round. As expected, when the pool corresponds to the ground truth bin, the epistemic value approaches 1, indicating a relatively precise aggregate. However, in rounds 3 and 4, a bin excluding p^* emerges victorious in the approval vote, while round 5 results in a tie, leading VfB to map to the empty set, thereby yielding an epistemic value of 0.

Despite the promising performance of VfB, it is important to highlight that it comes with three major drawbacks: (1) it necessitates additional information about the group of agents, particularly an estimate for Δp . Moreover, (2) for a given number of agents, there must be a certain minimal competence threshold satisfied, as elaborated in [9]. Without meeting this threshold, and for specific combinations of n , Δp , and P_{min} , fewer than two bins are permissible, rendering VfB inapplicable. Finally, (3) VfB assumes the agents to be independent.

4 Characterization

In the following section, we explore a set of desirable properties for pooling functions. Analogous to voting rules in the social choice literature, it has been demonstrated that not all desirable properties outlined in the literature can be simultaneously satisfied when pooling precise probabilistic beliefs [12]. This limitation also extends to pooling imprecise beliefs [17]. We adopt the presentation of a subset of central, desirable properties that can be simultaneously satisfied, as collected and generalized to imprecise pooling in the work of Quintana [15].

In this study, certain original properties have been adapted to fit into the imprecise pooling framework while preserving their essence from their original formulation in the precise setting. We proceed in a similar fashion. Initially, we acknowledge that these functions typically take as input only a profile. Considering *Voting for Bins* as a pooling method, we need to expand the function's input to include the number of bins the agents can vote for, which is determined by Δp , P_{min} , and n . We denote the set of bins by \mathbb{B} with $\mathbb{B} = \{B_1, \dots, B_m\}$ and include this in the input of the pooling function for each property presentation. Secondly, some of the original properties are not expected to hold when voting on probability intervals, but can be salvaged by slight modifications. Those properties will be highlighted by a superscript $*$.

The first property is known as the *Strong Setwise Function Property* (SSFP). It is defined over a common algebra \mathcal{A}^* for all profiles with $\mathcal{P}(\mathbb{P})$ denoting the power set of the set of all probability functions over \mathcal{A}^* [15] and where $[0, 1]^n$ is shorthand notation for taking n values from $[0, 1]$ [17]. It states that:

Definition 9 (SSFP). *There exists a function $\mathcal{G} : (\mathcal{P}([0, 1]^n), \mathbb{B}) \rightarrow \mathcal{P}([0, 1])$ such that for every event $A \in \mathcal{A}^*$, $\mathcal{F}(\mathcal{P}_1, \dots, \mathcal{P}_n, \mathbb{B})(A) = \mathcal{G}(\mathcal{P}_1(A), \dots, \mathcal{P}_n(A), \mathbb{B})$.*

Intuitively, this property states that pooling two profiles with respect to a proposition A depends only on the values these profiles assign to A . For certain pooling functions, it may occur that two profiles, identical regarding A , produce different outcomes when their values for A are pooled, simply because they differ in another event B . In such cases, no such function could exist [17]. We demonstrate that VfB satisfies SSFP.

Proof. This proof, analogous to the one in [15] for convex pooling, is straightforward. Let \mathcal{G} be VfB, and observe that VfB solely takes the probabilities for proposition A as input for a set of bins. Trivially, $\mathcal{F} = \mathcal{G}$ as \mathcal{F} equals VfB by assumption.

The next desirable property is known as the *Zero Preservation Property* (ZPP). Intuitively, it states that when every agent holds a precise belief equal to 0 for a proposition A to be true, the aggregated belief should also be equal to 0. However, ZPP does not hold for VfB, as the aggregate depends on the number and width of the bins, which in turn depend on Δp , P_{\min} , and n . For example, consider a scenario with only two bins $[0, 0.5)$ and $[0.5, 1]$. According to VfB, $[0, 0.5)$ wins the election when every agent holds a belief of 0 in that proposition, yet $[0, 0.5) \neq 0$. Although ZPP does not hold in general, we can reformulate this property to align with our voting framework.

Definition 10 (ZPP*). *For any $A \in \mathcal{A}^*$, if $\mathcal{P}_i(A) = 0$ for all $i = 1, \dots, n$, then $\mathcal{F}(\mathcal{P}_1, \dots, \mathcal{P}_n, \mathbb{B})(A) = [a, b]$ such that $0 \in [a, b]$.*

Proof. Let there be n agents, each agent holding a belief $\mathcal{P}_i(A) = 0$ for a given proposition A . Then, for any $\Delta p, P_{\min}$ and, thereby, any number of bins, there exists a leftmost bin of the form $B_1 = [0, b]$. By construction of the bins, $0 \notin B_k$ for $B_k \neq B_1$. Thus, for all i, k , $l(\mathcal{P}_i(A) \cap B_1) \geq l(\mathcal{P}_i(A) \cap B_k)$. As $0 \in B_1$ and since $n, \Delta p, P_{\min}$ as well as the width and number of bins B_k are arbitrarily chosen, VfB satisfies ZPP*.

