

Fusion of Probability Density Functions

This article provides useful insight into the problem of probability density function fusion by reviewing numerous approaches developed in different disciplines and providing a higher level perspective of the viewpoints and solutions.

By GÜNTHER KOLIANDER^{ID}, YOUSEF EL-LAHAM^{ID}, PETAR M. DJURIĆ^{ID}, *Fellow IEEE*, AND FRANZ HLAWATSCH^{ID}, *Fellow IEEE*

ABSTRACT | Fusing probabilistic information is a fundamental task in signal and data processing with relevance to many fields of technology and science. In this work, we investigate the fusion of multiple probability density functions (pdfs) of a continuous random variable or vector. Although the case of continuous random variables and the problem of pdf fusion frequently arise in multisensor signal processing, statistical inference, and machine learning, a universally accepted method for pdf fusion does not exist. The diversity of approaches, perspectives, and solutions related to pdf fusion motivates a unified presentation of the theory and methodology of the field. We discuss three different approaches to fusing pdfs. In the axiomatic approach, the fusion rule is defined indirectly by a set of properties (axioms). In the optimization approach, it is the result of minimizing an objective function that involves an information-theoretic divergence or a distance measure. In the supra-Bayesian approach, the fusion center interprets the pdfs to be fused as random observations. Our work is partly a survey, reviewing in a structured and coherent fashion many of the concepts and methods that have been developed in the literature. In addition, we present new results for each of the three approaches. Our original contributions include new fusion rules, axioms, and axiomatic and optimization-based

characterizations; a new formulation of supra-Bayesian fusion in terms of finite-dimensional parametrizations; and a study of supra-Bayesian fusion of posterior pdfs for linear Gaussian models.

KEYWORDS | α -divergence; Chernoff fusion; covariance intersection; forecasting; Hölder mean; information fusion; Kullback–Leibler divergence; linear Gaussian model; model averaging; multisensor signal processing; pooling function; probabilistic machine learning; probabilistic opinion pooling; sensor network; supra-Bayesian fusion; target tracking.

NOMENCLATURE

Probabilistic Opinion Pooling

$q_k(\theta)$	Pdf of agent k .
$q(\theta)$	Aggregate (fused) pdf.
$Q_k(\mathcal{A})$	Probability of event \mathcal{A} according to $q_k(\theta)$.
$Q(\mathcal{A})$	Probability of event \mathcal{A} according to $q(\theta)$.
μ_{q_k}	Mean associated with $q_k(\theta)$.
μ_q	Mean associated with $q(\theta)$.
Σ_{q_k}	Covariance matrix associated with $q_k(\theta)$.
Σ_q	Covariance matrix associated with $q(\theta)$.

Supra-Bayesian Framework

\mathbf{y}_k	Local observation vector of agent k .
\mathbf{y}	Global observation vector (stacking all \mathbf{y}_k).
\mathbf{t}_k	Local statistic of agent k .
\mathbf{t}	Stacked vector of all local statistics \mathbf{t}_k .
$p(\theta)$	Prior pdf.
$\ell_k(\theta)$	Local observation likelihood function of agent k .
$\ell(\theta)$	Global observation likelihood function.
$\lambda_k(\theta)$	Local \mathbf{t}_k -likelihood function of agent k .
$\lambda(\theta)$	Global \mathbf{t} -likelihood function.
$\pi_k(\theta)$	Local posterior pdf of agent k .

Manuscript received May 2, 2021; revised January 21, 2022; accepted February 16, 2022. Date of current version April 1, 2022. This work was supported in part by the Vienna Science and Technology Fund (WWTF) under Grant MA16-053, in part by the Austrian Science Fund (FWF) under Grant P32055-N31 and Grant Y1199, and in part by the Growing Convergence Research Program of the National Science Foundation under Grant 2021002. (Corresponding author: Petar M. Djurić.)

Günther Koliander is with the Acoustics Research Institute, Austrian Academy of Sciences, 1040 Vienna, Austria (e-mail: gkoliander@kfs.oewa.ac.at).

Yousef El-Laham and **Petar M. Djurić** are with the Department of Electrical and Computer Engineering, Stony Brook University, Stony Brook, NY 11794 USA (e-mail: yousef.ellaham@stonybrook.edu; petar.djurić@stonybrook.edu).

Franz Hlawatsch is with the Institute of Telecommunications, TU Wien, 1040 Vienna, Austria (e-mail: franz.hlawatsch@tuwien.ac.at).

Digital Object Identifier 10.1109/JPROC.2022.3154399

General Notation

g	Pooling function.
$g[q_1, \dots, q_K](\theta)$	Fused pdf resulting from application of pooling function g to pdfs $q_1(\theta), \dots, q_K(\theta)$.
$\mathbb{E}_\psi[\cdot]$	Expectation operator with respect to pdf $\psi(\theta)$.
$\mathbb{E}[\cdot]$	Expectation operator with respect to the joint pdf of all involved random variables.
$\mathcal{N}(\theta; \mu, \Sigma)$	Pdf of a Gaussian random vector θ with mean μ and covariance matrix Σ .

I. INTRODUCTION

The fusion of multiple probabilistic descriptions of a random quantity is a fundamental task with applications in many fields, including multisensor signal processing [1]–[8], machine learning [9]–[12], robotics [4], smart environments [13], medicine [14], transportation [15], precision agriculture [16], pharmacology [17], weather forecasting [18], [19], economics [20], [21], and financial engineering [22]. While this task has been studied for several decades, an in-depth treatment with a focus on continuous random variables and, accordingly, on the fusion of probability density functions (pdfs) appears to be lacking. This article attempts to fill this gap. Our focus on continuous random variables is motivated by the fact that continuously distributed quantities are the primary object of interest in many applications.

The fusion of pdfs can be considered in different contexts, and several different techniques for this task have been proposed in the literature. Our treatment is partly a survey of existing concepts and techniques with an emphasis on a structured and coherent presentation. In addition, we present numerous original contributions related to axiomatic, optimization-based, and Bayesian approaches to pdf fusion.

A. Motivation

The field of pdf fusion is multifaceted and somewhat fuzzy: there are many possible approaches to the problem of finding a pdf fusion rule, and there is no universally accepted measure of performance [23], [24]. An appropriate fusion rule and performance measure depend on the scenario and application. This situation can be aggravated by the fact that different fusion rules could lead to very different results.

Although, in specific applications, certain pdf fusion rules have been established and found to be useful, the rationales of these rules and their possible alternatives are not always obvious. Thus, it is both theoretically interesting and practically relevant to study the problem of pdf fusion and the existing viewpoints and solutions in a general way that abstracts from specific applications, and to put these viewpoints and solutions into a higher level perspective. Our hope is that this analysis will support an

informed choice of a pdf fusion rule for specific scenarios and applications. Accordingly, rather than considering a single framework or method for pdf fusion, this article reviews the different approaches that have been developed over several decades in different disciplines and by different communities. In addition, these approaches are categorized into three fundamental approaches to principled pdf fusion, which we term the axiomatic, optimization, and supra-Bayesian approaches.

Fusing pdfs is a special variant of the general task of “data fusion” or “information fusion,” and one may ask why it can be advantageous to perform data/information fusion at the level of pdfs. Possible answers include the following [20], [24], [25].

- 1) A pdf constitutes a *complete* probabilistic description of a continuous random variable or random vector. In addition to its mean or its mode (which can be used as point estimates of the random variable or vector), this description includes further important information, such as effective support, multimodality, tail decay, and a detailed characterization of the “dispersion” around the mean. Moreover, it enables the calculation of quantitative measures of the accuracy of point estimates.
- 2) A pdf provides a standardized and “genesis-agnostic” representation of the state of information of an agent or sensor, i.e., it abstracts from the intricacies of the processing employed by the agent or sensor to obtain it from the raw data. This “no questions asked” characteristic enables or facilitates an information fusion even between heterogeneous agents that employ different sensing modalities and/or different types of data preprocessing. Furthermore, the lack of a transparent relation to the raw data is a desirable feature in privacy-sensitive applications.
- 3) Because a pdf provides a standardized, genesis-agnostic representation, pdf fusion is well suited to a decentralized (peer-to-peer) network topology. In decentralized, possibly *ad hoc* networks, a distributed in-network type of processing is used where each agent communicates with a limited set of neighboring agents, and typically, little or no information about the characteristics of far-away agents is available locally. The pdf format here facilitates the dissemination of information through the network.
- 4) Computationally efficient pdf fusion algorithms based on parametric pdf representations are available. For example, the fusion of Gaussian pdfs reduces to fusing the corresponding means and variances or covariance matrices. More generally, there are efficient algorithms for fusing Gaussian mixture pdfs. In distributed implementations, parametric pdf representations enable pdf fusion with low or moderate communication cost. Thus, pdf fusion is attractive because detailed probabilistic information can be

fused with moderate complexity in terms of computation and communication.

B. Probabilistic Opinion Pooling

Consider K “agents,” “experts,” or “models,” each providing an “opinion” about an unknown random object that may be a scalar or vector. In the probabilistic setting studied in this work, the opinions provided by the agents are not point estimates of the random object but probability distributions. More specifically, we focus on the case of a continuous random variable or vector θ , where the opinion of agent k is expressed by a pdf $q_k(\theta)$.

The problem studied in this article is to combine, or fuse, the pdfs of the K agents, $q_k(\theta)$ for $k = 1, 2, \dots, K$, into an aggregate pdf $q(\theta)$. This problem is traditionally referred to as *probabilistic opinion pooling* although that term is also used for the fusion of discrete (categorical) distributions. We assume that the combination of the agent pdfs $q_k(\theta)$ is done by a central agent or unit, termed a “fusion center,” which has access to all the agent pdfs. The function employed by the fusion center to map the $q_k(\theta)$ into the aggregate (fused) pdf $q(\theta)$ is termed a *fusion rule* or a *pooling function*. Many different pooling functions have been proposed in the literature based on various models and considerations. Important examples include the linear pooling function (a weighted arithmetic mean, also known as arithmetic mean density) [26], [27] and the log-linear pooling function (a weighted geometric mean, also referred to as Chernoff fusion or geometric mean density) [27]–[31]. For Gaussian pdfs, the covariance intersection technique [29], [32] is an instance of a log-linear pooling function. These and several other pooling functions will be discussed in Sections III–VII.

An alternative to the centralized setting for probabilistic opinion pooling described above would be a decentralized network of agents without a dedicated fusion center [8], [30], [33], [34]. Here, the agents communicate their pdfs only locally, i.e., to neighboring agents, and each agent can be considered to act as a local fusion center. In this “in-network” or “network-centric” type of probabilistic opinion pooling, the agents use a distributed communication-and-fusion protocol, such as flooding, consensus, gossip, or diffusion, to disseminate their local pdfs through the network and emulate a given overall pooling function. This relies on a suitable pdf representation, such as a Gaussian, Gaussian mixture, or particle representation. The fusion methods that we discuss in this work are also relevant to decentralized probabilistic opinion pooling. We note, however, that there are numerous methods for in-network signal and information processing in which the local processing results that are being combined are not pdfs. For example, some methods combine local likelihood functions [1], [35], [36] or messages within a message passing algorithm, such as belief propagation [7], [37], or certain iterated quantities within a networkwide adaptation-diffusion procedure [38], to name a few.

C. Relevance and Applications

Probabilistic opinion pooling is a fundamental and elementary functionality with widespread applications. Historically, the first motivation was to combine expert opinions into an aggregate opinion [26]. Nowadays, it is more likely that the different probability distributions do not represent the opinions of multiple experts but originate from the use of multiple sensors, models, processors, or datasets. In particular, probabilistic opinion pooling is often formulated in a Bayesian setting as the fusion of local *posterior* pdfs that are produced by multiple agents using local implementations of Bayesian inference [39]. The ideal aggregate pdf here is the global posterior pdf, which takes into account all the data available to the agents. However, the calculation of the global posterior pdf generally requires additional knowledge besides the local posterior pdfs, such as the local likelihood functions, the prior pdfs used by the agents, and possible statistical dependencies between the agents. By contrast, probabilistic opinion pooling requires only the local posterior pdfs. In many settings, it is easily and widely applicable because it does not make any assumptions about the local inference methods, the types of the sensors, or the nature of the local data, which can all be different at different agents.

From the viewpoint of the processed data, there is a wide range of scenarios for probabilistic opinion pooling. Two extreme cases are particularly important: all the agents process different data, or they process exactly the same data. Furthermore, the processing may be carried out with completely unrelated models but with the same objective (e.g., predicting future observations or classifying observations).

Current applications of probabilistic opinion pooling include, but are not limited to, the following selection.

- 1) In multisensor signal processing applications of probabilistic opinion pooling, multiple sensors derive local pdfs based on local observations and either submit these pdfs (or finite-dimensional representations thereof) to a fusion center or fuse them in a distributed, peer-to-peer manner [1], [3], [27], [30]–[33], [39]–[42]. In particular, probabilistic opinion pooling plays an important role in multisensor target tracking [2], [3], [5], [6], [8], [30], [43]–[46]. For tracking an unknown number of targets, probabilistic opinion pooling has recently also been applied to the “multi-object” pdfs or to the probability hypothesis densities (PHDs) (i.e., the “densities” of the first moment measures) of finite point processes, also known as random finite sets [2], [5], [6], [8], [43]–[50]. Although, in this work, we do not consider finite point processes, much of our discussion is also relevant in that domain. The application of probabilistic opinion pooling to multisensor target tracking will be discussed in more detail in Section II-A.
- 2) In probabilistic machine learning, several scenarios suggest the combination of probability distributions.

For example, the concept of ensemble learning [9] is based on applying multiple learning algorithms whose outputs are combined to obtain an aggregate result that is more accurate than that of any of the individual learning algorithms in the ensemble. Furthermore, in federated learning [12], [51], [52], multiple edge devices learn statistical models individually from their local datasets without explicitly exchanging these datasets, and a fusion center aggregates the learned models without having access to the original data. This is attractive for privacy-sensitive applications since no private data have to be shared. More details on probabilistic machine learning are provided in Section II-B.

- 3) The main goal in the combination of forecasts [53], [54] is the estimation of a parameter by combining several different models. To this end, certain methods perform a fusion of pdfs and usually refer to it as “combining density forecasts” [20], [55], [56]. This application will be addressed in more detail in Section II-C.
- 4) In Bayesian model averaging, several different models are used to derive different posterior pdfs based on the same data [57], [58]. An aggregate pdf is derived as a weighted average of the individual pdfs, where the weights are given by the posterior probabilities of the models. The Bayesian model averaging has been widely used in phylogenetics [59], [60], economics [21], [61], ecology [62], and many other fields [63], [64].
- 5) Traditional implementations of Monte Carlo-based inference schemes do not easily scale to large datasets (“big data”). A common expedient then is to partition the dataset into subsets and obtain a partial posterior pdf approximation for each subset. The partial approximations are subsequently fused into an approximation of the overall posterior pdf, which, thereby, takes into account the full dataset [65]–[67]. More details on this application are given in Section II-B. Another approach [68], [69] directly fuses sample representations of distributions by interpreting these samples as a weighted sum of Dirac measures.

To focus the scope of this work, we assume for the most part that the fusion center does not have any additional data about the random vector θ beyond the pdfs provided by the agents. (Here, an exception is given by the suprabayesian setting studied in Sections VIII and IX, where we assume that the fusion center knows a statistical model related to θ .) In particular, the fusion center cannot access any training data that were used by the agents and it does not have any validation data that it could use to validate the agents’ pdfs. Thus, although the fundamental problems are similar, we will not consider several ensemble learning methods, such as stacking [70], [71], or many other machine learning settings related to probabilistic

fusion [72]–[74]. Furthermore, given our focus on pdfs rather than discrete probability distributions, we will not touch upon methods tailored to the combination of classifiers, another large and growing field [75]. Finally, we are interested in obtaining a pdf and not merely a point estimate of θ . This is motivated by the fact that the pdf of θ contains all the probabilistic information about θ and can, thus, be used to obtain point estimates or other types of statistics. Hence, certain works on multimodel inference [76] and on the combination of forecasts [53], [54], [77] share some ideas with this work but ultimately have a different focus.

D. Approaches to Probabilistic Opinion Pooling

Although the probabilistic opinion pooling problem may appear simple and elementary, no single pooling function is universally accepted or uniformly best. Generally speaking, we would like the pooling function to involve the agent pdfs $q_k(\theta)$ in a way that follows some rationale. This rationale and the resulting choice of a pooling function may depend on the overall problem setting, application-specific aspects, side constraints, additional information available to the fusion center, and other considerations. The probabilistic opinion pooling problem has been studied for many decades, and substantial research efforts have been dedicated to the definition or derivation of pooling functions. One of the earliest works is [26], where the linear pooling function was introduced. Several survey articles on probabilistic opinion pooling with detailed literature reviews have been published [23], [78]–[80]; however, they often focus on discrete random variables.

In this work, we consider three principled approaches to defining a pooling function for pdfs. In what we call the *axiomatic approach*, the pooling function is defined indirectly by a set of properties (axioms) that it is required to satisfy. For example, it may be reasonable to require that the aggregate pdf $q(\theta)$ does not depend on the indexing order of the agent pdfs $q_k(\theta)$, or that for equal $q_k(\theta)$ —i.e., unanimity among all the agents—the aggregate pdf $q(\theta)$ conforms to that unanimous opinion. Most of the early literature in the field was dedicated to the axiomatic approach [23], [80]. An axiomatic approach is also adopted in the literature based on *imprecise probabilities* [81]. There, the idea is to define pooling operators that map from the agents’ probability mass functions (pmfs) to a set of pmfs rather than a single pmf. To the best of our knowledge, the concept of imprecise probabilities has so far been considered only for discrete probability spaces [81], [82].

In the *optimization approach*, the pooling function is the result of an optimization, i.e., the minimization or maximization of an objective function. Usually, the idea is that the aggregate pdf $q(\theta)$ should be as close as possible to all the agent pdfs $q_k(\theta)$ simultaneously. This can be formulated as a minimization involving an information-theoretic divergence [83]–[85] or a distance measure [48],

[84], [86]. The resulting optimum $q(\theta)$ can typically be interpreted as an “average” of $q_k(\theta)$.

Finally, the *supra-Bayesian approach* considers the fusion center as a Bayesian observer that interprets the agent pdfs $q_k(\theta)$ as random observations. This Bayesian observer builds on additional information about the dependence of these pdfs on θ (represented by the conditional probability distribution $p(q_1, \dots, q_K | \theta)$ of the random functions q_1, \dots, q_K given θ) to calculate a posterior pdf, which then constitutes the fusion result [87], [88]. Most of the early literature [89]–[91] describes $p(q_1, \dots, q_K | \theta)$ implicitly by assuming that the joint distribution of the errors $\mu_k - \theta$ (where μ_k is the expectation of θ induced by the pdf q_k) is multivariate Gaussian. This reduces the fusion problem to the calculation of the posterior pdf for a simple Bayesian linear Gaussian model where the μ_k are treated as observations at the fusion center and the covariance structure is known. The practically most important scenario in the supra-Bayesian framework is where each agent has access to certain random observations that are statistically dependent on the random vector θ , and both the agents and the fusion center have knowledge of a prior distribution of θ and the local likelihood functions of the agents. The agent pdf $q_k(\theta)$ is here given by the agent’s local posterior pdf. The fusion center is also aware of any statistical dependencies between the observations of different agents, which are described by a global likelihood function.

E. Contributions and Article Organization

The diversity of approaches, perspectives, and solutions related to probabilistic opinion pooling motivates a survey that presents the theory and methodology of the field in a coherent manner. This article attempts to answer this call. In addition, it provides a number of original contributions and results, including the following.

- 1) A rigorous and coherent treatment of probabilistic opinion pooling for a *continuous* random vector θ and, accordingly, for the fusion of pdfs. In particular, for the first time, the axiomatic approach is rigorously and thoroughly discussed for pdfs (see Section IV). So far, the focus in the literature has mostly been on discrete probability distributions, and it has been claimed that analogous results hold for pdfs. Although this is indeed often the case, the nonatomic structure of the pdf setting sometimes allows for stronger or different results.
- 2) The definition of a new pooling function, referred to as “generalized multiplicative pooling function” (see Section III-B8).
- 3) Two new axioms for pooling functions, referred to as “factorization preservation” and “generalized Bayesianity” (see Axioms 9 and 12).
- 4) Several new theorems presenting axiomatic characterizations of pooling functions for pdfs and related results (see Theorems 1, 2, and 6–11 and

Appendices A–F). These theorems are partly adaptations of existing results formulated for discrete probability distributions and partly entirely new results.

- 5) Proofs of the following results: The pooling function minimizing the weighted sum of α -divergences is given by the weighted Hölder mean; the pooling function minimizing the weighted sum of Pearson χ^2 -divergences is given by the weighted harmonic mean; and the pooling function minimizing the weighted sum of L_2 distances is given by the weighted arithmetic mean (see Theorems 14 and 16 and Appendices G and I). Furthermore, we derive the solution to the problem of minimizing a general class of weighted symmetric distance functions (see Theorem 17 and Appendix J).
- 6) A new framework of supra-Bayesian fusion of posterior pdfs in terms of finite-dimensional “local statistics” (see Sections VIII-B–VIII-D). Our results include an explicit pooling function for the case of agents collecting conditionally independent observations (see Theorem 18), a formal definition of and result for finite-dimensional supra-Bayesian fusion (see Definition 1 and Theorem 19), and a general procedure for establishing a fusion rule for the case of agents collecting conditionally dependent observations (see Section VIII-D).
- 7) A detailed study of supra-Bayesian fusion of posterior pdfs for linear Gaussian models (see Section IX), including the derivation of explicit pooling functions and fusion rules (see Sections IX-C and IX-D and Appendices K and L).

The structure of this article is given as follows. In Section II, we illustrate the applicability and relevance of probabilistic opinion pooling by discussing three specific example applications. In Section III, we formulate the probabilistic opinion pooling problem for pdfs and present a collection of specific pooling functions. Section IV discusses the axiomatic approach to opinion pooling and provides several new characterization theorems. In Section V, we consider the optimization approach to opinion pooling. We describe various optimization criteria and show that they partly lead to the same pooling functions as the axiomatic approach and partly to different pooling functions, such as the family of Hölder means. The fusion of Gaussian distributions using the pooling functions from Sections III and V is considered in Section VI. Section VII addresses the choice of the weights involved in the two most prominent and popular pooling functions, namely, the linear and log-linear pooling functions, as well as the choice of the parameter α involved in the Hölder pooling function. In Section VIII, we present a new view of the supra-Bayesian pooling approach using finite-dimensional parametrizations. The results of Section VIII are specialized to linear Gaussian models in Section IX. The model of Section IX includes as a special case the supra-Bayesian setting presented in [89]–[91]. We broaden this

setting significantly and present detailed fusion rules. In Section X, we provide suggestions for future research, and in Section XI, we provide a summary of our main insights and results. Detailed proofs of our main results are provided in the appendices.

F. Notation

We will use the following basic notation. Vectors are denoted by boldface lower case letters (e.g., \mathbf{t} and $\boldsymbol{\theta}$), matrices by boldface upper case letters (e.g., \mathbf{H} and $\boldsymbol{\Sigma}$), and sets and events by calligraphic letters (e.g., \mathcal{A}). The transpose is written as $(\cdot)^\top$. We write \mathbf{I}_d for the identity matrix of dimension d , $\mathbf{0}_{d_1 \times d_2}$ for the $d_1 \times d_2$ zero matrix, $\mathbf{1}_d$ for the all-one vector of dimension d , and \otimes for the Kronecker product. The symbol \mathcal{P} denotes the set of all pdfs, and \mathcal{S}_K denotes the probability simplex on $[0, 1]^K$, i.e., the set of all $(w_1, \dots, w_K) \in [0, 1]^K$ with $\sum_{k=1}^K w_k = 1$. For a set or event \mathcal{A} , we denote the complement as \mathcal{A}^c , the indicator function as $\mathbb{1}_{\mathcal{A}}$, and the Lebesgue measure as $|\mathcal{A}|$. Further notation is listed in the Nomenclature at the beginning of this article.

II. ILLUSTRATIVE APPLICATIONS

To illustrate the broad applicability of probabilistic opinion pooling or, more concretely, of the fusion of pdfs, we consider three illustrative applications in more detail.

A. Target Tracking

Target tracking aims to estimate the time-varying state (e.g., position and velocity) of a “target” from a sequence of observations [92], [93]. Applications include aeronautical and maritime situational awareness, surveillance, autonomous driving, biomedical analytics, remote sensing, and robotics. The performance of target tracking can be enhanced by using multiple sensors. This can be done in an optimal manner if the multisensor observation model is completely known, including possible statistical dependencies between the observations. However, in many cases, a simplified approach to multisensor target tracking based on probabilistic opinion pooling is adopted. Each sensor node operates a Bayesian filter that, at each time step, calculates a local posterior pdf of the current state based solely on the observation of that sensor. Fig. 1 illustrates the local posterior pdfs of two sensor nodes at two different time steps. The local posterior pdfs of the various sensor nodes are then fused using, typically, log-linear pooling or its second-order version known as covariance intersection [3], [27], [29]–[33], [40], [41], [94] (see Sections III-B3 and VI-B). This approach is practically convenient because: 1) the multisensor fusion is decoupled from the filtering and 2) it works for any choice of Bayesian filter methods used at the sensor nodes and for any sensing modalities, even when they are different at different sensor nodes. These characteristics make the probabilistic opinion pooling approach well suited to heterogeneous and/or decentralized sensor networks.

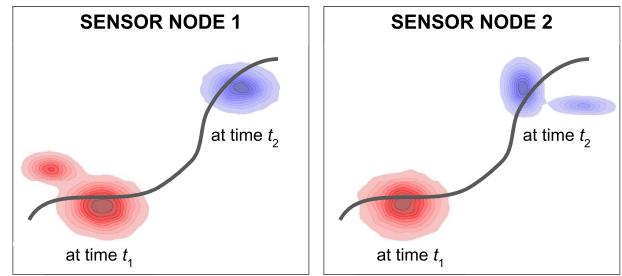


Fig. 1. Schematic illustration of the state trajectory of a target and the local posterior pdfs of two sensor nodes at two different time steps.

A nontrivial extension of target tracking is *multitarget* tracking, which involves an unknown time-varying number of targets and a more complicated observation model [95]–[102]. More specifically, targets can appear and disappear randomly, and there are missed detections (i.e., some sensors do not produce observations for some of the targets), clutter or false-alarm observations (which are not related to any target), and an observation-origin uncertainty (i.e., the sensor nodes do not know whether a given observation originated from a target, and from which target, or is clutter). Probabilistic opinion pooling can be used both for “vector-based” multitarget tracking methods, which describe the joint state of the targets by a random vector, and “set-based” methods, which describe it by a random finite set or equivalently a finite point process [97], [100], [103]. In the vector-based case, the target states are fused individually using, typically, log-linear pooling or covariance intersection. This presupposes an association of the target states across the sensors [104], [105].

In set-based methods, on the other hand, probabilistic opinion pooling is applied either to the posterior *multiobject pdfs* or the posterior PHDs of the sensor nodes, which provide two alternative *joint* descriptions of all the target states [97], [100]. Here, both log-linear pooling—also termed geometric average fusion, exponential mixture density, generalized covariance intersection, or Kullback–Leibler averaging [2], [5], [8], [43], [50], [85], [106]–[108]—and linear pooling (see Section III-B1)—also termed arithmetic average fusion and minimum information loss fusion [5], [43]–[46], [48], [85], [108]–[111]—have been used. Log-linear pooling is more sensitive to missed detections, whereas linear pooling is more sensitive to clutter. Regarding this sensitivity tradeoff, we note that pooling functions that are intermediate between the linear and log-linear ones are provided by the family of Hölder pooling functions to be presented and discussed in Sections III-B5 and V-C.

Finally, both log-linear pooling and linear pooling have recently been generalized to multitarget tracking methods based on *labeled* random finite sets, which track the identities of the targets in addition to their states [6],

[47], [49], [112]–[114]. Some of these methods require a label association step that is similar in spirit to the target association step required by vector-based methods [47], [49], [113], [114].

B. Probabilistic Machine Learning

Probabilistic machine learning [115], [116] has recently seen applications in many different areas, including quantum molecular dynamics [117], medical diagnosis [118], [119], scene understanding [120], and geotechnical engineering [121]. In machine learning, *uncertainty quantification* for predictive models is required for problems that involve risk assessment. Unfortunately, classical machine learning models do not account for parameter and prediction uncertainty, which makes them more susceptible to failure when dealing with unseen and/or unrelated data [122]. This is a prominent issue for deep learning models [123]. One way to account for predictive uncertainty in machine learning is to adopt a Bayesian framework: using training data, a prior pdf over the model parameters is updated to obtain a posterior pdf. This posterior pdf is then used to calculate a predictive pdf for unobserved data (test data). This pdf is often represented in a parametric form—e.g., a Gaussian pdf is parameterized by its mean and covariance matrix—or by a set of samples. Examples of Bayesian machine learning models include Bayesian linear regression, Bayesian neural networks [124], [125], Gaussian processes [126], and deep Gaussian processes [127].

In certain scenarios of probabilistic machine learning, probabilistic opinion pooling can be used to resolve practical challenges. For example, the choice of a model (or an architecture or a set of parameters) is frequently not obvious, and thus, there is a model uncertainty that has to be taken into account to ensure robustness and generalization. A class of methods dealing with this issue is known as *ensemble learning*. The learning is carried out by a collection of algorithms based on different models, and the final result of classification, regression, or clustering is obtained by combining the individual results [128]–[133]. The combination of the results of individual probabilistic learning algorithms can be implemented via probabilistic opinion pooling, i.e., by fusing the predictive pdfs produced by the individual algorithms. An example in the context of binary classification is shown in Fig. 2. Probabilistic opinion pooling in ensemble learning has been successfully applied, e.g., in the context of deep ensembles [9], neural network ensembles [10], and ensemble Gaussian processes [11]. We note that, in ensemble learning, unlike in multisensor signal processing and, in particular, target tracking as discussed in Section II-A, all the algorithms may operate on the same set of data.

Another practical challenge in machine learning is posed by privacy-sensitive scenarios. Here, local (private) data observed at individual nodes may not be disseminated across the nodes or to a fusion center and, thus, can be

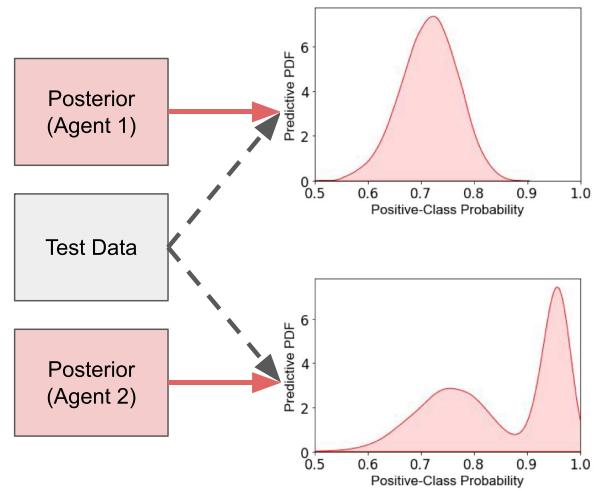


Fig. 2. Bayesian machine learning in the context of binary classification with two agents. Each agent obtains a posterior pdf from training data and uses it to derive its predictive pdf of the probability that test data belong to the positive class. These predictive pdfs are subsequently combined to obtain an aggregate predictive pdf.

used only to train local models at the respective nodes. This framework, often referred to as *federated learning*, requires the combination of local models at a fusion center [12], [51], [52], [134]. Although, in many instances of federated learning, updates are also communicated from the fusion center to the nodes, several works consider problem settings along the lines of probabilistic opinion pooling. For example, agnostic federated learning [135] combines sample representations of probability distributions trained on private data into an aggregate distribution.

Finally, the application of machine learning methods to “big data” scenarios calls for divide-and-conquer strategies that partition the data into much smaller sets, perform learning on each set, and combine the respective predictive or posterior distributions [65]–[67], [136]. Here, a focus has so far been on Markov chain Monte Carlo (MCMC) samplers for Bayesian inference [65]–[67]. For example, in [65], the idea is to generate a “subposterior” for each small dataset and combine the subposterior using the multiplicative pooling function (see Section III-B7). Each subposterior is initially represented by a set of samples produced by an MCMC sampler but is then converted into a continuous pdf given by a kernel density estimate. The different pdfs are finally fused to form an approximation to the overall posterior pdf. This approach can be motivated by the fact, to be shown in Section VIII-A, that, under a suitable conditional independence assumption, a multiplicative pooling function operating on the subposterior gives the overall posterior pdf.

The use of probabilistic machine learning has so far been restricted by the fact that many popular methods of machine learning do not provide probabilistic results. However, we expect that the outcomes of recent and

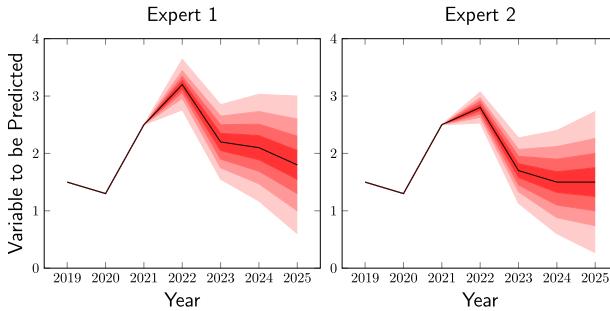


Fig. 3. Density forecasts of a variable (e.g., inflation) beyond 2021 made by two experts, visualized as fan charts. The values of past years are already observed and thus fixed, while predictions farther into the future become increasingly uncertain.

ongoing research will remove this limitation and thereby increase the successful application of probabilistic opinion pooling in this field.

C. Forecasting

The goal of forecasting is to predict future values of some variable of interest based on present and past observed data [137]. An issue that may limit the performance of forecasting is a lack of confidence in the underlying model. This issue can be addressed by the *combination of forecasts*, which fuses the forecasting results obtained with several different models [53], [54], [56]. While classical work has considered point forecasts, *probabilistic forecasting* uses a description of the variable of interest in terms of probability distributions. Here, for a long time, the focus was on discrete probability distributions [18], and accordingly, continuous random variables were approximated by discrete random variables through quantization. For example, in meteorology, the amount of precipitation was binned into a finite number of categories [138].

By contrast, the idea of *density forecasting* is to predict continuous random variables directly in terms of their pdfs [139]. This is visualized by Fig. 3, which shows a fan chart representation of density forecasts made by two experts. Density forecasting was suggested already more than 50 years ago [140], [141]. However, the combination of density forecasts [142]—which is a special setting of the fusion of pdfs—was considered only much later. Suggestions to combine density forecasts started with [20] and [55], which discussed the optimization of the weights in the linear pooling function based on training data. At about the same time, the use of Bayesian model averaging [57] in forecasting was proposed [143], again resulting in a linear pooling function. Also subsequent work focused on linear pooling [24]. Nonlinear pooling functions were mostly obtained by a preprocessing of the individual pdfs (e.g., in the spread-adjusted linear pool [144]) or by a postprocessing of the aggregate pdf (e.g., in the Beta-transformed linear pool [142]). Recently, the combination of density forecasts has also been studied in a

nonparametric Bayesian setting based on the Beta-transformed linear pool [145].

