Taming Dilation in Imprecise Pooling

Jonas Karge

Technische Universität Dresden, Germany
jonas.karge@tu-dresden.de

Abstract. If an agent's belief in a proposition is represented by *imprecise probabilities*, i.e. intervals of probability values, a phenomenon called *dilation* can occur, where updating the agent's belief with a new observation can only widen the probability interval, thus making the agent more uncertain, regardless of the observation acquired. Similar to standard updating, dilation can also occur in the context of *imprecise opinion pooling*, where the imprecise beliefs of multiple agents are aggregated. In this work, we provide the first formal investigation of dilation and its counterpart, contraction, in the context of imprecise opinion pooling. To this end, we use a recently defined voting rule, *Voting for Bins* (VfB), as a means to handle dilation and contraction, consistent with intuitions about the quality of additional opinions. VfB, inspired by the *Condorcet Jury Theorem* (CJT), is extended to account for correlation by an opinion leader. This model is further generalized to account for average correlation.

Keywords: Multi-Agent Opinion Pooling · Dilation · Jury Theorem

1 Introduction

Combining the probabilistic beliefs of multiple agents regarding a particular event is commonly known as *probabilistic opinion pooling*. In contexts where events entail significant uncertainty, such as forecasting the occurrence of critical climate tipping points [14], it is frequently presumed that the agents' beliefs are most accurately captured by *imprecise probabilities*. These imprecise probabilities denote intervals of probability values assigned to an event.

Aggregating probabilistic beliefs across a group of agents, and representing beliefs using imprecise probabilities, are central topics in the field of multi-agent systems. Recent studies on pooling explore the interplay between direct evidential updating and probability aggregation in multi-agent systems [16], develop a Bayesian approach to probability pooling, employing copulas to capture dependencies among agents [15], or delve into consensus formation for multi-agent systems where agents' beliefs are both vague and uncertain [6]. Additionally, researchers have demonstrated the utility of imprecise probabilities in model-checking multi-agent systems [23] and in defining a decidable multi-agent logic [7]. Furthermore, there is emerging interest in imprecise pooling itself within the context of multi-agent systems [16].

Arguably, the most significant concern associated with imprecise probabilities, also referred to as "the specter [that] is haunting the theory of imprecise probabilities [is] the specter of dilation" [21]. Intuitively, dilation occurs when the acquisition of a new observation unavoidably leads to an increase in uncertainty. Within the framework of imprecise probabilities, this implies that the interval representing an agent's belief necessarily expands when updating that belief with a new observation. This phenomenon stands in stark contrast to a fundamental principle in precise probability theory, namely, that it is always advantageous for an agent to seek free evidence [4, 10]. Departing from this principle, dilation is sometimes regarded as disconcerting [11] and serves as grounds for general objections against the imprecise framework as a whole [24].

While dilation has predominantly been examined within the realm of updating imprecise beliefs held by individual agents, it can also manifest in the context of imprecise opinion pooling. In this regard, it has been noted, albeit not explicitly termed as "dilation", that when incorporating the beliefs of additional agents into an existing group, the resulting aggregate—represented by a probability interval—may only broaden, thereby increasing uncertainty, under certain pooling functions [17].

To deal with dilation in this context, we resort to a recently developed voting method inspired by the *Condorcet Jury Theorem* and known as *Voting for Bins* [12], which operates on imprecise probabilities and can serve as an imprecise pooling function. Notably, it distinguishes itself through its close correlation between the overall quality of a group's expertise and the precision, or width, of the aggregated belief.

Contribution. In this paper, our primary objective is twofold: Firstly, we develop the first formal treatment of dilation within the realm of imprecise probabilistic pooling. Secondly, we harness the Voting for Bins framework to tame dilation in that we not only delineate precise conditions under which dilation is anticipated to occur but also derive a concrete bound that enables the computation of the extent of dilation based on estimations of the group's expertise.

To achieve this aim, we extend an existing generalization of the Condorcet Jury Theorem, which models dependence among agents by influence through an *Opinion Leader* (OL), to encompass *average OL influence.*

2 Preliminaries

In this section, we introduce the basic terminology necessary to treat dilation in our framework, and describe the assumptions necessary for our analysis. Because we study dilation within the framework of *Voting for Bins*, an imprecise pooling method rooted in the Condorcet Jury Theorem (CJT), our discussion requires incorporating elements from five different frameworks: (i) imprecise beliefs, (ii) imprecise opinion pooling, (iii) the CJT, (iv) Voting for Bins, and (v) dilation.

Imprecise Beliefs. The conventional approach to representing an agent's probabilistic belief involves employing a single probability function, denoted as $\mathbb{P}$, which adheres to the Kolmogorov axioms. This function 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, $\mathcal{A}$ signifies 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 [22]. The value assigned to an event represents the agent's degree of belief in that proposition [14].

However, for certain applications, it may be unrealistic to expect agents to maintain precise degrees of belief. Consider, for instance, the event predicting a global sea level rise of at least 1.5 meters by the year 2100 relative to the 2000 level. To address such events characterized by significant uncertainty, the standard representation has been extended to encompass what are known as *Imprecise Probabilities*:

Definition 1 (Imprecise Probabilities). *Imprecise probabilities are sets of probability functions [3].*

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 $\mathcal{P}(A) = \{\mathbb{P}(A) : \mathbb{P} \in \mathcal{P}\}$ [2].*

Throughout this work, and as is common in the imprecise pooling literature, we presume the imprecise degree of belief to be *convex*. In essence, this implies that if the set of probability functions assigns differing values to an event, all values between those are also within the imprecise degree of belief. Consequently, the belief in a proposition is denoted by an interval of values of the form $[a, b]$, where $a, b \in [0, 1]$, with a denoted as the *lower probability* of the interval and b as the *upper probability*.

Imprecise Pooling. An imprecise pooling function, denoted as $\mathcal{F}$, operates by taking as input a *profile* comprising sets of probability functions, wherein each agent is represented by a set of probability functions, denoted as $(\mathcal{P}_1, ..., \mathcal{P}_n)$. These sets collectively form a set of sets of probability functions. Subsequently, the pooling function maps this profile to a single set of probability functions, known as the aggregate or pool of $\mathcal{F}$. When profiles are pooled event-wise, i.e. with respect to a given proposition A, the input to $\mathcal{F}$ consists of sets of imprecise degrees of belief, i.e., intervals of probability values, with the output being a single such interval.

A plethora of pooling functions have been defined in the literature. In this work, we exclusively introduce *convex pooling* and *pooling by intersection* from among the standard pooling methods. These two methods will serve as our corner cases for our analysis of dilation. Convex pooling, in particular, has garnered recent support as a superior pooling method compared to alternative approaches, either due to its adherence to more desirable pooling properties [19] or its ability to prevent group members from making regretful bets [8]. To define these two pooling functions, and throughout this paper, we define a finite set $\mathcal{N} = \{a_1, \ldots, a_n\}$ consisting of 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 [19]:

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

In essence, for a given proposition A, convex pooling takes as input n imprecise degrees of belief and yields a single interval. The lower probability of the interval corresponds to the smallest endpoint from the input profile, while the upper probability corresponds to the greatest endpoint from the profile. *Pooling by Intersection*, on the other hand, produces as output interval precisely the set of values that exists in all input beliefs, which could potentially be empty and is defined as [19]:

Definition 4 (Pooling by Intersection). $\mathcal{F}(\mathcal{P}_1, ..., \mathcal{P}_n)(A) = \bigcap_{i=1}^{n} \mathcal{P}_i(A)$.