In a similar vein, the *Unanimity Preservation Property* (UP) states that if all agents have the same imprecise belief in a proposition, their pool should exactly reflect that belief:

Definition 11 (UP). *For all $(\mathcal{P}_1, \dots, \mathcal{P}_n) \in \mathcal{P}^n$, if $\mathcal{P}_i = \mathcal{P}_j$ for all $i, j = 1, \dots, n$, then $\mathcal{F}(\mathcal{P}_1, \dots, \mathcal{P}_n) = \mathcal{P}_i$.*

By a similar argument as for ZPP, it is straightforward to show that UP does not hold in general, as the aggregate depends on the number and width of bins. Moreover, we cannot define an analogous property to ZPP*, since not every value $\mathbb{P} \in \mathcal{P}_i$ is necessarily in the winning bin. For example, consider a setting with only two bins $[0, 0.5)$, $[0.5, 1]$ and $\mathcal{P}_i = [0.4, 1]$. In that case the pool is $[0.5, 1]$, but for all elements $\mathbb{P} \in [0.4, 0.5)$ it holds that they are not in aggregate. Nonetheless,

by the definition of VfB, it can be guaranteed that the largest part of the group belief will be preserved through aggregation. The following property is known as *Confirmational Irrelevance Preservation* (IP) and states intuitively that if updating each precise probability for an event A within an agent's imprecise belief by an event B that is independent of A, then the aggregate of our pooling function should be equal to conditionalizing the aggregate of our pooling on B. For this, let $F^B(\mathcal{P}_1, \dots, \mathcal{P}_n, \mathbb{B})(A)$ denote the result from conditionalizing each member of the aggregate on B [17]. In the original formalization of IP, it is assumed that $\mathbb{P}(A \mid B) = \mathbb{P}(A)$ for all $\mathbb{P} \in \bigcup_{i=1}^n \mathcal{P}_i$ [15]. We slightly alter, but preserve the essence of, IP by assuming that $\mathbb{P}(A \mid B) = \mathbb{P}(A)$ for all conceivable probability functions, not just the ones actually held by the agents:

Definition 12 (IP*). *If $\mathbb{P}(A \mid B) = \mathbb{P}(A)$ for all \mathbb{P} , then $\mathcal{F}(\mathcal{P}_1, \dots, \mathcal{P}_n, \mathbb{B})(A) = \mathcal{F}^B(\mathcal{P}_1, \dots, \mathcal{P}_n, \mathbb{B})(A)$.*

Proof. Let $\mathbb{P}(A \mid B) = \mathbb{P}(A)$ for all \mathbb{P} . We need to show that $VfB(\mathcal{P}_1, \dots, \mathcal{P}_n, \mathbb{B})(A) = VfB^B(\mathcal{P}_1, \dots, \mathcal{P}_n, \mathbb{B})(A)$. Let \mathbb{B}^* denote the winning bin with precise probabilities $\mathbb{P}_i \in \mathbb{B}^*$. As $VfB^B(\mathcal{P}_1, \dots, \mathcal{P}_n, \mathbb{B})(A)$ results from updating each $\mathbb{P}_i \in \mathbb{B}^*$ and since $\mathbb{P}(A \mid B) = \mathbb{P}(A)$ for all \mathbb{P} , the equality holds.

Observe that the original IP does not hold as the winning bin may contain values $\mathbb{P} \notin \bigcup_{i=1}^n \mathcal{P}_i$ such that $\mathbb{P}(A \mid B) \neq \mathbb{P}(A)$. When updating those values within the winning bin, we can construct a counterexample such that $\mathcal{F}(\mathcal{P}_1, \dots, \mathcal{P}_n, \mathbb{B})(A) \neq \mathcal{F}^B(\mathcal{P}_1, \dots, \mathcal{P}_n, \mathbb{B})(A)$. The next property, *External Bayesianity* (EB), states that when a probability function is updated by a likelihood function λ (instead of an event) before aggregation, it should yield the same result as the aggregate updated by the same likelihood function.