While the combination of density forecasts has the same goal as pdf fusion—namely, to fuse pdfs from different sources—there are two distinctive features. First, realizations of the random variable to be predicted are observed on a regular basis, which enables an evaluation of density forecasts and their combinations based on new data. A significant part of the literature focuses on this aspect. Although beyond the scope of our work, such an evaluation can obviously be performed also within the general setting of pdf fusion if the required data are available. Second, forecasts usually concern 1-D random variables. This implies that the combination of forecasts can be formulated in terms of the 1-D cumulative distribution function (cdf), and more specific properties, such as calibration [142], can be studied. Also the combination of forecasts—in particular, the choice of weights—is often based on new data and the evaluation of the fused 1-D cdf [146].

Probabilistic forecasting has been used in the broad domains of meteorology [18] and economics [20], [61], [139], [147] and, more specifically and more recently, in many disciplines, including wind forecasting [148], [149], electric load forecasting [150], electricity price forecasting [151], and solar power forecasting [152]. The combination of density forecasts has been considered, e.g., in [20], [149], and [152], and we conjecture that successful deployments of this variant of pdf fusion will emerge in many further applications of probabilistic forecasting.

III. PROBABILISTIC OPINION POOLING

A. Basic Framework

In probabilistic opinion pooling, we are interested in fusing the pdfs of K agents or “experts” into a single pdf. Let $\theta \in \Theta \subseteq \mathbb{R}^{d_\theta}$ be a continuous random variable or vector defined on some probability space.¹ Furthermore, let the pdf $q_k(\theta) \in \mathcal{P}$ represent the *opinion* of the k th agent. The sequence of all opinions $(q_1, q_2, \dots, q_K) \in \mathcal{P}^K$ is called the *opinion profile*. We consider events to be (measurable) subsets of Θ . The probability of an event $\mathcal{A} \subseteq \Theta$ according to the opinion of the k th agent is given by

$$Q_k(\mathcal{A}) = \int_{\mathcal{A}} q_k(\theta) d\theta.$$

Given an opinion profile (q_1, q_2, \dots, q_K) , a *pooling function* $g: \mathcal{P}^K \rightarrow \mathcal{P}$ is used to fuse the agents’ pdfs $q_k(\theta)$ into a single pdf

$$q(\theta) = g[q_1, \dots, q_K](\theta).$$

¹Our results extend to arbitrary probability measures that are absolutely continuous with respect to a σ -finite nonatomic measure. However, to keep the presentation more easily accessible, we present all results in the familiar setting of pdfs on \mathbb{R}^{d_θ} .

The probability of an event $\mathcal{A} \subseteq \Theta$ according to the fused pdf $q(\boldsymbol{\theta})$ is then given by

$$Q(\mathcal{A}) = \int_{\mathcal{A}} q(\boldsymbol{\theta}) d\boldsymbol{\theta}.$$

The fused pdf $q(\boldsymbol{\theta})$ summarizes the opinions of the K agents and will be referred to as the *aggregate pdf*. The fusion of the agent opinions via the pooling function is done (at least virtually) at a *fusion center*.

B. Pooling Functions

Over the years, many different pooling functions g have been proposed. We summarize some of them in the following. These pooling functions will be reconsidered in later sections.

1) *Linear Pooling*: A popular pooling function is the *linear pooling function*, which was introduced in [26]. Linear pooling aggregates the agent opinions through a weighted arithmetic average, i.e.,

$$g[q_1, \dots, q_K](\boldsymbol{\theta}) = \sum_{k=1}^K w_k q_k(\boldsymbol{\theta}) \quad (1)$$

where $(w_1, \dots, w_K) \in \mathcal{S}_K$.

One can establish a connection between linear opinion pooling and model averaging [57]. Let us consider the joint distribution $q(\boldsymbol{\theta}, M)$ of the unknown random vector $\boldsymbol{\theta}$ and a discrete “model” random variable $M \in \{M_1, \dots, M_K\}$. Furthermore, let $q(\boldsymbol{\theta}|M_k)$ denote the pdf of $\boldsymbol{\theta}$ conditioned on model M_k and let $P(M_k)$ denote the probability of M_k . Then, the marginal pdf of $\boldsymbol{\theta}$ is given by

$$q(\boldsymbol{\theta}) = \sum_{k=1}^K P(M_k) q(\boldsymbol{\theta}|M_k). \quad (2)$$

This is equivalent to the linear pooling operation (1), wherein the agent pdf $q_k(\boldsymbol{\theta})$ is interpreted as the pdf of $\boldsymbol{\theta}$ under model M_k , the weight w_k equals the probability of M_k , and the aggregate pdf $q(\boldsymbol{\theta})$ is the marginal pdf of $\boldsymbol{\theta}$.

2) *Generalized Linear Pooling*: The *generalized linear pooling function* defined in [153] includes an arbitrary pdf q_0 in the weighted arithmetic average (1), i.e.,

$$g[q_1, \dots, q_K](\boldsymbol{\theta}) = \sum_{k=0}^K w_k q_k(\boldsymbol{\theta}) \quad (3)$$

where $(w_0, \dots, w_K) \in \mathcal{S}_{K+1}$. We note that, in the general, measure-theoretic formulation of generalized linear opinion pooling in [153], some weights w_k are allowed to be negative. However, in the setting of fusing pdfs, this would result in a fusion rule g that does not give a valid (nonnegative) pdf for all possible opinion profiles

(q_1, \dots, q_K) . Thus, we restrict to nonnegative weights. One possible interpretation of the pdf q_0 is as the opinion of the fusion center. Alternatively, q_0 can be interpreted as a regularization.

3) *Log-Linear Pooling*: Another popular pooling function is the *log-linear pooling function* [28]. This function aggregates the agent opinions using a weighted geometric average, i.e.,

$$g[q_1, \dots, q_K](\boldsymbol{\theta}) = c \prod_{k=1}^K (q_k(\boldsymbol{\theta}))^{w_k} \quad (4)$$

where c is a normalization factor given by

$$c = \frac{1}{\int_{\Theta} \prod_{k=1}^K (q_k(\boldsymbol{\theta}))^{w_k} d\boldsymbol{\theta}} \quad (5)$$

and $(w_1, \dots, w_K) \in \mathcal{S}_K$. To avoid the possibility of the integral in (5) being zero and, thus, c being undefined, this pooling function is usually only defined for pdfs that are positive on the domain Θ . We will refer to opinion profiles (q_1, \dots, q_K) that satisfy

$$q_k(\boldsymbol{\theta}) > 0 \quad \text{for all } \boldsymbol{\theta} \in \Theta \quad (6)$$

as *positive opinion profiles*.

The pooling function is called “log-linear” because it is a linear function of the agent pdfs in the log-domain, i.e., the logarithm of the right-hand side of (4) is

$$\log \left(c \prod_{k=1}^K (q_k(\boldsymbol{\theta}))^{w_k} \right) = \log(c) + \sum_{k=1}^K w_k \log(q_k(\boldsymbol{\theta}))$$

which is a weighted arithmetic average [up to the additive constant $\log(c)$]. We will, therefore, refer to the powers w_1, \dots, w_K as “weights.”

4) *Generalized Log-Linear Pooling*: Similar to the generalized linear pooling function, a generalization of the log-linear pooling function can be obtained by including a function ξ_0 as an additional factor. However, in contrast to the generalized linear pooling function, ξ_0 is not necessarily a pdf. More specifically, the *generalized log-linear pooling function* [154] is defined as

$$g[q_1, \dots, q_K](\boldsymbol{\theta}) = c \xi_0(\boldsymbol{\theta}) \prod_{k=1}^K (q_k(\boldsymbol{\theta}))^{w_k} \quad (7)$$

where

$$c = \frac{1}{\int_{\Theta} \xi_0(\boldsymbol{\theta}) \prod_{k=1}^K (q_k(\boldsymbol{\theta}))^{w_k} d\boldsymbol{\theta}}$$

ξ_0 is a bounded, positive function, and $(w_1, \dots, w_K) \in \mathcal{S}_K$. Here, we again restrict to positive opinion profiles.

The function ξ_0 can be used, e.g., to include the opinion of the fusion center or regularize the fused density.

5) *Hölder Pooling*: The following pooling function was apparently first suggested in [78] as a generalization of the linear and log-linear pooling functions:

$$g[q_1, \dots, q_K](\boldsymbol{\theta}) = c \left(\sum_{k=1}^K w_k (q_k(\boldsymbol{\theta}))^\alpha \right)^{1/\alpha} \quad (8)$$

where

$$c = \frac{1}{\int_{\Theta} \left(\sum_{k=1}^K w_k (q_k(\boldsymbol{\theta}))^\alpha \right)^{1/\alpha} d\boldsymbol{\theta}}$$

and $\alpha \in \mathbb{R} \setminus \{0\}$. While, for $\alpha \geq 1$, it can be shown that c is defined for arbitrary opinion profiles, in the other cases, we have to restrict to opinion profiles such that c is defined. Because the pooling function in (8) is the weighted Hölder mean (also called the generalized average) [155] of the agent pdfs $q_k(\boldsymbol{\theta})$, we will refer to (8) as the *Hölder pooling function*. The linear and log-linear pooling functions are special cases of the Hölder pooling function for $\alpha = 1$ and $\alpha \rightarrow 0$, respectively.

6) *Inverse-Linear Pooling*: The *inverse-linear pooling function* (weighted harmonic average) is defined as

$$g[q_1, \dots, q_K](\boldsymbol{\theta}) = c \left(\sum_{k=1}^K \frac{w_k}{q_k(\boldsymbol{\theta})} \right)^{-1} \quad (9)$$

where

$$c = \frac{1}{\int_{\Theta} \left(\sum_{k=1}^K \frac{w_k}{q_k(\boldsymbol{\theta})} \right)^{-1} d\boldsymbol{\theta}}.$$

This is the special case of the Hölder pooling function for $\alpha = -1$.

7) *Multiplicative Pooling*: The *multiplicative pooling function*, as proposed in [80] for pmfs, is defined as

$$g[q_1, \dots, q_K](\boldsymbol{\theta}) = c (q_0(\boldsymbol{\theta}))^{1-K} \prod_{k=1}^K q_k(\boldsymbol{\theta}) \quad (10)$$

where

$$c = \frac{1}{\int_{\Theta} (q_0(\boldsymbol{\theta}))^{1-K} \prod_{k=1}^K q_k(\boldsymbol{\theta}) d\boldsymbol{\theta}}$$

and q_0 is a positive pdf called the *calibrating pdf*. Here, we restrict to positive opinion profiles and further assume that $q_k(\boldsymbol{\theta})/q_0(\boldsymbol{\theta})$ is bounded for all $k = 1, \dots, K$. These assumptions guarantee that the normalization constant c is well-defined and nonzero. In Section VIII-A, we will show that, within the supra-Bayesian framework, the multiplicative pooling function is the correct fusion rule for combining posterior pdfs in the case of conditionally independent observations. In that case, the calibrating pdf q_0 is the prior pdf used by the agents to form their posterior pdfs.

8) *Generalized Multiplicative Pooling*: We propose another pooling function that is a generalization of both the generalized log-linear pooling function and the multiplicative pooling function. In addition to a calibrating pdf q_0 , we also allow for arbitrary weights in the generalized log-linear pooling function (7). More specifically, we define the *generalized multiplicative pooling function* as

$$g[q_1, \dots, q_K](\boldsymbol{\theta}) = c (q_0(\boldsymbol{\theta}))^{1-\sum_{k=1}^K w_k} \prod_{k=1}^K (q_k(\boldsymbol{\theta}))^{w_k} \quad (11)$$

where

$$c = \frac{1}{\int_{\Theta} (q_0(\boldsymbol{\theta}))^{1-\sum_{k=1}^K w_k} \prod_{k=1}^K (q_k(\boldsymbol{\theta}))^{w_k} d\boldsymbol{\theta}}$$

q_0 is a positive calibrating pdf, and the weights $w_1, \dots, w_K \in \mathbb{R}$ are arbitrary real numbers. We again restrict to positive opinion profiles and assume that $(q_k(\boldsymbol{\theta})/q_0(\boldsymbol{\theta}))^{w_k}$ is bounded for all $k = 1, \dots, K$. In Section IX-C, we will show that, within the supra-Bayesian framework with a linear Gaussian model, the generalized multiplicative pooling function is the correct fusion rule for combining posterior pdfs.

9) *Dictatorship Pooling*: The *dictatorship pooling function* maps the opinion profile to a single agent opinion, i.e.,

$$g[q_1, \dots, q_K](\boldsymbol{\theta}) = q_k(\boldsymbol{\theta}) \quad (12)$$

for some fixed $k \in \{1, \dots, K\}$. Although this function is a valid pooling function, one would not normally expect it to be a good choice.

10) *Dogmatic Pooling*: The *dogmatic pooling function* enforces a fixed pdf q_0 independently of the opinion profile, i.e.,

$$g[q_1, \dots, q_K](\boldsymbol{\theta}) = q_0(\boldsymbol{\theta}). \quad (13)$$

Again, this pooling function will not be suitable in most applications.

IV. AXIOMATIC APPROACH

Fundamentally, we would like the pooling function $g[q_1, \dots, q_K]$ to depend directly on all the agent pdfs q_k in a way that follows some rationale. One principled approach to probabilistic opinion pooling is the axiomatic approach, which seeks to determine all pooling functions that satisfy a set of desirable properties (axioms). In this section, we first formulate some axioms and then rigorously analyze the relationships between these axioms and the pooling functions presented in Section III-B.

A. Axioms

To begin, one basic restriction we may impose on the pooling function is that it be a symmetric function,

i.e., a function whose arguments can be interchanged without altering the output of the function. This means that there is no “natural order” of the agents, and all agents are treated equally. This is formally stated in the following axiom.

Axiom 1 (Symmetry): For all permutations $\beta: \mathcal{K} \rightarrow \mathcal{K}$ of the set $\mathcal{K} = \{1, \dots, K\}$ and all opinion profiles (q_1, \dots, q_K) , the pooling function g satisfies

$$g[q_1, \dots, q_K](\theta) = g[q_{\beta(1)}, \dots, q_{\beta(K)}](\theta).$$

A symmetric pooling function seems to be desirable and natural since it treats the pdfs of the agents equally at the fusion center. However, if certain agents are known *a priori* to be more “reliable” or “informative” than other agents, then it may be reasonable to emphasize them in the pooling function. For example, in the linear or log-linear pooling function, we may assign larger weights w_k . If this is done in a fixed manner, the pooling function is no longer symmetric. On the other hand, if the weights are chosen adaptively such that each weight is an explicit function of the opinion profile and this adaptation rule involves each agent in the same way, then all agents are treated equally and the resulting pooling function is still symmetric. This will be further discussed in Section VII.

Another basic property for a pooling function is the preservation of agreement among agents. For instance, if each of the agents believes that a certain event $\mathcal{A} \subset \Theta$ is a null event, i.e., the probability of \mathcal{A} is 0 according to all the agents, then \mathcal{A} should also be a null event according to the aggregate pdf. This property is called the *zero preservation property* (ZPP) [156].

Axiom 2 (ZPP): For any event $\mathcal{A} \subset \Theta$, if $Q_k(\mathcal{A}) = 0$ for all k , then $Q(\mathcal{A}) = 0$.

The next property, termed *unanimity preservation* [80], asserts that, if the opinions of the agents are identical, then the aggregate pdf should conform to that unanimous opinion.

Axiom 3 (Unanimity Preservation): If, for all events $\mathcal{A} \subset \Theta$, the probabilities $Q_k(\mathcal{A}) = p_{\mathcal{A}}$ coincide for all k , then $Q(\mathcal{A}) = p_{\mathcal{A}}$. Equivalently, if² $q_k(\theta) = q_0(\theta)$ for all k and some pdf $q_0(\theta)$, then $q(\theta) = q_0(\theta)$.

Another property that may be desirable in a pooling function is the *strong setwise function property* (SSFP) [156]. The SSFP states that the probability of an event $\mathcal{A} \subset \Theta$ according to the aggregate pdf $q(\theta)$ can be expressed as a function of the probabilities of that event according to each agent, i.e., $Q_1(\mathcal{A}), \dots, Q_K(\mathcal{A})$.

Axiom 4 (SSFP): There exists a function $h: [0, 1]^K \rightarrow [0, 1]$ such that, for all opinion profiles (q_1, \dots, q_K) and for all events $\mathcal{A} \subseteq \Theta$,

$$Q(\mathcal{A}) = h(Q_1(\mathcal{A}), \dots, Q_K(\mathcal{A})). \quad (14)$$

²We consider two pdfs to be equal if they are equal almost everywhere with respect to the Lebesgue measure.

We note that this axiom is in general not equivalent to the property that there exists a function $\tilde{h}: [0, \infty)^K \rightarrow [0, \infty)$ such that, for all opinion profiles (q_1, \dots, q_K) and each point $\theta \in \Theta$,

$$q(\theta) = \tilde{h}(q_1(\theta), \dots, q_K(\theta)). \quad (15)$$

In particular, for the case that Θ has finite Lebesgue measure $|\Theta|$, the dogmatic pooling function $q(\theta) = 1/|\Theta|$ for $\theta \in \Theta$ trivially satisfies (15) but not (14) (as a simple consequence of Theorem 1 in Section IV-B).

A more relaxed criterion than the SSFP is the *weak setwise function property* (WSFP) [156]. The WSFP states that the probability of an event according to the aggregate pdf is a function of the probabilities of that event according to each agent *and* the event itself.

Axiom 5 (WSFP): For all events $\mathcal{A} \subseteq \Theta$, there exists a generally \mathcal{A} -dependent function $h_{\mathcal{A}}: [0, 1]^K \rightarrow [0, 1]$ such that, for all opinion profiles (q_1, \dots, q_K) ,

$$Q(\mathcal{A}) = h_{\mathcal{A}}(Q_1(\mathcal{A}), \dots, Q_K(\mathcal{A})). \quad (16)$$

The WSFP is also equivalent to the so-called *marginalization property*, which states that marginalization and fusion are commutative operations. Formulating the marginalization property requires a measure-theoretic language that is beyond the scope of this article. We, thus, omit a discussion of the marginalization property and refer the interested reader to [153] and [156].

Another relaxation of the SSFP is the *likelihood principle* [28]. Here, the value of the aggregate pdf $q(\theta)$ at some point θ may only depend on the values of all $q_k(\theta)$ at the same θ up to a normalization constant that can depend on the opinion profile.

Axiom 6 (Likelihood Principle): There exists a function $h: [0, \infty)^K \rightarrow [0, \infty)$ such that, for all opinion profiles (q_1, \dots, q_K) and each point $\theta \in \Theta$,

$$q(\theta) = \frac{h(q_1(\theta), \dots, q_K(\theta))}{\int_{\Theta} h(q_1(\theta'), \dots, q_K(\theta')) d\theta'}.$$

The name “likelihood principle” is motivated by viewing the pdfs as normalized likelihood functions: in this viewpoint, the idea is that the fused likelihood at θ should only depend on the local likelihoods at θ up to normalization [28]. Note that (15) is a significantly stronger assumption because the function \tilde{h} in (15) has to normalize to one.

We can also formulate a weak version of the likelihood principle, where the function h may depend on θ [28].

Axiom 7 (Weak Likelihood Principle): For all $\theta \in \Theta$, there exists a generally θ -dependent function $h_{\theta}: [0, \infty)^K \rightarrow [0, \infty)$ such that, for all opinion

profiles (q_1, \dots, q_K) ,

$$q(\boldsymbol{\theta}) = \frac{h_{\boldsymbol{\theta}}(q_1(\boldsymbol{\theta}), \dots, q_K(\boldsymbol{\theta}))}{\int_{\Theta} h_{\boldsymbol{\theta}}(q_1(\boldsymbol{\theta}'), \dots, q_K(\boldsymbol{\theta}')) d\boldsymbol{\theta}'}.$$

Another important axiom is *independence preservation*³ [157]. This axiom asserts that, if all the agents agree that two events $\mathcal{A}, \mathcal{B} \subseteq \Theta$ are independent, then these events should be independent also according to the aggregate pdf.

Axiom 8 (Independence Preservation): For any events $\mathcal{A}, \mathcal{B} \subseteq \Theta$, if

$$Q_k(\mathcal{A} \cap \mathcal{B}) = Q_k(\mathcal{A})Q_k(\mathcal{B})$$

for all $k \in \{1, \dots, K\}$, then $Q(\mathcal{A} \cap \mathcal{B}) = Q(\mathcal{A})Q(\mathcal{B})$.

A relaxation of independence preservation that, to the best of our knowledge, has not been considered before is to assume the preservation of a given factorization structure.

Axiom 9 (Factorization Preservation): For any functions $f_1: \Theta \rightarrow \mathbb{R}^{d_1}$ and $f_2: \Theta \rightarrow \mathbb{R}^{d_2}$, if there exist functions $q_{k,1}$ and $q_{k,2}$ such that

$$q_k(\boldsymbol{\theta}) = q_{k,1}(f_1(\boldsymbol{\theta}))q_{k,2}(f_2(\boldsymbol{\theta}))$$

for all $k \in \{1, \dots, K\}$, then there exist functions $q_{a,1}$ and $q_{a,2}$ such that

$$q(\boldsymbol{\theta}) = q_{a,1}(f_1(\boldsymbol{\theta}))q_{a,2}(f_2(\boldsymbol{\theta})).$$

This axiom expresses, in particular, preservation of the independence of components of $\boldsymbol{\theta}$. Assume that $\boldsymbol{\theta} = (\boldsymbol{\theta}_1, \boldsymbol{\theta}_2)$ and all agent pdfs factor according to $q_k(\boldsymbol{\theta}) = q_{k,1}(\boldsymbol{\theta}_1)q_{k,2}(\boldsymbol{\theta}_2)$. Then, choosing $f_1(\boldsymbol{\theta}) = \boldsymbol{\theta}_1$ and $f_2(\boldsymbol{\theta}) = \boldsymbol{\theta}_2$, factorization preservation implies that also the aggregate pdf preserves the independence of $\boldsymbol{\theta}_1$ and $\boldsymbol{\theta}_2$, i.e., $q(\boldsymbol{\theta}) = q_{a,1}(\boldsymbol{\theta}_1)q_{a,2}(\boldsymbol{\theta}_2)$.

The final axioms that we consider are motivated by Bayesian updating of probabilities. More specifically, we interpret each agent pdf $q_k(\boldsymbol{\theta})$ as the agent's belief about an unknown quantity $\boldsymbol{\theta}$ after observing some data. When observing new (additional) data, $q_k(\boldsymbol{\theta})$ is updated by multiplying it by a likelihood function $\ell: \Theta \rightarrow [0, \infty)$, which relates the agent's new data to $\boldsymbol{\theta}$. The updated belief of the k th agent, $q_k^{(\ell)}(\boldsymbol{\theta})$, is, thus, given as

$$q_k^{(\ell)}(\boldsymbol{\theta}) = \frac{\ell(\boldsymbol{\theta})q_k(\boldsymbol{\theta})}{\int_{\Theta} \ell(\boldsymbol{\theta}')q_k(\boldsymbol{\theta}') d\boldsymbol{\theta}'}.$$
 (17)

To avoid degenerate cases, one usually assumes in the following axioms that all pdfs are positive on the domain Θ . Thus, we restrict the statements of the axioms to positive

³Independence preservation should not be confused with the WSFP, which is sometimes referred to as the independence or eventwise independence property (e.g., [80]).

opinion profiles. The first axiom related to the Bayesian framework is known as *external Bayesianity* [28], [158].

Axiom 10 (External Bayesianity): For all functions $\ell: \Theta \rightarrow [0, \infty)$ and all positive opinion profiles (q_1, \dots, q_K) satisfying $0 < \int_{\Theta} \ell(\boldsymbol{\theta})q_k(\boldsymbol{\theta}) d\boldsymbol{\theta} < \infty$ for all $k \in \{1, \dots, K\}$, we have

$$q^{(\ell)}(\boldsymbol{\theta}) = g\left[q_1^{(\ell)}, \dots, q_K^{(\ell)}\right](\boldsymbol{\theta})$$

where $q_k^{(\ell)}$ is given in (17) and

$$q^{(\ell)}(\boldsymbol{\theta}) \triangleq \frac{\ell(\boldsymbol{\theta})q(\boldsymbol{\theta})}{\int_{\Theta} \ell(\boldsymbol{\theta}')q(\boldsymbol{\theta}') d\boldsymbol{\theta}'} \quad (18)$$

with $q(\boldsymbol{\theta}) = g[q_1, \dots, q_K](\boldsymbol{\theta})$.

This axiom is motivated by the following Bayesian scenario. Assume that q_1, \dots, q_K are prior pdfs of K agents. Some data are observed, and the resulting likelihood function ℓ is provided to all agents. Then, a pooling function g satisfying external Bayesianity gives the same fusion result if it first aggregates the priors q_k into a fused prior q , and then q is updated according to (18), or if it aggregates the posterior pdfs $q_k^{(\ell)}$ resulting from all agents updating their priors according to (17). Thus, external Bayesianity states that pdf updating and fusion are commutative operations. Such a property is desirable in applications where the agents share identical data (i.e., a global likelihood function) but have distinct prior distributions [159].

The second axiom related to the Bayesian framework is known as *individualized Bayesianity* [80]. This axiom is motivated by the idea of combining posterior probabilities, where each agent's posterior probability is based on private data (i.e., a local likelihood function) in contrast to all agents sharing identical data.

Axiom 11 (Individualized Bayesianity): For all $k \in \{1, \dots, K\}$, all bounded, positive⁴ functions $\ell: \Theta \rightarrow [0, \infty)$, and all positive opinion profiles (q_1, \dots, q_K) , we have

$$q^{(\ell)}(\boldsymbol{\theta}) = g\left[q_1, \dots, q_{k-1}, q_k^{(\ell)}, q_{k+1}, \dots, q_K\right](\boldsymbol{\theta}) \quad (19)$$

where $q_k^{(\ell)}$ and $q^{(\ell)}$ are defined by (17) and (18), respectively.

This axiom is motivated by a scenario that is partly different from the scenario motivating external Bayesianity. We again assume that q_1, \dots, q_K are prior pdfs of the agents. For some arbitrary but fixed k , the k th agent observes (private) data in terms of a likelihood function ℓ . Then, a pooling function g satisfying individualized Bayesianity gives the same fusion result if it first aggregates the priors q_k into a fused prior q , and then q is updated according to (18), or if it aggregates the priors of all but the k th agent and the posterior pdf $q_k^{(\ell)}$ resulting from the k th agent updating its prior according

⁴The assumption of boundedness and positivity is needed to obtain the characterization theorems involving individualized Bayesianity in Section IV-B.

Table 1 Axioms Satisfied by the Pooling Functions Presented in Section III-B. (*: Satisfied If and Only If All Weights Are Equal)

Pooling Function	Axiom											
	1	2	3	4	5	6	7	8	9	10	11	12
Linear	*	✓	✓	✓	✓	✓	✓					
Generalized Linear	*				✓		✓					
Log-linear	*	n.a.	✓			✓	✓		✓	✓		✓
Generalized Log-linear	*	n.a.				✓	✓		✓	✓		✓
Hölder	*	n.a.	✓			✓	✓					
Inverse-linear	*	n.a.	✓			✓	✓					
Multiplicative	✓	n.a.				✓		✓		✓		✓
Generalized Multiplicative	*	n.a.				✓		✓				✓
Dictatorship		✓	✓	✓	✓	✓	✓	✓	✓	✓		✓
Dogmatic	✓				✓		✓					✓

to (17). Thus, individualized Bayesianity states that pdf updating at a *single* agent and fusion are commutative operations.

Finally, we state a novel axiom that generalizes individualized Bayesianity. We, thus, call it *generalized Bayesianity*.

Axiom 12 (Generalized Bayesianity): For all bounded, positive functions $\ell_k: \Theta \rightarrow [0, \infty)$, $k \in \{1, \dots, K\}$, there exists a fused likelihood function $h[\ell_1, \dots, \ell_K]$ such that, for all positive opinion profiles (q_1, \dots, q_K) , we have

$$q^{(h[\ell_1, \dots, \ell_K])}(\theta) = g[q_1^{(\ell_1)}, \dots, q_K^{(\ell_K)}](\theta) \quad (20)$$

where $q_k^{(\ell_k)}$ and $q^{(h[\ell_1, \dots, \ell_K])}$ are defined by (17) and (18), respectively.

This axiom states that fusing $q_1^{(\ell_1)}, \dots, q_K^{(\ell_K)}$, i.e., the result of updating q_1, \dots, q_K , is equivalent to updating q , i.e., the result of fusing q_1, \dots, q_K , by a “fused likelihood function” $h[\ell_1, \dots, \ell_K]$. Note that the fused likelihood function is not allowed to depend on the opinion profile (q_1, \dots, q_K) .

The axioms related to the Bayesian framework presented above are not directly related to the supra-Bayesian approach presented in Sections VIII and IX. More specifically, in the supra-Bayesian framework, we have explicit likelihood functions, and thus, the pooling function does not necessarily satisfy properties that relate to arbitrary likelihood functions as in the axioms above.

B. Relations Between Axioms and Pooling Functions

Having presented various pooling functions in Section III-B and various axioms in Section IV-A, we next analyze which pooling functions satisfy which axioms and, conversely, which axioms imply which pooling functions. Our results are summarized in Table 1. In what follows, we will abbreviate the various axioms as A1, A2, and so on.

Theorem 1: The linear pooling function in (1) satisfies the ZPP (A2), unanimity preservation (A3), the SSFP (A4), the WSFP (A5), the likelihood principle (A6), and the

weak likelihood principle (A7). In addition, it satisfies the symmetry axiom (A1) if and only if all weights are equal, i.e., $w_1 = w_2 = \dots = w_K = 1/K$. Furthermore, for a pooling function g , the following statements are equivalent.

- 1) g is a linear pooling function.
- 2) g satisfies the SSFP (A4).
- 3) g satisfies the WSFP (A5) and the ZPP (A2).
- 4) g satisfies the WSFP (A5) and unanimity preservation (A3).

The equivalence of 1–3 was first proven in [156] for pmfs and in [153] for arbitrary probability measures. However, to the best of our knowledge, a proof for pdfs has not been provided so far.⁵ In [80], the equivalence of 4 and 3 was presented for pmfs. In Appendix A, we give a proof of Theorem 1 for pdfs.

Theorem 2: The generalized linear pooling function in (3) satisfies the WSFP (A5) and the weak likelihood principle (A7). Conversely, any pooling function that satisfies the WSFP (A5) is a generalized linear pooling function. In addition, the generalized linear pooling function satisfies the symmetry axiom (A1) if and only if all weights except w_0 are equal, i.e., $w_1 = w_2 = \dots = w_K$.

The measure-theoretic equivalence of generalized linear pooling functions with possibly negative weights and pooling functions satisfying the WSFP (A5) was proven in [153]. However, in the case of the fusion of pdfs considered here, the generalized linear pooling functions cannot have negative weights. Accordingly, we present a proof with the necessary adaptations in Appendix B.

We next turn to pooling functions that include multiplication of pdfs or of powers of pdfs. In this context, we restrict to positive opinion profiles, i.e., we assume that (6) is satisfied. Note that, in this setting, the ZPP (A2) is not

⁵Note that the proof for arbitrary probability measures in [153] does not imply the result for pdfs. Indeed, in our pdf framework, only probability measures that are absolutely continuous with respect to a fixed reference measure (usually the Lebesgue measure) are considered. This implicates the following difference from the framework of [153]: whereas we only assume that an axiom holds for all pdfs, in [153] it is assumed to hold also for other probability measures such as, e.g., a Dirac measure. Therefore, if in [153] it is stated that, e.g., the assumption 2 implies 1, then this refers to a stronger version of 2.

applicable since $Q_K(\mathcal{A}) = 0$ is not possible except for sets \mathcal{A} of Lebesgue measure zero; therefore, we will disregard the ZPP in the following considerations.

Theorem 3: The log-linear pooling function in (4) satisfies unanimity preservation (A3), the likelihood principle (A6), the weak likelihood principle (A7), factorization preservation (A9), external Bayesianity (A10), and generalized Bayesianity (A12). In addition, it satisfies the symmetry axiom (A1) if and only if all weights are equal, i.e., $w_1 = w_2 = \dots = w_K = 1/K$. Furthermore, for a pooling function g , the following statements are equivalent.

- 1) g is a log-linear pooling function.
- 2) g satisfies the likelihood principle (A6) and external Bayesianity (A10).
- 3) g satisfies unanimity preservation (A3), the weak likelihood principle (A7), and external Bayesianity (A10).

The equivalence of 1 and 2 was proven in [28] and the equivalence of 1 and 3 in [154]. The remaining claimed axioms follow straightforwardly from the definition of the log-linear pooling function in (4).

Theorem 4: The generalized log-linear pooling function in (7) satisfies the weak likelihood principle (A7), factorization preservation (A9), external Bayesianity (A10), and generalized Bayesianity (A12). In addition, it satisfies the symmetry axiom (A1) if and only if all weights except w_0 are equal, i.e., $w_1 = w_2 = \dots = w_K$. Furthermore, for a pooling function g , the following statements are equivalent.

- 1) g is a generalized log-linear pooling function.
- 2) g satisfies the weak likelihood principle (A7) and external Bayesianity (A10).

This characterization theorem was proven in [154]. We note that the assumption of fusing pdfs (rather than general measures) is essential. In particular, for pmfs, axioms A7 and A10 would imply only a modified generalized log-linear pooling function that may contain negative weights [154]. In [154], one can also find a characterization of all pooling functions that satisfy external Bayesianity (A10). However, these pooling functions do not have a simple structure.

Theorem 5: The Hölder pooling function in (8) satisfies unanimity preservation (A3), the likelihood principle (A6), and the weak likelihood principle (A7). In addition, it satisfies the symmetry axiom (A1) if and only if all weights are equal, i.e., $w_1 = w_2 = \dots = w_K$.

The proof of this theorem is straightforward and, thus, omitted. Because the inverse-linear pooling function (9) is a special case of the Hölder pooling function, it follows that it also satisfies A3, A6, and A7.

Theorem 6: The multiplicative pooling function in (10) satisfies the symmetry axiom (A1), the weak likelihood principle (A7), factorization preservation (A9), individualized Bayesianity (A11), and generalized Bayesianity (A12). Furthermore, for a pooling function g , the following statements are equivalent.

- 1) g is a multiplicative pooling function with calibrating pdf q_0 .

- 2) g satisfies individualized Bayesianity (A11), and there exists a pdf $q_0(\theta)$ such that $g[q_0, \dots, q_0](\theta) = q_0(\theta)$.

The claimed axioms follow straightforwardly from the definition of the pooling function. A result similar to the equivalence of 1 and 2 was proven for pmfs in [80]. We provide a proof for pdfs in Appendix C.

Theorem 7: The generalized multiplicative pooling function in (11) satisfies the weak likelihood principle (A7), factorization preservation (A9), and generalized Bayesianity (A12). In addition, it satisfies the symmetry axiom (A1) if and only if all weights are equal, i.e., $w_1 = w_2 = \dots = w_K$.

Again, the claimed axioms follow straightforwardly from the definition of the pooling function.

Theorem 8: The dictatorship pooling function in (12) satisfies the ZPP (A2), unanimity preservation (A3), the SSFP (A4), the WSFP (A5), the likelihood principle (A6), the weak likelihood principle (A7), independence preservation (A8), factorization preservation (A9), external Bayesianity (A10), and generalized Bayesianity (A12). Furthermore, for a pooling function g , the following statements are equivalent.