The Condorcet Jury Theorem. As a foundational theorem in voting theory, the *Condorcet Jury Theorem* (CJT) offers probabilistic assurances for determining the presumed correct alternative among a set of alternatives under specific conditions. Originally, the CJT assumes agents to be equally competent *(homogeneity)*, to be more likely to vote for the correct alternative than for a competitor *(reliability)*, to not influence one another, or to be influenced by an external factor, in the voting process *(independence)*, and to choose exactly one *(completeness)* from two alternatives only *(dichotomy)* under majority voting. With that, the classical CJT [5] states the following.

Theorem 1. *For odd-numbered, homogeneous 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.*

Since most real-world applications cannot guarantee adherence to these ideal conditions, it is essential in CJT research to seek generalizations that maintain the asymptotic part under weakened assumptions. However, the monotonic increase fails to hold as soon as heterogeneously competent agents are allowed [18].

Recently, a novel generalization of the CJT has been proposed that successfully relaxes all original assumptions simultaneously [13]. In this generalization, agents are permitted to vote for any finite number of alternatives while accommodating heterogeneous competence levels and a degree of correlation among the electorate, which is modeled through an *opinion leader* (OL), the classical dependence model in the CJT literature [1]. The opinion leader (OL) serves as an abstract external influence in the voting process, without actively participating in it. Instead, the OL approves or disapproves presented alternatives based on her own competency $\hat{p}$, representing the probability of approving the correct alternative. Subsequently, her choice influences the agents' votes: each agent votes according to the OL's preference rather than their own "inner voice" with a certain probability π. In the standard OL model, this probability π is identical across all agents, i.e. it is uniformly distributed. For convenience, we refer to this case as *uniform* π. In this paper, we harness this specific generalization of the CJT to address dilation within the Voting for Bins framework.

For this purpose, we now introduce the underlying voting and probabilistic framework, following the presentation by Karge et al. (2024) [13]. Let $\mathcal{W} = \{\omega_1, \dots, \omega_m\}$ denote a finite set of m alternatives. With the previously defined set of agents, $\mathcal{N}$, we can then represent a single *approval voting (instance)* by $V \subseteq \mathcal{N} \times \mathcal{W}$ where $(a_i, \omega_j) \in V$ means that agent a_i approves choice ω_j. Subsequently, we define the *score* $\#_V\, \omega$ of some choice $\omega \in \mathcal{W}$ as: $\#_V\, \omega = |\{a_i \in \mathcal{N}_n \mid (a_i, \omega) \in V\}|$. Finally, the winner of V is defined to be the alternative that receives a strictly higher score than any alternative: $\#_V\, \omega > \max_{\omega' \in \mathcal{W} \setminus \{\omega\}} \#_V\, \omega'$.

The described voting scenario is modeled by a random process that generates the correct alternative, ω_*, the OL's choice as well as V and is govenerd by a joint probability distribution $\mathbb{P}$ over the Bernoulli (i.e., $\{0,1\}$-valued) random variables $X_*^{\omega_1}, ..., X_*^{\omega_m}$, $X_o^{\omega_1}, ..., X_o^{\omega_m}$ as well as $X_i^{\omega_1}, ..., X_i^{\omega_m}$ for all agents $1, ..., i, ..., n$ and all alternatives $1, ..., j, ..., m$ such that 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 true world state (i.e., $\omega_j = \omega_*$), else 0,
- $X_o^{\omega_j}$ is 1 if the OL approves ω_j, and 0 otherwise,
- $X_i^{\omega_j}$ represents the private signal of the ith agent regarding his approval of the jth world state: it is 1 if a_i privately approves ω_j and otherwise 0.

Given this joint distribution, we introduce the random variable $V_i^{\omega_j}$ representing the final outcome of an agent's vote, i.e. after the OL potentially exerted influence. According to our assumption, $V_i^{\omega_j}$ is the probabilistic mixture of $X_o^{\omega_j}$ with probability π and of $X_i^{\omega_j}$ with probability $1 - \pi$. From this, we obtain for any $x \in \{0, 1\}$ that

$$\mathbb{P}(V_i^{\omega_j} = x) = \pi \mathbb{P}(X_o^{\omega_j} = x) + (1 - \pi)\mathbb{P}(X_i^{\omega_j} = x).$$

We denote by $p_1^{\omega}, \dots, p_n^{\omega}$ the Bernoulli parameters of the "inner voice" random variables $X_1^{\omega}, \dots, X_n^{\omega}$, for all $\omega \in \mathcal{W}$, that is, $p_i^{\omega_j} = \mathbb{P}(X_i^{\omega_j} = 1)$. In a similar vein, for every $\omega \in \mathcal{W}$, we let $\hat{p}^{\omega_1}, \dots, \hat{p}^{\omega_m}$ denote the Bernoulli parameters of the random variables $X_o^{\omega_1}, \dots, X_o^{\omega_m}$. Whether the OL approves the correct alternative, i.e. whether $X_o^{\omega_*} = 1$, is governed by the parameter, $\hat{p} = \mathbb{P}(X_o^{\omega_*} = 1)$. For convenience, the choice of the correct alternative, or correct world state, being unknown to the agents, will be abbreviated by $[\omega_* = \omega_j]$.

In the following, we define the two central assumptions regarding the joint distribution. Conditioning upon the actual world state, we may define *private agent approval independence* as follows:

Definition 5. *A joint distribution satisfies* private agent approval independence *if, conditioned on the actual world state, the private 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 $v_1, ..., v_n$ of values from $\{0,1\}$ the following holds:*

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

A further central assumption deals with the "internal competency" p_k^{ω} of the kth agent regarding his capacity to identify the true world state among any number of alternatives if no influence is exerted. This assumption can be formalized as follows, where we denote the average over these "internal competencies" with $\bar{p}^{\omega} = \frac{1}{n}\sum_{k=1}^{n} p_k^{\omega}$.

Definition 6. *A joint probability distribution satisfies* Δp-group reliability *for some $\Delta p > 0$, if the probability, with respect to the agent's inner voice, 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 n and $\omega_\dagger \in \mathcal{W} \setminus \{\omega_*\}$ holds $\bar{p}^{\omega_*} \geq \Delta p + \bar{p}^{\omega_\dagger}$.*

From these assumptions, an implicit bound on the number of agents necessary to achieve a minimal success probability, $P_{\min}$, can be derived [13]. This bound will later be utilized to derive a bound on the achievable precision in pooling by Voting for Bins. Intuitively, throughout the paper, we interpret an increase in Δp as an increase in average competency, and an increase in π as an increase in correlation among the agents.