Definition 13 (EB). *For every profile $(\mathcal{P}_1, \dots, \mathcal{P}_n)$ in the domain of F and every likelihood function λ such that $(\mathcal{P}_1^\lambda, \dots, \mathcal{P}_n^\lambda)$ remains in the domain of F ,*

$$F(\mathcal{P}_1^\lambda, \dots, \mathcal{P}_n^\lambda, \mathbb{B}) = F^\lambda(\mathcal{P}_1, \dots, \mathcal{P}_n, \mathbb{B}).$$

Here $\mathcal{P}_i^\lambda = \{\mathbb{P}_j^\lambda : \mathbb{P}_j \in \mathcal{P}_i\}$, where updating an imprecise belief by a likelihood function involves updating all precise members of that belief according to the formula $\mathbb{P}^\lambda(\cdot) = \frac{\mathbb{P}(\cdot)\lambda(\cdot)}{\sum_{\omega' \in \Omega} \mathbb{P}(\omega')\lambda(\omega')}$ provided $\sum_{\omega} \mathbb{P}(\omega)\lambda(\omega) > 0$. Updating a profile of imprecise beliefs by a likelihood function can then be defined as $F^\lambda(\mathcal{P}_1, \dots, \mathcal{P}_n, \mathbb{B}) = \{\mathbb{P}^\lambda : \mathbb{P} \in \mathcal{F}(\mathcal{P}_1, \dots, \mathcal{P}_n, \mathbb{B})\}$ [15]. However, EB does not hold for VfB.

Counterexample. Suppose, we have a single agent holding the belief $\mathcal{P}_1(A) = 0.5$ for the event $A = \omega_1$ and that $\Omega = \{\omega_1, \omega_2\}$ with $\omega_2 = \neg A$. Moreover, assume that $\lambda(\omega_1) = 0.6$ and $\lambda(\omega_2) = 0.4$. We choose Δp and P_{min} in such a way that there exists only two bins. Consider $F(\mathcal{P}_1^\lambda, \dots, \mathcal{P}_n^\lambda, \mathbb{B}) = VfB(\mathcal{P}_1^\lambda(A), \mathbb{B})$. By definition, $\mathcal{P}_1^\lambda(A)$ represents the updated belief after applying the likelihood function λ to all possible outcomes with $\mathcal{P}_1(A) = 0.6$. This yields $VfB(0.6, \mathbb{B}) = [0.5, 1]$ as the winning bin. Now, consider $\mathcal{F}^\lambda(\mathcal{P}_1, \dots, \mathcal{P}_n, \mathbb{B}) = VfB^\lambda(\mathcal{P}_1(A), \mathbb{B})$. Since $\mathcal{P}_1(A) = 0.5$ by assumption, this is equal to $VfB^\lambda(0.5, \mathbb{B})$. Recall that

$\mathcal{F}^\lambda(\mathcal{P}_1, \dots, \mathcal{P}_n, \mathbb{B}) = \{\mathbb{P}^\lambda : \mathbb{P} \in \mathcal{F}(\mathcal{P}_1, \dots, \mathcal{P}_n, \mathbb{B})\}$. Since $VfB(0.5, \mathbb{B}) = [0.5, 1]$, we have that $VfB^\lambda(\mathcal{P}_1(A), \mathbb{B}) = [0.6, 1]$ by updating each member of the winning bin. As $[0.5, 1] \neq [0.6, 1]$, we conclude that $\mathcal{F}(\mathcal{P}_1^\lambda, \dots, \mathcal{P}_n^\lambda, \mathbb{B}) \neq \mathcal{F}^\lambda(\mathcal{P}_1, \dots, \mathcal{P}_n, \mathbb{B})$.

Two additional properties merit consideration in evaluating imprecise pooling functions [15]: the Weak Setwise Function Property (WSFP) and the Marginalization Property (MP). We delve into these properties on an intuitive level, as they naturally stem from preceding properties. WSFP, akin to SSFP but incorporating an additional input event A , ensures consistency in mapping profiles of imprecise beliefs to their respective values, regardless of the associated event. Given that SSFP implies WSFP [17], it follows that VfB adheres to WSFP. Additionally, MP dictates that when the relevant algebra for aggregation is reduced by eliminating irrelevant events, the result of the aggregation remains unchanged. It has previously been shown that any imprecise pooling function satisfies MP iff it satisfies WSFP [17]. As VfB upholds WSFP, it inherently satisfies MP.

The following table offers a comparison of satisfied and unsatisfied properties among pooling methods, including VfB, Linear Pooling, Logarithmic Pooling (also known as geometric pooling [5]), and Convex Pooling. It extends the previously known characterization of the standard pooling methods [17] by VfB. Instances where VfB only fulfills modified versions of certain properties are indicated by *.