- 1) g is a dictatorship pooling function.
- 2) g satisfies the SSFP (A4) and independence preservation (A8).
- 3) g satisfies the WSFP (A5) and independence preservation (A8).
- 4) g satisfies the SSFP (A4) and external Bayesianity (A10).
- 5) g satisfies the WSFP (A5) and external Bayesianity (A10).
- 6) g satisfies the SSFP (A4) and generalized Bayesianity (A12).

Our statements regarding the satisfied axioms follow easily from the definition of the dictatorship pooling function. The equivalence of 1 and 2 was proven in [153, Th. 3.1]. In Appendix D, we strengthen this result and show that the WSFP—instead of the (stronger) SSFP—in combination with independence preservation suffices to axiomatically define the dictatorship pooling function, i.e., that 3 implies 2. The equivalence of 1 and 4 was proven in [160]. In fact, in [160], even the equivalence of 1 and 5 is shown by proving that the version of external Bayesianity considered in [160] implies the ZPP. However, our formulation of external Bayesianity assumes positive opinion profiles, and thus, the ZPP cannot be proven. To close this gap, we further show in Appendix D that 5 implies 4. Finally, we also show in Appendix D that 6 implies 1.

We note that the dictatorship pooling function is a special case of both the linear and log-linear pooling functions when one of the weights is 1 and all the others are 0. The fact that the dictatorship pooling function satisfies ten axioms shows that a pooling function that satisfies many axioms is not necessarily a useful pooling function.

Turning to the dogmatic pooling function, we first present a preliminary result that is proven in Appendix E.

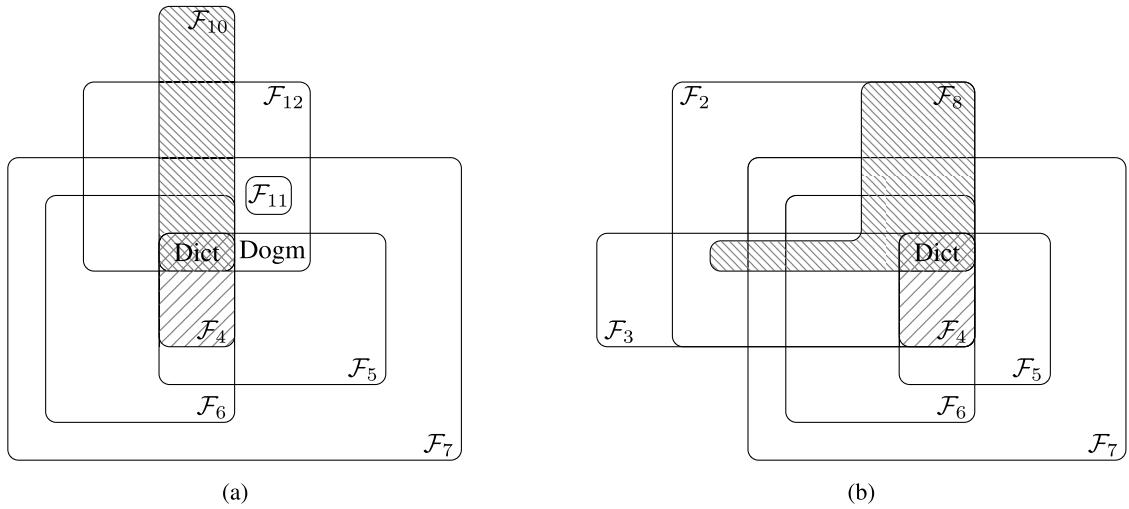


Fig. 4. Venn diagrams representing the implication structure for the axioms from Section IV-A. (a) A4–A7 and A10–A12, as well as intersections resulting in dictatorship (Dict) or dogmatic (Dogm) pooling functions. (b) A2–A8, as well as intersections resulting in dictatorship (Dict) pooling functions. Note that the diagrams illustrate the currently known implications, and some regions that appear nonempty in the diagrams may actually be empty sets. For better visibility, the sets \mathcal{F}_4 , \mathcal{F}_8 , and \mathcal{F}_{10} are highlighted by different line patterns.

Lemma 9: Assume that a pooling function g satisfies the WSFP (A5) and generalized Bayesianity (A12). Then, g is either a dogmatic pooling function or a dictatorship pooling function.

The following characterization of the dogmatic pooling function now follows easily.

Theorem 10: The dogmatic pooling function in (13) satisfies the symmetry axiom (A1), the WSFP (A5), the weak likelihood principle (A7), and generalized Bayesianity (A12). Conversely, a pooling function that satisfies the symmetry axiom (A1), the WSFP (A5), and generalized Bayesianity (A12) is a dogmatic pooling function.

It is obvious that the dogmatic pooling function satisfies the stated axioms. The converse follows because, by Lemma 9, the pooling function must be either a dogmatic pooling function or a dictatorship pooling function, but of these only the dogmatic pooling function is symmetric.

Based on the theorems above, we can establish an implication structure for the different axioms from Section IV-A, which indicates which axioms imply which other axioms. To formalize this structure, we will designate the set of all pooling functions that satisfy Axiom i as \mathcal{F}_i . The next theorem states the currently known implications. Venn diagrams representing the implication structure are presented in Fig. 4.

Theorem 11: For the axioms introduced in Section IV-A, the following implications hold.

- 1) The SSFP (A4) implies the ZPP (A2), unanimity preservation (A3), the WSFP (A5), the likelihood principle (A6), and the weak likelihood principle (A7), i.e., $\mathcal{F}_4 \subseteq \mathcal{F}_2 \cap \mathcal{F}_3 \cap \mathcal{F}_5 \cap \mathcal{F}_6 \cap \mathcal{F}_7$. Furthermore, $\mathcal{F}_4 = \mathcal{F}_2 \cap \mathcal{F}_5 = \mathcal{F}_3 \cap \mathcal{F}_5$.
- 2) The WSFP (A5) implies the weak likelihood principle (A7), i.e., $\mathcal{F}_5 \subseteq \mathcal{F}_7$.

- 3) The likelihood principle (A6) implies the weak likelihood principle (A7), i.e., $\mathcal{F}_6 \subseteq \mathcal{F}_7$.
- 4) Independence preservation (A8) implies the ZPP (A2), i.e., $\mathcal{F}_8 \subseteq \mathcal{F}_2$.
- 5) Individualized Bayesianity (A11) implies generalized Bayesianity (A12), i.e., $\mathcal{F}_{11} \subseteq \mathcal{F}_{12}$.

Most of these implications follow from our earlier theorems. For completeness, we provide a proof of, or references for, all implications in Appendix F.

V. OPTIMIZATION APPROACH

In Section IV-B, we identified pooling functions that satisfy certain axioms. An alternative approach to establishing pooling functions for probabilistic opinion pooling is the optimization approach. Here, a pooling function is obtained by minimizing the weighted average of some discrepancy measure between the pdfs of the K agents, $q_1(\theta), \dots, q_K(\theta)$, and the aggregate pdf $q(\theta)$. The underlying idea is to make the aggregate pdf as similar as possible to all the agent pdfs simultaneously. As we will see, the obtained $q(\theta)$ turns out to be some sort of average of the agent pdfs $q_1(\theta), \dots, q_K(\theta)$.

One class of discrepancy measures that can be considered is f -divergences. For a convex function $f: \mathbb{R}^+ \rightarrow \mathbb{R}$ with $f(1) = 0$, the f -divergence between two pdfs $q_k(\theta)$ and $\varphi(\theta)$ with common domain Θ is defined as [161]–[164]

$$\mathcal{D}_f(q_k \parallel \varphi) = \int_{\Theta} \varphi(\theta) f\left(\frac{q_k(\theta)}{\varphi(\theta)}\right) d\theta. \quad (21)$$

The fusion of the agent pdfs $q_1(\theta), \dots, q_K(\theta)$ can then be based on defining the aggregate pdf $q(\theta)$ as the pdf that

Table 2 Optimization-Based Definition of Pooling Functions: Some Pooling Functions Along With the Underlying f -Divergence $\mathcal{D}_f(q_k\|\varphi)$ and Squared Distance Function $d^2(q_k, \varphi) = \|\chi(q_k) - \chi(\varphi)\|_2^2$ Used in the Optimization Problems in (22) and (32), Respectively

Pooling function	$f(x)$	$\mathcal{D}_f(q_k\ \varphi)$	$\chi(x)$	$\ \chi(q_k) - \chi(\varphi)\ _2^2$
Linear: $q(\boldsymbol{\theta}) = \sum_{k=1}^K w_k q_k(\boldsymbol{\theta})$	$x \log x$	$\int_{\Theta} q_k(\boldsymbol{\theta}) \log \left(\frac{q_k(\boldsymbol{\theta})}{\varphi(\boldsymbol{\theta})} \right) d\boldsymbol{\theta}$	x	$\int_{\Theta} (q_k(\boldsymbol{\theta}) - \varphi(\boldsymbol{\theta}))^2 d\boldsymbol{\theta}$
Log-linear: $q(\boldsymbol{\theta}) \propto \prod_{k=1}^K (q_k(\boldsymbol{\theta}))^{w_k}$	$-\log x$	$\int_{\Theta} \varphi(\boldsymbol{\theta}) \log \left(\frac{\varphi(\boldsymbol{\theta})}{q_k(\boldsymbol{\theta})} \right) d\boldsymbol{\theta}$	$\log x$	$\int_{\Theta} (\log q_k(\boldsymbol{\theta}) - \log \varphi(\boldsymbol{\theta}))^2 d\boldsymbol{\theta}$
Inverse-linear: $q(\boldsymbol{\theta}) \propto \left(\sum_{k=1}^K \frac{w_k}{q_k(\boldsymbol{\theta})} \right)^{-1}$	$\frac{1-x}{2x}$	$\int_{\Theta} \frac{(q_k(\boldsymbol{\theta}) - \varphi(\boldsymbol{\theta}))^2}{q_k(\boldsymbol{\theta})} d\boldsymbol{\theta}$	$\frac{1}{x}$	$\int_{\Theta} \left(\frac{1}{q_k(\boldsymbol{\theta})} - \frac{1}{\varphi(\boldsymbol{\theta})} \right)^2 d\boldsymbol{\theta}$
Hölder: $q(\boldsymbol{\theta}) \propto \left(\sum_{k=1}^K w_k (q_k(\boldsymbol{\theta}))^{\alpha} \right)^{1/\alpha}$	$\frac{x^{\alpha}-1}{\alpha(\alpha-1)}$	$\frac{1}{\alpha(\alpha-1)} \int_{\Theta} \varphi(\boldsymbol{\theta}) \frac{(q_k(\boldsymbol{\theta}))^{\alpha} - (\varphi(\boldsymbol{\theta}))^{\alpha}}{(\varphi(\boldsymbol{\theta}))^{\alpha}} d\boldsymbol{\theta}$	x^{α}	$\int_{\Theta} ((q_k(\boldsymbol{\theta}))^{\alpha} - (\varphi(\boldsymbol{\theta}))^{\alpha})^2 d\boldsymbol{\theta}$

minimizes a weighted average of f -divergences⁶:

$$q = \arg \min_{\varphi \in \mathcal{P}} \sum_{k=1}^K w_k \mathcal{D}_f(q_k\|\varphi) \quad (22)$$

where the weights satisfy $(w_1, \dots, w_K) \in \mathcal{S}_K$. In what follows, we consider some specific f -divergences and derive the associated pooling functions defined by (22). These results are summarized in Table 2.

A. Kullback–Leibler Divergence

For $f(x) = x \log x$, the f -divergence is the Kullback–Leibler divergence (KLD) [167]

$$\mathcal{D}_{\text{KL}}(q_k\|\varphi) = \int_{\Theta} q_k(\boldsymbol{\theta}) \log \left(\frac{q_k(\boldsymbol{\theta})}{\varphi(\boldsymbol{\theta})} \right) d\boldsymbol{\theta}. \quad (23)$$

Under this choice of divergence, the pooling function that solves the optimization problem in (22) is the linear pooling function in (1):

Theorem 12: Let $f(x) = x \log x$ (i.e., $\mathcal{D}_f(q_k\|\varphi) = \mathcal{D}_{\text{KL}}(q_k\|\varphi)$) and $(w_1, \dots, w_K) \in \mathcal{S}_K$. Then, the solution to the optimization problem in (22) is

$$q(\boldsymbol{\theta}) = \sum_{k=1}^K w_k q_k(\boldsymbol{\theta}).$$

A proof of this theorem can be found in [83]. The proof is based on the fact that minimizing the weighted average of KLDs is equivalent to minimizing the cross-entropy

$$H(q_{\text{mix}}, \varphi) = - \int_{\Theta} q_{\text{mix}}(\boldsymbol{\theta}) \log (\varphi(\boldsymbol{\theta})) d\boldsymbol{\theta}$$

⁶This minimization establishes a conceptual link to a central problem in the field of robust hypothesis testing, namely, the identification of a vector of “least favorable” pdfs within a given set of hypothesized pdfs. For two pdfs, this problem can be shown to be equivalent to the joint minimization of all f -divergences (21) for all twice differentiable convex functions f [165], [166]. The solution to this minimization can be interpreted as the pdfs that are maximally similar within the set of hypothesized pdfs, which means that a statistical test between the respective pdfs is “as hard as possible.” It is interesting that an interpretation as a maximally similar pdf holds for both the optimization approach to pdf fusion and robust hypothesis testing.

between the mixture pdf $q_{\text{mix}}(\boldsymbol{\theta}) = \sum_{k=1}^K w_k q_k(\boldsymbol{\theta})$ and the pdf $\varphi \in \mathcal{P}$. That is,

$$\arg \min_{\varphi \in \mathcal{P}} \sum_{k=1}^K w_k \mathcal{D}_{\text{KL}}(q_k\|\varphi) = \arg \min_{\varphi \in \mathcal{P}} H(q_{\text{mix}}, \varphi).$$

The cross-entropy $H(q_{\text{mix}}, \varphi)$ is minimized if and only if $q_{\text{mix}}(\boldsymbol{\theta})$ and $\varphi(\boldsymbol{\theta})$ are equal. This follows from the fact that $H(q_{\text{mix}}, \varphi)$ is equal to the sum of the KLD between $q_{\text{mix}}(\boldsymbol{\theta})$ and $\varphi(\boldsymbol{\theta})$ and the differential entropy of $q_{\text{mix}}(\boldsymbol{\theta})$ [129, Ch. 2], i.e.,

$$H(q_{\text{mix}}, \varphi) = \mathcal{D}_{\text{KL}}(q_{\text{mix}}\|\varphi) - \int_{\Theta} q_{\text{mix}}(\boldsymbol{\theta}) \log(q_{\text{mix}}(\boldsymbol{\theta})) d\boldsymbol{\theta}.$$

Hence, $H(q_{\text{mix}}, \varphi)$ is minimized if and only if $\mathcal{D}_{\text{KL}}(q_{\text{mix}}\|\varphi)$ is minimized, which implies that $\varphi(\boldsymbol{\theta}) = q_{\text{mix}}(\boldsymbol{\theta})$.

B. Reverse Kullback–Leibler Divergence

Next, consider $f(x) = -\log x$. In this case, the f -divergence equals the KLD whose arguments are reversed with respect to (23) [168], i.e.,

$$\mathcal{D}_{\text{KL}}(\varphi\|q_k) = \int_{\Theta} \varphi(\boldsymbol{\theta}) \log \left(\frac{\varphi(\boldsymbol{\theta})}{q_k(\boldsymbol{\theta})} \right) d\boldsymbol{\theta}.$$

We refer to $\mathcal{D}_{\text{KL}}(\varphi\|q_k)$ as the *reverse KLD*. For the reverse KLD, the solution to the optimization problem in (22) is the log-linear pooling function in (4):

Theorem 13: Let $f(x) = -\log x$ (i.e., $\mathcal{D}_f(q_k\|\varphi) = \mathcal{D}_{\text{KL}}(\varphi\|q_k)$) and $(w_1, \dots, w_K) \in \mathcal{S}_K$. Then, the solution to the optimization problem in (22) is

$$q(\boldsymbol{\theta}) = c \prod_{k=1}^K (q_k(\boldsymbol{\theta}))^{w_k}$$

where $c = 1 / \int_{\Theta} \prod_{k=1}^K (q_k(\boldsymbol{\theta}))^{w_k} d\boldsymbol{\theta}$.

A proof of this theorem can be found in [83] and [169]. The idea behind the proof is to derive a lower bound on the weighted average of reverse KLDs using Jensen’s inequality and then to show that the lower bound is achieved if and only if (4) is satisfied.

C. α -Divergences

We have shown that both the linear and log-linear pooling functions can be derived using the optimization approach involving the KLD or reverse KLD, respectively. These two results can be extended to an entire family of divergences and a corresponding family of pooling functions that are both parameterized by a real parameter α . Indeed, let us consider the f -divergence $\mathcal{D}_f(q_k\|\varphi)$ induced by

$$f(x) = f_\alpha(x) \triangleq \frac{x^\alpha - 1}{\alpha(\alpha - 1)}$$

where $x > 0$ and $\alpha \in \mathbb{R} \setminus \{0, 1\}$. This yields the family of α -divergences defined as [170]–[172]

$$\begin{aligned} \mathcal{D}_\alpha(q_k\|\varphi) &\triangleq \mathcal{D}_{f_\alpha}(q_k\|\varphi) & (24) \\ &= \frac{1}{\alpha(\alpha - 1)} \int_{\Theta} \varphi(\boldsymbol{\theta}) \frac{(q_k(\boldsymbol{\theta}))^\alpha - (\varphi(\boldsymbol{\theta}))^\alpha}{(\varphi(\boldsymbol{\theta}))^\alpha} d\boldsymbol{\theta}. & (25) \end{aligned}$$

We remark that the α -divergence equals the so-called Hellinger divergence up to a scaling factor and is also a one-to-one transformation of the Rényi divergence [164]. Using the optimization approach for the α -divergences, we obtain the α -parameterized family of Hölder pooling functions in (8). As noted earlier, this family comprises the linear, log-linear, and inverse-linear pooling functions as special cases.

Theorem 14: Let $f(x) = f_\alpha(x) = (x^\alpha - 1)/(\alpha(\alpha - 1))$ [i.e., $\mathcal{D}_f(q_k\|\varphi) = \mathcal{D}_\alpha(q_k\|\varphi)$] with $\alpha \in \mathbb{R} \setminus \{0, 1\}$ and $(w_1, \dots, w_K) \in \mathcal{S}_K$. Then, the solution to the optimization problem in (22) is

$$q(\boldsymbol{\theta}) = c \left(\sum_{k=1}^K w_k (q_k(\boldsymbol{\theta}))^\alpha \right)^{1/\alpha} \quad (26)$$

where $c = 1/\int_{\Theta} \left(\sum_{k=1}^K w_k (q_k(\boldsymbol{\theta}))^\alpha \right)^{1/\alpha} d\boldsymbol{\theta}$.

Although this result was mentioned in [84, Fig. 1], to the best of our knowledge, a proof does not exist in the literature. We provide a proof in Appendix G.

In the limiting case $\alpha \rightarrow 0$, the Hölder pooling function (26) becomes the log-linear pooling function (weighted geometric average) in (4), while for $\alpha = 1$, it equals the linear pooling function (weighted arithmetic average) in (1). These results are consistent with the fact that $\lim_{\alpha \rightarrow 0} \mathcal{D}_\alpha(q_k\|\varphi) = \mathcal{D}_{\text{KL}}(\varphi\|q_k)$ and $\lim_{\alpha \rightarrow 1} \mathcal{D}_\alpha(q_k\|\varphi) = \mathcal{D}_{\text{KL}}(q_k\|\varphi)$ [171]. For $\alpha = -1$, the Hölder pooling function (26) becomes the inverse-linear pooling function (9). Furthermore, the α -divergence in the case $\alpha = 2$ is (up to a scaling factor 2) equal to the Pearson χ^2 -divergence [164], [173]

$$\chi^2(q_k, \varphi) \triangleq \int_{\Theta} \frac{(q_k(\boldsymbol{\theta}) - \varphi(\boldsymbol{\theta}))^2}{\varphi(\boldsymbol{\theta})} d\boldsymbol{\theta}$$

$$= \int_{\Theta} \varphi(\boldsymbol{\theta}) \frac{(q_k(\boldsymbol{\theta}))^2 - (\varphi(\boldsymbol{\theta}))^2}{(\varphi(\boldsymbol{\theta}))^2} d\boldsymbol{\theta}.$$

The corresponding Hölder pooling function (26) is

$$q(\boldsymbol{\theta}) = c \left(\sum_{k=1}^K w_k (q_k(\boldsymbol{\theta}))^2 \right)^{1/2}$$

where $c = 1/\int_{\Theta} \left(\sum_{k=1}^K w_k (q_k(\boldsymbol{\theta}))^2 \right)^{1/2} d\boldsymbol{\theta}$.

D. Reverse α -Divergences

As for the KLD, one can exchange the order of q_k and φ in the α -divergence in (25). Again, this is equivalent to changing to a different f -divergence. More precisely, it is stated in [162, eq. (1.13)] (see also [164, Proposition 2]) that

$$\mathcal{D}_f(\varphi\|q_k) = \mathcal{D}_{f^*}(q_k\|\varphi) \quad (27)$$

where $f^*(x) = xf(1/x)$. Based on this result, we show in Appendix H that

$$\mathcal{D}_\alpha(\varphi\|q_k) = \mathcal{D}_{\alpha^*}(q_k\|\varphi)$$

where $\alpha^* = 1 - \alpha$. Thus, Theorem 14 implies the following result.

Corollary 15: The solution to the optimization problem

$$q = \arg \min_{\varphi \in \mathcal{P}} \sum_{k=1}^K w_k \mathcal{D}_\alpha(\varphi\|q_k) \quad (28)$$

is

$$q(\boldsymbol{\theta}) = c \left(\sum_{k=1}^K w_k (q_k(\boldsymbol{\theta}))^{\alpha^*} \right)^{1/\alpha^*} \quad (29)$$

where $c = 1/\int_{\Theta} \left(\sum_{k=1}^K w_k (q_k(\boldsymbol{\theta}))^{\alpha^*} \right)^{1/\alpha^*} d\boldsymbol{\theta}$ and $\alpha^* = 1 - \alpha$.

In particular, the reverse α -divergence for $\alpha = 2$ corresponds to the Pearson χ^2 -divergence in the reverse direction, i.e., $\chi^2(\varphi, q_k) = \int_{\Theta} (q_k(\boldsymbol{\theta}) - \varphi(\boldsymbol{\theta}))^2/q_k(\boldsymbol{\theta}) d\boldsymbol{\theta}$. In this case, $\alpha^* = 1 - 2 = -1$, and the corresponding Hölder pooling function (29) is the inverse-linear pooling function (9).

E. Symmetric Discrepancy Measures

As mentioned at the beginning of this section, the optimization approach defines pooling functions by minimizing a weighted average of discrepancy measures between the agent pdfs and the aggregate pdf. So far, our focus has been on minimizing a weighted average of f -divergences, where our choices of f yielded *asymmetric* discrepancy measures. Through this approach, we derived pooling functions that are the weighted arithmetic, geometric, harmonic, and Hölder averages of the agent pdfs. Interestingly, these fusion rules can also be derived using an

alternative formulation, where the goal is to minimize a weighted average of *symmetric* discrepancy measures (distance functions). Let $d(q_k, \varphi)$ be a symmetric function expressing a distance between the k th agent pdf $q_k(\theta)$ and the pdf $\varphi(\theta)$, where symmetric means that $d(q_k, \varphi) = d(\varphi, q_k)$. Then, we can define the aggregate pdf to be the solution to the following optimization problem:

$$q(\theta) = \arg \min_{\varphi \in \mathcal{P}} \sum_{k=1}^K w_k d^2(q_k, \varphi) \quad (30)$$

where $(w_1, \dots, w_K) \in \mathcal{S}_K$. The resulting $q(\theta)$ has been referred to as *Fréchet mean* [44].

An important distance function is the L_2 distance function defined as

$$\|q_k - \varphi\|_2 = \sqrt{\int_{\Theta} (q_k(\theta) - \varphi(\theta))^2 d\theta}. \quad (31)$$

The linear pooling function can be obtained alternatively by minimizing a weighted average of squared L_2 distances.

Theorem 16: Let $d(q_k, \varphi) = \|q_k - \varphi\|_2$ and $(w_1, \dots, w_K) \in \mathcal{S}_K$. Then, the solution to the optimization problem in (30) is

$$q(\theta) = \sum_{k=1}^K w_k q_k(\theta).$$

This result was mentioned without proof in [84, Fig. 1]. We provide a proof in Appendix I.

Unfortunately, for arbitrary distance functions $d(q_k, \varphi)$, an analytical solution to the optimization problem in (30) does not exist. This is due to the difficulty in satisfying the constraint $\varphi \in \mathcal{P}$, which ensures that the obtained aggregate pdf $q(\theta)$ is a valid pdf. To overcome this difficulty, following [44], we can instead solve the unconstrained version of the optimization problem in (30), i.e.,

$$\tilde{q}(\theta) = \arg \min_{\varphi} \sum_{k=1}^K w_k d^2(q_k, \varphi) \quad (32)$$

and then normalize the result, i.e.,

$$q(\theta) = \frac{\tilde{q}(\theta)}{\int_{\Theta} \tilde{q}(\theta') d\theta'}.$$

However, we emphasize that the obtained $q(\theta)$ is generally different from the solution of the constrained optimization problem in (30), and it is not guaranteed to be nonnegative.

Using this unconstrained approach, the minimization for the L_2 distance function (31) results again in the linear pooling function [44]. Here, the solution satisfies the constraint $q \in \mathcal{P}$ without explicitly enforcing it. Furthermore,

the log-linear [44], inverse-linear, and Hölder pooling functions can be derived in an analogous manner using suitable distance functions. We can arrive at all of these results and many more in a unified manner by considering the general class of distance functions $d(q_k, \varphi)$ defined as

$$\|\chi(q_k) - \chi(\varphi)\|_2 = \sqrt{\int_{\Theta} (\chi(q_k(\theta)) - \chi(\varphi(\theta)))^2 d\theta} \quad (33)$$

where $\chi: (0, \infty) \rightarrow (a, b)$ with $a \in \mathbb{R} \cup \{-\infty\}$ and $b \in \mathbb{R} \cup \{\infty\}$ is an invertible function. Solving the optimization problem (32) for the distance functions (33) leads to the rich class of pooling functions stated by the following result.

Theorem 17: Let $d(q_k, \varphi) = \|\chi(q_k) - \chi(\varphi)\|_2$ and $(w_1, \dots, w_K) \in \mathcal{S}_K$. Then, the solution to the optimization problem in (32) is

$$\tilde{q}(\theta) = \chi^{-1} \left(\sum_{k=1}^K w_k \chi(q_k(\theta)) \right). \quad (34)$$

A proof is provided in Appendix J, and the functions χ leading to the linear, log-linear, inverse-linear, and Hölder pooling functions are listed in Table 2. Note that the solution $\tilde{q}(\theta)$ in (34) is always nonnegative because the domain of χ is $(0, \infty)$.

VI. GAUSSIAN DENSITIES

In Sections III–V, we discussed a variety of pooling functions that can be used to fuse the pdfs of several agents into a single aggregate pdf. We now consider the practically important special case where the opinions of the agents are represented by Gaussian pdfs. That is, we assume that

$$q_k(\theta) = \mathcal{N}(\theta; \mu_{q_k}, \Sigma_{q_k}), \quad k = 1, \dots, K \quad (35)$$

where $\mathcal{N}(\theta; \mu_{q_k}, \Sigma_{q_k})$ denotes a multivariate Gaussian pdf with mean $\mu_{q_k} = \mathbb{E}_{q_k}[\theta]$ and covariance matrix $\Sigma_{q_k} = \mathbb{E}_{q_k}[(\theta - \mu_{q_k})(\theta - \mu_{q_k})^T]$. An important aspect of the Gaussian case is the fact that each agent pdf $q_k(\theta)$ is completely characterized by its first- and second-order moments μ_{q_k} and Σ_{q_k} .

A. Linear Pooling

The fusion of Gaussian pdfs using the linear pooling function in (1) results in an aggregate pdf that is a mixture of Gaussians, i.e.,

$$q(\theta) = \sum_{k=1}^K w_k \mathcal{N}(\theta; \mu_{q_k}, \Sigma_{q_k}). \quad (36)$$

A convenient property in this context is that the expected value of a function $h(\theta)$ with respect to the pdf $q(\theta)$ in (1) is the weighted average of the expected values of

$h(\theta)$ with respect to the agent pdfs $q_1(\theta), \dots, q_K(\theta)$, i.e., $\mathbb{E}_q[h(\theta)] = \sum_{k=1}^K w_k \mathbb{E}_{q_k}[h(\theta)]$. This implies that the mean of the aggregate pdf in (36), $\mu_q = \mathbb{E}_q[\theta]$, is simply the weighted average of the agent means, i.e.,

$$\mu_q = \sum_{k=1}^K w_k \mu_{q_k}. \quad (37)$$

Similarly, the covariance matrix of the aggregate pdf in (36), $\Sigma_q = \mathbb{E}_q[(\theta - \mu_q)(\theta - \mu_q)^\top]$, is obtained as [174]

$$\Sigma_q = \sum_{k=1}^K w_k (\Sigma_{q_k} + (\mu_{q_k} - \mu_q)(\mu_{q_k} - \mu_q)^\top). \quad (38)$$

Thus, the mean and covariance matrix of the aggregate pdf $q(\theta)$ can be calculated easily from the agent means and covariance matrices. This is useful from a practical perspective because it provides a way for obtaining an estimate of the parameters (e.g., mean) and a measure of uncertainty for that estimate (e.g., covariance matrix). It is important to note, however, that, since $q(\theta)$ is a mixture of Gaussians and, therefore, is non-Gaussian, it is not fully characterized by its mean and covariance matrix. Indeed, a mixture of Gaussians can have properties that a Gaussian cannot have, including heavy tails, multiple modes, and nonzero skewness [175].

In the case that the agent pdfs are Gaussian, the connection of linear opinion pooling to model averaging established in Section III-B1 extends to an estimation technique in the Kalman filtering literature called *multiple model adaptive estimation* (MMAE) [176]. MMAE uses a bank of Kalman filters to estimate an unknown state (time-varying parameter), where each Kalman filter assumes a distinct model describing the state's time evolution and its relation to the observed data. In this context, μ_{q_k} is the local state estimate provided by the k th Kalman filter at a given time, while Σ_{q_k} is the covariance of that estimate. The local state estimates are then combined according to (37) to obtain a final state estimate μ_q , whose covariance Σ_q is determined by (38). Here, the weight w_k equals the posterior probability of the model assumed by the k th Kalman filter.

B. Log-Linear Pooling

The fusion of the Gaussian pdfs $q_k(\theta)$ in (35) by the log-linear pooling function in (4) results in an aggregate pdf that is also Gaussian, i.e.,

$$q(\theta) = \mathcal{N}(\theta; \mu_q, \Sigma_q)$$

with the mean vector

$$\mu_q = \left(\sum_{k=1}^K w_k \Sigma_{q_k}^{-1} \right)^{-1} \sum_{j=1}^K w_j \Sigma_{q_j}^{-1} \mu_{q_j} \quad (39)$$

and the covariance matrix

$$\Sigma_q = \left(\sum_{k=1}^K w_k \Sigma_{q_k}^{-1} \right)^{-1}. \quad (40)$$

Unlike the case of linear pooling, since the aggregate pdf $q(\theta)$ is Gaussian, it is unimodal and symmetric about the mean μ_q , and it is, moreover, fully characterized by the mean μ_q and covariance Σ_q .

There is a strong link between log-linear pooling of Gaussian pdfs and a second-order fusion method called *covariance intersection* [29], [32], which is often employed in distributed (decentralized) Kalman filter implementations [3], [41], [94]. In the covariance intersection context, there are K agents, each of which uses its own local observations to form a local estimate of an unknown quantity θ . The goal of covariance intersection is to fuse the local estimates in a way that does not underestimate the overall covariance of the fused estimate. Let μ_{q_k} be the local estimate of the k th agent, whose covariance is denoted by Σ_{q_k} . The fused state estimate μ_q is determined according to (39), while the corresponding covariance matrix Σ_q is given by (40). The weights w_1, \dots, w_K used in (39) and (40) are typically chosen to minimize the determinant or the trace of Σ_q [29].

C. Hölder Pooling

Finally, we consider the Hölder pooling functions. The normalization factor c in the Hölder pooling function in (8) for general $\alpha \in \mathbb{R} \setminus \{0, 1\}$ involves an intractable integral and cannot be evaluated, even if the agent pdfs are Gaussian. Therefore, typically, the aggregate pdf $q(\theta)$ resulting from the Hölder pooling function is only known up to a normalization factor. Computing expected values with respect to $q(\theta)$ would require the use of numerical integration techniques, such as the trapezoidal quadrature rule or Monte Carlo methods [177]. Because numerical integration techniques are plagued by the *curse of dimensionality* [178], computing expectations with respect to $q(\theta)$ under the Hölder pooling function becomes challenging when the dimension of θ is large.