Theorem 2. *Consider an approval voting setting with $m > 1$ alternatives, satisfying private agent approval independence and Δp-group reliability for some $\Delta p \in (0,1]$, influenced by an opinion leader with $\pi \in [0, \frac{\Delta p}{\Delta p+1})$ and $\hat{p} \in [0,1]$. Then, given a probability $P_{\min} < 1$, it is guaranteed that the success probability of the approval voting process is greater than $P_{\min}$ if the number n of agents obeys the following condition*

$$\hat{p}e^{-\frac{n}{2}\Delta p^2(1-\pi)^2} + (1{-}\hat{p})e^{-\frac{n}{2}(\Delta p(1-\pi)-\pi)^2} \leq \frac{1-P_{\min}}{m{-}1}.$$

Voting for Bins. *Voting for Bins* (VfB), a recently introduced voting method [12], can also be viewed as an imprecise pooling function. The overarching idea is to incorporate the imprecise beliefs of a given number of agents and to interpret the set of alternatives, within the framework of the CJT as described above, as partitions of the unit interval called *bins*. This approach allows to exploit the probabilistic guarantees provided by the CJT. An underlying crucial assumption is that for a given proposition A about which beliefs are to be pooled, there exists a presumed true probability p^* for A to occur, such that p^* falls into exactly one of these bins. That particular bin then represents the correct alternative, or ground truth, in the voting process. The number of bins directly correlates with the precision achievable in the voting process, thereby translating straightforwardly into the precision attainable in opinion pooling by VfB. To the best of our knowledge, VfB is the only imprecise pooling function where the pooled belief is directly dependent on parameters that quantify the quality of the opinions to be aggregated. This makes it an ideal pooling mechanism for taming dilation. Consequently, it is crucial to formally define the set of bins and elucidate how agents vote based on their imprecise beliefs, following the exposition by Karge (2023) [12].

Consider an approval voting setting as described in Theorem 2. 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 7 (Bin). *Each* $\omega_k \in \mathcal{W} = \{\omega_1, \dots, \omega_m\}$ *represents a subinterval (bin) of the form* $[a_1, a_2)$*, 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 and denote a particular set of bins as $\mathbb{B} = \{B_1, ..., B_m\}$. Intuitively, each agent votes for the set of bins they are predominantly confident in:

Definition 8 (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 Lebesque 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 9 (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 1. 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 $\mathcal{P}(A)_1 \cap B_2 = [0.5, 0.9]$ 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 if there is a tie among alternatives with no bin having strictly more votes than any other, there is no winning bin in the approval vote.

3 Contributions

This section outlines our contributions to dilation in imprecise pooling. Firstly, we utilize Theorem 2 to establish a bound on the number of alternatives permissible for VfB, akin to the work by Karge (2023) [12], but with the integration of OL influence. Secondly, we formally define the problem of dilation, along with its counterpart, contraction, within the imprecise pooling framework. Subsequently, we expand upon the generalization proposed by Karge et al. (2024) [13] to incorporate average OL influence, allowing for a meaningful discussion regarding the potential increase in correlation within a group of agents when adding an additional agent. Specifically, while it may not be meaningful to discuss an increase in correlation when employing a uniform distribution of π, extending the result to average π-values provides a context where such discussions are relevant and informative. We then present a series of corollaries, each delineating conditions under which dilation or contraction is anticipated, and illustrate the method for computing the resultant change in dilation or contraction.

3.1 Voting for Bins and the OL

By establishing a direct correspondence between the alternatives and subintervals of the unit interval, the bins, we can determine a bound on the maximum number of bins permissible for a given set of input parameters [12]. This approach is extended to accommodate an opinion leader by solving the bound in Theorem 2 for the number of alternatives.

Corollary 1. *Consider an approval voting setting as described in Theorem 2. Then, given a probability $P_{\min} < 1$, it is guaranteed that the success probability of the approval voting process is greater than $P_{\min}$ if the number m of bins is at most*

$$m \leq \frac{1-P_{min}}{(\hat{p}(e^{-\frac{1}{2}n(1-\pi)^2\Delta p^2})+(1-\hat{p})e^{-\frac{1}{2}n(\Delta p(1-\pi)-\pi)^2}} + 1.$$

Finally, this translates directly into the maximum allowed precision in percent, denoted by C, which is the fraction of the unit interval covered by a subinterval given by the formula: $C = \frac{100}{m}$ [12].

In real-world scenarios involving the aggregation of imprecise probabilistic beliefs, it is not uncommon to have large expert panels comprising, for example, 50 climate scientists [14] or 140 epidemiologists [20].

Example 2. Suppose $\Delta p = 0.4$, $P_{\min} = 0.9, \hat{p} = 0.7, \pi = 0.1$, and $n = 150$. Then, we may allow for 52 bins of equal Lebesque measure. This translates to a precision of 1.9%.

With these definitions in hand, we can now interpret VfB as a pooling method. Initially, we compute the maximal permissible number of bins, denoted as m, based on parameters $n, P_{min}, \Delta p, \pi$, and $\hat{p}$. Subsequently, we partition the unit interval into a maximum of m subintervals of equal width, known as bins. In conclusion, we utilize the input profile of imprecise beliefs, $(\mathcal{P}_1(A), ..., \mathcal{P}_n(A))$, along with $\mathbb{B}$, to compute $\mathrm{VfB}(\mathcal{P}_1, ..., \mathcal{P}_n)$ based on Definition 9.

3.2 Dilation and Contraction in Imprecise Pooling

Dilation, within the context of an agent's imprecise belief in a proposition A, denoted as $\mathcal{P}(A)$, refers to the phenomenon where updating the belief by conditionalizing on an outcome B from the event space results in $\mathcal{P}(A \mid B)$, whereby the belief interval $\mathcal{P}(A)$ is necessarily contained within $\mathcal{P}(A \mid B)$ [9]. In other words, regardless of the outcome learned from the event space, the belief interval expands, leading to increased uncertainty regarding the proposition. Updating rules that allow for dilation are often viewed as unsettling, as one would expect that learning new evidence would, at least in some instances, decrease an agent's uncertainty. Similarly, an opposing effect, known as *contraction*, can also occur in a manner analogous to that of dilation. Contraction entails that, irrespective of the new observation acquired, the belief interval can only diminish, thereby consistently reducing uncertainty [9]. Just as with dilation, contraction intuitively should not occur, regardless of the evidence presented to the agent. Interpreting a smaller aggregated belief interval as a reduction in uncertainty, and reduced uncertainty as an increase in knowledge, suggests that not every observation should augment the agent's knowledge. In this paper, we explore the phenomena of dilation and contraction within the context of imprecise opinion pooling.

Considering the two most extreme pooling functions, we examine convex pooling, where the pooled interval expands solely by incorporating more opinions, thereby becoming less informative, and pooling by intersection, where the pooled interval contracts exclusively, becoming more informative [17]. Given these extreme points of imprecise pooling, we proceed to define the process of updating an imprecise pooled belief within our framework, followed by a formal definition of dilation and contraction in this context.