	Linear Convex Logarithmic VfB			
SSFP	✓	✓	×	✓
WSFP	✓	✓	×	✓
ZPP	✓	✓	✓	*
UP	✓	✓	✓	×
IP	×	✓	×	*
EB	×	✓	✓	×
MP	✓	✓	×	✓

5 Summary and Future Work

In this work, we conducted a comparative analysis between VfB and three standard pooling functions for imprecise probabilistic beliefs using multi-agent simulations. Assessing their performance through epistemic values, we consistently observed VfB outperforming the standard pooling functions across all considered parameter settings. Secondly, we integrated VfB into the imprecise pooling framework, revealing that its superior performance is accompanied by a trade-off—it satisfies fewer properties than linear and convex pooling, especially. Additionally, VfB necessitates an estimate of agent competencies and assumes an independent electorate. Moving forward, we plan to extend the comparative analysis to VfB with weighted approval voting. In this setting, we aim at assigning optimal weights to each agents based on their beliefs and proceed analogously for linear and logarithmic pooling.

Acknowledgments. This work is partly supported by BMBF in DAAD project 57616814 (SECAI).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Bradley, S.: How to choose among choice functions. *Episteme* **11**, 277–289 (2014)
2. Bradley, S., Steele, K.: Should subjective probabilities be sharp? *Episteme* **11**(3), 277–289 (2014)
3. Marquis de Condorcet, M.J.A.N.C.: *Essai sur l’application de l’analyse à la probabilité des décisions rendues à la pluralité des voix*. Imprimerie Royale, Paris (1785)
4. Crosscombe, M., Lawry, J.: A model of multi-agent consensus for vague and uncertain beliefs. *Adaptive Behavior* **24**(4), 249–260 (2016)
5. Dietrich, F., List, C.: Probabilistic Opinion Pooling. In: *The Oxford Handbook of Probability and Philosophy*. Oxford University Press (2016)
6. Doder, D., Savić, N., Ognjanović, Z.: A decidable multi-agent logic with iterations of upper and lower probability operators. In: Ferrarotti, F., Woltran, S. (eds.) *Foundations of Information and Knowledge Systems*. pp. 170–185 (2018)
7. Karge, J.: A modified supervaluationist framework for decision-making. *Logos & Episteme* **12**(2), 175–191 (2021)
8. Karge, J.: Second-order confidence in supervaluationism. *Journal for General Philosophy of Science* pp. 1–16 (2023)
9. Karge, J.: Voting for bins: Integrating imprecise probabilistic beliefs into the condorcet jury theorem. In: *Proceedings of the The 2nd International Workshop on Knowledge Diversity (KoDis23)* (2023)
10. Karge, J., Rudolph, S.: The More the Worst-Case-Merrier: A Generalized Condorcet Jury Theorem for Belief Fusion. In: *Proceedings of the 19th International Conference on Principles of Knowledge Representation and Reasoning*. pp. 205–214 (8 2022)
11. Konek, J.: Ip scoring rules: foundations and applications. In: *International Symposium on Imprecise Probabilities: Theories and Applications*. pp. 256–264 (2019)
12. Kriegler, E., Hall, J.W., Held, H., Dawson, R., Schellnhuber, H.J.: Imprecise probability assessment of tipping points in the climate system. *Proceedings of the National Academy of Sciences* **106**(13), 5041–5046 (2009)
13. Lawry, J., Lee, C.: Probability pooling for dependent agents in collective learning. *Artificial Intelligence* **288**, 103371 (2020)
14. Lee, C., Lawry, J., Winfield, A.: Combining opinion pooling and evidential updating for multi-agent consensus. In: *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*. pp. 347–353 (7 2018)
15. Quintana, I.O.: Radical pooling and imprecise probabilities. *Erkenntnis* pp. 1–28 (2022)
16. Recchia, G., Freeman, A.L., Spiegelhalter, D.: How well did experts and laypeople forecast the size of the covid-19 pandemic? *PloS one* **16**(5), e0250935 (2021)
17. Stewart, R.T., Quintana, I.O.: Probabilistic opinion pooling with imprecise probabilities. *Journal of Philosophical Logic* **47**, 17–45 (2018)
18. Termine, A., Antonucci, A., Primiero, G., Facchini, A.: Logic and model checking by imprecise probabilistic interpreted systems. In: Rosenfeld, A., Talmon, N. (eds.) *Multi-Agent Systems*. pp. 211–227. Springer International Publishing (2021)
19. Termine, A., Antonucci, A., Primiero, G., Facchini, A.: Imprecise probabilistic model checking for stochastic multi-agent systems. *SN Computer Science* **4**(5), 443 (2023)