To illustrate the behavior of the linear and log-linear pooling functions, and to demonstrate the effect of different choices of α on the Hölder pooling function, we present in Fig. 5 simulation results for two different sets of $K = 2$ Gaussian agent pdfs $q_k(\theta)$ with $\theta \in \mathbb{R}$. We used the trapezoidal quadrature rule to compute the normalization factor of the aggregate pdf. Fig. 5(a) shows the fusion of two Gaussian pdfs with different means but the same variance. In this case, the value of α in the Hölder pooling function controls the multimodality of the aggregate pdf, in the sense that smaller (larger) values of α attenuate (enhance) the modes of the agent pdfs in the aggregate pdf. Fig. 5(b) shows the fusion of two Gaussian pdfs with the same mean but different variances. In this case, the value of α controls the shape of the tails of the aggregate

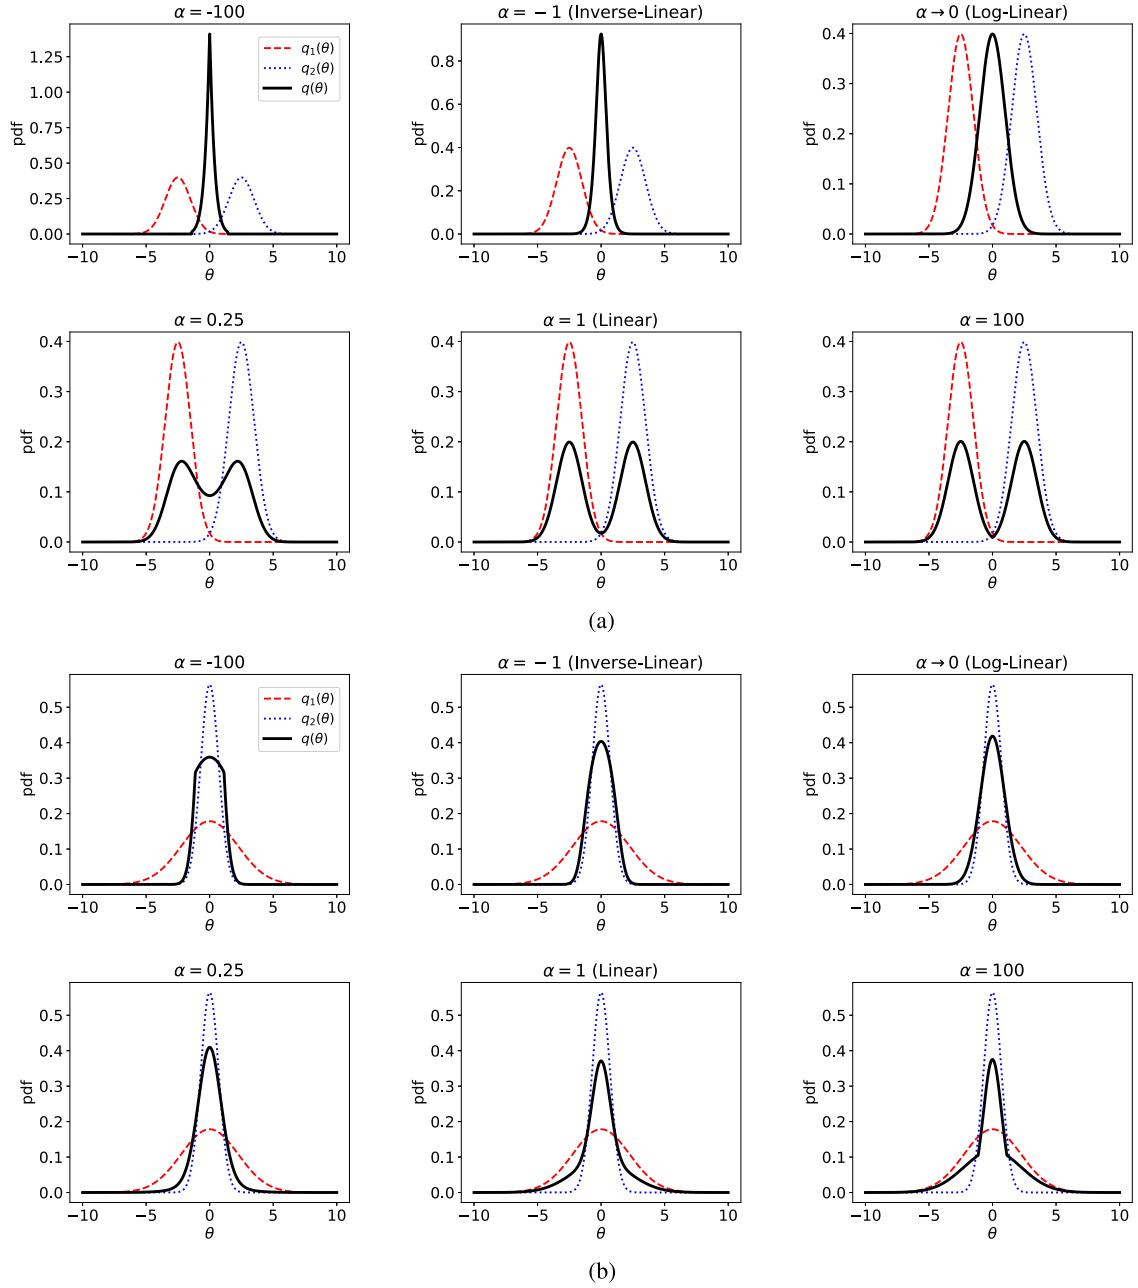


Fig. 5. Results of Hölder pooling of two pdfs $q_1(\theta)$ and $q_2(\theta)$ using weights $w_1 = w_2 = 0.5$ and different values of α . The pdfs are defined as follows. (a) $q_1(\theta) = \mathcal{N}(\theta; -2.5, 1)$ and $q_2(\theta) = \mathcal{N}(\theta; 2.5, 1)$ (different means, same variance). (b) $q_1(\theta) = \mathcal{N}(\theta; 0.5, 0.5)$ and $q_2(\theta) = \mathcal{N}(\theta; -0.5, 0.5)$ (same mean, different variances). Note that $\alpha = -1$, $\alpha \rightarrow 0$, and $\alpha = 1$ correspond to the inverse-linear, log-linear, and linear pooling functions, respectively.

pdf, in the sense that smaller (larger) values of α lead to less heavy (heavier) tails.

VII. CHOOSING THE POOLING PARAMETERS

An important consideration in opinion pooling is the choice of the parameters involved in the various pooling functions. While most of our discussion will concern the weights w_1, \dots, w_K , we also provide some insight on the choice of the parameter α in the Hölder pooling function.

The problem of choosing the weights in probabilistic opinion pooling is well researched. The simplest approach

is to assign equal weights to all agents, i.e., $w_k = 1/K$ for all k [55]. However, alternative strategies for assigning weights have been proposed for linear pooling [179]–[181] and log-linear pooling [159], [182], [183]. These strategies are usually based on solving some optimization problem, where the definition of the objective function depends on how the weights are interpreted by the fusion center. In some instances, the optimization of the weights solely depends on the agent pdfs. In other scenarios, weight assignment takes into consideration data that are observed at the fusion center and is based on a Bayesian interpretation involving likelihood functions

or posterior distributions. These data-dependent methods have also been extended to the sequential case, where observed data are streamed and the weights are updated when new data become available [179].

In the following, we describe several options for choosing the weights in linear pooling and log-linear pooling. We focus on methods that do not assume that the fusion center has observed any data. At this point, it is important to emphasize that, in both the axiomatic and optimization approaches to probabilistic opinion pooling, the weights w_k were assumed fixed, i.e., not dependent on the agent pdfs $q_k(\theta)$. If, on the other hand, the weights are chosen adaptively according to an additional optimization procedure that involves the agent pdfs $q_k(\theta)$, then this implies a deviation from the strict mathematical framework established by both the axiomatic and optimization approaches. For example, the linear pooling function with adaptively chosen weights is no longer linear in the agent pdfs $q_k(\theta)$.

A. Linear Pooling

The problem of assigning the weights in the linear pooling function has been considered in many works (see [179] for a review). One approach is based on interpreting the weight w_k as a *veridical probability*, i.e., as the probability that the true pdf of θ is $q_k(\theta)$ [184]. Accordingly, w_k is chosen to equal a prior or posterior estimate of that probability. This approach is connected to the model-averaging view of linear opinion pooling mentioned in Section III-B1 since, in (2), $P(M_k)$ equals the probability that the model of the k th agent, M_k , is the correct one. When data are considered, the weights w_k equal the posterior probabilities of the models M_k , and this is exactly how they are assigned in the MMAE algorithm mentioned in Section VI-A [174], [176].

Alternatively, the weights can be assigned according to the predictive performance of each agent by viewing the weights as *outranking probabilities* [185]. In this view, w_k is the probability that predictions made based on $q_k(\theta)$ will outperform the predictions based on the pdfs of the other agents. This rationale for choosing the weights requires consideration of data and a mechanism for assessing the predictive performance of the agents.

Another idea is to interpret the weights as a measure of distance [186]. Based on this interpretation, agents that have “middle of the road” opinions are assigned higher weights, while those that have more extreme (controversial) opinions are assigned lower weights. The opposite strategy would in principle also be possible, namely, giving more weight to controversial opinions. Such weight assignments can be achieved by assigning a nonnegative score γ_k to each agent pdf $q_k(\theta)$. For example, one can choose the score γ_k to be inversely related to the *maximum discrepancy* between agent k and the other agents, i.e.,

$$\gamma_k = \frac{1}{\max_{j \in \{1, \dots, K\}} \mathcal{D}_{\text{KL}}(q_k \| q_j)} \geq 0, \quad k = 1, \dots, K. \quad (41)$$

Here, the KLD is used to measure the discrepancy between agents although other divergences can be used instead. The weight of each agent is then obtained as a normalized version of γ_k , i.e.,

$$w_k = \frac{\gamma_k}{\sum_{j=1}^K \gamma_j}, \quad k = 1, \dots, K.$$

Finally, there are also iterative schemes for weight assignment, where each agent considers itself to be a fusion center and assigns weights to all the other agents. The weights are iteratively updated until a consensus is reached. In [187], the weight vector of each agent is updated by multiplying it by a transition matrix, and under some conditions, a consensus is reached asymptotically. The work [188] builds on this idea but updates the weights according to how closely the agent pdfs agree using a scoring function similar to (41).

B. Log-Linear Pooling

The choice of the weights in the log-linear pooling function has been considered less intensely in the literature. Some of the aforementioned methods for linear opinion pooling can also be applied to log-linear opinion pooling; for example, the scoring rule in (41) is still reasonable. Moreover, as mentioned in Section VI-B, for Gaussian agent pdfs, log-linear pooling corresponds to the covariance intersection fusion method. Here, the weights can be chosen using schemes proposed in the covariance intersection literature, such as minimizing the trace or determinant of the covariance matrix in (40) [29].

One criterion proposed in the literature that does not require the consideration of data is the *minimum KLD* criterion [183]. If there is no basis for determining the reliability of each agent, one can choose the weights such that the aggregate pdf is maximally close to all the agent pdfs simultaneously. This is the criterion that was used in Section V to find an optimal pooling function for given weights w_k . Similar to Section V-A, the criterion can be formulated as a minimization of the average of the KLDs between the agent pdfs $q_k(\theta)$ and the aggregate pdf $q(\theta)$. Introducing the weight vector $\mathbf{w} \triangleq (w_1, \dots, w_K)$, the optimal weights are defined as

$$\mathbf{w}^* = \arg \min_{\mathbf{w} \in \mathcal{S}_K} L(\mathbf{w})$$

with

$$\begin{aligned} L(\mathbf{w}) &\triangleq \frac{1}{K} \sum_{k=1}^K \mathcal{D}_{\text{KL}}(q_k \| q) \\ &= \frac{1}{K} \sum_{k=1}^K \mathcal{D}_{\text{KL}}\left(q_k \left\| c \prod_{\ell=1}^K (q_\ell(\theta))^{w_\ell}\right.\right) \end{aligned}$$

where expression (4) was inserted for $q(\theta)$. Using the KLD definition (23), one can obtain [183]

$$L(\mathbf{w}) = -\log c(\mathbf{w}) + \frac{1}{K} \sum_{k=1}^K \sum_{j \neq k} w_j \mathcal{D}_{\text{KL}}(q_k \| q_j) \quad (42)$$

where $c(\mathbf{w})$ is the normalization factor in (5), which depends on \mathbf{w} . The objective function $L(\mathbf{w})$ is convex, since the first term $-\log c(\mathbf{w})$ is convex [159] and the second term is a linear function of \mathbf{w} . Therefore, tools from convex optimization can be used to compute the optimal weight vector \mathbf{w}^* . We note that the minimum KLD criterion would also be a reasonable criterion for use with other pooling functions; however, the expression for $L(\mathbf{w})$ in (42) applies specifically to the log-linear pooling function.

Furthermore, we remark that, if the average of the reverse KLDs, i.e.,

$$\begin{aligned} \tilde{L}(\mathbf{w}) &\triangleq \frac{1}{K} \sum_{k=1}^K \mathcal{D}_{\text{KL}}(q \| q_k) \\ &= \frac{1}{K} \sum_{k=1}^K \mathcal{D}_{\text{KL}} \left(c \prod_{\ell=1}^K (q_\ell(\theta))^{w_\ell} \middle\| q_k \right) \end{aligned} \quad (43)$$

was chosen as the objective function to be minimized, the optimal weights would be given by

$$\arg \min_{\mathbf{w} \in \mathcal{S}_K} \tilde{L}(\mathbf{w}) = \left(\frac{1}{K}, \dots, \frac{1}{K} \right). \quad (44)$$

Indeed, let $q^*(\theta)$ be defined by (4) with weights $\mathbf{w} = (1/K, \dots, 1/K)$, i.e.,

$$q^*(\theta) \triangleq c \prod_{k=1}^K (q_k(\theta))^{1/K}. \quad (45)$$

By Theorem 13, $q^*(\theta)$ minimizes the objective function in (43) over all pdfs φ , i.e.,

$$q^* = \arg \min_{\varphi \in \mathcal{P}} \frac{1}{K} \sum_{k=1}^K \mathcal{D}_{\text{KL}}(\varphi \| q_k). \quad (46)$$

Thus, we have

$$\begin{aligned} \tilde{L}(\mathbf{w}) &\stackrel{(43)}{=} \frac{1}{K} \sum_{k=1}^K \mathcal{D}_{\text{KL}} \left(c \prod_{\ell=1}^K (q_\ell(\theta))^{w_\ell} \middle\| q_k \right) \\ &\stackrel{(46)}{\geq} \frac{1}{K} \sum_{k=1}^K \mathcal{D}_{\text{KL}}(q^* \| q_k) \\ &\stackrel{(45)}{=} \frac{1}{K} \sum_{k=1}^K \mathcal{D}_{\text{KL}} \left(c \prod_{\ell=1}^K (q_\ell(\theta))^{1/K} \middle\| q_k \right) \\ &\stackrel{(43)}{=} \tilde{L} \left(\left(\frac{1}{K}, \dots, \frac{1}{K} \right) \right). \end{aligned}$$

Thus, for any \mathbf{w} , $\tilde{L}(\mathbf{w})$ is lower bounded by $\tilde{L}((1/K, \dots, 1/K))$. This proves (44).

Other approaches minimize an alternative KLD criterion [182], [189] or take a Bayesian approach by specifying a prior distribution over the weights [183]. However, these approaches require data to be available and usually lead to closed-form solutions only if the prior pdfs take the form of conjugate priors for the considered likelihood functions.

C. Hölder Pooling

In addition to the weights, the parameter α involved in the Hölder pooling function in (8) strongly impacts the resulting aggregate pdf, as was demonstrated in Fig. 5. An appropriate choice of α depends on the application at hand. For example, in risk assessment, the choice of α is relevant to quantification of uncertainty. In a risk-averse scenario, one may opt to choose a larger value of α or at least a positive α . Indeed, for any $\alpha > 0$, the supports of the agent pdfs are preserved by the fusion in the sense that the support of the aggregate pdf equals the union of the supports of all the agent pdfs. Furthermore, a larger α tends to yield a larger uncertainty in the aggregate pdf. This latter characteristic is related to the fact, shown in Fig. 5, that a larger α tends to promote multimodality and/or heavy-tailed properties in the aggregate pdf.

If one instead chooses a small positive value of α , then components of different agent pdfs that occur at different θ locations will have substantially less influence on the aggregate pdf. This means, in particular, that an “outlier behavior” of one agent will tend to be attenuated in the fusion process. Furthermore, for $\alpha = 0$ (log-linear pooling), if the pdf of any agent k is zero for some θ_0 , i.e., $q_k(\theta_0) = 0$, this implies that the aggregate pdf is also zero at θ_0 irrespectively of the values of the other agent pdfs. This “veto property” can be problematic in certain situations. Finally, for $\alpha < 0$, Hölder pooling is restricted to positive opinion profiles, which implies that all agents have to agree on the support Θ of θ .

Hölder pooling appears to be practically relevant mostly for values of α in $[0, 1]$. Here, we recall that $\alpha = 0$ and $\alpha = 1$ correspond to the log-linear pooling function and the linear pooling function, respectively; furthermore, values of α between 0 and 1 correspond to pooling functions whose characteristics—e.g., with regard to multimodality and tail decay—are intermediate between those of the linear and log-linear pooling functions, as demonstrated in Fig. 5. An application where this observation is potentially relevant was considered in Section II-A.

VIII. SUPRA-BAYESIAN FRAMEWORK

The supra-Bayesian framework is fundamentally different from the approaches discussed so far. In this section, we consider θ to be a random variable with prior pdf $p(\theta)$ and assume that the fusion center follows a Bayesian update rule to derive a posterior pdf. Our focus will be

on scenarios where observations (data) that depend on θ are obtained by the agents but are not known to the fusion center. We will start this section with a formulation using conditionally independent observations and extend from there to the general supra-Bayesian framework.

A. Agents Collecting Conditionally Independent Observations

Let us consider a scenario with K agents where each agent $k \in \{1, \dots, K\}$ obtains observations $\mathbf{y}_k \in \mathbb{R}^{d_{y_k}}$. These observations are statistically related to the random vector $\theta \in \mathbb{R}^{d_\theta}$ according to the “local” likelihood functions $p(\mathbf{y}_k | \theta)$. We consider the observations fixed (i.e., already observed) and emphasize the dependence of $p(\mathbf{y}_k | \theta)$ on θ by writing the local likelihood functions as $\ell_k(\theta) \triangleq p(\mathbf{y}_k | \theta)$. Furthermore, each agent has access to the prior pdf $p(\theta)$ and is, thus, able to calculate its local posterior $\pi_k(\theta) \triangleq p(\theta | \mathbf{y}_k)$ according to Bayes’ rule, i.e.,

$$\pi_k(\theta) = p(\theta | \mathbf{y}_k) = \frac{\ell_k(\theta)p(\theta)}{\int_{\Theta} \ell_k(\theta')p(\theta') d\theta'}. \quad (47)$$

We further assume that the local observations \mathbf{y}_k are conditionally independent given θ for all $k \in \{1, \dots, K\}$. This implies that the “global” likelihood function $\ell(\theta) \triangleq p(\mathbf{y} | \theta)$ for $\mathbf{y} \triangleq [\mathbf{y}_1^\top, \dots, \mathbf{y}_K^\top]^\top$ factors into the local likelihood functions $\ell_k(\theta) = p(\mathbf{y}_k | \theta)$, i.e.,

$$\ell(\theta) = \prod_{k=1}^K \ell_k(\theta). \quad (48)$$

The task of the fusion center is to fuse the local posteriors $\pi_k(\theta)$ provided by the agents into an aggregate (fused) pdf $g[\pi_1, \dots, \pi_K](\theta)$. We assume that the fusion center is aware of the statistical properties of all the observations (i.e., the conditional pdfs $p(\mathbf{y}_k | \theta)$) and the prior $p(\theta)$ but does not have access to the observations \mathbf{y}_k directly. From a Bayesian viewpoint, the best possible fusion result is the posterior pdf of θ using the observations from all the agents as represented by the total observation vector \mathbf{y} , i.e., $p(\theta | \mathbf{y})$. We will refer to $p(\theta | \mathbf{y})$ as *oracle posterior* because the fusion center does not know the observations \mathbf{y} explicitly. Nevertheless, the following result shows that the fusion center is still able to fuse the $\pi_k(\theta)$ into the oracle posterior $p(\theta | \mathbf{y})$.

Theorem 18: Let θ be a random vector with prior $p(\theta)$. Furthermore, let the local observations $\mathbf{y}_1, \dots, \mathbf{y}_K$ given θ be mutually independent and distributed according to $p(\mathbf{y}_k | \theta)$. Then, the global posterior $p(\theta | \mathbf{y})$ with $\mathbf{y} = [\mathbf{y}_1^\top, \dots, \mathbf{y}_K^\top]^\top$ is given by

$$p(\theta | \mathbf{y}) = g[\pi_1, \dots, \pi_K](\theta) = c(p(\theta))^{1-K} \prod_{k=1}^K \pi_k(\theta) \quad (49)$$

where $c = 1 / \int_{\Theta} (p(\theta))^{1-K} (\prod_{k=1}^K \pi_k(\theta)) d\theta$ is a normalization factor and the local posteriors $\pi_k(\theta)$ are given by (47).

Proof: We recall that $\ell(\theta) = p(\mathbf{y} | \theta)$, $\ell_k(\theta) = p(\mathbf{y}_k | \theta)$, and $\pi_k(\theta) = p(\theta | \mathbf{y}_k)$. We have by Bayes’ rule that

$$p(\theta | \mathbf{y}) \propto p(\theta) \ell(\theta)$$

$$\begin{aligned} &\stackrel{(48)}{=} p(\theta) \prod_{k=1}^K \ell_k(\theta) \\ &\propto p(\theta) \prod_{k=1}^K \frac{p(\theta | \mathbf{y}_k)}{p(\theta)} \\ &= (p(\theta))^{1-K} \prod_{k=1}^K \pi_k(\theta). \end{aligned} \quad (50)$$

Since $p(\theta | \mathbf{y})$ is a conditional pdf, normalizing the function in (50) gives (49). \square

The fusion rule in (49) is recognized to be an instance of the multiplicative pooling function in (10), where the calibrating pdf $q_0(\theta)$ is given by the prior $p(\theta)$. Thus, Theorem 18 states that the multiplicative pooling function applied to the local posteriors $\pi_k(\theta)$ provides the oracle posterior $p(\theta | \mathbf{y})$ in the case of conditionally independent local observations \mathbf{y}_k .

We note that the fusion center could calculate $p(\theta | \mathbf{y})$ equally well from the local likelihood functions $\ell_k(\theta) = p(\mathbf{y}_k | \theta)$, rather than from the local posteriors $\pi_k(\theta)$. Indeed, the fusion rule (49) can be interpreted as first dividing each local posterior $\pi_k(\theta)$ by the prior $p(\theta)$ to obtain the local likelihood function $p(\mathbf{y}_k | \theta)$, then fusing (multiplying) the local likelihood functions into the global likelihood function $p(\mathbf{y} | \theta)$, and, finally, multiplying by the prior to obtain the oracle posterior $p(\theta | \mathbf{y})$. (This corresponds to reading the proof of Theorem 18 bottom up.) Thus, in the present scenario of conditionally independent observations \mathbf{y}_k , the agents may also communicate their local likelihood functions $\ell_k(\theta)$ to the fusion center, rather than their posteriors $\pi_k(\theta)$.

B. Supra-Bayesian Framework and Local Statistics

To generalize the scenario considered in Section VIII-A, we take the perspective of the fusion center. In our Bayesian setting, the fusion center aims to calculate the posterior distribution of θ , given all the information it has access to. However, in more general settings than the case of conditionally independent observations discussed in Section VIII-A, we cannot expect that the fusion center is able to calculate the oracle posterior $p(\theta | \mathbf{y})$. This is because the fusion center does not have direct access to the observations \mathbf{y}_k ; rather, it observes the effect of the \mathbf{y}_k only indirectly through the local posteriors $\pi_k(\theta) = p(\theta | \mathbf{y}_k)$. In addition to knowing the local posteriors $\pi_k(\theta)$, the fusion center is aware of the prior $p(\theta)$ and the conditional distribution $p(\mathbf{y} | \theta)$ (as a function of \mathbf{y} and θ , not for the

fixed, observed \mathbf{y}). Finally, the fusion center knows how the agents derive their local posteriors $\pi_k(\theta) = p(\theta | \mathbf{y}_k)$ given their local observations \mathbf{y}_k , i.e., it is aware that each π_k depends on θ in a well-defined probabilistic way, namely, by the two-step process of first generating a random \mathbf{y}_k given θ according to the conditional pdf $p(\mathbf{y}_k | \theta)$ and then deriving π_k from \mathbf{y}_k using (47).

This setup can be formulated generically via an abstract “observation model” $p(\pi_1, \dots, \pi_K | \theta)$ in which the local posteriors π_K are considered as “observations.” This approach is known in the literature as the supra-Bayesian model [87], [88]. In this abstract setting, we no longer have to consider the intermediate step of generating the observations \mathbf{y}_k given θ , and we no longer have to assume that the local pdfs π_k are generated as posteriors. Instead, we directly define an observation model by specifying a probability distribution over the local pdfs π_k given θ . Thus, at the fusion center, the local pdfs of all agents are considered as observations, i.e., as random objects whose statistical relation to θ is described by the “likelihood function” $p(\pi_1, \dots, \pi_K | \theta)$. As always in Bayesian settings, we need in addition some prior $p(\theta)$. By Bayes’ theorem, we can then express the posterior distribution of θ given the local pdfs π_k as

$$p(\theta | \pi_1, \dots, \pi_K) = \frac{p(\pi_1, \dots, \pi_K | \theta)p(\theta)}{\int_{\Theta} p(\pi_1, \dots, \pi_K | \theta')p(\theta')d\theta'} \quad (51)$$

which is considered to be the supra-Bayesian fusion result, also to be referred to as “supra-Bayesian posterior.”

For any given θ , $p(\pi_1, \dots, \pi_K | \theta)$ is a probability distribution over the infinite-dimensional space of functions that is given by the K -fold Cartesian product of the space of all pdfs \mathcal{P} . It is both mathematically and practically convenient to restrict to a finite-dimensional subset of this space. Indeed, a finite-dimensional parameterization is very often used in practical applications. In particular, if π_k depends deterministically on some finite-dimensional observation \mathbf{y}_k , then π_k is obviously restricted to a finite-dimensional subset. Thus, we will hereafter assume that each π_k depends deterministically and in a one-to-one manner on a finite-dimensional random vector $\mathbf{t}_k \in \mathbb{R}^{d_{t_k}}$. Then, the probability distribution $p(\pi_1, \dots, \pi_K | \theta)$ simplifies to a conventional conditional pdf $p(\mathbf{t}_1, \dots, \mathbf{t}_K | \theta)$. This finite-dimensional setting is formalized by the following definition.

Definition 1: A finite-dimensional supra-Bayesian model for a parameter $\theta \in \Theta \subseteq \mathbb{R}^{d_{\theta}}$ consists of the following:

- 1) a prior pdf $p(\theta)$;
- 2) a conditional pdf $p(\mathbf{t} | \theta)$, where $\mathbf{t} = [\mathbf{t}_1^T, \dots, \mathbf{t}_K^T]^T$ with $\mathbf{t}_k \in \mathbb{R}^{d_{t_k}}$ for $k = 1, \dots, K$;
- 3) for each $k \in \{1, \dots, K\}$, a one-to-one mapping $\psi_k: \mathbb{R}^{d_{t_k}} \rightarrow \mathcal{P}$.

The vectors \mathbf{t}_k are referred to as *local statistics* and the functions $\pi_k(\theta) = \psi_k[\mathbf{t}_k](\theta)$ as *local pdfs*.

In a finite-dimensional supra-Bayesian model, each local pdf π_k is uniquely defined by a corresponding local statistic \mathbf{t}_k . As a consequence, the conditional distribution $p(\pi_1, \dots, \pi_K | \theta)$ is implicitly given by the conditional pdf $p(\mathbf{t} | \theta)$ with $\mathbf{t} = [\mathbf{t}_1^T, \dots, \mathbf{t}_K^T]^T$, and we will refer to $\lambda(\theta) \triangleq p(\mathbf{t} | \theta)$ as *global likelihood function*. The function ψ_k specifies which family of distributions π_k belongs to. For example, if we want to model the fact that $\pi_k(\theta)$ belongs to the family of Gaussian distributions with fixed and known covariance matrix Σ , then we define $\psi_k[\mu_k](\theta) = \mathcal{N}(\theta; \mu_k, \Sigma)$. In this example, then, $\mathbf{t}_k = \mu_k$.

In Definition 1, we further assumed that there is a one-to-one relation between the local pdf π_k and \mathbf{t}_k , i.e., different vectors \mathbf{t}_k and $\tilde{\mathbf{t}}_k$ correspond to different pdfs π_k and $\tilde{\pi}_k$. In addition to the fact that the pdf π_k is uniquely specified by the vector \mathbf{t}_k , this assumption also implies that we can uniquely determine \mathbf{t}_k from π_k , i.e., \mathbf{t}_k is a function of π_k , and we can, thus, interpret it as a statistic of π_k . This justifies the designation of the vectors \mathbf{t}_k as *local statistics*. In summary, the local statistic \mathbf{t}_k represents the information provided by the pdf π_k of agent k in a more accessible, finite-dimensional way.

The following result is an immediate consequence of our definition of a finite-dimensional supra-Bayesian model (see Definition 1) and Bayes’ theorem: the one-to-one relationship between π_k and \mathbf{t}_k for each $k \in \{1, \dots, K\}$ implies that $p(\theta | \pi_1, \dots, \pi_K) = p(\theta | \mathbf{t})$, and Bayes’ theorem implies that $p(\theta | \mathbf{t}) \propto p(\mathbf{t} | \theta)p(\theta)$.

Theorem 19: In a finite-dimensional supra-Bayesian model, the supra-Bayesian fusion result (or supra-Bayesian posterior) is given by

$$p(\theta | \pi_1, \dots, \pi_K) = p(\theta | \mathbf{t}) = \frac{\lambda(\theta)p(\theta)}{\int_{\Theta} \lambda(\theta')p(\theta')d\theta'} \quad (52)$$

where $\lambda(\theta) = p(\mathbf{t} | \theta)$.

Since the fusion center knows \mathbf{t} , $p(\mathbf{t} | \theta)$, and $p(\theta)$, it is able to calculate (52). However, in general, (52) does not provide an explicit rule for fusing the pdfs $\pi_k(\theta)$ into the supra-Bayesian posterior $p(\theta | \mathbf{t})$, i.e., it does not specify a pooling function g such that $p(\theta | \mathbf{t}) = g[\pi_1, \dots, \pi_K](\theta)$. Nevertheless, we can already deduce an interesting fact from the structure of (52): The supra-Bayesian posterior is proportional to the product of the prior $p(\theta)$ and the global likelihood function $\lambda(\theta)$ and, thus, depends on the pdfs $\pi_k(\theta)$ only indirectly via the global likelihood function $\lambda(\theta)$. Hence, the actual task in supra-Bayesian fusion is to establish a rule for obtaining the global likelihood function $\lambda(\theta) = p(\mathbf{t} | \theta)$ from the local posteriors $\pi_k(\theta)$ or, equivalently, from the vector of local statistics $\mathbf{t} = [\mathbf{t}_1^T, \dots, \mathbf{t}_K^T]^T$. In what follows, we will see that this approach can result in interesting fusion rules for specific scenarios. In particular, we will consider conditionally independent agents in Section VIII-C and dependent agents in Section VIII-D. Furthermore, the special case given by the linear Gaussian model will be studied in Section IX.

C. Supra-Bayesian Fusion for Conditionally Independent Agents

Generalizing the scenario in Section VIII-A, we assume that, given θ , the information provided by each agent to the fusion center is conditionally independent of the information provided by the other agents. In our finite-dimensional supra-Bayesian model, this means that the \mathbf{t}_k are conditionally independent given θ , i.e., the global likelihood function $\lambda(\theta)$ factors according to

$$\lambda(\theta) = p(\mathbf{t} | \theta) = \prod_{k=1}^K p(\mathbf{t}_k | \theta) = \prod_{k=1}^K \lambda_k(\theta) \quad (53)$$

where we introduced the local likelihood functions $\lambda_k(\theta) \triangleq p(\mathbf{t}_k | \theta)$. Because conditional independence of the \mathbf{t}_k is equivalent to conditional independence of the random local pdfs π_k , we immediately obtain the following corollary by inserting (53) into (52).

Corollary 20: In a finite-dimensional supra-Bayesian model where the local pdfs π_k are conditionally independent given θ , the supra-Bayesian fusion result (or supra-Bayesian posterior) is given by

$$p(\theta | \pi_1, \dots, \pi_K) = p(\theta | \mathbf{t}) = \frac{(\prod_{k=1}^K \lambda_k(\theta)) p(\theta)}{\int_{\Theta} (\prod_{k=1}^K \lambda_k(\theta')) p(\theta') d\theta'} \quad (54)$$

where $\lambda_k(\theta) = p(\mathbf{t}_k | \theta)$.

To establish a link to the scenario of Section VIII-A, let us consider the local statistics \mathbf{t}_k and the global likelihood function $\lambda(\theta) = p(\mathbf{t} | \theta)$ in that scenario. Recall that, in Section VIII-A, we assumed that each agent has observations $\mathbf{y}_k \in \mathbb{R}^{d_{y_k}}$ related to θ according to the local observation likelihood function $\ell_k(\theta) = p(\mathbf{y}_k | \theta)$, and these observations are conditionally independent given θ . The local pdfs $\pi_k(\theta)$ —which, in this scenario, are the local posteriors $p(\theta | \mathbf{y}_k)$ —are given by (47), and they are, thus, parametrized by the local observations \mathbf{y}_k . However, in common observation models, the observations \mathbf{y}_k cannot be uniquely reconstructed from the posterior pdf $\pi_k(\theta)$. Indeed, local statistics \mathbf{t}_k that parametrize the local posteriors $\pi_k(\theta)$ in a one-to-one manner are usually obtained as some function $T_k(\mathbf{y}_k)$ of the observations, where $T_k: \mathbb{R}^{d_{y_k}} \rightarrow \mathbb{R}^{d_{t_k}}$ with $d_{t_k} \leq d_{y_k}$ is, in general, not invertible. The random variable $\mathbf{t}_k = T_k(\mathbf{y}_k)$ is then a sufficient statistic [190, Sec. 6.2] of \mathbf{y}_k for θ , i.e.,

$$p(\theta | \mathbf{y}_k) = p(\theta | \mathbf{t}_k). \quad (55)$$

Thus, our local statistic \mathbf{t}_k uniquely parametrizing the local posterior $\pi_k(\theta)$ is given by $\mathbf{t}_k = T_k(\mathbf{y}_k)$, with a non-invertible, possibly dimension-reducing function T_k . The local statistics \mathbf{t}_k given θ are conditionally independent for $k = 1, \dots, K$ because they are deterministic functions

of the conditionally independent observations \mathbf{y}_k . Hence, the factorization (53) holds, and indeed, we have a finite-dimensional supra-Bayesian model with a prior $p(\theta)$, a likelihood function $p(\mathbf{t} | \theta)$, and local pdfs π_k that are given by $\pi_k(\theta) = p(\theta | \mathbf{t}_k)$, i.e., $\psi_k[\mathbf{t}_k](\theta) = p(\theta | \mathbf{t}_k)$. Thus, the supra-Bayesian fusion result $p(\theta | \mathbf{t})$ is given by the expression in (54). We will now demonstrate that $p(\theta | \mathbf{t})$ coincides with the fusion result given in (49). Recalling that $\lambda_k(\theta) = p(\mathbf{t}_k | \theta)$, the supra-Bayesian fusion result (54) becomes

$$\begin{aligned} p(\theta | \mathbf{t}) &\propto \left(\prod_{k=1}^K p(\mathbf{t}_k | \theta) \right) p(\theta) \\ &\propto \left(\prod_{k=1}^K \frac{p(\theta | \mathbf{t}_k)}{p(\theta)} \right) p(\theta) \\ &= (p(\theta))^{1-K} \prod_{k=1}^K p(\theta | \mathbf{t}_k) \end{aligned}$$

where we used Bayes' theorem. By (55), we further have

$$\begin{aligned} p(\theta | \mathbf{t}) &\propto (p(\theta))^{1-K} \prod_{k=1}^K p(\theta | \mathbf{y}_k) \\ &= (p(\theta))^{1-K} \prod_{k=1}^K \pi_k(\theta) \end{aligned} \quad (56)$$

which, indeed, equals the fusion rule (49). In particular, a comparison with (49) shows that, for conditionally independent \mathbf{y}_k , the supra-Bayesian posterior $p(\theta | \mathbf{t})$ coincides with the oracle posterior $p(\theta | \mathbf{y})$. Thus, in this case, \mathbf{t} is a sufficient statistic of \mathbf{y} for θ .

Example 1 (Exponential Families): A convenient and versatile class of likelihood functions is given by exponential families [191]. Let us specialize the results discussed above to these models. A local observation likelihood function of the exponential family type can be written as

$$p(\mathbf{y}_k | \theta) = h_k(\mathbf{y}_k) \exp(\boldsymbol{\eta}(\theta)^\top T_k(\mathbf{y}_k) - A_k(\theta)) \quad (57)$$

with some functions $h_k(\mathbf{y}_k) \geq 0$, $\boldsymbol{\eta}(\theta) \in \mathbb{R}^{d_\theta}$, and $T_k(\mathbf{y}_k) \in \mathbb{R}^{d_\theta}$. The function $A_k(\theta)$ is determined by the other functions via the fact that $p(\mathbf{y}_k | \theta)$ is normalized. We assume that the observations \mathbf{y}_k are conditionally independent given θ . Furthermore, the fusion center is supposed to know the conditional pdfs $p(\mathbf{y}_k | \theta)$ in terms of the functions $\boldsymbol{\eta}$, h_k , T_k , and A_k for all k (but, as always, it does not know the \mathbf{y}_k) and to be also aware of the prior $p(\theta)$.