Definition 10 (Updated Opinion Pool). *Suppose we are given the imprecise beliefs $\mathcal{P}_1(A), ..., \mathcal{P}_n(A)$ of n agents for some proposition A as well as their pooled belief $\mathcal{F}(\mathcal{P}_1(A), ..., \mathcal{P}_n(A))$ according to some pooling function $\mathcal{F}$. Upon receiving an additional opinion $\mathcal{P}_{n+1}(A)$, the updated opinion pool is simply $\mathcal{F}(\mathcal{P}_1(A), ..., \mathcal{P}_{n+1}(A))$.*

We define dilation in imprecise pooling as follows:

Definition 11 (Dilation in Imprecise Pooling). *Given $\mathcal{P}_1(A), ..., \mathcal{P}_n(A)$ and $\mathcal{F}(\mathcal{P}_1(A), ..., \mathcal{P}_n(A))$, dilation occurs if for every possible $\mathcal{P}_{n+1}(A)$, the Lebesgue measure of the updated pool is greater than that of the original pool, i.e. $l(\mathcal{F}(\mathcal{P}_1(A), ..., \mathcal{P}_{n+1}(A))) \geq l(\mathcal{F}(\mathcal{P}_1(A), ..., \mathcal{P}_n(A)))$.*

Contraction in imprecise pooling is defined analogously:

Definition 12 (Contraction in Imprecise Pooling). *Given $\mathcal{P}_1(A), ..., \mathcal{P}_n(A)$ and $\mathcal{F}(\mathcal{P}_1(A), ..., \mathcal{P}_n(A))$, contraction occurs if for every possible $\mathcal{P}_{n+1}(A)$, $l(\mathcal{F}(\mathcal{P}_1(A), ..., \mathcal{P}_{n+1}(A))) \leq l(\mathcal{F}(\mathcal{P}_1(A), ..., \mathcal{P}_n(A)))$.*

With these definitions in hand, we are now ready to formulate the central problem addressed in this paper:

Given: $\mathcal{P}_1(A), ..., \mathcal{P}_n(A), \mathcal{P}_{n+1}(A)$ and $\mathcal{F}(\mathcal{P}_1(A), ..., \mathcal{P}_n(A))$.
Question: Under what conditions and to what extent should $\mathcal{F}(\mathcal{P}_1(A), ..., \mathcal{P}_{n+1}(A))$ dilate or contract compared to $\mathcal{F}(\mathcal{P}_1(A), ..., \mathcal{P}_n(A))$?

Intuitively, dilation is expected to occur when the additional opinion is of poor quality, consequently increasing uncertainty. Conversely, contraction is anticipated when the opinion is of high quality, leading to an increase in knowledge and a decrease in uncertainty among the group of agents.

As a general strategy to address this central problem, we demonstrate how Voting for Bins (VfB) aligns with this intuition. To illustrate this, we fix a minimal success probability, P_{min}, for VfB to identify the bin containing the underlying probability of the event in question. By fixing P_{min} and striving for the maximum number of permissible bins, we illustrate that if the additional agent's opinion is of poor quality, we anticipate dilation of the pooled belief. Conversely, if the opinion is of good quality, we expect the pooled belief to contract. In our framework, we define a 'good quality opinion' as one that either elevates the average competence level of the group of agents or diminishes correlation among them. Similarly, a 'bad quality opinion' is characterized by a decrease in average competence or an increase in correlation.

Notice that the current opinion leader (OL) model exclusively accounts for a uniform probability π for each agent to adhere to the OL. However, if we introduce an additional agent to an established group, it is conceptually unclear how correlation could increase for a uniform π. It is not reasonable to expect that every agent follows the OL with a higher probability simply because an additional agent is added. To address this limitation, we extend the OL model to accommodate average correlation.

3.3 Average Opinion Leader Influence

As a preliminary step towards incorporating average OL influence, we introduce the allowance for a finite number of distinct π-values $(\pi_1, ..., \pi_k, ...\pi_s)$, where $1 \leq s \leq n$, instead of a single uniform value. This means that the maximum number of π-values corresponds to having one unique value for each agent. It is worth noting that if there are multiple π_k values within a given electorate, we can delineate the group of voters into distinct subgroups. Each subgroup comprises individuals who share the same π-value, meaning they follow the OL with an equal probability.

We denote each subgroup as $\mathcal{G}_{\pi_k}$ where π_k represents the precise π-value of that subgroup, such that

$$\mathcal{G}_{\pi_1} \cup ...\mathcal{G}_{\pi_k} ... \cup \mathcal{G}_{\pi_s} = \mathcal{N} = \{a_1, ..., a_n\}.$$

Finally, for each subgroup $\mathcal{G}_{\pi_k}$, we refer to the number of agents in that subgroup as $|\mathcal{G}_{\pi_k}|$. From this, we define the average π-value, $\bar{\pi}$ as follows:

$$\bar{\pi} = \sum_{k=1}^{s} \frac{|\mathcal{G}_{\pi_k}|}{n} \pi_k.$$

Now that we have defined a notion for average correlation based on the OL influence, we proceed to derive a bound on the maximal number of alternatives in the next step.

Theorem 3. *Consider an approval voting setting as described in Theorem 2, but allowing for different π_k-values as defined above. Then, given a probability $P_{\min} < 1$, it is guaranteed that the success probability of the approval voting process is greater than $P_{\min}$ if the number m of bins is at most*

$$1 + \frac{1-P_{min}}{\hat{p}e^{-\frac{2}{4n}(\sum\limits_{k=1}^{s}|\mathcal{G}_{\pi_k}|(1-\pi_k)\Delta p)^2} + (1-\hat{p})e^{-\frac{2}{4n}(\sum\limits_{k=1}^{s}|\mathcal{G}_{\pi_k}|(\Delta p(1-\pi_k)-\pi_k))^2}}.$$

Proof. The proof parallels that of Karge et al. (2024) [13] for uniformly distributed π and can be found in the appendix. In essence, the proof proceeds by distinguishing between two cases: one where the OL is correct and another where the OL is incorrect. Within each case, the underlying marginal probabilities may be treated as independent once conditioned on the OL's selection. Utilizing Hoeffding's inequality, we aggregate both cases to derive the aforementioned bound.

We proceed to demonstrate that having a uniform π is equivalent to having the average OL influence $\bar{\pi}$, in the sense that the worst-case bound on the number of bins remains the same if $\pi = \bar{\pi}$. Let m_π denote the maximal number of permissible bins under uniform π, as derived by Corollary 1, and let $m_{\bar{\pi}}$ refer to the maximal number of bins derived from Theorem 3.

Theorem 4. *If $\pi = \bar{\pi}$, then $m_\pi = m_{\bar{\pi}}$.*

Proof. Assume, $\pi = \bar{\pi}$. Observe that the denominator of Theorem 3 and Corollary 1 respectively consist of two parts: one where the OL endorses the correct alternative, and one where this is not the case. Consider the first case in Theorem 3, $(e^{-\frac{2}{4n}(\sum_{k=1}^{s}|\mathcal{G}_{\pi_k}|(1-\pi_k)\Delta p)^2})$. By definition of $\sum_{k=1}^{s}|\mathcal{G}_{\pi_k}|$, $\sum_{k=1}^{s}|\mathcal{G}_{\pi_k}| = n$. Moreover, by definition of $\bar{\pi}$, $\sum_{k=1}^{s}\frac{|\mathcal{G}_{\pi_k}|}{n}\pi_k = \bar{\pi}$. Thus, as $\pi = \bar{\pi}$ by assumption, $\sum_{k=1}^{s}\frac{|\mathcal{G}_{\pi_k}|}{n}\pi_k = \pi$. Finally, $\sum_{k=1}^{s}|\mathcal{G}_{\pi_k}|(1-\pi_k) = n(1-\pi)$ when $\sum_{k=1}^{s}\frac{|\mathcal{G}_{\pi_k}|}{n}\pi_k = \pi$ and $\sum_{k=1}^{s}|\mathcal{G}_{\pi_k}| = n$. As Δp is fixed, we have that $(\sum_{k=1}^{s}|\mathcal{G}_{\pi_k}|(1-\pi_k)\Delta p)^2 = (n(1-\pi)\Delta p)^2$.