For the likelihood function (57), it is known that the local statistic $\mathbf{t}_k = T_k(\mathbf{y}_k)$ is a sufficient statistic of \mathbf{y}_k for θ [191, Proposition 1.5]. To verify that there is a one-to-one relation between the local posterior π_k and \mathbf{t}_k , we have to

show that \mathbf{t}_k can be recovered from π_k . We have

$$\pi_k(\boldsymbol{\theta}) \propto p(\boldsymbol{\theta})p(\mathbf{y}_k | \boldsymbol{\theta}) \propto p(\boldsymbol{\theta}) \exp(\boldsymbol{\eta}(\boldsymbol{\theta})^\top \mathbf{t}_k - A_k(\boldsymbol{\theta})). \quad (58)$$

Then,

$$\log \left(\frac{\pi_k(\boldsymbol{\theta})}{p(\boldsymbol{\theta})} \exp(A_k(\boldsymbol{\theta})) \right) = \boldsymbol{\eta}(\boldsymbol{\theta})^\top \mathbf{t}_k + C \quad (59)$$

where C is a constant that does not depend on $\boldsymbol{\theta}$. To be able to solve (59) for \mathbf{t}_k and C , we make the technical assumption that there exist $d_\theta + 1$ different $\boldsymbol{\theta}_j$ such that the matrix

$$\mathbf{B} \triangleq \begin{pmatrix} \boldsymbol{\eta}(\boldsymbol{\theta}_1)^\top & 1 \\ \vdots & \vdots \\ \boldsymbol{\eta}(\boldsymbol{\theta}_{d_\theta+1})^\top & 1 \end{pmatrix} \in \mathbb{R}^{(d_\theta+1) \times (d_\theta+1)} \quad (60)$$

is nonsingular. Then, evaluating (59) at $\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_{d_\theta+1}$ gives a system of $d_\theta + 1$ equations that can be written as

$$\mathbf{B} \begin{pmatrix} \mathbf{t}_k \\ C \end{pmatrix} = \begin{pmatrix} \log \left(\frac{\pi_k(\boldsymbol{\theta}_1)}{p(\boldsymbol{\theta}_1)} \exp(A_k(\boldsymbol{\theta}_1)) \right) \\ \vdots \\ \log \left(\frac{\pi_k(\boldsymbol{\theta}_{d_\theta+1})}{p(\boldsymbol{\theta}_{d_\theta+1})} \exp(A_k(\boldsymbol{\theta}_{d_\theta+1})) \right) \end{pmatrix}.$$

Because \mathbf{B} is nonsingular, this equation can be solved for \mathbf{t}_k and C . Thus, we are able to recover \mathbf{t}_k from π_k . We conclude that our exponential family model is a finite-dimensional supra-Bayesian model.

Using (58) in (56), the supra-Bayesian fusion result is obtained as

$$\begin{aligned} p(\boldsymbol{\theta} | \mathbf{t}) &\propto (p(\boldsymbol{\theta}))^{1-K} \prod_{k=1}^K p(\boldsymbol{\theta}) \exp(\boldsymbol{\eta}(\boldsymbol{\theta})^\top \mathbf{t}_k - A_k(\boldsymbol{\theta})) \\ &= p(\boldsymbol{\theta}) \exp(\boldsymbol{\eta}(\boldsymbol{\theta})^\top \bar{\mathbf{t}} - \bar{A}(\boldsymbol{\theta})) \end{aligned} \quad (61)$$

with

$$\bar{\mathbf{t}} = \sum_{k=1}^K \mathbf{t}_k, \quad \bar{A}(\boldsymbol{\theta}) = \sum_{k=1}^K A_k(\boldsymbol{\theta}).$$

We see that, for conditionally independent observations \mathbf{y}_k , $p(\boldsymbol{\theta} | \mathbf{t})$ depends on the observations \mathbf{y}_k only via the local statistics $\mathbf{t}_k = T_k(\mathbf{y}_k)$, and furthermore, supra-Bayesian fusion essentially amounts to the summation of the local statistics \mathbf{t}_k and the normalization functions $A_k(\boldsymbol{\theta})$.

This simple summation rule is augmented when the prior $p(\boldsymbol{\theta})$ is chosen as

$$p(\boldsymbol{\theta}) \propto \exp(\boldsymbol{\eta}(\boldsymbol{\theta})^\top \mathbf{t}_0 - A_0(\boldsymbol{\theta})) \quad (62)$$

for some vector \mathbf{t}_0 and function $A_0(\boldsymbol{\theta})$. Inserting (62) into (61), we obtain

$$p(\boldsymbol{\theta} | \mathbf{t}) \propto \exp(\boldsymbol{\eta}(\boldsymbol{\theta})^\top \mathbf{t}_{\text{post}} - A_{\text{post}}(\boldsymbol{\theta})) \quad (63)$$

with

$$\mathbf{t}_{\text{post}} = \bar{\mathbf{t}} + \mathbf{t}_0 = \sum_{k=0}^K \mathbf{t}_k \quad (64)$$

and

$$A_{\text{post}}(\boldsymbol{\theta}) = \bar{A}(\boldsymbol{\theta}) + A_0(\boldsymbol{\theta}) = \sum_{k=0}^K A_k(\boldsymbol{\theta}). \quad (65)$$

In particular, when all $A_k(\boldsymbol{\theta})$ for $k = 1, \dots, K$ are equal to the same $A(\boldsymbol{\theta})$ and $A_0(\boldsymbol{\theta}) = a_0 A(\boldsymbol{\theta})$, then the prior becomes the conjugate prior [191, Definition 4.18]

$$p(\boldsymbol{\theta}) \propto \exp(\boldsymbol{\eta}(\boldsymbol{\theta})^\top \mathbf{t}_0 - a_0 A(\boldsymbol{\theta}))$$

with the two hyperparameters \mathbf{t}_0 and $a_0 > 0$. Here, the supra-Bayesian fusion result simplifies to

$$p(\boldsymbol{\theta} | \mathbf{t}) \propto \exp(\boldsymbol{\eta}(\boldsymbol{\theta})^\top \mathbf{t}_{\text{post}} - (K + a_0) A(\boldsymbol{\theta})).$$

We see that $p(\boldsymbol{\theta} | \mathbf{t})$ has the same form as the prior $p(\boldsymbol{\theta})$, while the hyperparameters \mathbf{t}_0 and a_0 are replaced by $\mathbf{t}_{\text{post}} = \mathbf{t}_0 + \bar{\mathbf{t}}$ and $a_0 + K$, respectively.

An important special case of the exponential family setting is given by linear Gaussian observations. This case will be considered in Section IX, both for conditionally dependent and independent observations (see, in particular, Example 2).

D. Supra-Bayesian Fusion for Agents Collecting Dependent Observations

Similar to the setting of independent agents studied in Section VIII-C, we consider K agents that obtain observations $\mathbf{y}_k \in \mathbb{R}^{d_{y_k}}$ distributed according to the local observation likelihood functions $p(\mathbf{y}_k | \boldsymbol{\theta})$ with $k \in \{1, \dots, K\}$. Again, each agent has access also to the prior pdf $p(\boldsymbol{\theta})$, and the local posterior pdfs $\pi_k(\boldsymbol{\theta})$ are still given by (47). However, in contrast to Section VIII-C, we do not assume that the observations are conditionally independent. We assume that the fusion center is aware of the conditional pdf⁷ $p(\mathbf{y} | \boldsymbol{\theta})$ of all observations $\mathbf{y} = [\mathbf{y}_1^\top, \dots, \mathbf{y}_K^\top]^\top$ given $\boldsymbol{\theta}$, the prior pdf $p(\boldsymbol{\theta})$, and the local posterior pdfs $\pi_k(\boldsymbol{\theta}) = p(\boldsymbol{\theta} | \mathbf{y}_k)$. We emphasize that, although the fusion center has access to $p(\mathbf{y} | \boldsymbol{\theta})$ as a function of \mathbf{y} and $\boldsymbol{\theta}$, it does not know the global observation \mathbf{y} and, thus, cannot use $p(\mathbf{y} | \boldsymbol{\theta})$ as a global likelihood function.

⁷Note that the conditional pdfs $p(\mathbf{y}_k | \boldsymbol{\theta})$ are marginals of the conditional pdf $p(\mathbf{y} | \boldsymbol{\theta})$.

To establish a supra-Bayesian fusion scheme for this scenario, we again consider a finite-dimensional supra-Bayesian model, i.e., for each agent k , there exists a local statistic \mathbf{t}_k such that $\pi_k(\boldsymbol{\theta}) = \psi_k[\mathbf{t}_k](\boldsymbol{\theta})$, and there is a one-to-one relation between \mathbf{t}_k and the local posterior π_k . Because $\pi_k(\boldsymbol{\theta}) = p(\boldsymbol{\theta} | \mathbf{y}_k)$, the local pdf π_k is also uniquely determined by \mathbf{y}_k , and thus, the one-to-one relation between π_k and \mathbf{t}_k implies that there exists a function $T_k: \mathbb{R}^{d_{y_k}} \rightarrow \mathbb{R}^{d_{t_k}}$ such that $\mathbf{t}_k = T_k(\mathbf{y}_k)$. As before, the function T_k is not one-to-one in general, i.e., it is not possible to recover \mathbf{y}_k from \mathbf{t}_k . However, \mathbf{t}_k is again a sufficient statistic of \mathbf{y}_k for $\boldsymbol{\theta}$, i.e., $p(\boldsymbol{\theta} | \mathbf{y}_k) = p(\boldsymbol{\theta} | \mathbf{t}_k)$.

Because the local observations \mathbf{y}_k are subvectors of the global observation $\mathbf{y} = [\mathbf{y}_1^\top, \dots, \mathbf{y}_K^\top]^\top \in \mathbb{R}^{d_y}$, we can introduce $T: \mathbb{R}^{d_y} \rightarrow \mathbb{R}^{\sum_{k=1}^K d_{t_k}}$ as

$$T(\mathbf{y}) = [T_1(\mathbf{y}_1)^\top, \dots, T_K(\mathbf{y}_K)^\top]^\top$$

and, thus, we have

$$\mathbf{t} = [\mathbf{t}_1^\top, \dots, \mathbf{t}_K^\top]^\top = T(\mathbf{y}).$$

The random vector \mathbf{t} summarizes all the information that the agents communicate to the fusion center, and it is thus known to the fusion center (whereas \mathbf{y} is not). Note that, although each \mathbf{t}_k is a sufficient statistic of \mathbf{y}_k for $\boldsymbol{\theta}$, the global statistic \mathbf{t} is, in general, not a sufficient statistic of \mathbf{y} . This is due to the fact that \mathbf{t} generally does not capture all the dependencies between the individual \mathbf{y}_k .

Because $\mathbf{t} = T(\mathbf{y})$, we can use the general change-of-variables formula [192, Sec. 3.4.3] to calculate the conditional pdf $p(\mathbf{t} | \boldsymbol{\theta})$ from the conditional pdf $p(\mathbf{y} | \boldsymbol{\theta})$, provided that the function T is differentiable. Since \mathbf{t} summarizes the information communicated by the agents to the fusion center, $\lambda(\boldsymbol{\theta}) = p(\mathbf{t} | \boldsymbol{\theta})$ is the global likelihood function that the fusion center has to use in the calculation of the supra-Bayesian posterior $p(\boldsymbol{\theta} | \mathbf{t})$ according to (52). Therefore, to obtain the supra-Bayesian fusion rule $g[\pi_1, \dots, \pi_K]$, based on (52), we have to perform the following three steps.

- 1) Identify the local statistics \mathbf{t}_k that uniquely represent the local posterior pdfs π_k within the given statistical model.
- 2) Apply the general change-of-variables formula to transform the (known) conditional pdf $p(\mathbf{y} | \boldsymbol{\theta})$ into the global likelihood function $\lambda(\boldsymbol{\theta}) = p(\mathbf{t} | \boldsymbol{\theta})$.
- 3) Calculate the supra-Bayesian posterior $p(\boldsymbol{\theta} | \mathbf{t})$ according to (52).

While this three-step process can, in principle, be performed in any setting satisfying our assumptions, an explicit characterization of the resulting supra-Bayesian fusion rule (pooling function) $g[\pi_1, \dots, \pi_K]$ can only be derived for special cases. The important case of a linear Gaussian model will be explored in the following.

IX. SUPRA-BAYESIAN FUSION FOR THE LINEAR GAUSSIAN MODEL

We consider the supra-Bayesian pdf fusion for the linear observation model

$$\mathbf{y} = \mathbf{H}\boldsymbol{\theta} + \mathbf{n} \quad (66)$$

where $\mathbf{H} \in \mathbb{R}^{d_y \times d_\theta}$ is a known observation matrix and $\mathbf{n} \in \mathbb{R}^{d_y}$ is additive zero-mean Gaussian noise with a known covariance matrix $\boldsymbol{\Sigma}$, i.e., $p(\mathbf{n}) = \mathcal{N}(\mathbf{n}; \mathbf{0}, \boldsymbol{\Sigma})$. Thus, \mathbf{y} given $\boldsymbol{\theta}$ is Gaussian distributed with mean $\mathbf{H}\boldsymbol{\theta}$ and covariance matrix $\boldsymbol{\Sigma}$, i.e.,

$$p(\mathbf{y} | \boldsymbol{\theta}) = \mathcal{N}(\mathbf{y}; \mathbf{H}\boldsymbol{\theta}, \boldsymbol{\Sigma}). \quad (67)$$

The local observation at agent k is given as $\mathbf{y}_k = \mathbf{H}_k\boldsymbol{\theta} + \mathbf{n}_k \in \mathbb{R}^{d_{y_k}}$, where

$$\mathbf{H} = \begin{pmatrix} \mathbf{H}_1 \\ \mathbf{H}_2 \\ \vdots \\ \mathbf{H}_K \end{pmatrix} \quad (68)$$

with $\mathbf{H}_k \in \mathbb{R}^{d_{y_k} \times d_\theta}$ and $\mathbf{n} = [\mathbf{n}_1^\top, \dots, \mathbf{n}_K^\top]^\top$. Thus, each local observation \mathbf{y}_k given $\boldsymbol{\theta}$ is again Gaussian with mean $\mathbf{H}_k\boldsymbol{\theta}$ and covariance matrix $\boldsymbol{\Sigma}_{kk} \in \mathbb{R}^{d_{y_k} \times d_{y_k}}$. We note that the overall covariance matrix $\boldsymbol{\Sigma}$ is block-structured according to

$$\boldsymbol{\Sigma} = \begin{pmatrix} \boldsymbol{\Sigma}_{11} & \cdots & \boldsymbol{\Sigma}_{1K} \\ \vdots & \ddots & \vdots \\ \boldsymbol{\Sigma}_{K1} & \cdots & \boldsymbol{\Sigma}_{KK} \end{pmatrix} \quad (69)$$

where the off-diagonal cross-covariance matrices $\boldsymbol{\Sigma}_{kk'}$ for $k \neq k'$ describe the conditional dependence between the observations of different agents. The case of conditionally independent observations \mathbf{y}_k is obtained for $\boldsymbol{\Sigma}_{kk'} = \mathbf{0}$ for all $k \neq k'$. For simplicity, we further assume that, for all $k = 1, \dots, K$, $d_{y_k} \geq d_\theta$, \mathbf{H}_k has full rank, and $\boldsymbol{\Sigma}_{kk}$ is positive definite. The local observation likelihood functions are here given by

$$\begin{aligned} \ell_k(\boldsymbol{\theta}) &= p(\mathbf{y}_k | \boldsymbol{\theta}) \\ &= \mathcal{N}(\mathbf{y}_k; \mathbf{H}_k\boldsymbol{\theta}, \boldsymbol{\Sigma}_{kk}) \\ &\propto \exp\left(-\frac{(\mathbf{y}_k - \mathbf{H}_k\boldsymbol{\theta})^\top \boldsymbol{\Sigma}_{kk}^{-1} (\mathbf{y}_k - \mathbf{H}_k\boldsymbol{\theta})}{2}\right). \end{aligned} \quad (70)$$

A. Local Statistics

We can rewrite (70) as

$$\begin{aligned} \ell_k(\boldsymbol{\theta}) &\propto \exp\left(-\frac{(\boldsymbol{\theta} - \mathbf{V}_k \mathbf{y}_k)^\top \mathbf{H}_k^\top \boldsymbol{\Sigma}_{kk}^{-1} \mathbf{H}_k (\boldsymbol{\theta} - \mathbf{V}_k \mathbf{y}_k)}{2}\right) \\ &= \exp\left(-\frac{(\boldsymbol{\theta} - \mathbf{t}_k)^\top \mathbf{H}_k^\top \boldsymbol{\Sigma}_{kk}^{-1} \mathbf{H}_k (\boldsymbol{\theta} - \mathbf{t}_k)}{2}\right) \end{aligned} \quad (71)$$

where

$$\mathbf{V}_k = (\mathbf{H}_k^\top \boldsymbol{\Sigma}_{kk}^{-1} \mathbf{H}_k)^{-1} \mathbf{H}_k^\top \boldsymbol{\Sigma}_{kk}^{-1} \quad (72)$$

and

$$\mathbf{t}_k = \mathbf{V}_k \mathbf{y}_k = (\mathbf{H}_k^\top \boldsymbol{\Sigma}_{kk}^{-1} \mathbf{H}_k)^{-1} \mathbf{H}_k^\top \boldsymbol{\Sigma}_{kk}^{-1} \mathbf{y}_k. \quad (73)$$

The proportionality in (71) is as a function of θ , i.e., the proportionality constant will depend on \mathbf{y}_k .

We claim that \mathbf{t}_k in (73) qualifies as a local statistic in a finite-dimensional supra-Bayesian model. For a proof, we note that the local posteriors are again given as $\pi_k(\theta) = p(\theta | \mathbf{y}_k) \propto \ell_k(\theta)p(\theta)$. To see that there is a one-to-one relation between the local posterior π_k and the finite-dimensional parameter $\mathbf{t}_k \in \mathbb{R}^{d_\theta}$, recall that the fusion center is aware of the prior $p(\theta)$ and the matrices \mathbf{H} and $\boldsymbol{\Sigma}$. In particular, the fusion center is aware of \mathbf{H}_k and $\boldsymbol{\Sigma}_{kk}$, and thus, it is able to recover from \mathbf{t}_k the local observation likelihood function $\ell_k(\theta)$ in (71) and, in turn, the local posterior $\pi_k(\theta) \propto \ell_k(\theta)p(\theta)$. Conversely, the fusion center is able to obtain \mathbf{t}_k from the local posterior $\pi_k(\theta)$ by first dividing by the prior $p(\theta)$, normalizing as a function of θ (to obtain a function proportional to $\ell_k(\theta)$), and, finally, calculating the mean of the resulting pdf in θ [which is \mathbf{t}_k according to (71)]. Thus, \mathbf{t}_k is related to π_k in a one-to-one manner, and hence, it is a local statistic.

Example 2 (Conditionally Independent Agents): In the case of conditionally independent agents, i.e., the observations \mathbf{y}_k are conditionally independent given θ , we can easily calculate the supra-Bayesian posterior. Indeed, the structure of the local likelihood function in (71) shows that we are in the exponential family setting of Example 1. More specifically, we can rewrite (71) as

$$\ell_k(\theta) \propto \exp\left(\theta^\top \tilde{\mathbf{t}}_k - \frac{\theta^\top \mathbf{H}_k^\top \boldsymbol{\Sigma}_{kk}^{-1} \mathbf{H}_k \theta}{2}\right) \quad (74)$$

where

$$\tilde{\mathbf{t}}_k = \mathbf{H}_k^\top \boldsymbol{\Sigma}_{kk}^{-1} \mathbf{H}_k \mathbf{t}_k = \mathbf{H}_k^\top \boldsymbol{\Sigma}_{kk}^{-1} \mathbf{y}_k$$

is a bijective transformation of \mathbf{t}_k and, thus, also a valid choice for a local statistic. Considering a Gaussian prior $p(\theta)$ with mean μ_0 and covariance matrix $\boldsymbol{\Sigma}_0$, we can rewrite $p(\theta)$ as

$$p(\theta) \propto \exp\left(\theta^\top \tilde{\mathbf{t}}_0 - \frac{\theta^\top \boldsymbol{\Sigma}_0^{-1} \theta}{2}\right) \quad (75)$$

where $\tilde{\mathbf{t}}_0 = \boldsymbol{\Sigma}_0^{-1} \mu_0$. Comparing (74) with (57) and (75) with (62), we see that $\ell_k(\theta) = p(\mathbf{y}_k | \theta)$ belongs to the exponential family (57) with \mathbf{t}_k formally replaced by $\tilde{\mathbf{t}}_k$ and $A_k(\theta) = \theta^\top \mathbf{H}_k^\top \boldsymbol{\Sigma}_{kk}^{-1} \mathbf{H}_k \theta / 2$. Furthermore, $p(\theta)$ conforms to (62) with $A_0(\theta) = \theta^\top \boldsymbol{\Sigma}_0^{-1} \theta / 2$. With our assumption of conditionally independent agents, we can use the result (63)–(65) and obtain for the supra-Bayesian

fusion result

$$\begin{aligned} p(\theta | \mathbf{t}) &\propto \exp\left(\theta^\top \left(\sum_{k=0}^K \tilde{\mathbf{t}}_k\right) - \frac{\theta^\top (\boldsymbol{\Sigma}_0^{-1} + \sum_{k=1}^K \mathbf{H}_k^\top \boldsymbol{\Sigma}_{kk}^{-1} \mathbf{H}_k) \theta}{2}\right). \end{aligned} \quad (76)$$

This is again a Gaussian pdf, with mean

$$\boldsymbol{\mu}_1 = \boldsymbol{\Sigma}_1 \sum_{k=0}^K \tilde{\mathbf{t}}_k = \boldsymbol{\Sigma}_1 \left(\boldsymbol{\Sigma}_0^{-1} \boldsymbol{\mu}_0 + \sum_{k=1}^K \mathbf{H}_k^\top \boldsymbol{\Sigma}_{kk}^{-1} \mathbf{y}_k \right)$$

and covariance matrix

$$\boldsymbol{\Sigma}_1 = \left(\boldsymbol{\Sigma}_0^{-1} + \sum_{k=1}^K \mathbf{H}_k^\top \boldsymbol{\Sigma}_{kk}^{-1} \mathbf{H}_k \right)^{-1}.$$

It is straightforward to verify that (76) is equal to the oracle posterior $p(\theta | \mathbf{y})$. Thus, we see once again (cf. Section VIII-C) that, although the supra-Bayesian fusion result depends on the observations \mathbf{y}_k only via the local statistics $\tilde{\mathbf{t}}_k$, it still equals the oracle posterior $p(\theta | \mathbf{y})$, as if the fusion center had access to all observations \mathbf{y}_k directly. As we will see below, this crucially depends on our assumption of conditionally independent agents and is no longer true if we assume conditional dependencies between the observations.

B. Global Likelihood Function

In Section IX-A, for the general linear Gaussian model with conditionally dependent \mathbf{y}_k , we identified local statistics $\mathbf{t}_k = T_k(\mathbf{y}_k) = \mathbf{V}_k \mathbf{y}_k$ that are related in a one-to-one manner to the local posteriors π_k . The next step according to our three-step program from Section VIII-D is to calculate the global likelihood function $\lambda(\theta) = p(\mathbf{t} | \theta)$ by transforming the conditional pdf $p(\mathbf{y} | \theta)$ into the conditional pdf $p(\mathbf{t} | \theta)$. According to (67), the conditional pdf of \mathbf{y} given θ is⁸

$$p(\mathbf{y} | \theta) \propto \exp\left(-\frac{(\mathbf{y} - \mathbf{H}\theta)^\top \boldsymbol{\Sigma}^{-1} (\mathbf{y} - \mathbf{H}\theta)}{2}\right). \quad (77)$$

We further have that

$$\mathbf{t} = [\mathbf{t}_1^\top, \dots, \mathbf{t}_K^\top]^\top = \mathbf{V}\mathbf{y} \quad (78)$$

where $\mathbf{V} = \text{diag}(\mathbf{V}_1, \dots, \mathbf{V}_K)$ denotes the block-diagonal matrix with block entries \mathbf{V}_k on the diagonal. Thus, \mathbf{t} is a linear function of \mathbf{y} , and hence, \mathbf{t} given θ is Gaussian and

⁸This conditional pdf only exists if the covariance matrix $\boldsymbol{\Sigma}$ is positive definite. However, the derivations that follow do not require the existence of a pdf and are also valid if $\boldsymbol{\Sigma}$ is positive semidefinite.

has mean $\mathbf{V}\mathbf{H}\boldsymbol{\theta}$ and covariance matrix

$$\tilde{\Sigma} = \mathbf{V}\Sigma\mathbf{V}^\top. \quad (79)$$

We assume that $\tilde{\Sigma}$ is nonsingular. The mean can be simplified to

$$\mathbf{V}\mathbf{H}\boldsymbol{\theta} = \begin{pmatrix} \mathbf{V}_1\mathbf{H}_1 \\ \mathbf{V}_2\mathbf{H}_2 \\ \vdots \\ \mathbf{V}_K\mathbf{H}_K \end{pmatrix} \boldsymbol{\theta} = \begin{pmatrix} \mathbf{I}_{d_\theta} \\ \mathbf{I}_{d_\theta} \\ \vdots \\ \mathbf{I}_{d_\theta} \end{pmatrix} \boldsymbol{\theta} = \begin{pmatrix} \boldsymbol{\theta} \\ \boldsymbol{\theta} \\ \vdots \\ \boldsymbol{\theta} \end{pmatrix} = \mathbf{1}_K \otimes \boldsymbol{\theta}$$

where we used (68) and the fact that, by (72),

$$\mathbf{V}_k\mathbf{H}_k = \mathbf{I}_{d_\theta}. \quad (80)$$

The global likelihood function $\lambda(\boldsymbol{\theta})$ is, thus, obtained as

$$\begin{aligned} \lambda(\boldsymbol{\theta}) &= p(\mathbf{t} \mid \boldsymbol{\theta}) \\ &= \mathcal{N}(\mathbf{t}; \mathbf{1}_K \otimes \boldsymbol{\theta}, \tilde{\Sigma}) \\ &\propto \exp\left(-\frac{(\mathbf{t} - \mathbf{1}_K \otimes \boldsymbol{\theta})^\top \tilde{\Sigma}^{-1} (\mathbf{t} - \mathbf{1}_K \otimes \boldsymbol{\theta})}{2}\right). \end{aligned} \quad (81)$$

To summarize, for the linear Gaussian model, local statistics \mathbf{t}_k characterizing the local posteriors π_k are given by (73), and the corresponding global likelihood function $\lambda(\boldsymbol{\theta}) = p(\mathbf{t} \mid \boldsymbol{\theta})$ is given by (81).

C. Supra-Bayesian Fusion Rule for a Scalar θ

After identifying local statistics \mathbf{t}_k and calculating the global likelihood function $\lambda(\boldsymbol{\theta}) = p(\mathbf{t} \mid \boldsymbol{\theta})$, the final step in the derivation of the supra-Bayesian fusion rule is to calculate the supra-Bayesian posterior $p(\boldsymbol{\theta} \mid \mathbf{t})$ according to (52). We first develop the supra-Bayesian fusion rule for the case that $d_\theta = 1$, i.e., for a scalar random variable $\theta \in \mathbb{R}$. Here, the observation matrix \mathbf{H} reduces to a vector $\mathbf{h} \in \mathbb{R}^{d_y}$, and the observation model (66) is given by

$$\mathbf{y} = \mathbf{h}\boldsymbol{\theta} + \mathbf{n}.$$

Similarly, the local observation at agent k is given as $\mathbf{y}_k = \mathbf{h}_k\boldsymbol{\theta} + \mathbf{n}_k$ with $\mathbf{h}_k \in \mathbb{R}^{d_{y_k}}$, and the local statistic at agent k follows from (73) as

$$t_k = \mathbf{v}_k^\top \mathbf{y}_k \in \mathbb{R} \quad (82)$$

where \mathbf{V}_k reduces to the (row) vector

$$\mathbf{v}_k^\top = \frac{1}{\mathbf{h}_k^\top \Sigma_{kk}^{-1} \mathbf{h}_k} \mathbf{h}_k^\top \Sigma_{kk}^{-1}. \quad (83)$$

Note that $\mathbf{V} = \text{diag}(\mathbf{v}_1^\top, \dots, \mathbf{v}_K^\top)$ is still a matrix. In this case, we can give the following explicit fusion rule, which is derived in Appendix K-A.

Theorem 21: For $d_\theta = 1$, let $\ell_k(\boldsymbol{\theta}) = p(\mathbf{y}_k \mid \boldsymbol{\theta})$ denote the local observation likelihood functions given by (70) for $k = 1, \dots, K$, and let $\lambda(\boldsymbol{\theta}) = p(\mathbf{t} \mid \boldsymbol{\theta})$ be the global likelihood function given by (81). Then,

$$\lambda(\boldsymbol{\theta}) \propto \prod_{k=1}^K (\ell_k(\boldsymbol{\theta}))^{w_k} \quad (84)$$

where

$$w_k = \frac{\mathbf{1}_K^\top \tilde{\Sigma}^{-1} \mathbf{e}_k}{\mathbf{h}_k^\top \Sigma_{kk}^{-1} \mathbf{h}_k} \quad (85)$$

with $\tilde{\Sigma} = \mathbf{V}\Sigma\mathbf{V}^\top$ and \mathbf{e}_k denoting the k th unit vector in \mathbb{R}^K . Furthermore, for a given prior $p(\boldsymbol{\theta})$ and local posteriors $\pi_k(\boldsymbol{\theta}) = p(\boldsymbol{\theta} \mid \mathbf{y}_k) \propto p(\boldsymbol{\theta})\ell_k(\boldsymbol{\theta})$, the supra-Bayesian fusion result $g[\pi_1, \dots, \pi_K](\boldsymbol{\theta}) = p(\boldsymbol{\theta} \mid \mathbf{t}) \propto p(\boldsymbol{\theta})\lambda(\boldsymbol{\theta})$ is given by

$$g[\pi_1, \dots, \pi_K](\boldsymbol{\theta}) \propto (p(\boldsymbol{\theta}))^{1 - \sum_{k=1}^K w_k} \prod_{k=1}^K (\pi_k(\boldsymbol{\theta}))^{w_k}. \quad (86)$$

We emphasize that, in this theorem, we do not assume that the observations \mathbf{y}_k are conditionally independent given $\boldsymbol{\theta}$. Furthermore, it should be noted that the weights w_k in (85) do not generally sum to one, and they may be negative. Thus, the fusion rule (86) is an instance of the generalized multiplicative pooling function in (11).

Finally, if the prior $p(\boldsymbol{\theta})$ is Gaussian, we can show that the supra-Bayesian fusion result $p(\boldsymbol{\theta} \mid \mathbf{t})$ is again Gaussian and reduce the fusion rule (86) to a second-order rule involving only the mean and variance.

Corollary 22: Under the assumptions of Theorem 21, let the prior $p(\boldsymbol{\theta})$ be Gaussian with mean μ_0 and variance σ_0^2 , i.e., $p(\boldsymbol{\theta}) = \mathcal{N}(\boldsymbol{\theta}; \mu_0, \sigma_0^2)$. Then, the supra-Bayesian fusion result $p(\boldsymbol{\theta} \mid \mathbf{t})$ is again Gaussian, i.e., $p(\boldsymbol{\theta} \mid \mathbf{t}) = \mathcal{N}(\boldsymbol{\theta}; \mu_1, \sigma_1^2)$, with mean

$$\mu_1 = \frac{\hat{\sigma}^2}{\hat{\sigma}^2 + \sigma_0^2} (\sigma_0^2 \mathbf{1}_K^\top \tilde{\Sigma}^{-1} \mathbf{t} + \mu_0) \quad (87)$$

and variance

$$\sigma_1^2 = \frac{\hat{\sigma}^2 \sigma_0^2}{\hat{\sigma}^2 + \sigma_0^2}$$

where $\tilde{\Sigma} = \mathbf{V}\Sigma\mathbf{V}^\top$,

$$\hat{\sigma}^2 = \frac{1}{\mathbf{1}_K^\top \tilde{\Sigma}^{-1} \mathbf{1}_K} \quad (88)$$

and $\mathbf{t} = [t_1, \dots, t_K]^\top$ is given by (82) and (83).

As mentioned before, the supra-Bayesian fusion result $p(\boldsymbol{\theta} \mid \mathbf{t})$ is, in general, different from the oracle posterior $p(\boldsymbol{\theta} \mid \mathbf{y})$. Indeed, the oracle posterior is proportional to the product of the prior $p(\boldsymbol{\theta})$ and the global observation likelihood function $p(\mathbf{y} \mid \boldsymbol{\theta})$ in (77). It can then easily be

seen that the oracle posterior $p(\theta | \mathbf{y})$ is also Gaussian but with mean

$$\mu_2 = \frac{\hat{\sigma}_2^2}{\hat{\sigma}_2^2 + \sigma_0^2} (\sigma_0^2 \mathbf{h}^\top \Sigma^{-1} \mathbf{y} + \mu_0) \quad (89)$$

and variance

$$\sigma_2^2 = \frac{\hat{\sigma}_2^2 \sigma_0^2}{\hat{\sigma}_2^2 + \sigma_0^2}$$

where

$$\hat{\sigma}_2^2 = \frac{1}{\mathbf{h}^\top \Sigma^{-1} \mathbf{h}}. \quad (90)$$

To better understand the difference, we note that, in (87),

$$\mathbf{1}_K^\top \tilde{\Sigma}^{-1} \mathbf{t} = \mathbf{h}^\top \mathbf{V}^\top (\mathbf{V} \Sigma \mathbf{V}^\top)^{-1} \mathbf{V} \mathbf{y}$$

and in (88),

$$\mathbf{1}_K^\top \tilde{\Sigma}^{-1} \mathbf{1}_K = \mathbf{h}^\top \mathbf{V}^\top (\mathbf{V} \Sigma \mathbf{V}^\top)^{-1} \mathbf{V} \mathbf{h}$$

where we used (78)–(80). Comparing with $\mathbf{h}^\top \Sigma^{-1} \mathbf{y}$ and $\mathbf{h}^\top \Sigma^{-1} \mathbf{h}$ arising in (89) and (90), respectively, we conclude that the difference between the oracle posterior and the supra-Bayesian posterior is that the matrix Σ^{-1} involved in the former is replaced by $\mathbf{V}^\top (\mathbf{V} \Sigma \mathbf{V}^\top)^{-1} \mathbf{V}$.

A simplified version of Theorem 21 has been shown in [91] and is the setting of the early supra-Bayesian approaches. More specifically, in [91], it is assumed that a fusion center obtains from K agents estimates μ_k of a scalar random variable θ . These estimates can be interpreted as our local statistics t_k . Furthermore, the fusion center has a Gaussian prior for θ and knows that the vector of the estimation errors of all agents, $\mathbf{u} = [u_1, \dots, u_K]^\top$ with $u_k = \mu_k - \theta$, also follows a Gaussian distribution with zero mean and some covariance matrix $\tilde{\Sigma}$ (in general, the errors may be correlated). Equivalently, conditionally on θ , the estimates $\mu = [\mu_1, \dots, \mu_K]^\top$ follow a Gaussian distribution with mean $\mathbf{1}_K \theta$ and the same covariance matrix $\tilde{\Sigma}$. Thus, the setting in [91] directly assumes the conditional distribution of \mathbf{t} given θ without starting from any detailed observation model.