By algebra, we obtain $(e^{-\frac{2}{4n}(\sum_{k=1}^{s}|\mathcal{G}_{\pi_k}|(1-\pi_k)\Delta p)^2}) = (e^{-\frac{1}{2}n(1-\pi)^2\Delta p^2})$. An analogous argument can be made for the subcase where the OL is wrong. Hence, we conclude that $m_\pi = m_{\bar{\pi}}$.

This result suggests that irrespective of how competence levels and π-values are distributed among agents, in the worst-case, having a uniform π is equivalent to having different π-values across subgroups. Moreover, this finding enables us to leverage the bound for uniform π, which is slightly easier to work with, when discussing an increase in $\bar{\pi}$ as long as only the worst-case bound is of interest.

3.4 Taming Dilation and Contraction

With the direct correspondence between the number of alternatives and the permissible width of the aggregate interval computed by VfB, we can deduce a series of corollaries from Corollary 1. These corollaries not only elucidate how dilation and contraction align with our intuitions regarding incoming observations but also enable us to directly compute the change in dilation or contraction that may occur. For each corollary, we examine a specific scenario involving the addition of an extra agent to the group of agents, along with the corresponding updated opinion pool.

Corollary 2. *If the addition of agent a_{n+1} does not alter the average correlation or average competence, the maximal number of bins increases, resulting in the contraction of the pooled interval.*

Proof. Let P_{min}, Δp, $\hat{p}$, and $\bar{\pi}$ be fixed. Recall Corollary 1. As the denominator is of the form $\hat{p}e^{-nx} + (1-\hat{p}e^{-ny})$, it decreases for any fixed set of input parameters as n, causing nx and ny to become larger. Consequently, as the denominator decreases, the overall expression increases. Therefore, a larger number of maximal alternatives can be accommodated. Translated to VfB, this implies that the aggregate interval becomes smaller, thereby becoming more informative.

Intuitively, this effect stems from the *wisdom of the crowds* effect underlying the Condorcet Jury Theorem. As long as the group of agents is Δp-group reliable and all parameters remain fixed, the addition of an extra agent with the same expertise enhances the group's ability to identify the ground truth.

Example 3. Consider the following parameters: $P_{min} = 0.9$, $\Delta p = 0.4$, $\hat{p} = 0.5$, $\bar{\pi} = 0.1$, and $n = 200$. This yields a maximal number of 173 bins. Now, if we increase n to 201, we obtain 179 bins. Consequently, the aggregate interval contracts.

Corollary 3. *If agent a_{n+1} raises the average competence level, the maximal number of bins increases, resulting in a contraction of the pooled interval.*

Proof. Consider adding an agent a_{n+1} such that $a_{n+1} \in \mathcal{G}_{\pi_j}$. Let Δp and $\Delta p'$ be specific values such that $\Delta p < \Delta p'$. We claim that $m_n < m_{n+1}$. By a similar argument as above, it is evident that the denominator, of the form $\hat{p}e^{-\Delta px} + (1-\hat{p})e^{-\Delta py})$, decreases for any fixed set of input parameters and increasing Δp, causing Δpx and Δpy to become larger. With decreased denominator, the overall expression increases.

In scenarios where the average correlation increases upon adding an additional agent, the underlying intuition is less straightforward. Typically, an increase in correlation suggests that the electorate may be less effective in identifying the correct bin. However, as we've observed, simply adding an extra agent enhances the group's capabilities. Therefore, the informativeness, regarded as the counterpart of uncertainty, of the aggregate should only decrease when the increase in correlation outweighs the impact of adding a new member. Let's consider the following two examples:

Example 4 (Increase in Informativeness). Consider the parameters $P_{min} = 0.2$, $\Delta p = 0.9$, $\hat{p} = 0.5$, $\bar{\pi} = 0.1$, and $n = 20$. With these parameters, we obtain 204 bins. Now, by slightly increasing π to $\pi' = 0.11$ and considering $n = 21$, we obtain 205 bins. Thus, informativeness increases despite the increase in average correlation.

Example 5 (Decrease in Informativeness). Let $P_{min} = 0.2$, $\Delta p = 0.9$, $\hat{p} = 0.5$, $\bar{\pi} = 0.2$, and $n = 40$. With these parameters, we obtain 355 bins. Now, with a slight increase in π to $\pi' = 0.21$ and considering $n = 41$, we obtain 274 bins. Thus, informativeness decreases despite the addition of an extra agent.

Next, we aim to determine the threshold for increased average correlation that disrupts the enhanced group capabilities due to the new agent. To determine this threshold we introduce the parameter ϵ, which represents the amount by which the average correlation increases by adding the new agent. We claim that adding a new agent is outweighed by increased correlation as soon as the group of agents without the new one allows for more alternatives than the larger group. That is, as soon as:

$$\frac{1-P_{min}}{(\hat{p}(e^{-\frac{1}{2}n(1-\pi)^2\Delta p^2})+(1-\hat{p})e^{-\frac{1}{2}n(\Delta p(1-\pi)-\pi)^2}} + 1 > 1 + \frac{1-P_{min}}{(\hat{p}(e^{-\frac{1}{2}(n+1)(1-(\pi+\epsilon))^2\Delta p^2})+(1-\hat{p})e^{-\frac{1}{2}(n+1)(\Delta p(1-(\pi+\epsilon))-(\pi+\epsilon))^2}}.$$

Instead of determining the threshold for the general case, we observe that increasing π has the most detrimental effect on the group's capacity to identify the correct bin when $\hat{p}$ is low. In the worst-case scenario, $\hat{p} = 0$. Therefore, to derive the threshold at which increased correlation outweighs adding a new member, we consider this worst case and assume $\hat{p} = 0$. In this scenario, the above expression simplifies to

$$\frac{1-P_{min}}{e^{-\frac{1}{2}n(\Delta p(1-\pi)-\pi)^2}} + 1 > \frac{1-P_{min}}{e^{-\frac{1}{2}(n+1)(\Delta p(1-(\pi+\epsilon))-(\pi+\epsilon))^2}} + 1.$$

We simplify this problem further by noting that both expressions only differ in the exponent of their denominators. As the denominators are of the form e^{-x} and e^{-y} respectively, they increase as x and y decrease. With an increased denominator, the overall number of bins, and thereby, informativeness, decreases. Therefore, to derive the exact threshold, we investigate under what conditions

$$n(\Delta p(1-\pi)-\pi)^2 > (n+1)(\Delta p(1-(\pi+\epsilon))-(\pi+\epsilon))^2.$$

By solving for ϵ, we can derive the following bound:

$$\epsilon > \frac{\Delta p(1-\pi)-\pi}{\Delta p+1} - \sqrt{\frac{n(\Delta p(\pi-1)+\pi)^2}{(\Delta p+1)^2(n+1)}}. \quad (1)$$

This leads directly to the following corollary:

Corollary 4. *If agent a_{n+1} increases the average correlation level and inequality 1 is true, then the maximal number of bins decreases, resulting in a dilation of the pooled interval.*

We illustrate the change in number of bins in subfigure (a) of Figure 1 where dilation occurs as soon as the threshold for ϵ is surpassed.

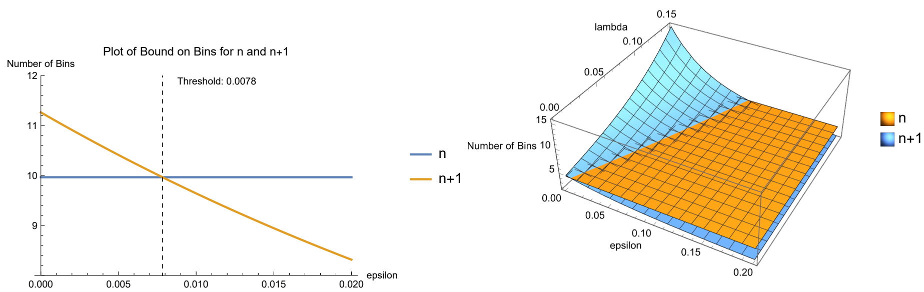

(a) Threshold for corollary 3.

(b) Threshold for corollary 5.

Fig. 1: Number of Bins depending on ϵ and for $\Delta p = 0.6, P_{min} = 0.4, n = 20, \pi = 0.05$ (left), and depending on ϵ and λ for $P_{min} = 0.2, \Delta p = 0.4, n = 25, \pi = 0.05$ (right).

Corollary 5. *If the additional agent decreases the average correlation, the maximal number of bins increases. This results in a contraction of the aggregate.*

In this case, the terms e^{-x} and e^{-y} in the denominator increase as π decreases, resulting in a larger overall expression and thus increasing the number of bins.

Finally, we address the scenario in which both average competency and average correlation increase. We derive a threshold indicating when the increase in average competency outweighs the increase in correlation. Introducing λ as a parameter for the increase in average competency, we determine the point at which an increase in competency leads to a larger number of bins, despite the simultaneous increase in average correlation:

$$\frac{1-P_{min}}{e^{-\frac{1}{2}n(\Delta p(1-\pi)-\pi)^2}} + 1 \frac{1-P_{min}}{e^{-\frac{1}{2}(n+1)((\Delta p+\lambda)(1-(\pi+\epsilon))-(\pi+\epsilon))^2}} + 1.$$

As before, we further simplify this problem and demonstrate the conditions under which:

$$n(\Delta p(1-\pi)-\pi)^2 < (n+1)((\Delta p+\lambda)(1-(\pi+\epsilon))-(\pi+\epsilon))^2.$$

By solving for ϵ, we obtain a bound for how much increase in correlation can be tolerated for a given increase in competency:

$$\epsilon < \frac{(\Delta p(-\pi)+\Delta p-\pi(\lambda+1)+\lambda)}{(\Delta p+\lambda+1)} - \sqrt{\frac{(n(\Delta p(\pi-1)+\pi)^2)}{((n+1)(\Delta p+\lambda+1)^2)}}. \quad (2)$$

This yields the following corollary:

Corollary 6. *If the average correlation increases and Inequality 2 is true, the overall number of bins increases. In this case, the pooled interval contracts.*

We illustrate the maximal permissible number of bins for n and $n+1$ agents with $P_{min} = 0.2$, $\Delta p = 0.4$, $n = 25$, and $\pi = 0.05$ for different λ and ϵ values in subfigure (b) of Figure 1. The orange surface area illustrates the combinations of values where the threshold is surpassed, indicating where the increase in competence does not outweigh the increase in correlation.

In this subsection, we discussed a series of corollaries, each illustrating how pooling by Voting for Bins manages dilation and contraction precisely as one would intuitively expect, depending on the quality of the additional opinion. Furthermore, it is now clear how to compute the extent to which the updated aggregate dilates or contracts compared to the original pool: it is simply the difference in the width of both pools, assuming we always choose the maximal number of bins according to Corollary 1.

4 Conclusion

In this work, we delved into the phenomena of *dilation* and *contraction* within the framework of imprecise opinion pooling. Through our analysis, we presented the first formal treatment of both effects within imprecise pooling and demonstrated how they can be managed using a specific pooling technique known as *Voting for Bins*. This approach enabled us to establish precise conditions under which dilation or contraction is expected to occur, aligning with our intuitive expectations based on the quality of the additional opinion being incorporated. Moreover, we illustrated how to compute the extent to which dilation or contraction is expected to occur. In addition, we extended the OL model to account for the average correlation among agents, allowing us to study the interplay between the increase in average correlation and the quality of additional opinions.

Looking ahead, we aim to derive translations between the opinion leader model and standard correlation measures used in multi-agent systems such as covariance, the correlation coefficient, and mutual information. These translations would provide a bridge between the abstract notion of opinion leader influence and more familiar correlation metrics, enhancing the applicability to real-world scenarios with complex dependence structures among agents.

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

References

1. Boland, P.J., Proschan, F., Tong, Y.L.: Modelling dependence in simple and indirect majority systems. Journal of Applied Probability **26**(1), 81–88 (1989)
2. Bradley, S.: How to choose among choice functions. Episteme **11**, 277–289 (2014)
3. Bradley, S., Steele, K.: Should subjective probabilities be sharp? Episteme **11**(3), 277–289 (2014)
4. Bradley, S., Steele, K.: Can free evidence be bad? value of information for the imprecise probabilist. Philosophy of Science **83**(1), 1–28 (2016)
5. 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)
6. Crosscombe, M., Lawry, J.: A model of multi-agent consensus for vague and uncertain beliefs. Adaptive Behavior **24**(4), 249–260 (2016)
7. 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)
8. Elkin, L.: Regret averse opinion aggregation. Ergo **8**(16) (2021)
9. Gong, R., Meng, X.L.: Judicious judgment meets unsettling updating: Dilation, sure loss and simpson's paradox. Statistical Science **36**(2), 169–190 (2021)
10. Good, I.J.: On the principle of total evidence. The British Journal for the Philosophy of Science **17**(4), 319–321 (1967)
11. Joyce, J.M.: A defense of imprecise credences in inference and decision making1. Philosophical Perspectives **24**(1), 281–323 (2010)
12. 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)
13. Karge, J., Burkhardt, J.M., Rudolph, S., Rusovac, D.: To lead or to be led: A generalized condorcet jury theorem under dependence. In: AAMAS '24: Proceedings of the 2024 International Conference on Autonomous Agents and Multiagent Systems (2024)
14. 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)
15. Lawry, J., Lee, C.: Probability pooling for dependent agents in collective learning. Artificial Intelligence **288**, 103371 (2020)
16. 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)
17. Nau, R.F.: The aggregation of imprecise probabilities. Journal of Statistical Planning and Inference **105**(1), 265–282 (2002)
18. Owen, G., Grofman, B., Feld, S.L.: Proving a distribution-free generalization of the Condorcet Jury Theorem. Mathematical Social Sciences **17**(1), 1–16 (1989)
19. Quintana, I.O.: Radical pooling and imprecise probabilities. Erkenntnis (2022)
20. 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)
21. Stewart, R.T., Nielsen, M.: Distention for sets of probabilities. Philosophy of Science **89**(3), 604–620 (2022)
22. Stewart, R.T., Quintana, I.O.: Probabilistic opinion pooling with imprecise probabilities. Journal of Philosophical Logic **47**, 17–45 (2018)

23. 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)
24. White, R.: Evidential symmetry and mushy credence. Oxford Studies in Epistemology **3**, 161–186 (2009)

5 Appendix

Proof, Theorem 3.