To get a better intuition about the role of the weights w_k and the meaning of negative weights in the setting of Theorem 21, we will consider a specific example.

Example 3 (Private and Shared Observations): We assume that agent k has r_k private observations, i.e., observations that no other agent observes, and r_0 shared observations, i.e., observations that all agents know jointly. The resulting total number of observations is, thus, $d_y = \sum_{k=1}^K (r_0 + r_k)$. However, there are only $r_0 + \sum_{k=1}^K r_k$ different observations. We assume that these different observations given θ are independent and have variance one and mean θ . To embed this scenario into our linear model, we choose $\mathbf{h}_k = \mathbf{1}_{r_0+r_k}$ and the submatrices of the

covariance matrix Σ in (69) as

$$\Sigma_{kk'} = \begin{pmatrix} \mathbf{I}_{r_0} & \mathbf{0}_{r_0 \times r_{k'}} \\ \mathbf{0}_{r_k \times r_0} & \mathbf{0}_{r_k \times r_{k'}} \end{pmatrix} \in \mathbb{R}^{(r_0+r_k) \times (r_0+r_{k'})} \quad (91)$$

for $k \neq k'$ and

$$\Sigma_{kk} = \mathbf{I}_{r_0+r_k}.$$

Thus, we have that

$$\mathbf{y}_k = \mathbf{1}_{r_0+r_k} \theta + \mathbf{n}_k$$

where \mathbf{n}_k is a vector of independent identically distributed standard Gaussian random variables, i.e., $p(\mathbf{n}_k) = \mathcal{N}(\mathbf{n}_k; \mathbf{0}_{(r_0+r_k) \times 1}, \mathbf{I}_{r_0+r_k})$. The covariance structure (91) between the \mathbf{n}_k , for $k \in \{1, \dots, K\}$, implies that, for $i \in \{1, \dots, r_0\}$, the i th entry of \mathbf{n}_k and the i th entry of $\mathbf{n}_{k'}$ with $k' \neq k$ coincide with probability one, i.e.,

$$\mathbb{E}[(n_{k,i} - n_{k',i})^2] = \underbrace{\mathbb{E}[n_{k,i}^2]}_{=1} + \underbrace{\mathbb{E}[n_{k',i}^2]}_{=1} - 2 \underbrace{\mathbb{E}[n_{k,i} n_{k',i}]}_{=1} = 0.$$

Thus, the first r_0 observations are the same for all agents.

With these choices and assuming that $r_k > 0$ and $r_0 > 0$, a tedious but straightforward calculation (for details see Appendix K-B) shows that the weights w_k in (85) simplify to

$$w_k = 1 - \frac{K-1}{r_k} \left(\sum_{k'=0}^K \frac{1}{r_{k'}} \right)^{-1}. \quad (92)$$

In particular, we see that all weights are upper bounded by 1 and are emphasized according to their amount of independent information as given by r_k . More surprising is the possibility of negative weights for agents with few private observations (e.g., the setting $K = 3$, $r_1 = 1$, and $r_0 = r_2 = r_3 = 4$ gives $w_1 = -1/7$). An explanation for this result is that negatively weighting agents with few private observations can counteract the repeated counting of the shared observations that are part of all agents' posteriors. More generally, it follows from (92) that $w_k \geq 0$ if and only if

$$r_k \geq (K-1) \left(\sum_{k'=0}^K \frac{1}{r_{k'}} \right)^{-1}$$

or, equivalently,

$$\sum_{k'=0}^K \frac{r_k}{r_{k'}} \geq K-1.$$

The sum of all weights is given by

$$\sum_{k=1}^K w_k = K - (K-1) \left(\sum_{k'=0}^K \frac{1}{r_{k'}} \right)^{-1} \sum_{k=1}^K \frac{1}{r_k}. \quad (93)$$

From this expression, we readily conclude that

$$1 \leq \sum_{k=1}^K w_k \leq K. \quad (94) \quad \text{where}$$

Indeed, this follows from the fact that the second term on the right-hand side of (93), $(K-1) \left(\sum_{k'=0}^K (1/r_{k'}) \right)^{-1} \times \sum_{k=1}^K (1/r_k)$, is nonnegative and upper bounded by $K-1$ since $\left(\sum_{k'=0}^K (1/r_{k'}) \right)^{-1} \sum_{k=1}^K (1/r_k) \leq 1$. The twofold bound (94) shows that, although some weights may be negative, the sum of all weights is always between the sum of all weights in the log-linear pooling function in (4) (there, the sum was 1) and the sum of all weights in the multiplicative pooling function in (10) (there, all weights were 1, and hence, the sum was K).

Another conclusion that we can draw is that varying the number of shared observations r_0 —while keeping the number of private observations r_k fixed—corresponds to an “interpolation” between the multiplicative pooling function and the log-linear pooling function. First, consider the case that the agents have the same number of private observations, i.e., $r_1 = \dots = r_K$. When $r_0 = 0$, a derivation similar to that in Appendix K-B gives $w_k = 1$. This implies that, when the agents do not share any observations, the pooling function in (86) corresponds exactly to the standard multiplicative pooling function in (10). On the other hand, as the number of shared observations r_0 increases, the pooling function behaves closer to a symmetric log-linear pooling function (i.e., using $w_k = 1/K$). Indeed, it follows from (92) that

$$\lim_{r_0 \rightarrow \infty} w_k = \frac{1}{K}.$$

If we remove the restriction that $r_1 = \dots = r_K$, the connection to multiplicative pooling still holds; however, the connection to log-linear pooling only holds under the condition of nonnegative weights, i.e., $w_k \geq 0$ for all k , which may be violated if some agents hold only few private observations compared to the total number of observations.

D. Supra-Bayesian Fusion Rule for a Vector θ

We can generalize Theorem 21 to a vector $\theta \in \mathbb{R}^{d_\theta}$ with $d_\theta > 1$. However, formally, the weights w_k in (85) become matrices \mathbf{W}_k and, thus, cannot be used as powers in a fusion rule. Hence, the following fusion result is more complicated, and the relation to the 1-D case is not obvious. A proof is provided in Appendix L-A.

Theorem 23: Let $\ell_k(\theta) = p(\mathbf{y}_k | \theta)$ denote the local observation likelihood functions given by (70) for $k = 1, \dots, K$, and let $\lambda(\theta) = p(\mathbf{t} | \theta)$ be the global likelihood function given by (81). Then,

$$\lambda(\theta) \propto \xi_0(\theta) \prod_{k=1}^K \ell_k(\mathbf{W}_k \theta) \quad (95)$$

$$\mathbf{W}_k = (\mathbf{H}_k^\top \Sigma_{kk}^{-1} \mathbf{H}_k)^{-1} (\mathbf{e}_k \otimes \mathbf{I}_{d_\theta})^\top \tilde{\Sigma}^{-1} (\mathbf{1}_K \otimes \mathbf{I}_{d_\theta}) \quad (96)$$

with \mathbf{e}_k denoting the k th unit vector in \mathbb{R}^K and $\tilde{\Sigma} = \mathbf{V} \Sigma \mathbf{V}^\top$, and

$$\xi_0(\theta) = \exp \left(-\frac{\theta^\top \mathbf{G} \theta}{2} \right). \quad (97)$$

Here,

$$\mathbf{G} = \hat{\Sigma}^{-1} - \sum_{k=1}^K \mathbf{W}_k^\top \mathbf{H}_k^\top \Sigma_{kk}^{-1} \mathbf{H}_k \mathbf{W}_k \quad (98)$$

with

$$\hat{\Sigma}^{-1} = (\mathbf{1}_K \otimes \mathbf{I}_{d_\theta})^\top \tilde{\Sigma}^{-1} (\mathbf{1}_K \otimes \mathbf{I}_{d_\theta}). \quad (99)$$

Furthermore, for a given prior $p(\theta)$ and local posteriors $\pi_k(\theta) = p(\theta | \mathbf{y}_k) \propto p(\theta) \ell_k(\theta)$, the supra-Bayesian fusion result $g[\pi_1, \dots, \pi_K](\theta) = p(\theta | \mathbf{t}) \propto p(\theta) \lambda(\theta)$ is given by

$$g[\pi_1, \dots, \pi_K](\theta) \propto p(\theta) \xi_0(\theta) \prod_{k=1}^K \frac{\pi_k(\mathbf{W}_k \theta)}{p(\mathbf{W}_k \theta)}. \quad (100)$$

Finally, if the prior $p(\theta)$ is Gaussian, then the supra-Bayesian fusion result $p(\theta | \mathbf{t})$ is again Gaussian, and the fusion rule (100) can be reduced to a second-order rule involving only the mean and covariance matrix.

Corollary 24: Under the assumptions of Theorem 23, let the prior $p(\theta)$ be Gaussian with mean μ_0 and covariance matrix Σ_0 , i.e., $p(\theta) = \mathcal{N}(\theta; \mu_0, \Sigma_0)$. Then, the supra-Bayesian fusion result $p(\theta | \mathbf{t})$ is again Gaussian, i.e., $p(\theta | \mathbf{t}) = \mathcal{N}(\theta; \mu_1, \Sigma_1)$, with mean

$$\mu_1 = (\hat{\Sigma}^{-1} + \Sigma_0^{-1})^{-1} ((\mathbf{1}_K \otimes \mathbf{I}_{d_\theta})^\top \tilde{\Sigma}^{-1} \mathbf{t} + \Sigma_0^{-1} \mu_0) \quad (101)$$

and covariance matrix

$$\Sigma_1 = (\hat{\Sigma}^{-1} + \Sigma_0^{-1})^{-1} \quad (102)$$

where $\hat{\Sigma}^{-1}$ is given by (99).

Here, we recall that $\mathbf{t} = [\mathbf{t}_1^\top, \dots, \mathbf{t}_K^\top]^\top$ with \mathbf{t}_k given by (73). A proof of Corollary 24 is provided in Appendix L-B.

The supra-Bayesian fusion result in (100) has an intriguing structure: the agent pdfs are first preprocessed by a multiplication in the argument and then combined via a generalized multiplicative pooling function. The relevance of this fusion rule beyond the linear Gaussian setting, especially for approximately linear Gaussian observation models, is an open issue.

As in the scalar case, the supra-Bayesian fusion result $p(\theta | \mathbf{t})$ is, in general, different from the oracle posterior $p(\theta | \mathbf{y})$. Again, the oracle posterior is proportional to the product of the prior $p(\theta)$ and the global observation likelihood function $p(\mathbf{y} | \theta)$ in (77); it is easily seen that $p(\theta | \mathbf{y})$ is also Gaussian but with mean

$$\boldsymbol{\mu}_2 = (\hat{\boldsymbol{\Sigma}}_2^{-1} + \boldsymbol{\Sigma}_0^{-1})^{-1} (\mathbf{H}^\top \boldsymbol{\Sigma}^{-1} \mathbf{y} + \boldsymbol{\Sigma}_0^{-1} \boldsymbol{\mu}_0) \quad (103)$$

and covariance matrix

$$\boldsymbol{\Sigma}_2 = (\hat{\boldsymbol{\Sigma}}_2^{-1} + \boldsymbol{\Sigma}_0^{-1})^{-1}$$

where

$$\hat{\boldsymbol{\Sigma}}_2^{-1} = \mathbf{H}^\top \boldsymbol{\Sigma}^{-1} \mathbf{H}. \quad (104)$$

The difference can be better understood by noting that in (101)

$$(\mathbf{1}_K \otimes \mathbf{I}_{d_\theta})^\top \tilde{\boldsymbol{\Sigma}}^{-1} \mathbf{t} = \mathbf{H}^\top \mathbf{V}^\top (\mathbf{V} \boldsymbol{\Sigma} \mathbf{V}^\top)^{-1} \mathbf{V} \mathbf{y}$$

and in (99)

$$(\mathbf{1}_K \otimes \mathbf{I}_{d_\theta})^\top \tilde{\boldsymbol{\Sigma}}^{-1} (\mathbf{1}_K \otimes \mathbf{I}_{d_\theta}) = \mathbf{H}^\top \mathbf{V}^\top (\mathbf{V} \boldsymbol{\Sigma} \mathbf{V}^\top)^{-1} \mathbf{V} \mathbf{H}$$

where we used (78)–(80). Comparing with $\mathbf{H}^\top \boldsymbol{\Sigma}^{-1} \mathbf{y}$ and $\mathbf{H}^\top \boldsymbol{\Sigma}^{-1} \mathbf{H}$ in (103) and (104), respectively, we conclude that, as in the scalar case considered earlier, the difference between the oracle posterior and the supra-Bayesian posterior is that $\boldsymbol{\Sigma}^{-1}$ involved in the former is replaced by $\mathbf{V}^\top (\mathbf{V} \boldsymbol{\Sigma} \mathbf{V}^\top)^{-1} \mathbf{V}$.

X. OUTLOOK

The fusion of pdfs presents numerous interesting aspects beyond those considered in our treatment. Moreover, certain extensions can be envisioned. In what follows, we suggest some related directions of future research.

- 1) Our discussion of pdf fusion emphasized theoretical considerations. In practical implementations, a finite-dimensional representation or parametrization of the agent pdfs $q_k(\theta)$ is required. Popular examples are Gaussian, Gaussian mixture, and particle representations [31], [68], [69]. Since these representations are usually approximations of the true pdfs, a relevant issue is the tradeoff between low representation complexity (small number of parameters) and high accuracy of the approximation. Furthermore,

algorithms implementing a given pooling function for a given type of parametric representation are required. Examples of finite-dimensional parametric fusion rules were considered in Sections VI and IX.

- 2) In the case of a centralized agent network where each agent pdf $q_k(\theta)$ is transmitted to the fusion center via a channel, communication cost is another practical issue. Although a low-dimensional parametric representation of the agent pdfs may be used to achieve a low communication cost, the reduction of communication cost is ultimately a source coding (rate-distortion) problem.
- 3) In many cases, the aggregate pdf $q(\theta) = g[q_1, \dots, q_K](\theta)$ is not used as the final result but arises as part of a method performing a statistical inference task, such as estimation, detection, classification, or clustering. In this setting, the pooling function (or certain parameters within a given family of pooling functions) should be chosen or optimized such that the performance of the statistical inference method is maximized. Note that this is different from the optimization approach considered in Section V.
- 4) Our discussion assumed the existence of a fusion center that has access to all pdfs $q_k(\theta)$. In a decentralized agent network, there is no fusion center and each agent is able to communicate only with certain neighboring agents. Besides the basic necessity of using a distributed communication-and-fusion protocol, challenging aspects in the decentralized setting include communication cost, efficient representation of pdfs, and repeated counting of information along cycles in the network graph.
- 5) In many scenarios, the agent pdfs $q_k(\theta)$ are time-varying, and a temporal sequence $q_k^{(n)}(\theta)$, where $n = 1, 2, \dots$ is a discrete time index, is available at the k th agent. This serial setting suggests a sequential variant of pdf fusion in which, at each time n , the fused pdf is not calculated from scratch, but the previous fusion result is updated using the new set of $q_k^{(n)}(\theta)$. Practical implementations of sequential updating can be based on both parametric and nonparametric representations of the pdfs.
- 6) The fusion of multiobject pdfs or PHDs of finite point processes (random finite sets), especially in the context of multitarget tracking, is a topic of active research [2], [5], [6], [8], [43]–[50], [85], [106]–[114]. While the current focus is on the finite point process counterparts of the linear and log-linear pooling functions, it would also be interesting to investigate the applicability of the other pooling functions considered in Sections III and V. In particular, the fact that the family of Hölder pooling functions offers fusion characteristics that are intermediate between those of the linear and log-linear pooling functions may be relevant to multitarget tracking. Furthermore, it may be rewarding to reformulate and develop our

results on supra-Bayesian pdf fusion in the context of finite point processes.

7) Big data problems allow a natural application of pdf fusion. When the data to be processed are so large in size that they exceed the capacity of a single computer, it is logical to partition them and process the different parts separately. Furthermore, data related to some quantity of interest may be available in heterogeneous form, so that all of the data cannot be processed within a single framework, and hence, different parts have to be processed separately. In either case, the individual processing results can be represented as summaries, which then need to be fused into one overall summary. The concepts and techniques presented in this article provide suggestions regarding the construction and fusion of the summaries. This is of particular interest in the context of modern machine learning methods [65]–[67], [136], [193].

8) Ensemble learning [128], [130], i.e., the combination of the results of multiple learning algorithms, is currently one of the most successful learning paradigms. At the same time, there is a growing demand for probabilistic machine learning methods that provide along with a point estimate also a measure of reliability. Up until now, only few works have considered ensembles of probabilistic machine learning methods. We conjecture that the success of the ensemble learning paradigm will soon lead to its increased use also in probabilistic machine learning. At that point, it is likely that probabilistic opinion pooling will outperform the simple linear voting rules that are currently used to combine point estimates.

9) With a collaborative machine learning methodology known as federated learning, a learning algorithm is trained across multiple decentralized edge devices or servers that hold local data, which are not exchanged [134]. In other words, model parameters are learned collectively by many interconnected devices without sharing or disclosing local training data. The devices send summaries instead of raw data to a server for fusion. Here, again, fusion plays a central role. The fusion process can be challenging in the case of a large number of heterogeneous devices with different constraints. Using pdfs to represent the local summaries enables the use of different pdf representations at the individual devices, from simple parametric models to complex kernel density estimates, which can still be combined in a meaningful way. Moreover, different levels of quality of the local data can be taken into account by using appropriate weights in the pooling function used for pdf fusion.

10) A potential theoretical basis of pdf fusion that has not been explored in this work is information geometry, which studies probability theory and statistics using tools from differential geometry [42]. The focus of information geometry is on statistical manifolds whose points correspond to probability distributions. This theoretical framework can be exploited for fusion by assuming that the agent pdfs belong to a parametric family with the structure of a Riemannian manifold [42]. One can then formulate pdf fusion, e.g., by considering the fused pdf to be an informative barycenter of the manifold [194].

11) Within the finite-dimensional supra-Bayesian setting, an explicit fusion rule was obtained for linear Gaussian observation models (see Section IX). This fusion rule can formally be used also for nonlinear/non-Gaussian models with known first and second moments. However, it is here unclear how close the obtained fusion result will be to the true supra-Bayesian fusion result. Characterizing the error for approximately linear Gaussian observation models is an interesting topic for future research. Another interesting topic is the derivation of explicit supra-Bayesian fusion rules for simple nonlinear/non-Gaussian observation models.

12) Our supra-Bayesian framework is currently limited to a finite-dimensional setting. Although this is the setting most frequently encountered in practical applications, it would be interesting to find a definition of a likelihood function for random pdfs that do not admit a finite-dimensional parameterization. For this, nonparametric Bayesian models [195] appear to be a feasible starting point. The challenge is to model a useful and nontrivial dependence on the parameter θ that accounts for the constraint that random pdfs must be nonnegative and integrate to one with probability one.

XI. CONCLUSION

The problem of fusing multiple pdfs $q_k(\theta)$, $k = 1, \dots, K$ of a continuous random vector θ into an aggregate pdf $q(\theta) = g[q_1, \dots, q_K](\theta)$ has many possible solutions, and indeed, several different approaches to this fusion problem have been developed in the past decades. We have attempted to survey and study these approaches and the related solutions in a structured and coherent manner. Our discussion has emphasized a first basic distinction between the axiomatic approach, the optimization approach, and the conceptually more complex supra-Bayesian framework.

Regarding the axiomatic approach, we formulated a set of axioms and determined the axioms satisfied by each considered pooling function. This analysis demonstrated the prominent roles of the linear, log-linear, and multiplicative pooling functions within the axiomatic framework. However, it also revealed that several desirable axioms are effectively incompatible, and postulating those simultaneously implies a dictatorship pooling function.

Regarding the optimization approach, besides other results, we proved that the minimization of the weighted sum of α -divergences yields the family of Hölder mean pooling functions. This family contains the two most popular pooling functions—the linear and log-linear pooling

functions—as special cases. Moreover, it offers an infinite number of further interesting pooling functions with different multimodality and tail decay characteristics depending on the choice of a single parameter.

The supra-Bayesian framework is different from the classical probabilistic opinion pooling framework in that the pdfs $q_k(\theta)$ are modeled as random observations, and additional information regarding the statistical structure of θ is available to the fusion center. In this framework, the optimal aggregate pdf $q(\theta)$ is the global posterior pdf of θ given the pdfs $q_k(\theta)$. Since random functions are difficult to work with, we introduced the finite-dimensional supra-Bayesian model based on random “local statistics.” Using this framework, we formulated a general procedure for obtaining the supra-Bayesian posterior pdf conditioned on all the local statistics, and we derived explicit fusion rules for special cases.

While the theory of pdf fusion appears mature, interesting directions of future work are related to implementation and application aspects. We provided some suggestions including implementations using parametric representations, integration into probabilistic methods for multiagent signal processing, machine learning, and forecasting, and extensions to decentralized scenarios and point processes. \square

APPENDIX A

PROOF OF THEOREM 1

A. Axioms Satisfied by the Linear Pooling Function

We first show that all the mentioned axioms are satisfied by the linear pooling function. Let $g[q_1, \dots, q_K](\theta) = \sum_{k=1}^K w_k q_k(\theta)$ with $(w_1, \dots, w_K) \in \mathcal{S}_K$. We first show the ZPP (A2). Assume that, for some event \mathcal{A} , we have $Q_k(\mathcal{A}) = 0$ for all $k = 1, \dots, K$. Because $Q_k(\mathcal{A}) = \int_{\mathcal{A}} q_k(\theta) d\theta$ and $q_k(\theta)$ is nonnegative, this implies that $q_k(\theta) = 0$ for almost all $\theta \in \mathcal{A}$ and all $k = 1, \dots, K$. Thus,

$$q(\theta) = g[q_1, \dots, q_K](\theta) = \sum_{k=1}^K w_k q_k(\theta) = 0$$

for almost all $\theta \in \mathcal{A}$. Hence, $Q(\mathcal{A}) = \int_{\mathcal{A}} q(\theta) d\theta = 0$, which concludes the proof of the ZPP.

We next show unanimity preservation (A3). To this end, assume that $q_k(\theta) = q_0(\theta)$ for all $k = 1, \dots, K$. Then,

$$\begin{aligned} q(\theta) &= g[q_1, \dots, q_K](\theta) \\ &= \sum_{k=1}^K w_k q_k(\theta) \\ &= q_0(\theta) \sum_{k=1}^K w_k \\ &= q_0(\theta) \end{aligned}$$

which shows unanimity preservation.

To show the SSFP (A4), we define $h: [0, 1]^K \rightarrow [0, 1]$ as

$$h(p_1, \dots, p_K) \triangleq \sum_{k=1}^K w_k p_k. \quad (105)$$

For an arbitrary set $\mathcal{A} \subseteq \Theta$ and any opinion profile (q_1, \dots, q_K) , we have that

$$\begin{aligned} Q(\mathcal{A}) &= \int_{\mathcal{A}} q(\theta) d\theta \\ &= \int_{\mathcal{A}} \sum_{k=1}^K w_k q_k(\theta) d\theta \\ &= \sum_{k=1}^K w_k \int_{\mathcal{A}} q_k(\theta) d\theta \\ &= \sum_{k=1}^K w_k Q_k(\mathcal{A}) \\ &= h(Q_1(\mathcal{A}), \dots, Q_K(\mathcal{A})) \end{aligned}$$

i.e., h satisfies the condition stated in A4.

The WSFP (A5) follows by setting $h_{\mathcal{A}} = h$ with h given in (105). The likelihood principles (A6 and A7) are obviously satisfied with

$$h(t_1, \dots, t_K) \triangleq \sum_{k=1}^K w_k t_k$$

and $h_{\theta} = h$, respectively, as is the symmetry statement.

B. Equivalence Statement

We now prove the other direction, namely, that each of the assumptions 2–4 stated in Theorem 1 implies that g is a linear pooling function. More specifically, we will show the chain of implications $4 \Rightarrow 3 \Rightarrow 2 \Rightarrow 1$. Because we already showed $1 \Rightarrow 4$, this implies that 1–4 are equivalent and, thus, concludes the proof.

1) 4 Implies 3: We assume that g satisfies 4, i.e., the WSFP (A5) and unanimity preservation (A3). We will show that this implies that g satisfies the ZPP (A2), i.e., 4 implies 3. Let $h_{\mathcal{A}}: [0, 1]^K \rightarrow [0, 1]$ denote the function satisfying (16) for all opinion profiles. For any set \mathcal{A} that satisfies $|\mathcal{A}^c| > 0$, let us choose $q_k(\theta) = \mathbb{1}_{\mathcal{A}^c}(\theta)/|\mathcal{A}^c|$ for all $k = 1, \dots, K$. Then, $Q_k(\mathcal{A}) = \int_{\mathcal{A}} q_k(\theta) d\theta = 0$ for all k . By unanimity preservation, this implies that $Q(\mathcal{A}) = 0$. On the other hand, we have

$$Q(\mathcal{A}) \stackrel{(16)}{=} h_{\mathcal{A}}(Q_1(\mathcal{A}), \dots, Q_K(\mathcal{A})) = h_{\mathcal{A}}(0, \dots, 0)$$

and, hence,

$$h_{\mathcal{A}}(0, \dots, 0) = 0 \quad (106)$$

for any set \mathcal{A} such that $|\mathcal{A}^c| > 0$.

To show the ZPP, assume that, for a given opinion profile (q_1, \dots, q_K) , we have $Q_k(\mathcal{A}) = 0$ for all $k = 1, \dots, K$. Note that this is only possible if $|\mathcal{A}^c| > 0$ as otherwise $Q_k(\Theta) = Q_k(\mathcal{A}) + Q_k(\mathcal{A}^c) = 0$. Thus, we can calculate $Q(\mathcal{A})$ as

$$Q(\mathcal{A}) \stackrel{(16)}{=} h_{\mathcal{A}}(Q_1(\mathcal{A}), \dots, Q_K(\mathcal{A})) = h_{\mathcal{A}}(0, \dots, 0) \stackrel{(106)}{=} 0$$

which shows that the ZPP (A2) is satisfied.

2) 3 Implies 2: Next, we show that 3, i.e., the WSFP (A5) and the ZPP (A2), implies 2, i.e., the SSFP (A4). Let again $h_{\mathcal{A}} : [0, 1]^K \rightarrow [0, 1]$ denote the function satisfying (16) for all opinion profiles. Our proof consists of three steps.

- 1) Show that, for two nontrivial events \mathcal{A} and \mathcal{B} (i.e., $|\mathcal{A}|, |\mathcal{A}^c|, |\mathcal{B}|, |\mathcal{B}^c| > 0$) that have a nontrivial intersection and a nontrivial union, we have $h_{\mathcal{A}} = h_{\mathcal{B}}$.
- 2) Show that, for any nontrivial events \mathcal{A} and \mathcal{B} , there exists a nontrivial event \mathcal{C} such that $\mathcal{A} \cap \mathcal{C}$, $\mathcal{A} \cup \mathcal{C}$, $\mathcal{B} \cap \mathcal{C}$, and $\mathcal{B} \cup \mathcal{C}$ are all nontrivial. This implies by step 1 that $h_{\mathcal{A}} = h_{\mathcal{C}}$ and $h_{\mathcal{B}} = h_{\mathcal{C}}$, and thus, $h_{\mathcal{A}} = h_{\mathcal{B}}$. Therefore, setting $h \triangleq h_{\mathcal{A}}$, we have $h_{\mathcal{A}'} = h$ for all nontrivial events \mathcal{A}' , and hence, the same function h satisfies (14) for all nontrivial events.
- 3) Show that the function h satisfies (14) also for trivial events.

To show step 1, we consider two nontrivial events \mathcal{A} and \mathcal{B} that have a nontrivial intersection, in particular, $|\mathcal{A} \cap \mathcal{B}| > 0$, and a nontrivial union, in particular, $|\mathcal{A} \cup \mathcal{B}| > 0$. We fix arbitrary $(p_1, \dots, p_K) \in [0, 1]^K$ and will show that $h_{\mathcal{A}}(p_1, \dots, p_K) = h_{\mathcal{B}}(p_1, \dots, p_K)$. Because $|\mathcal{A} \cap \mathcal{B}| > 0$ and $|\mathcal{A} \cup \mathcal{B}| > 0$, there exists an opinion profile (q_1, \dots, q_K) such that

$$Q_k(\mathcal{A} \cap \mathcal{B}) = p_k \quad (107)$$

and

$$Q_k((\mathcal{A} \cup \mathcal{B})^c) = 1 - p_k \quad (108)$$

for all $k = 1, \dots, K$. Because $\Theta = (\mathcal{A} \cup \mathcal{B}) \cup (\mathcal{A} \cup \mathcal{B})^c$ is a disjoint union and $Q_k(\Theta) = 1$, (108) implies that $Q_k(\mathcal{A} \cup \mathcal{B}) = p_k$. Hence, as also $Q_k(\mathcal{A} \cap \mathcal{B}) = p_k$ by (107), the difference set $(\mathcal{A} \cup \mathcal{B}) \setminus (\mathcal{A} \cap \mathcal{B})$ satisfies

$$Q_k((\mathcal{A} \cup \mathcal{B}) \setminus (\mathcal{A} \cap \mathcal{B})) = 0. \quad (109)$$

Because $(\mathcal{A} \cup \mathcal{B}) \setminus (\mathcal{A} \cap \mathcal{B}) = (\mathcal{A} \setminus \mathcal{B}) \cup (\mathcal{B} \setminus \mathcal{A})$, this implies that $Q_k(\mathcal{A} \setminus \mathcal{B}) = 0$ and $Q_k(\mathcal{B} \setminus \mathcal{A}) = 0$. Thus,

$$\begin{aligned} Q_k(\mathcal{A}) &= Q_k((\mathcal{A} \setminus \mathcal{B}) \cup (\mathcal{A} \cap \mathcal{B})) \\ &= Q_k(\mathcal{A} \setminus \mathcal{B}) + Q_k(\mathcal{A} \cap \mathcal{B}) \\ &= p_k \end{aligned} \quad (110)$$

and, similarly,

$$Q_k(\mathcal{B}) = p_k. \quad (111)$$

Furthermore,

$$Q_k((\mathcal{A} \setminus \mathcal{B}) \cup (\mathcal{B} \setminus \mathcal{A})) = Q_k((\mathcal{A} \cup \mathcal{B}) \setminus (\mathcal{A} \cap \mathcal{B})) \stackrel{(109)}{=} 0 \quad (112)$$

for all $k = 1, \dots, K$. By the ZPP (112) implies that $Q((\mathcal{A} \setminus \mathcal{B}) \cup (\mathcal{B} \setminus \mathcal{A})) = 0$, and in turn, $Q(\mathcal{A} \setminus \mathcal{B}) = Q(\mathcal{B} \setminus \mathcal{A}) = 0$. Thus,

$$\begin{aligned} Q(\mathcal{A}) &= Q(\mathcal{A} \setminus \mathcal{B}) + Q(\mathcal{A} \cap \mathcal{B}) \\ &= Q(\mathcal{A} \cap \mathcal{B}) \\ &= Q(\mathcal{B} \setminus \mathcal{A}) + Q(\mathcal{A} \cap \mathcal{B}) \\ &= Q(\mathcal{B}). \end{aligned} \quad (113)$$

For the functions $h_{\mathcal{A}}$ and $h_{\mathcal{B}}$, these properties imply that

$$\begin{aligned} h_{\mathcal{A}}(p_1, \dots, p_K) &\stackrel{(110)}{=} h_{\mathcal{A}}(Q_1(\mathcal{A}), \dots, Q_K(\mathcal{A})) \\ &\stackrel{(16)}{=} Q(\mathcal{A}) \\ &\stackrel{(113)}{=} Q(\mathcal{B}) \\ &\stackrel{(16)}{=} h_{\mathcal{B}}(Q_1(\mathcal{B}), \dots, Q_K(\mathcal{B})) \\ &\stackrel{(111)}{=} h_{\mathcal{B}}(p_1, \dots, p_K) \end{aligned}$$

i.e.,

$$h_{\mathcal{A}}(p_1, \dots, p_K) = h_{\mathcal{B}}(p_1, \dots, p_K) \quad (114)$$

for any nontrivial events $\mathcal{A}, \mathcal{B} \subseteq \Theta$ that have a nontrivial intersection and a nontrivial union.

To show step 2, we first construct a set $\mathcal{C} \subseteq \mathcal{A} \cup \mathcal{B}$ such that $\mathcal{A} \cap \mathcal{C}$, $\mathcal{B} \cap \mathcal{C}$, $\mathcal{A} \cup \mathcal{C}$, and $\mathcal{B} \cup \mathcal{C}$ are nontrivial. If $|\mathcal{A} \cap \mathcal{B}| > 0$, then $\mathcal{C} = \mathcal{A} \cap \mathcal{B}$ can easily be seen to satisfy these assumptions. If $|\mathcal{A} \cap \mathcal{B}| = 0$, we choose $\mathcal{C} = \mathcal{C}_{\mathcal{A}} \cup \mathcal{C}_{\mathcal{B}}$, where $\mathcal{C}_{\mathcal{A}} \subseteq \mathcal{A}$ with $|\mathcal{C}_{\mathcal{A}}|, |\mathcal{A} \setminus \mathcal{C}_{\mathcal{A}}| > 0$ and $\mathcal{C}_{\mathcal{B}} \subseteq \mathcal{B}$ with $|\mathcal{C}_{\mathcal{B}}|, |\mathcal{B} \setminus \mathcal{C}_{\mathcal{B}}| > 0$. The separations $\mathcal{A} = \mathcal{C}_{\mathcal{A}} \cup (\mathcal{A} \setminus \mathcal{C}_{\mathcal{A}})$ and $\mathcal{B} = \mathcal{C}_{\mathcal{B}} \cup (\mathcal{B} \setminus \mathcal{C}_{\mathcal{B}})$ are possible because the Lebesgue measure is nonatomic, i.e., any set of positive Lebesgue measure can be separated into two disjoint sets of positive Lebesgue measure.

Now, we choose

$$h(p_1, \dots, p_K) = h_{\mathcal{A}}(p_1, \dots, p_K) \quad (115)$$

for any nontrivial set \mathcal{A} . Then, for any nontrivial set $\mathcal{B} \subseteq \Theta$, we construct \mathcal{C} as above and obtain

$$\begin{aligned} h(Q_1(\mathcal{B}), \dots, Q_K(\mathcal{B})) &\stackrel{(115)}{=} h_{\mathcal{A}}(Q_1(\mathcal{B}), \dots, Q_K(\mathcal{B})) \\ &\stackrel{(114)}{=} h_{\mathcal{C}}(Q_1(\mathcal{B}), \dots, Q_K(\mathcal{B})) \\ &\stackrel{(114)}{=} h_{\mathcal{B}}(Q_1(\mathcal{B}), \dots, Q_K(\mathcal{B})) \\ &\stackrel{(16)}{=} Q(\mathcal{B}) \end{aligned} \quad (116)$$

i.e., (14) is satisfied for any nontrivial set \mathcal{B} .