We now derive a lower bound on the minimum probability of success for identifying the correct alternative for an electorate across different subgroups.

Recall the aggregated random variable $V^{\omega_*-\omega_\dagger}$ defined by

$$V^{\omega_*-\omega_\dagger} = \sum_{k=1}^{n}(V_k^{\omega_*} - V_k^{\omega_\dagger}) = \sum_{k=1}^{n} V_k^{\omega_*} - \sum_{k=1}^{n} V_k^{\omega_\dagger} = V^{\omega_*} - V^{\omega_\dagger}$$

and observe that ω_* wins against $\omega_\dagger$ exactly if $V^{\omega_*-\omega_\dagger} > 0$.

For the derivation, we extend the definition of $V^{\omega_*-\omega_\dagger}$ to include the subgroups. Suppose the subgroups are ordered from the lowest $\pi - value$, denoted by π_1, to the highest $\pi - value$, denoted by π_s. Moreover, we refer to $V^{\omega_j}_{|\mathcal{G}_{\pi_k}|}$ as the score ω_j received from the subgroup $\mathcal{G}_{\pi_k}$ with size $|\mathcal{G}_{\pi_k}|$.

$$\begin{aligned} V^{\omega_*-\omega_\dagger} &= \sum_{k=1}^{n}(V_k^{\omega_*} - V_k^{\omega_\dagger}) \\ &= \sum_{k=1}^{n} V_k^{\omega_*} - \sum_{k=1}^{n} V_k^{\omega_\dagger} \\ &= \sum_{k=1}^{s} V^{\omega_*}_{|\mathcal{G}_{\pi_k}|} - V^{\omega_\dagger}_{|\mathcal{G}_{\pi_k}|} \\ &= V^{\omega_*} - V^{\omega_\dagger} \end{aligned}$$

To find good probability estimates, we use Hoeffding's inequality, which provides a tail estimate for the sum of independent random variables with the property of having zero probability outside a finite interval.

Lemma 1 (Hoeffding 1963). *Let $X_1, \ldots, X_n$ be independent random variables satisfying $\mathbb{P}(l_i \le X_i \le u_i) = 1$ 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*

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

We recall that the agent-wise distributions of $V_i^{\omega_*} - V_i^{\omega_\dagger}$ discussed above are of this type with $l_i = -1$ and $u_i = 1$.

The OL is right. We start by conditioning on the OL approving the correct alternative. To correctly reflect the worst case, we assume that for any number of alternatives, the OL always approves all competitors of ω_*. In the following, we consider an arbitrarily chosen but fixed competing alternative $\omega_\dagger \in \mathcal{W} \setminus \{\omega_*\}$ in the approval vote. In a first step, we derive a lower bound on the probability that ω_* wins against this fixed competitor, $\omega_\dagger$. In a next step, we can derive a lower bound for the probability of winning against all competitors.

$$
\begin{aligned}
&\mathbb{P}(V^{\omega_*} > V^{\omega_\dagger} \mid X_o^{\omega_*} = 1) \\
&\quad = \mathbb{P}(\sum_{k=1}^{s} V^{\omega_*}_{|\mathcal{G}_{\pi_k}|} - V^{\omega_\dagger}_{|\mathcal{G}_{\pi_k}|} > 0 \mid X_o^{\omega_*} = 1) \\
&\quad = 1 - \mathbb{P}(\sum_{k=1}^{s} V^{\omega_*}_{|\mathcal{G}_{\pi_k}|} - V^{\omega_\dagger}_{|\mathcal{G}_{\pi_k}|} \leq 0 \mid X_o^{\omega_*} = 1) \\
&\quad = 1 - \mathbb{P}(\sum_{k=1}^{s} V^{\omega_*}_{|\mathcal{G}_{\pi_k}|} - V^{\omega_\dagger}_{|\mathcal{G}_{\pi_k}|} - \mathbb{E}(\sum_{k=1}^{s} V^{\omega_*}_{|\mathcal{G}_{\pi_k}|} - V^{\omega_\dagger}_{|\mathcal{G}_{\pi_k}|} \mid X_o^{\omega_*} = 1) \leq \\
&\qquad - \mathbb{E}(\sum_{k=1}^{s} V^{\omega_*}_{|\mathcal{G}_{\pi_k}|} - V^{\omega_\dagger}_{|\mathcal{G}_{\pi_k}|} \mid X_o^{\omega_*} = 1) \mid X_o^{\omega_*} = 1) \\
&\quad = 1 - \mathbb{P}(\sum_{k=1}^{s} V^{\omega_\dagger}_{|\mathcal{G}_{\pi_k}|} - V^{\omega_*}_{|\mathcal{G}_{\pi_k}|} - \mathbb{E}(\sum_{k=1}^{s} V^{\omega_\dagger}_{|\mathcal{G}_{\pi_k}|} - V^{\omega_*}_{|\mathcal{G}_{\pi_k}|} \mid X_o^{\omega_*} = 1) \geq \\
&\qquad \mathbb{E}(\sum_{k=1}^{s} (V^{\omega_\dagger}_{\mathcal{G}_{\pi_k}} - V^{\omega_*}_{\mathcal{G}_{\pi_k}}) \mid X_o^{\omega_*} = 1) \mid X_o^{\omega_*} = 1) \\
&\quad = 1 - \mathbb{P}(\sum_{k=1}^{s} V^{\omega_\dagger}_{|\mathcal{G}_{\pi_k}|} - V^{\omega_*}_{|\mathcal{G}_{\pi_k}|} - \mathbb{E}(\sum_{k=1}^{s} V^{\omega_\dagger}_{|\mathcal{G}_{\pi_k}|} - V^{\omega_*}_{|\mathcal{G}_{\pi_k}|} \mid X_o^{\omega_*} = 1) \geq \\
&\qquad \mathbb{E}(\sum_{k=1}^{s} V^{\omega_*}_{|\mathcal{G}_{\pi_k}|} \mid X_o^{\omega_*} = 1) - \mathbb{E}(\sum_{k=1}^{s} V^{\omega_\dagger}_{|\mathcal{G}_{\pi_k}|} \mid X_o^{\omega_*} = 1) \mid X_o^{\omega_*} = 1)
\end{aligned}
$$

In the worst case, $\bar{p}^{\omega_*} = \bar{p}^{\omega_\dagger} + \Delta p$.