It remains to show step 3, i.e., with this choice of h , (14) is also satisfied by trivial sets. For trivial sets \mathcal{A} , i.e., such that $|\mathcal{A}| = 0$ or $|\mathcal{A}^c| = 0$, we have $Q_k(\mathcal{A}) = 0$ or $Q_k(\mathcal{A}) = 1$ for all $k = 1, \dots, K$, respectively. Also, the fused result must satisfy $Q(\mathcal{A}) = 0$ or $Q(\mathcal{A}) = 1$, respectively. Thus, we have to show $h(0, \dots, 0) = 0$ and $h(1, \dots, 1) = 1$ for our choice of h in (115). To this end, let $\mathcal{B} \subseteq \Theta$ be any nontrivial set and choose an opinion profile (q_1, \dots, q_K) such that $Q_k(\mathcal{B}) = 0$ for all $k = 1, \dots, K$. Then, the ZPP implies that $Q(\mathcal{B}) = 0$. On the other hand, since \mathcal{B} is a nontrivial set and, thus, (116) is satisfied, we have

$$h(0, \dots, 0) = h(Q_1(\mathcal{B}), \dots, Q_K(\mathcal{B})) \stackrel{(116)}{=} Q(\mathcal{B}).$$

Thus, $h(0, \dots, 0) = 0$. Furthermore, $Q_k(\mathcal{B}) = 0$ and $Q(\mathcal{B}) = 0$ imply that $Q_k(\mathcal{B}^c) = 1$ and $Q(\mathcal{B}^c) = 1$, respectively. Hence,

$$h(1, \dots, 1) = h(Q_1(\mathcal{B}^c), \dots, Q_K(\mathcal{B}^c)) \stackrel{(116)}{=} Q(\mathcal{B}^c).$$

Thus, $h(1, \dots, 1) = 1$. Hence, we identified a function h such that (14) holds for all sets $\mathcal{A} \subseteq \Theta$. This concludes the proof that 3 implies 2.

3) 2 Implies 1: Finally, we show that 2, i.e., the SSFP (A4), implies 1, i.e., g is a linear pooling function. Let h denote the function satisfying (14). Furthermore, let $\mathcal{A}, \mathcal{B}, \mathcal{C} \subseteq \Theta$ be disjoint events of positive Lebesgue measure. For arbitrary $p_1, \tilde{p}_1, \dots, p_K, \tilde{p}_K \in [0, 1]$ satisfying $p_k + \tilde{p}_k \leq 1$ for all $k = 1, \dots, K$, we define an opinion profile (q_1, \dots, q_K) such that $Q_k(\mathcal{A}) = p_k$, $Q_k(\mathcal{B}) = \tilde{p}_k$, and $Q_k(\mathcal{C}) = 1 - p_k - \tilde{p}_k$. Because \mathcal{A} and \mathcal{B} are disjoint, $Q_k(\mathcal{A} \cup \mathcal{B}) = p_k + \tilde{p}_k$ and $Q(\mathcal{A} \cup \mathcal{B}) = Q(\mathcal{A}) + Q(\mathcal{B})$. Thus,

$$\begin{aligned} h(p_1 + \tilde{p}_1, \dots, p_K + \tilde{p}_K) &= h(Q_1(\mathcal{A} \cup \mathcal{B}), \dots, Q_K(\mathcal{A} \cup \mathcal{B})) \\ &\stackrel{(14)}{=} Q(\mathcal{A} \cup \mathcal{B}) \\ &= Q(\mathcal{A}) + Q(\mathcal{B}) \\ &\stackrel{(14)}{=} h(Q_1(\mathcal{A}), \dots, Q_K(\mathcal{A})) + h(Q_1(\mathcal{B}), \dots, Q_K(\mathcal{B})) \\ &= h(p_1, \dots, p_K) + h(\tilde{p}_1, \dots, \tilde{p}_K) \end{aligned}$$

i.e., h is an additive function on its domain $[0, 1]^K$. It can, moreover, be extended to an additive function on \mathbb{R}^K . Because h is also bounded by 1 on $[0, 1]^K$, it must be linear according to [196, p. 215, Th. 1], i.e.,

$$h(p_1, \dots, p_K) = \sum_{k=1}^K w_k p_k. \quad (117)$$

Here, the weights w_k must be in $[0, 1]$ because $h(p_1, \dots, p_K) \in [0, 1]$ for all $(p_1, \dots, p_K) \in [0, 1]^K$. Furthermore, because $1 = Q(\Theta) = h(Q_1(\Theta), \dots,$

$Q_K(\Theta)) = h(1, \dots, 1) = \sum_{k=1}^K w_k$, the weights must sum to one. Thus, we have for any event $\mathcal{A} \subseteq \Theta$

$$\begin{aligned} \int_{\mathcal{A}} q(\boldsymbol{\theta}) d\boldsymbol{\theta} &= Q(\mathcal{A}) \\ &\stackrel{(14)}{=} h(Q_1(\mathcal{A}), \dots, Q_K(\mathcal{A})) \\ &\stackrel{(117)}{=} \sum_{k=1}^K w_k Q_k(\mathcal{A}) \\ &= \int_{\mathcal{A}} \sum_{k=1}^K w_k q_k(\boldsymbol{\theta}) d\boldsymbol{\theta} \end{aligned}$$

which implies that $q(\boldsymbol{\theta}) = \sum_{k=1}^K w_k q_k(\boldsymbol{\theta})$.

APPENDIX B PROOF OF THEOREM 2

A. Axioms Satisfied by the Generalized Linear Pooling Function

First, we show that all the mentioned axioms are satisfied by the generalized linear pooling function. Let $g[q_1, \dots, q_K](\boldsymbol{\theta}) = \sum_{k=0}^K w_k q_k(\boldsymbol{\theta})$ with $(w_0, \dots, w_K) \in \mathcal{S}_{K+1}$. To show the WSFP (A5), we define for an event $\mathcal{A} \subseteq \Theta$

$$h_{\mathcal{A}}(p_1, \dots, p_K) = w_0 \int_{\mathcal{A}} q_0(\boldsymbol{\theta}) d\boldsymbol{\theta} + \sum_{k=1}^K w_k p_k \quad (118)$$

for all $(p_1, \dots, p_K) \in [0, 1]^K$. For any opinion profile (q_1, \dots, q_K) , we then have that

$$\begin{aligned} Q(\mathcal{A}) &= \int_{\mathcal{A}} q(\boldsymbol{\theta}) d\boldsymbol{\theta} \\ &= \int_{\mathcal{A}} \sum_{k=0}^K w_k q_k(\boldsymbol{\theta}) d\boldsymbol{\theta} \\ &= \sum_{k=0}^K w_k \int_{\mathcal{A}} q_k(\boldsymbol{\theta}) d\boldsymbol{\theta} \\ &= w_0 \int_{\mathcal{A}} q_0(\boldsymbol{\theta}) d\boldsymbol{\theta} + \sum_{k=1}^K w_k Q_k(\mathcal{A}) \\ &\stackrel{(118)}{=} h_{\mathcal{A}}(Q_1(\mathcal{A}), \dots, Q_K(\mathcal{A})) \end{aligned}$$

i.e., $h_{\mathcal{A}}$ satisfies the condition stated in A5. The weak likelihood principle (A7) is obviously satisfied with

$$h_{\boldsymbol{\theta}}(t_1, \dots, t_K) = w_0 q_0(\boldsymbol{\theta}) + \sum_{k=1}^K w_k t_k$$

as is the symmetry statement in Theorem 2.

B. Converse Statement

Now, we prove the converse statement in Theorem 2, i.e., any pooling function g that satisfies the WSFP (A5) is a

generalized linear pooling function. For each event $\mathcal{A} \subseteq \Theta$, let $h_{\mathcal{A}}: [0, 1]^K \rightarrow [0, 1]$ denote the function satisfying (16) for all opinion profiles. Our proof consists of three steps. First, we construct the pdf q_0 and the corresponding weight $w_0 \leq 1$. In the second step, we show that, by adapting each function $h_{\mathcal{A}}$ to $\tilde{h}_{\mathcal{A}} = (h_{\mathcal{A}} - w_0 \int_{\mathcal{A}} q_0(\boldsymbol{\theta}) d\boldsymbol{\theta}) / (1 - w_0)$, we obtain a linear pooling function. Finally, we show that this implies that g is a generalized linear pooling function.

1) *Step 1: Construct q_0 and w_0 :* We define

$$Q_0(\mathcal{A}) = h_{\mathcal{A}}(0, \dots, 0) \quad (119)$$

for all nontrivial (i.e., $|\mathcal{A}|, |\mathcal{A}^c| > 0$) events $\mathcal{A} \subseteq \Theta$. The pdf q_0 will be a weighted version of a density associated with Q_0 . Thus, we first show that Q_0 can be expressed as an integral $Q_0(\mathcal{A}) = \int_{\mathcal{A}} \tilde{q}_0(\boldsymbol{\theta}) d\boldsymbol{\theta}$.

Let \mathcal{A}_0 be a fixed nontrivial event. Because $|\mathcal{A}_0^c| > 0$, there exists an opinion profile $(q_1^{(\mathcal{A}_0)}, \dots, q_K^{(\mathcal{A}_0)})$ such that $Q_k^{(\mathcal{A}_0)}(\mathcal{A}_0) = \int_{\mathcal{A}_0} q_k^{(\mathcal{A}_0)}(\boldsymbol{\theta}) d\boldsymbol{\theta} = 0$ for all $k = 1, \dots, K$. We denote the fused pdf of this particular profile as $q^{(\mathcal{A}_0)}(\boldsymbol{\theta})$ and the resulting probability measure as

$$Q^{(\mathcal{A}_0)}(\mathcal{A}) = \int_{\mathcal{A}_0} q^{(\mathcal{A}_0)}(\boldsymbol{\theta}) d\boldsymbol{\theta}. \quad (120)$$

Then,

$$Q^{(\mathcal{A}_0)}(\mathcal{A}) \stackrel{(16)}{=} h_{\mathcal{A}}(Q_1^{(\mathcal{A}_0)}(\mathcal{A}), \dots, Q_K^{(\mathcal{A}_0)}(\mathcal{A})).$$

In particular, for any event $\mathcal{A} \subseteq \mathcal{A}_0$, we have $Q_k^{(\mathcal{A}_0)}(\mathcal{A}) = 0$ for all $k = 1, \dots, K$ (because $Q_k^{(\mathcal{A}_0)}(\mathcal{A}) \leq Q_k^{(\mathcal{A}_0)}(\mathcal{A}_0) = 0$), and thus, we obtain further

$$Q^{(\mathcal{A}_0)}(\mathcal{A}) = h_{\mathcal{A}}(0, \dots, 0) \stackrel{(119)}{=} Q_0(\mathcal{A}).$$

Recalling (120), we conclude that the fused pdf $q^{(\mathcal{A}_0)}(\boldsymbol{\theta})$ satisfies

$$Q_0(\mathcal{A}) = \int_{\mathcal{A}} q^{(\mathcal{A}_0)}(\boldsymbol{\theta}) d\boldsymbol{\theta} \quad \text{for any event } \mathcal{A} \subseteq \mathcal{A}_0. \quad (121)$$

Following the same steps with \mathcal{A}_0 replaced by \mathcal{A}_0^c , we obtain a pdf $q^{(\mathcal{A}_0^c)}(\boldsymbol{\theta})$ such that

$$Q_0(\mathcal{A}) = \int_{\mathcal{A}} q^{(\mathcal{A}_0^c)}(\boldsymbol{\theta}) d\boldsymbol{\theta} \quad \text{for any event } \mathcal{A} \subseteq \mathcal{A}_0^c. \quad (122)$$

Now, for an arbitrary nontrivial event $\mathcal{B} \subseteq \Theta$, there exists an opinion profile $(q_1^{(\mathcal{B})}, \dots, q_K^{(\mathcal{B})})$ such that $Q_k^{(\mathcal{B})}(\mathcal{B}) = \int_{\mathcal{B}} q_k^{(\mathcal{B})}(\boldsymbol{\theta}) d\boldsymbol{\theta} = 0$ for all $k = 1, \dots, K$. Again, we denote the fused probability measure as $Q^{(\mathcal{B})}$. We, thus, obtain for

$Q_0(\mathcal{B})$ as defined by (119)

$$\begin{aligned} Q_0(\mathcal{B}) &= h_{\mathcal{B}}(0, \dots, 0) \\ &= h_{\mathcal{B}}(Q_1^{(\mathcal{B})}(\mathcal{B}), \dots, Q_K^{(\mathcal{B})}(\mathcal{B})) \\ &\stackrel{(16)}{=} Q^{(\mathcal{B})}(\mathcal{B}). \end{aligned} \quad (123)$$

Because \mathcal{B} can be decomposed into disjoint subsets according to $\mathcal{B} = (\mathcal{B} \cap \mathcal{A}_0) \cup (\mathcal{B} \cap \mathcal{A}_0^c)$, we further obtain from (123)

$$\begin{aligned} Q_0(\mathcal{B}) &= Q^{(\mathcal{B})}(\mathcal{B} \cap \mathcal{A}_0) + Q^{(\mathcal{B})}(\mathcal{B} \cap \mathcal{A}_0^c) \\ &\stackrel{(16)}{=} h_{\mathcal{B} \cap \mathcal{A}_0}(Q_1^{(\mathcal{B})}(\mathcal{B} \cap \mathcal{A}_0), \dots, Q_K^{(\mathcal{B})}(\mathcal{B} \cap \mathcal{A}_0)) \\ &\quad + h_{\mathcal{B} \cap \mathcal{A}_0^c}(Q_1^{(\mathcal{B})}(\mathcal{B} \cap \mathcal{A}_0^c), \dots, Q_K^{(\mathcal{B})}(\mathcal{B} \cap \mathcal{A}_0^c)) \\ &\stackrel{(a)}{=} h_{\mathcal{B} \cap \mathcal{A}_0}(0, \dots, 0) + h_{\mathcal{B} \cap \mathcal{A}_0^c}(0, \dots, 0) \\ &\stackrel{(119)}{=} Q_0(\mathcal{B} \cap \mathcal{A}_0) + Q_0(\mathcal{B} \cap \mathcal{A}_0^c) \end{aligned}$$

where we used $Q_k^{(\mathcal{B})}(\mathcal{B}) = 0$ in (a). Using (121) with $\mathcal{A} = \mathcal{B} \cap \mathcal{A}_0$ and (122) with $\mathcal{A} = \mathcal{B} \cap \mathcal{A}_0^c$, this implies that

$$\begin{aligned} Q_0(\mathcal{B}) &= \int_{\mathcal{B} \cap \mathcal{A}_0} q^{(\mathcal{A}_0)}(\boldsymbol{\theta}) d\boldsymbol{\theta} + \int_{\mathcal{B} \cap \mathcal{A}_0^c} q^{(\mathcal{A}_0^c)}(\boldsymbol{\theta}) d\boldsymbol{\theta} \\ &= \int_{\mathcal{B}} \tilde{q}_0(\boldsymbol{\theta}) d\boldsymbol{\theta} \end{aligned}$$

where we defined

$$\tilde{q}_0(\boldsymbol{\theta}) \triangleq \begin{cases} q^{(\mathcal{A}_0)}(\boldsymbol{\theta}), & \text{if } \boldsymbol{\theta} \in \mathcal{A}_0 \\ q^{(\mathcal{A}_0^c)}(\boldsymbol{\theta}), & \text{if } \boldsymbol{\theta} \in \mathcal{A}_0^c. \end{cases}$$

Thus, we found an integral representation for Q_0 and can define

$$Q_0(\mathcal{B}) = \int_{\mathcal{B}} \tilde{q}_0(\boldsymbol{\theta}) d\boldsymbol{\theta} \quad (124)$$

also for trivial events \mathcal{B} . The nonnegativity of $q^{(\mathcal{A}_0)}(\boldsymbol{\theta})$ and $q^{(\mathcal{A}_0^c)}(\boldsymbol{\theta})$ implies that $\tilde{q}_0(\boldsymbol{\theta})$ is nonnegative, and in turn, Q_0 is a measure. However, $\tilde{q}_0(\boldsymbol{\theta})$ is not a pdf in general.

We define

$$w_0 \triangleq Q_0(\Theta) \quad (125)$$

(note that this implies that $w_0 \geq 0$) and

$$q_0(\boldsymbol{\theta}) \triangleq \frac{\tilde{q}_0(\boldsymbol{\theta})}{w_0} \quad (126)$$

provided $w_0 \neq 0$. If $w_0 = 0$, we choose $q_0(\boldsymbol{\theta})$ as an arbitrary pdf. We claim that $w_0 \leq 1$. To prove this claim, let \mathcal{B}_n be a sequence of nontrivial events such that $\mathcal{B}_n \subseteq \mathcal{B}_{n+1}$ and $\lim_{n \rightarrow \infty} \mathcal{B}_n = \Theta$. For each \mathcal{B}_n , there exists an opinion profile $(q_1^{(\mathcal{B}_n)}, \dots, q_K^{(\mathcal{B}_n)})$ such that $Q_k^{(\mathcal{B}_n)}(\mathcal{B}_n) = \int_{\mathcal{B}_n} q_k^{(\mathcal{B}_n)}(\boldsymbol{\theta}) d\boldsymbol{\theta} = 0$ for all $k = 1, \dots, K$. Again, we denote the sequence of fused probability measures as $Q^{(\mathcal{B}_n)}$.

Following the steps in (123), we have that

$$Q_0(\mathcal{B}_n) = Q^{(\mathcal{B}_n)}(\mathcal{B}_n) \leq 1$$

because $Q^{(\mathcal{B}_n)}$ is a probability measure. The continuity from below of measures [197, Lem. 3.4] implies that $w_0 = Q_0(\Theta) = \lim_{n \rightarrow \infty} Q_0(\mathcal{B}_n) \leq 1$.

A similar argument can be employed to show (for later use) that, for any nontrivial event $\mathcal{A} \subseteq \Theta$ and arbitrary probabilities p_k ,

$$h_{\mathcal{A}}(p_1, \dots, p_K) \geq Q_0(\mathcal{A}). \quad (127)$$

Indeed, for any nontrivial event $\mathcal{A} \subseteq \Theta$, let $\mathcal{B}_n \subseteq \mathcal{A}$ be a sequence satisfying $|\mathcal{A} \setminus \mathcal{B}_n| > 0$, $\mathcal{B}_n \subseteq \mathcal{B}_{n+1}$ for all $n \in \mathbb{N}$, and $\lim_{n \rightarrow \infty} \mathcal{B}_n = \mathcal{A}$. Then, for each $n \in \mathbb{N}$, there exists an opinion profile $(q_1^{(\mathcal{B}_n)}, \dots, q_K^{(\mathcal{B}_n)})$ satisfying $Q_k^{(\mathcal{B}_n)}(\mathcal{B}_n) = 0$, $Q_k^{(\mathcal{B}_n)}(\mathcal{A} \setminus \mathcal{B}_n) = p_k$, and, in turn, $Q_k^{(\mathcal{B}_n)}(\mathcal{A}) = p_k$. Again, we denote the sequence of fused probability measures as $Q^{(\mathcal{B}_n)}$. Following the steps in (123), $Q_0(\mathcal{B}_n) = Q^{(\mathcal{B}_n)}(\mathcal{B}_n)$. Thus, we have

$$\begin{aligned} h_{\mathcal{A}}(p_1, \dots, p_K) &= h_{\mathcal{A}}\left(Q_1^{(\mathcal{B}_n)}(\mathcal{A}), \dots, Q_K^{(\mathcal{B}_n)}(\mathcal{A})\right) \\ &\stackrel{(16)}{=} Q^{(\mathcal{B}_n)}(\mathcal{A}) \\ &\geq Q^{(\mathcal{B}_n)}(\mathcal{B}_n) \\ &= Q_0(\mathcal{B}_n). \end{aligned} \quad (128)$$

Here, $h_{\mathcal{A}}(p_1, \dots, p_K)$ does not depend on n . Hence, we can take the limit on the right-hand side of (128) and obtain

$$h_{\mathcal{A}}(p_1, \dots, p_K) \geq \lim_{n \rightarrow \infty} Q_0(\mathcal{B}_n) = Q_0(\mathcal{A})$$

using again the continuity from below of Q_0 .

2) Step 2: Define $\tilde{h}_{\mathcal{A}}$ and Prove That It Defines a Linear Pooling Function: We define

$$\tilde{h}_{\mathcal{A}}(p_1, \dots, p_K) \triangleq \frac{h_{\mathcal{A}}(p_1, \dots, p_K) - Q_0(\mathcal{A})}{1 - w_0}. \quad (129)$$

Here, we have to assume that $w_0 < 1$. Thus, we first show that g is a generalized linear pooling function in the case $w_0 = 1$. In this case, for any nontrivial event $\mathcal{A} \subseteq \Theta$ and arbitrary probabilities p_k , we choose an opinion profile that satisfies $Q_k(\mathcal{A}) = p_k$ and, hence, $Q_k(\mathcal{A}^c) = 1 - p_k$ for all $k = 1, \dots, K$. We then have

$$\begin{aligned} 1 &= Q(\mathcal{A}) + Q(\mathcal{A}^c) \\ &\stackrel{(16)}{=} h_{\mathcal{A}}(p_1, \dots, p_K) + h_{\mathcal{A}^c}(1 - p_1, \dots, 1 - p_K) \\ &\stackrel{(127)}{\geq} Q_0(\mathcal{A}) + Q_0(\mathcal{A}^c) \\ &\stackrel{(125)}{=} w_0 \\ &= 1. \end{aligned}$$

Thus, the inequality in the third line is actually an equality, which is only possible if $h_{\mathcal{A}}(p_1, \dots, p_K) = Q_0(\mathcal{A})$. Because \mathcal{A} and p_k were chosen arbitrarily, we have $h_{\mathcal{A}}(p_1, \dots, p_K) = Q_0(\mathcal{A})$ independently of the probabilities p_k . By (16), this further implies, for any opinion profile (q_1, \dots, q_K) , that the aggregate pdf q satisfies

$$\begin{aligned} \int_{\mathcal{A}} q(\boldsymbol{\theta}) d\boldsymbol{\theta} &= Q(\mathcal{A}) \\ &\stackrel{(16)}{=} h_{\mathcal{A}}(Q_1(\mathcal{A}), \dots, Q_K(\mathcal{A})) \\ &= Q_0(\mathcal{A}) \\ &\stackrel{(124)}{=} \int_{\mathcal{A}} \tilde{q}_0(\boldsymbol{\theta}) d\boldsymbol{\theta} \end{aligned}$$

for all events \mathcal{A} . Hence, $q(\boldsymbol{\theta}) = \tilde{q}_0(\boldsymbol{\theta})$, which implies that g is a dogmatic pooling function (which is a special case of a generalized linear pooling function with weights $w_0 = 1$, $w_k = 0$ for $k = 1, \dots, K$). This concludes the proof for the special case $w_0 = 1$, and thus, we can assume $w_0 < 1$ in what follows.

We define a new fusion rule \tilde{g} by

$$\tilde{g}[q_1, \dots, q_K](\boldsymbol{\theta}) \triangleq \frac{g[q_1, \dots, q_K](\boldsymbol{\theta}) - w_0 q_0(\boldsymbol{\theta})}{1 - w_0} \quad (130)$$

and claim that it satisfies the WSFP with the functions $\tilde{h}_{\mathcal{A}}$ defined by (129). Indeed, we have, for any opinion profile (q_1, \dots, q_K) and any event $\mathcal{A} \subseteq \Theta$, that

$$\begin{aligned} \tilde{g}(\mathcal{A}) &= \int_{\mathcal{A}} \tilde{g}[q_1, \dots, q_K](\boldsymbol{\theta}) d\boldsymbol{\theta} \\ &= \int_{\mathcal{A}} \frac{g[q_1, \dots, q_K](\boldsymbol{\theta}) - w_0 q_0(\boldsymbol{\theta})}{1 - w_0} d\boldsymbol{\theta} \\ &\stackrel{(a)}{=} \frac{h_{\mathcal{A}}(Q_1(\mathcal{A}), \dots, Q_K(\mathcal{A})) - \int_{\mathcal{A}} \tilde{q}_0(\boldsymbol{\theta}) d\boldsymbol{\theta}}{1 - w_0} \\ &\stackrel{(124)}{=} \frac{h_{\mathcal{A}}(Q_1(\mathcal{A}), \dots, Q_K(\mathcal{A})) - Q_0(\mathcal{A})}{1 - w_0} \\ &\stackrel{(129)}{=} \tilde{h}_{\mathcal{A}}(Q_1(\mathcal{A}), \dots, Q_K(\mathcal{A})) \end{aligned} \quad (131)$$

where we used in (a) that, by (16), $h_{\mathcal{A}}(Q_1(\mathcal{A}), \dots, Q_K(\mathcal{A})) = Q(\mathcal{A}) = \int_{\mathcal{A}} g[q_1, \dots, q_K](\boldsymbol{\theta}) d\boldsymbol{\theta}$ and, by (126), $w_0 q_0(\boldsymbol{\theta}) = \tilde{q}_0(\boldsymbol{\theta})$. Furthermore, we claim that \tilde{g} satisfies the ZPP. To prove this, let (q_1, \dots, q_K) be an opinion profile and \mathcal{A} a nontrivial event such that $Q_k(\mathcal{A}) = 0$ for all $k = 1, \dots, K$. Because $h_{\mathcal{A}}(0, \dots, 0) = Q_0(\mathcal{A})$, we have

$$\begin{aligned} \tilde{g}(\mathcal{A}) &\stackrel{(131)}{=} \tilde{h}_{\mathcal{A}}(Q_1(\mathcal{A}), \dots, Q_K(\mathcal{A})) \\ &= \tilde{h}_{\mathcal{A}}(0, \dots, 0) \\ &\stackrel{(129)}{=} \frac{h_{\mathcal{A}}(0, \dots, 0) - Q_0(\mathcal{A})}{1 - w_0} \\ &= \frac{Q_0(\mathcal{A}) - Q_0(\mathcal{A})}{1 - w_0} \\ &= 0 \end{aligned}$$

proving the ZPP.

Finally, to see that \tilde{g} is a valid pooling function, we first show that, for any $(p_1, \dots, p_K) \in [0, 1]^K$, the function $\tilde{h}_{\mathcal{A}}(p_1, \dots, p_K)$ is nonnegative. This follows from (129), (127), and our assumption $w_0 < 1$. Hence, the measure \tilde{Q} is nonnegative, and thus, also the associated density $\tilde{g}[q_1, \dots, q_K](\theta)$ must be nonnegative. The fact that $\tilde{g}[q_1, \dots, q_K](\theta)$ integrates to one follows directly from the definition (130) and the fact that $g[q_1, \dots, q_K](\theta)$ and q_0 are pdfs.

Because \tilde{g} is a pooling function that satisfies the WSFP and the ZPP, Theorem 1 implies that it is a linear pooling function, i.e.,

$$\tilde{g}[q_1, \dots, q_K](\theta) = \sum_{k=1}^K w_k q_k(\theta) \quad (132)$$

with $(w_1, \dots, w_K) \in \mathcal{S}_K$.

3) *Step 3: Conclude That \tilde{g} Is a Generalized Linear Pooling Function:* Combining (130) and (132), we obtain

$$\sum_{k=1}^K w_k q_k(\theta) = \frac{g[q_1, \dots, q_K](\theta) - w_0 q_0(\theta)}{1 - w_0}$$

or, equivalently,

$$g[q_1, \dots, q_K](\theta) = w_0 q_0(\theta) + \sum_{k=1}^K (1 - w_0) w_k q_k(\theta).$$

From $\sum_{k=1}^K w_k = 1$, it follows that $w_0 + \sum_{k=1}^K (1 - w_0) w_k$ is one. Thus, g is a generalized linear pooling function.

APPENDIX C PROOF OF THE EQUIVALENCE STATEMENT IN THEOREM 6

We only show that 2, i.e., g satisfies individualized Bayesianity (A11) and $g[q_0, \dots, q_0](\theta) = q_0(\theta)$ for some pdf q_0 , implies 1, i.e., g is a multiplicative pooling function. The other direction is obvious.

Thus, let us assume that $g[q_0, \dots, q_0](\theta) = q_0(\theta)$ for some pdf q_0 . We have to show that, for any opinion profile (q_1, \dots, q_K) such that q_k/q_0 is bounded for all $k = 1, \dots, K$ (recall that we only consider those opinion profiles in the multiplicative pooling function), g is of the form (10), i.e.,

$$g[q_1, \dots, q_K](\theta) \propto (q_0(\theta))^{1-K} \prod_{k=1}^K q_k(\theta).$$

To this end, we first note that $q_k = q_0^{(\ell_k)}$ [see (17)] with $\ell_k = q_k/q_0$ for all $k = 1, \dots, K$. Thus,

$$g[q_1, \dots, q_K](\theta) = g[q_0^{(q_1/q_0)}, \dots, q_0^{(q_K/q_0)}](\theta).$$

By iteratively using individualized Bayesianity (19) with $\ell = q_k/q_0$ for each $k = 1, \dots, K$, we obtain further

$$\begin{aligned} g[q_1, \dots, q_K](\theta) &\propto g[q_0, q_0^{(q_2/q_0)}, \dots, q_0^{(q_K/q_0)}](\theta) \frac{q_1(\theta)}{q_0(\theta)} \\ &\propto g[q_0, \dots, q_0](\theta) \prod_{k=1}^K \frac{q_k(\theta)}{q_0(\theta)} \\ &= (q_0(\theta))^{1-K} \prod_{k=1}^K q_k(\theta) \end{aligned}$$

which is (10) and, thus, concludes the proof.

APPENDIX D PARTIAL PROOF OF THEOREM 8

A. 3 Implies 2

We first show that 3, i.e., the WSFP (A5) and independence preservation (A8), implies 2, i.e., the SSFP (A4) and independence preservation (A8). To this end, we show that independence preservation implies the ZPP (A2). The ZPP and the assumed WSFP, in turn, imply the SSFP by Theorem 1.

To show that independence preservation implies the ZPP, assume that, for some event \mathcal{A} , we have $Q_k(\mathcal{A}) = 0$ for all $k = 1, \dots, K$. This implies that

$$Q_k(\mathcal{A} \cap \mathcal{A}) = Q_k(\mathcal{A}) = 0 = Q_k(\mathcal{A})Q_k(\mathcal{A}).$$

Independence preservation now implies that also Q must satisfy $Q(\mathcal{A} \cap \mathcal{A}) = Q(\mathcal{A})Q(\mathcal{A})$, and thus, either $Q(\mathcal{A}) = 0$ or $Q(\mathcal{A}) = 1$. In the first case, the proof of the ZPP is finished. In the second case, i.e., $Q(\mathcal{A}) = 1$, there must exist a subset $\mathcal{B} \subseteq \mathcal{A}$ such that $Q(\mathcal{B}) = 1/2$. However, because $\mathcal{B} \subseteq \mathcal{A}$ and $Q_k(\mathcal{A}) = 0$, we have that also $Q_k(\mathcal{B}) = 0$, and thus, we again have that $Q_k(\mathcal{B} \cap \mathcal{B}) = 0 = Q_k(\mathcal{B})Q_k(\mathcal{B})$. This implies that $Q(\mathcal{B})$ is either 0 or 1, which is a contradiction to $Q(\mathcal{B}) = 1/2$. Thus, $Q(\mathcal{A}) = 0$ is the only valid conclusion, which proves that the ZPP is satisfied.

B. 5 Implies 4

We next show that 5, i.e., the WSFP (A5) and external Bayesianity (A10), implies 4, i.e., the SSFP (A4) and external Bayesianity (A10). Thus, we have to show that the WSFP and external Bayesianity imply the SSFP.

By Theorem 2, the WSFP implies the weak likelihood principle (A7). Furthermore, by Theorem 4, the weak likelihood principle and external Bayesianity imply that g is a generalized log-linear pooling function, i.e.,

$$g[q_1, \dots, q_K](\theta) = c \xi_0(\theta) \prod_{k=1}^K (q_k(\theta))^{w_k} \quad (133)$$

for all positive opinion profiles. Finally, by Theorem 2, the WSFP implies that g is also a generalized linear pooling

function, i.e.,

$$g[q_1, \dots, q_K](\boldsymbol{\theta}) = \sum_{k=0}^K w'_k q_k(\boldsymbol{\theta}) \quad (134)$$

for all opinion profiles. Thus, combining (133) and (134), we have

$$c \xi_0(\boldsymbol{\theta}) \prod_{k=1}^K (q_k(\boldsymbol{\theta}))^{w_k} = \sum_{k=0}^K w'_k q_k(\boldsymbol{\theta}) \quad (135)$$

for all positive opinion profiles. Note that $q_0(\boldsymbol{\theta})$ is not necessarily positive, i.e., it may be zero for certain values of $\boldsymbol{\theta}$.