$$
\geq 1 - \mathbb{P}((V^{\omega_\dagger} - V^{\omega_*}) - \mathbb{E}(V^{\omega_\dagger} - V^{\omega_*} \mid X_o^{\omega_*} = 1) \geq \sum_{k=1}^{s} \Delta p(|\mathcal{G}_{\pi_k}|(1-\pi_k) \mid X_o^{\omega_*} = 1)
$$

Hoeffding noting that $u_i - l_i = 2$ for all i

$$
\geq 1 - e^{-\frac{2}{4n}(\sum_{k=1}^{s} \Delta p|\mathcal{G}_{\pi_k}|(1-\pi_k))^2}
$$

Then we obtain for the winning against all competitors:

$$
\begin{aligned}
\mathbb{P}(\bigwedge \omega_\dagger \in \mathcal{W} \setminus \{\omega_*\} V^{\omega_*} &> V^{\omega_\dagger} \mid X_o^{\omega_*} = 1) \\
&\geq 1 - \textstyle\sum_{i=1}^{m-1}(1 - \mathbb{P}(V^{\omega_*} > V^{\omega_i} \mid X_o^{\omega_*} = 1)) \\
&= 1 - \textstyle\sum_{i=1}^{m-1}(1 - (1 - e^{-\frac{2}{4n}(\sum_{k=1}^{s} \Delta p|\mathcal{G}_{\pi_k}|(1-\pi_k))^2})) \\
&= 1 - (m-1)e^{-\frac{2}{4n}(\sum_{k=1}^{s} \Delta p|\mathcal{G}_{\pi_k}|(1-\pi_k)))^2}
\end{aligned}
$$

The OL is wrong. In a similar vein, we consider the case of $X_o^{\omega_*} = 0$.

$$
\begin{aligned}
&\mathbb{P}(V^{\omega_*} > V^{\omega_\dagger} \mid X_o^{\omega_*} = 0) \\
&\quad = 1 - \mathbb{P}((V^{\omega_\dagger} - V^{\omega_*}) - \mathbb{E}(V^{\omega_\dagger} - V^{\omega_*} \mid X_o^{\omega_*} = 0) \geq \\
&\qquad \sum_{k=1}^{s}(|\mathcal{G}_{\pi_k}|(1-\pi_k)\bar{p}^{\omega_*}) - \sum_{k=1}^{s}(|\mathcal{G}_{\pi_k}|(\pi_k + (1-\pi_k)\bar{p}^{\omega_\dagger})) \mid X_o^{\omega_*} = 0)
\end{aligned}
$$

In the worst case, $\bar{p}^{\omega_*} = \bar{p}^{\omega_\dagger} + \Delta p$.

$$
\begin{aligned}
&\quad \geq 1 - \mathbb{P}((V^{\omega_\dagger} - V^{\omega_*}) - \mathbb{E}(V^{\omega_\dagger} - V^{\omega_*} \mid X_o^{\omega_*} = 0) \geq \\
&\qquad \sum_{k=1}^{s}(-\Delta p - 1)|\mathcal{G}_{\pi_k}|\pi_k + |\mathcal{G}_{\pi_k}|\Delta p \mid X_o^{\omega_*} = 0) \\
&\quad = 1 - \mathbb{P}((V^{\omega_\dagger} - V^{\omega_*}) - \mathbb{E}(V^{\omega_\dagger} - V^{\omega_*} \mid X_o^{\omega_*} = 0) \geq \\
&\qquad \sum_{k=1}^{s}|\mathcal{G}_{\pi_k}|(\Delta p(1-\pi_k) - \pi_k) \mid X_o^{\omega_*} = 0)
\end{aligned}
$$

Hoeffding with $u_i - l_i = 2$ for all i, assuming $\pi \leq \frac{\Delta p}{\Delta p + 1}$

$$
= 1 - e^{-\frac{2}{4n}\left(\sum_{k=1}^{s}|\mathcal{G}_{\pi_k}|(\Delta p(1-\pi_k)-\pi_k)\right)^2}
$$

Then we obtain for the winning against all competitors:

$$
\begin{aligned}
&\mathbb{P}(\bigwedge \omega_\dagger \in \mathcal{W} \setminus \{\omega_*\} V^{\omega_*} > V^{\omega_\dagger} \mid X_o^{\omega_*} = 0) \\
&\quad = 1 - (m-1)e^{-\frac{2}{4n}\left(\sum_{k=1}^{s}|\mathcal{G}_{\pi_k}|(\Delta p(1-\pi_k)-\pi_k)\right)^2}
\end{aligned}
$$

Aggregating the Cases. Finally, we aggregate the cases where the OL is right and where the OL is wrong.

$$
\begin{aligned}
&\mathbb{P}(\bigwedge \omega_\dagger \in \mathcal{W} \setminus \{\omega_*\} X^{\omega_*} > X^{\omega_\dagger}) \\
&= \mathbb{P}(X_o^{\omega_*} = 1) \cdot \mathbb{P}(\bigwedge \omega_\dagger \in \mathcal{W} \setminus \{\omega_*\} V^{\omega_*} > V^{\omega_\dagger} \mid X_o^{\omega_*} = 1) \\
&\quad + \mathbb{P}(X_o^{\omega_*} = 0) \cdot \mathbb{P}(\bigwedge \omega_\dagger \in \mathcal{W} \setminus \{\omega_*\} V^{\omega_*} > V^{\omega_\dagger} \mid X_o^{\omega_*} = 0) \\
&\geq \hat{p}\big(1 - (m-1)e^{-\frac{2}{4n}(\sum_{k=1}^{s}\Delta p(|\mathcal{G}_{\pi_k}|(1-\pi_k))^2}\big) \\
&\quad + (1-\hat{p})\big(1 - (m-1)e^{-\frac{2}{4n}\left(\sum_{k=1}^{s}|\mathcal{G}_{\pi_k}|(\Delta p(1-\pi_k)-\pi_k)\right)^2}\big) \\
&= 1 - (m-1)(\hat{p}(e^{-\frac{2}{4n}(\Delta p(\sum_{k=1}^{s}|\mathcal{G}_{\pi_k}|(1-\pi_k)))^2}) + (1-\hat{p})e^{-\frac{2}{4n}\left(\sum_{k=1}^{s}|\mathcal{G}_{\pi_k}|(\Delta p(1-\pi_k)-\pi_k)\right)^2})
\end{aligned}
$$

Solving for for the number of alternatives (m):

$$
\begin{aligned}
&P_{min} \leq 1 - (m-1)(\hat{p}(e^{-\frac{2}{4n}(\Delta p(\sum_{k=1}^{s}|\mathcal{G}_{\pi_k}|(1-\pi_k)))^2}) + (1-\hat{p})e^{-\frac{2}{4n}\left(\sum_{k=1}^{s}|\mathcal{G}_{\pi_k}|(\Delta p(1-\pi_k)-\pi_k)\right)^2}) \\
&\quad m \leq \frac{1-P_{min}}{(\hat{p}(e^{-\frac{2}{4n}(\Delta p(\sum_{k=1}^{s}|\mathcal{G}_{\pi_k}|(1-\pi_k)))^2}) + (1-\hat{p})e^{-\frac{2}{4n}\left(\sum_{k=1}^{s}|\mathcal{G}_{\pi_k}|(\Delta p(1-\pi_k)-\pi_k)\right)^2})} + 1
\end{aligned}
$$