We choose an arbitrary positive pdf $\tilde{q}_0(\boldsymbol{\theta})$ and $\varepsilon \in (0, 1)$ and consider the opinion profile $(\varepsilon \tilde{q}_0 + (1 - \varepsilon)q_0, \dots, \varepsilon \tilde{q}_0 + (1 - \varepsilon)q_0)$. Since $\tilde{q}_0(\boldsymbol{\theta})$ is positive, this is a positive opinion profile for any $\varepsilon \in (0, 1)$. Using it in (135) gives

$$\begin{aligned} & c \xi_0(\boldsymbol{\theta}) (\varepsilon \tilde{q}_0(\boldsymbol{\theta}) + (1 - \varepsilon)q_0(\boldsymbol{\theta})) \\ &= w'_0 q_0(\boldsymbol{\theta}) + (1 - w'_0) (\varepsilon \tilde{q}_0(\boldsymbol{\theta}) + (1 - \varepsilon)q_0(\boldsymbol{\theta})) \end{aligned}$$

where $\sum_{k=1}^K w_k = 1$ and $\sum_{k=0}^K w'_k = 1$ were used, or equivalently,

$$c \xi_0(\boldsymbol{\theta}) = w'_0 \frac{q_0(\boldsymbol{\theta})}{\varepsilon \tilde{q}_0(\boldsymbol{\theta}) + (1 - \varepsilon)q_0(\boldsymbol{\theta})} + 1 - w'_0. \quad (136)$$

Taking the limit $\varepsilon \rightarrow 0$ in (136), we obtain

$$c \xi_0(\boldsymbol{\theta}) = \begin{cases} 1, & \text{if } q_0(\boldsymbol{\theta}) > 0 \\ 1 - w'_0, & \text{if } q_0(\boldsymbol{\theta}) = 0. \end{cases} \quad (137)$$

Inserting into (133) and evaluating (133) for the opinion profile $(\tilde{q}_0, \dots, \tilde{q}_0)$ yield

$$g[\tilde{q}_0, \dots, \tilde{q}_0](\boldsymbol{\theta}) = \begin{cases} \tilde{q}_0(\boldsymbol{\theta}), & \text{if } q_0(\boldsymbol{\theta}) > 0 \\ (1 - w'_0) \tilde{q}_0(\boldsymbol{\theta}), & \text{if } q_0(\boldsymbol{\theta}) = 0. \end{cases}$$

Because $g[\tilde{q}_0, \dots, \tilde{q}_0](\boldsymbol{\theta})$ is a pdf, this implies that

$$\begin{aligned} 1 &= \int_{\Theta} g[\tilde{q}_0, \dots, \tilde{q}_0](\boldsymbol{\theta}) d\boldsymbol{\theta} \\ &= \int_{\{\boldsymbol{\theta} \in \Theta : q_0(\boldsymbol{\theta}) > 0\}} \tilde{q}_0(\boldsymbol{\theta}) d\boldsymbol{\theta} \\ &\quad + (1 - w'_0) \int_{\{\boldsymbol{\theta} \in \Theta : q_0(\boldsymbol{\theta}) = 0\}} \tilde{q}_0(\boldsymbol{\theta}) d\boldsymbol{\theta}. \end{aligned} \quad (138)$$

On the other hand, because $\tilde{q}_0(\boldsymbol{\theta})$ is a pdf, we have

$$1 = \int_{\Theta} \tilde{q}_0(\boldsymbol{\theta}) d\boldsymbol{\theta}$$

$$= \int_{\{\boldsymbol{\theta} \in \Theta : q_0(\boldsymbol{\theta}) > 0\}} \tilde{q}_0(\boldsymbol{\theta}) d\boldsymbol{\theta} + \int_{\{\boldsymbol{\theta} \in \Theta : q_0(\boldsymbol{\theta}) = 0\}} \tilde{q}_0(\boldsymbol{\theta}) d\boldsymbol{\theta}. \quad (139)$$

Combining (138) and (139), we obtain

$$(1 - w'_0) \int_{\{\boldsymbol{\theta} \in \Theta : q_0(\boldsymbol{\theta}) = 0\}} \tilde{q}_0(\boldsymbol{\theta}) d\boldsymbol{\theta} = \int_{\{\boldsymbol{\theta} \in \Theta : q_0(\boldsymbol{\theta}) = 0\}} \tilde{q}_0(\boldsymbol{\theta}) d\boldsymbol{\theta}$$

or, equivalently,

$$w'_0 \int_{\{\boldsymbol{\theta} \in \Theta : q_0(\boldsymbol{\theta}) = 0\}} \tilde{q}_0(\boldsymbol{\theta}) d\boldsymbol{\theta} = 0.$$

Since $\tilde{q}_0(\boldsymbol{\theta})$ is a positive pdf on Θ , this can only hold if either $w'_0 = 0$ or $|\{\boldsymbol{\theta} \in \Theta : q_0(\boldsymbol{\theta}) = 0\}| = 0$. In the first case, (134) implies that g is actually a linear pooling function, and thus, by Theorem 1, g satisfies the SSFP. In the second case, $q_0(\boldsymbol{\theta}) > 0$ almost everywhere, and thus, (137) states that $c \xi_0(\boldsymbol{\theta}) = 1$. Using the opinion profile $(q_1, \dots, q_K) = (\tilde{q}_0, \dots, \tilde{q}_0)$ in (135) gives

$$\tilde{q}_0(\boldsymbol{\theta}) = w'_0 q_0(\boldsymbol{\theta}) + (1 - w'_0) \tilde{q}_0(\boldsymbol{\theta}). \quad (140)$$

In particular, let us partition Θ into disjoint sets \mathcal{A}_1 and \mathcal{A}_2 satisfying $\int_{\mathcal{A}_1} q_0(\boldsymbol{\theta}) d\boldsymbol{\theta} = \int_{\mathcal{A}_2} q_0(\boldsymbol{\theta}) d\boldsymbol{\theta} = 1/2$, and let us choose

$$\tilde{q}_0(\boldsymbol{\theta}) = \begin{cases} \frac{3}{2} q_0(\boldsymbol{\theta}), & \text{if } \boldsymbol{\theta} \in \mathcal{A}_1 \\ \frac{1}{2} q_0(\boldsymbol{\theta}), & \text{if } \boldsymbol{\theta} \in \mathcal{A}_2. \end{cases}$$

Then, (140) yields for all $\boldsymbol{\theta} \in \mathcal{A}_1$

$$\frac{3}{2} q_0(\boldsymbol{\theta}) = w'_0 q_0(\boldsymbol{\theta}) + (1 - w'_0) \frac{3}{2} q_0(\boldsymbol{\theta}) = \left(\frac{3}{2} - \frac{1}{2} w'_0 \right) q_0(\boldsymbol{\theta}).$$

This implies $w'_0 = 0$, and hence, we again conclude from (134) that g is actually a linear pooling function, and thus, by Theorem 1, g satisfies the SSFP.

C. 6 Implies 1

Finally, we prove that 6, i.e., the SSFP (A4) and generalized Bayesianity (A12), implies 1, i.e., g is a dictatorship pooling function. By Theorem 1, the SSFP implies that g is a linear pooling function, i.e.,

$$g[q_1, \dots, q_K](\boldsymbol{\theta}) = \sum_{k=1}^K w_k q_k(\boldsymbol{\theta}) \quad (141)$$

with $(w_1, \dots, w_K) \in \mathcal{S}_K$. We will show that, for an arbitrary k , the weight w_k is either 0 or 1, which is equivalent to g being a dictatorship pooling function.

We first choose a positive function f and two disjoint sets \mathcal{A} and \mathcal{B} such that $\Theta = \mathcal{A} \cup \mathcal{B}$ and $\int_{\mathcal{A}} f(\boldsymbol{\theta}) d\boldsymbol{\theta} = \int_{\mathcal{B}} f(\boldsymbol{\theta}) d\boldsymbol{\theta} = 1$. We fix an arbitrary k and define an opinion

profile (q_1, \dots, q_K) by setting

$$q_k(\boldsymbol{\theta}) = \begin{cases} \frac{1}{3}f(\boldsymbol{\theta}), & \text{if } \boldsymbol{\theta} \in \mathcal{A} \\ \frac{2}{3}f(\boldsymbol{\theta}), & \text{if } \boldsymbol{\theta} \in \mathcal{B} \end{cases}$$

and $q_{k'} = q_0$ for all $k' \neq k$, where

$$q_0(\boldsymbol{\theta}) = \frac{1}{2}f(\boldsymbol{\theta}). \quad (142)$$

Inserting this opinion profile into the fusion rule (141) and using $\sum_{k' \neq k} w_{k'} = 1 - w_k$ gives

$$\begin{aligned} g[q_1, \dots, q_K](\boldsymbol{\theta}) &= \begin{cases} \left(\frac{1}{3}w_k + \frac{1}{2}(1-w_k)\right)f(\boldsymbol{\theta}), & \text{if } \boldsymbol{\theta} \in \mathcal{A} \\ \left(\frac{2}{3}w_k + \frac{1}{2}(1-w_k)\right)f(\boldsymbol{\theta}), & \text{if } \boldsymbol{\theta} \in \mathcal{B} \end{cases} \\ &= \begin{cases} \left(\frac{1}{2} - \frac{1}{6}w_k\right)f(\boldsymbol{\theta}), & \text{if } \boldsymbol{\theta} \in \mathcal{A} \\ \left(\frac{1}{2} + \frac{1}{6}w_k\right)f(\boldsymbol{\theta}), & \text{if } \boldsymbol{\theta} \in \mathcal{B}. \end{cases} \quad (143) \end{aligned}$$

Next, we use generalized Bayesianity with $\ell_{k'} = \ell$ for all $k' = 1, \dots, K$, where

$$\ell(\boldsymbol{\theta}) = \begin{cases} 1, & \text{if } \boldsymbol{\theta} \in \mathcal{A} \\ 2, & \text{if } \boldsymbol{\theta} \in \mathcal{B}. \end{cases}$$

We easily obtain [see (17)]

$$q_k^{(\ell)}(\boldsymbol{\theta}) = \begin{cases} \frac{1}{5}f(\boldsymbol{\theta}), & \text{if } \boldsymbol{\theta} \in \mathcal{A} \\ \frac{4}{5}f(\boldsymbol{\theta}), & \text{if } \boldsymbol{\theta} \in \mathcal{B} \end{cases}$$

and

$$q_0^{(\ell)}(\boldsymbol{\theta}) = \begin{cases} \frac{1}{3}f(\boldsymbol{\theta}), & \text{if } \boldsymbol{\theta} \in \mathcal{A} \\ \frac{2}{3}f(\boldsymbol{\theta}), & \text{if } \boldsymbol{\theta} \in \mathcal{B}. \end{cases} \quad (144)$$

Now, (141) gives

$$\begin{aligned} g[q_1^{(\ell)}, \dots, q_K^{(\ell)}](\boldsymbol{\theta}) &= \begin{cases} \left(\frac{1}{5}w_k + \frac{1}{3}(1-w_k)\right)f(\boldsymbol{\theta}), & \text{if } \boldsymbol{\theta} \in \mathcal{A} \\ \left(\frac{4}{5}w_k + \frac{2}{3}(1-w_k)\right)f(\boldsymbol{\theta}), & \text{if } \boldsymbol{\theta} \in \mathcal{B} \end{cases} \\ &= \begin{cases} \left(\frac{1}{3} - \frac{2}{15}w_k\right)f(\boldsymbol{\theta}), & \text{if } \boldsymbol{\theta} \in \mathcal{A} \\ \left(\frac{2}{3} + \frac{2}{15}w_k\right)f(\boldsymbol{\theta}), & \text{if } \boldsymbol{\theta} \in \mathcal{B}. \end{cases} \quad (145) \end{aligned}$$

On the other hand, because g satisfies generalized Bayesianity, there exists a function $h[\ell, \dots, \ell]$ such that

$$g[q_1^{(\ell)}, \dots, q_K^{(\ell)}](\boldsymbol{\theta}) = \frac{g[q_1, \dots, q_K](\boldsymbol{\theta})h[\ell, \dots, \ell](\boldsymbol{\theta})}{c_\ell} \quad (146)$$

where $c_\ell = \int_{\Theta} g[q_1, \dots, q_K](\boldsymbol{\theta})h[\ell, \dots, \ell](\boldsymbol{\theta}) d\boldsymbol{\theta}$. Inserting (143) and (145) into (146) gives

$$\frac{1}{3} - \frac{2}{15}w_k = \frac{\left(\frac{1}{2} - \frac{1}{6}w_k\right)h[\ell, \dots, \ell](\boldsymbol{\theta})}{c_\ell} \quad (147)$$

for all $\boldsymbol{\theta} \in \mathcal{A}$ and

$$\frac{2}{3} + \frac{2}{15}w_k = \frac{\left(\frac{1}{2} + \frac{1}{6}w_k\right)h[\ell, \dots, \ell](\boldsymbol{\theta})}{c_\ell} \quad (148)$$

for all $\boldsymbol{\theta} \in \mathcal{B}$.

Using again the generalized Bayesianity of g , we also have

$$g[q_0^{(\ell)}, \dots, q_0^{(\ell)}](\boldsymbol{\theta}) = \frac{g[q_0, \dots, q_0](\boldsymbol{\theta})h[\ell, \dots, \ell](\boldsymbol{\theta})}{c_{0,\ell}} \quad (149)$$

where $c_{0,\ell} = \int_{\Theta} g[q_0, \dots, q_0](\boldsymbol{\theta})h[\ell, \dots, \ell](\boldsymbol{\theta}) d\boldsymbol{\theta}$. Because linear pooling functions are unanimity preserving (see Theorem 1), we have $g[q_0^{(\ell)}, \dots, q_0^{(\ell)}](\boldsymbol{\theta}) = q_0^{(\ell)}(\boldsymbol{\theta})$ and $g[q_0, \dots, q_0](\boldsymbol{\theta}) = q_0(\boldsymbol{\theta})$, and thus, (149) is equivalent to

$$q_0^{(\ell)}(\boldsymbol{\theta}) = \frac{q_0(\boldsymbol{\theta})h[\ell, \dots, \ell](\boldsymbol{\theta})}{c_{0,\ell}}$$

or, inserting (144) and (142),

$$\frac{1}{3} = \frac{\frac{1}{2}h[\ell, \dots, \ell](\boldsymbol{\theta})}{c_{0,\ell}}$$

for all $\boldsymbol{\theta} \in \mathcal{A}$ and

$$\frac{2}{3} = \frac{\frac{1}{2}h[\ell, \dots, \ell](\boldsymbol{\theta})}{c_{0,\ell}}$$

for all $\boldsymbol{\theta} \in \mathcal{B}$. Thus, we obtain

$$h[\ell, \dots, \ell](\boldsymbol{\theta}) = \begin{cases} \frac{2}{3}c_{0,\ell}, & \text{if } \boldsymbol{\theta} \in \mathcal{A} \\ \frac{4}{3}c_{0,\ell}, & \text{if } \boldsymbol{\theta} \in \mathcal{B}. \end{cases}$$

Inserting this into (147) and (148) yields

$$\frac{1}{3} - \frac{2}{15}w_k = \frac{\left(\frac{1}{2} - \frac{1}{6}w_k\right)\frac{2}{3}c_{0,\ell}}{c_\ell}$$

and

$$\frac{2}{3} + \frac{2}{15}w_k = \frac{\left(\frac{1}{2} + \frac{1}{6}w_k\right)\frac{4}{3}c_{0,\ell}}{c_\ell}$$

or, equivalently,

$$\frac{\frac{1}{3} - \frac{2}{15}w_k}{\frac{1}{3} - \frac{1}{9}w_k} = \frac{c_{0,\ell}}{c_\ell} = \frac{\frac{2}{3} + \frac{2}{15}w_k}{\frac{2}{3} + \frac{2}{9}w_k}.$$

This amounts to the quadratic equation $w_k^2 - w_k = 0$, which has the solutions $w_k = 0$ and $w_k = 1$. Since k was arbitrary, this concludes the proof.

APPENDIX E PROOF OF LEMMA 9

By Theorem 2, the WSFP implies that g is a generalized linear pooling function, i.e.,

$$g[q_1, \dots, q_K](\boldsymbol{\theta}) = \sum_{k=0}^K w_k q_k(\boldsymbol{\theta}) \quad (150)$$

with $(w_0, \dots, w_K) \in \mathcal{S}_{K+1}$. We will show that w_0 is either 0 or 1, which is equivalent to g being either a linear pooling function or a dogmatic pooling function.

We first choose two disjoint sets \mathcal{A} and \mathcal{B} such that $\Theta = \mathcal{A} \cup \mathcal{B}$ and $\int_{\mathcal{A}} q_0(\boldsymbol{\theta}) d\boldsymbol{\theta} = \int_{\mathcal{B}} q_0(\boldsymbol{\theta}) d\boldsymbol{\theta} = 1/2$. Furthermore, we choose an opinion profile (q_1, \dots, q_K) as

$$q_k(\boldsymbol{\theta}) = \begin{cases} \frac{2}{3}q_0(\boldsymbol{\theta}), & \text{if } \boldsymbol{\theta} \in \mathcal{A} \\ \frac{4}{3}q_0(\boldsymbol{\theta}), & \text{if } \boldsymbol{\theta} \in \mathcal{B} \end{cases}$$

for all $k = 1, \dots, K$. Inserting this opinion profile into (150) and using $\sum_{k=1}^K w_k = 1 - w_0$ gives

$$\begin{aligned} g[q_1, \dots, q_K](\boldsymbol{\theta}) &= \begin{cases} (w_0 + \frac{2}{3}(1 - w_0))q_0(\boldsymbol{\theta}), & \text{if } \boldsymbol{\theta} \in \mathcal{A} \\ (w_0 + \frac{4}{3}(1 - w_0))q_0(\boldsymbol{\theta}), & \text{if } \boldsymbol{\theta} \in \mathcal{B} \end{cases} \\ &= \begin{cases} (\frac{2}{3} + \frac{1}{3}w_0)q_0(\boldsymbol{\theta}), & \text{if } \boldsymbol{\theta} \in \mathcal{A} \\ (\frac{4}{3} - \frac{1}{3}w_0)q_0(\boldsymbol{\theta}), & \text{if } \boldsymbol{\theta} \in \mathcal{B}. \end{cases} \end{aligned} \quad (151)$$

Next, we use generalized Bayesianity with $\ell_k = \ell$ for all $k = 1, \dots, K$, where

$$\ell(\boldsymbol{\theta}) = \begin{cases} 1, & \text{if } \boldsymbol{\theta} \in \mathcal{A} \\ 2, & \text{if } \boldsymbol{\theta} \in \mathcal{B}. \end{cases} \quad (152)$$

We easily obtain [see (17)]

$$q_k^{(\ell)}(\boldsymbol{\theta}) = \begin{cases} \frac{2}{5}q_0(\boldsymbol{\theta}), & \text{if } \boldsymbol{\theta} \in \mathcal{A} \\ \frac{8}{5}q_0(\boldsymbol{\theta}), & \text{if } \boldsymbol{\theta} \in \mathcal{B} \end{cases}$$

and, then, (150) gives

$$\begin{aligned} g[q_1^{(\ell)}, \dots, q_K^{(\ell)}](\boldsymbol{\theta}) &= \begin{cases} (w_0 + \frac{2}{5}(1 - w_0))q_0(\boldsymbol{\theta}), & \text{if } \boldsymbol{\theta} \in \mathcal{A} \\ (w_0 + \frac{8}{5}(1 - w_0))q_0(\boldsymbol{\theta}), & \text{if } \boldsymbol{\theta} \in \mathcal{B} \end{cases} \\ &= \begin{cases} (\frac{2}{5} + \frac{3}{5}w_0)q_0(\boldsymbol{\theta}), & \text{if } \boldsymbol{\theta} \in \mathcal{A} \\ (\frac{8}{5} - \frac{3}{5}w_0)q_0(\boldsymbol{\theta}), & \text{if } \boldsymbol{\theta} \in \mathcal{B}. \end{cases} \end{aligned} \quad (153)$$

Because g satisfies generalized Bayesianity, we have that there exists a function $h[\ell, \dots, \ell]$ such that

$$g[q_1^{(\ell)}, \dots, q_K^{(\ell)}](\boldsymbol{\theta}) = \frac{g[q_1, \dots, q_K](\boldsymbol{\theta})h[\ell, \dots, \ell](\boldsymbol{\theta})}{c_{\ell}} \quad (154)$$

where $c_{\ell} = \int_{\Theta} g[q_1, \dots, q_K](\boldsymbol{\theta})h[\ell, \dots, \ell](\boldsymbol{\theta})d\boldsymbol{\theta}$. Inserting (151) and (153) into (154) gives

$$\frac{2}{5} + \frac{3}{5}w_0 = \frac{(\frac{2}{3} + \frac{1}{3}w_0)h[\ell, \dots, \ell](\boldsymbol{\theta})}{c_{\ell}} \quad (155)$$

for all $\boldsymbol{\theta} \in \mathcal{A}$ and

$$\frac{8}{5} - \frac{3}{5}w_0 = \frac{(\frac{4}{3} - \frac{1}{3}w_0)h[\ell, \dots, \ell](\boldsymbol{\theta})}{c_{\ell}} \quad (156)$$

for all $\boldsymbol{\theta} \in \mathcal{B}$.

Again using the generalized Bayesianity of g , we also have

$$g[q_0^{(\ell)}, \dots, q_0^{(\ell)}](\boldsymbol{\theta}) = \frac{g[q_0, \dots, q_0](\boldsymbol{\theta})h[\ell, \dots, \ell](\boldsymbol{\theta})}{c_{0,\ell}} \quad (157)$$

where $c_{0,\ell} = \int_{\Theta} g[q_0, \dots, q_0](\boldsymbol{\theta})h[\ell, \dots, \ell](\boldsymbol{\theta})d\boldsymbol{\theta}$. Using (152) and (17), we obtain

$$q_0^{(\ell)}(\boldsymbol{\theta}) = \begin{cases} \frac{2}{3}q_0(\boldsymbol{\theta}), & \text{if } \boldsymbol{\theta} \in \mathcal{A} \\ \frac{4}{3}q_0(\boldsymbol{\theta}), & \text{if } \boldsymbol{\theta} \in \mathcal{B}. \end{cases}$$

Inserting into (150) yields

$$\begin{aligned} g[q_0^{(\ell)}, \dots, q_0^{(\ell)}](\boldsymbol{\theta}) &= \begin{cases} (w_0 + \frac{2}{3}(1 - w_0))q_0(\boldsymbol{\theta}), & \text{if } \boldsymbol{\theta} \in \mathcal{A} \\ (w_0 + \frac{4}{3}(1 - w_0))q_0(\boldsymbol{\theta}), & \text{if } \boldsymbol{\theta} \in \mathcal{B} \end{cases} \\ &= \begin{cases} (\frac{2}{3} + \frac{1}{3}w_0)q_0(\boldsymbol{\theta}), & \text{if } \boldsymbol{\theta} \in \mathcal{A} \\ (\frac{4}{3} - \frac{1}{3}w_0)q_0(\boldsymbol{\theta}), & \text{if } \boldsymbol{\theta} \in \mathcal{B}. \end{cases} \end{aligned} \quad (158)$$

Furthermore, again by (150), $g[q_0, \dots, q_0](\boldsymbol{\theta}) = q_0(\boldsymbol{\theta})$. Inserting this and (158) into (157), we obtain

$$\frac{2}{3} + \frac{1}{3}w_0 = \frac{h[\ell, \dots, \ell](\boldsymbol{\theta})}{c_{0,\ell}}$$

for all $\boldsymbol{\theta} \in \mathcal{A}$ and

$$\frac{4}{3} - \frac{1}{3}w_0 = \frac{h[\ell, \dots, \ell](\boldsymbol{\theta})}{c_{0,\ell}}$$

for all $\boldsymbol{\theta} \in \mathcal{B}$. Thus,

$$h[\ell, \dots, \ell](\boldsymbol{\theta}) = \begin{cases} (\frac{2}{3} + \frac{1}{3}w_0)c_{0,\ell}, & \text{if } \boldsymbol{\theta} \in \mathcal{A} \\ (\frac{4}{3} - \frac{1}{3}w_0)c_{0,\ell}, & \text{if } \boldsymbol{\theta} \in \mathcal{B}. \end{cases}$$

Inserting this into (155) and (156) gives

$$\frac{2}{5} + \frac{3}{5}w_0 = \frac{(\frac{2}{3} + \frac{1}{3}w_0)^2 c_{0,\ell}}{c_{\ell}}$$

and

$$\frac{8}{5} - \frac{3}{5}w_0 = \frac{\left(\frac{4}{3} - \frac{1}{3}w_0\right)^2 c_{0,\ell}}{c_\ell}$$

or, equivalently,

$$\frac{\frac{2}{5} + \frac{3}{5}w_0}{\left(\frac{2}{3} + \frac{1}{3}w_0\right)^2} = \frac{c_{0,\ell}}{c_\ell} = \frac{\frac{8}{5} - \frac{3}{5}w_0}{\left(\frac{4}{3} - \frac{1}{3}w_0\right)^2}.$$

This amounts to the cubic equation $w_0^3 - 3w_0^2 + 2w_0 = 0$, which has the solutions $w_0 = 0$, $w_0 = 1$, and $w_0 = 2$. Since w_0 cannot be larger than one, only the solutions $w_0 = 0$ and $w_0 = 1$ remain. In the first case, g is a linear pooling function, which satisfies the SSFP by Theorem 1. Hence, since g satisfies both the SSFP and generalized Bayesianity, it reduces to a dictatorship pooling function by Theorem 8. In the second case, g is a dogmatic pooling function.

APPENDIX F PROOF OF THEOREM 11

The implications in 1 follow from Theorem 1 because the SSFP implies that g is a linear pooling function and, in turn, satisfies the ZPP (A2), unanimity preservation (A3), the WSFP (A5), the likelihood principle (A6), and the weak likelihood principle (A7). Similarly, the implications in 2 follow from Theorem 2. Implication 3 follows directly from the concerned axioms. Implication 4 is shown in the first part of the proof of Theorem 8 in Appendix D. It remains to show implication 5, i.e., individualized Bayesianity implies generalized Bayesianity. This can easily be seen by defining

$$h[\ell_1, \dots, \ell_K](\boldsymbol{\theta}) \triangleq \prod_{k=1}^K \ell_k(\boldsymbol{\theta}). \quad (159)$$

Indeed, because g satisfies individualized Bayesianity, iterative application of (19) implies that

$$\begin{aligned} g[q_1^{(\ell_1)}, \dots, q_K^{(\ell_K)}](\boldsymbol{\theta}) &\propto g[q_1, q_2^{(\ell_2)}, \dots, q_K^{(\ell_K)}](\boldsymbol{\theta}) \ell_1(\boldsymbol{\theta}) \\ &\propto g[q_1, \dots, q_K](\boldsymbol{\theta}) \prod_{k=1}^K \ell_k(\boldsymbol{\theta}) \\ &\propto g[q_1, \dots, q_K](\boldsymbol{\theta}) \prod_{k=1}^K \ell_k(\boldsymbol{\theta}). \end{aligned}$$

Thus, (20) is satisfied by h defined in (159).

APPENDIX G PROOF OF THEOREM 14 (CONSTRAINED MINIMIZATION OF THE WEIGHTED AVERAGE OF α -DIVERGENCES)

Let $f_\alpha(x) = (x^\alpha - 1)/(\alpha(\alpha - 1))$. The inverse function is given by

$$f_\alpha^{-1}(x) = (x\alpha(\alpha - 1) + 1)^{1/\alpha}. \quad (160)$$

Furthermore, we have that, for two functions $p_1(\boldsymbol{\theta})$ and $p_2(\boldsymbol{\theta})$,

$$\begin{aligned} f_\alpha\left(\frac{p_1(\boldsymbol{\theta})}{p_2(\boldsymbol{\theta})}\right) &= \frac{\left(\frac{p_1(\boldsymbol{\theta})}{p_2(\boldsymbol{\theta})}\right)^\alpha - 1}{\alpha(\alpha - 1)} \\ &= \frac{(p_1(\boldsymbol{\theta}))^\alpha - (p_2(\boldsymbol{\theta}))^\alpha}{(p_2(\boldsymbol{\theta}))^\alpha \alpha(\alpha - 1)} \\ &= \frac{(p_1(\boldsymbol{\theta}))^\alpha - 1}{(p_2(\boldsymbol{\theta}))^\alpha \alpha(\alpha - 1)} - \frac{(p_2(\boldsymbol{\theta}))^\alpha - 1}{(p_2(\boldsymbol{\theta}))^\alpha \alpha(\alpha - 1)} \\ &= \frac{f_\alpha(p_1(\boldsymbol{\theta})) - f_\alpha(p_2(\boldsymbol{\theta}))}{(p_2(\boldsymbol{\theta}))^\alpha}. \end{aligned} \quad (161)$$

Therefore, the objective function in (22) for $f(x) = f_\alpha(x)$ can be written as

$$\begin{aligned} &\sum_{k=1}^K w_k \mathcal{D}_\alpha(q_k \parallel \varphi) \\ &= \sum_{k=1}^K w_k \int_{\Theta} \varphi(\boldsymbol{\theta}) f_\alpha\left(\frac{q_k(\boldsymbol{\theta})}{\varphi(\boldsymbol{\theta})}\right) d\boldsymbol{\theta} \\ &\stackrel{(161)}{=} \sum_{k=1}^K w_k \int_{\Theta} \varphi(\boldsymbol{\theta}) \frac{f_\alpha(q_k(\boldsymbol{\theta})) - f_\alpha(\varphi(\boldsymbol{\theta}))}{(\varphi(\boldsymbol{\theta}))^\alpha} d\boldsymbol{\theta}. \end{aligned}$$

Interchanging the summation and the integral gives

$$\begin{aligned} &\sum_{k=1}^K w_k \mathcal{D}_\alpha(q_k \parallel \varphi) \\ &= \int_{\Theta} \varphi(\boldsymbol{\theta}) \sum_{k=1}^K w_k \frac{f_\alpha(q_k(\boldsymbol{\theta})) - f_\alpha(\varphi(\boldsymbol{\theta}))}{(\varphi(\boldsymbol{\theta}))^\alpha} d\boldsymbol{\theta} \\ &\stackrel{(a)}{=} \int_{\Theta} \varphi(\boldsymbol{\theta}) \frac{\left(\sum_{k=1}^K w_k f_\alpha(q_k(\boldsymbol{\theta}))\right) - f_\alpha(\varphi(\boldsymbol{\theta}))}{(\varphi(\boldsymbol{\theta}))^\alpha} d\boldsymbol{\theta} \\ &= \int_{\Theta} \varphi(\boldsymbol{\theta}) \frac{f_\alpha^{-1}\left(\sum_{k=1}^K w_k f_\alpha(q_k(\boldsymbol{\theta}))\right) - f_\alpha(\varphi(\boldsymbol{\theta}))}{(\varphi(\boldsymbol{\theta}))^\alpha} d\boldsymbol{\theta} \\ &\stackrel{(161)}{=} \int_{\Theta} \varphi(\boldsymbol{\theta}) f_\alpha\left(\frac{f_\alpha^{-1}\left(\sum_{k=1}^K w_k f_\alpha(q_k(\boldsymbol{\theta}))\right)}{\varphi(\boldsymbol{\theta})}\right) d\boldsymbol{\theta} \end{aligned}$$

where we used in (a) that $\sum_{k=1}^K w_k = 1$. Since φ is a pdf and $f_\alpha(x) = (x^\alpha - 1)/(\alpha(\alpha - 1))$ is a convex function for $\alpha \in \mathbb{R} \setminus \{0, 1\}$, we can apply Jensen's inequality⁹ to obtain the following lower bound on the objective function:

$$\begin{aligned} &\int_{\Theta} \varphi(\boldsymbol{\theta}) f_\alpha\left(\frac{f_\alpha^{-1}\left(\sum_{k=1}^K w_k f_\alpha(q_k(\boldsymbol{\theta}))\right)}{\varphi(\boldsymbol{\theta})}\right) d\boldsymbol{\theta} \\ &\geq f_\alpha\left(\int_{\Theta} \varphi(\boldsymbol{\theta}) \frac{f_\alpha^{-1}\left(\sum_{k=1}^K w_k f_\alpha(q_k(\boldsymbol{\theta}))\right)}{\varphi(\boldsymbol{\theta})} d\boldsymbol{\theta}\right) \end{aligned}$$

⁹Jensen's inequality [198, Th. 3.3] asserts that, for a pdf $\varphi(\cdot)$, a measurable function $\zeta(\cdot)$, and a convex function $\psi(\cdot)$, we have that $\int \psi(\zeta(\boldsymbol{\theta})) \varphi(\boldsymbol{\theta}) d\boldsymbol{\theta} \geq \psi\left(\int \zeta(\boldsymbol{\theta}) \varphi(\boldsymbol{\theta}) d\boldsymbol{\theta}\right)$ with equality if and only if the function ζ is constant almost everywhere.

$$= f_\alpha \left(\int_{\Theta} f_\alpha^{-1} \left(\sum_{k=1}^K w_k f_\alpha(q_k(\theta)) \right) d\theta \right) \quad (162)$$

with equality if and only if the function

$$\zeta(\theta) \triangleq \frac{f_\alpha^{-1} \left(\sum_{k=1}^K w_k f_\alpha(q_k(\theta)) \right)}{\varphi(\theta)}$$

is constant almost everywhere. Note that this is equivalent to $\varphi(\theta) \propto f_\alpha^{-1}(\sum_{k=1}^K w_k f_\alpha(q_k(\theta)))$. Since the right-hand side of (162) is independent of φ , it is a lower bound for any choice of φ , and hence, the function $\varphi(\theta)$ minimizing the objective function (which is the desired solution $q(\theta)$ in (22)) is the one for which this lower bound is achieved with equality, i.e.,

$$\begin{aligned} q(\theta) &\propto f_\alpha^{-1} \left(\sum_{k=1}^K w_k f_\alpha(q_k(\theta)) \right) \\ &\stackrel{(160)}{=} \left(\left(\sum_{k=1}^K w_k f_\alpha(q_k(\theta)) \right) \alpha(\alpha-1) + 1 \right)^{1/\alpha} \\ &= \left(\sum_{k=1}^K w_k (q_k(\theta))^\alpha - \sum_{k=1}^K w_k + 1 \right)^{1/\alpha} \\ &= \left(\sum_{k=1}^K w_k (q_k(\theta))^\alpha \right)^{1/\alpha}. \end{aligned}$$

We conclude that the solution to (22) when $f(x) = f_\alpha(x)$ is $q(\theta) = c \left(\sum_{k=1}^K w_k (q_k(\theta))^\alpha \right)^{1/\alpha}$, where $c = 1 / \int_{\Theta} \left(\sum_{k=1}^K w_k (q_k(\theta))^\alpha \right)^{1/\alpha} d\theta$.

APPENDIX H CHARACTERIZATION OF THE REVERSE α -DIVERGENCE

We will show that $\mathcal{D}_\alpha(\varphi \| q_k) = \mathcal{D}_{\alpha^*}(q_k \| \varphi)$, where $\alpha^* = 1 - \alpha$. To this end, we will use (27) with $f(x) = f_\alpha(x) = (x^\alpha - 1) / (\alpha(\alpha - 1))$. Using $f^*(x) = xf(1/x)$, we have

$$\begin{aligned} f_\alpha^*(x) &= x \frac{x^{-\alpha} - 1}{\alpha(\alpha - 1)} \\ &= \frac{x^{-\alpha+1} - x}{\alpha(\alpha - 1)} \\ &= \frac{x^{-(\alpha-1)} - 1}{\alpha(\alpha - 1)} - \frac{1}{\alpha(\alpha - 1)}(x - 1) \\ &= f_{\alpha^*}(x) - \frac{1}{\alpha(\alpha - 1)}(x - 1). \end{aligned}$$

Thus, up to the additive term $-(1/(\alpha(\alpha - 1)))(x - 1)$, the function $f_\alpha^*(x)$ is equal to $f_{\alpha^*}(x)$. Now, by [164, Proposition 1], an f -divergence does not change if $f(x)$ is replaced by $f(x) + c(x - 1)$ for an arbitrary $c \in \mathbb{R}$. Hence, f_α^* and f_{α^*} result in the same f -divergence, and (27) together with

(24) implies that

$$\begin{aligned} \mathcal{D}_\alpha(\varphi \| q_k) &= \mathcal{D}_{f_\alpha}(\varphi \| q_k) \\ &= \mathcal{D}_{f_\alpha^*}(q_k \| \varphi) \\ &= \mathcal{D}_{f_{\alpha^*}}(q_k \| \varphi) \\ &= \mathcal{D}_{\alpha^*}(q_k \| \varphi). \end{aligned}$$

APPENDIX I PROOF OF THEOREM 16 (CONSTRAINED MINIMIZATION OF THE WEIGHTED AVERAGE OF SQUARED L_2 DISTANCES)

We want to find

$$q = \arg \min_{\varphi \in \mathcal{P}} \sum_{k=1}^K w_k \|q_k - \varphi\|_2^2. \quad (163)$$

To this end, we note that

$$\begin{aligned} &\min_{\varphi \in \mathcal{P}} \sum_{k=1}^K w_k \|q_k - \varphi\|_2^2 \\ &= \min_{\varphi \in \mathcal{P}} \int_{\Theta} \sum_{k=1}^K w_k (q_k(\theta) - \varphi(\theta))^2 d\theta \\ &\geq \int_{\Theta} \min_{\varphi(\theta) \geq 0} \left\{ \sum_{k=1}^K w_k (q_k(\theta) - \varphi(\theta))^2 \right\} d\theta. \end{aligned} \quad (164)$$

For each fixed θ , the function value $\varphi(\theta)$ that achieves the minimum $\min_{\varphi(\theta) \geq 0} \sum_{k=1}^K w_k (q_k(\theta) - \varphi(\theta))^2$ is easily seen to be

$$\varphi^*(\theta) = \sum_{k=1}^K w_k q_k(\theta).$$

Because $\varphi^* \in \mathcal{P}$ (due to $(w_1, \dots, w_K) \in \mathcal{S}_K$), we have that

$$\begin{aligned} \sum_{k=1}^K w_k \|q_k - \varphi^*\|_2^2 &\geq \min_{\varphi \in \mathcal{P}} \sum_{k=1}^K w_k \|q_k - \varphi\|_2^2 \\ &\stackrel{(164)}{\geq} \int_{\Theta} \sum_{k=1}^K w_k (q_k(\theta) - \varphi^*(\theta))^2 d\theta \\ &= \sum_{k=1}^K w_k \|q_k - \varphi^*\|_2^2. \end{aligned} \quad (165)$$

Thus, all inequalities in (165) are actually equalities. In particular,

$$\sum_{k=1}^K w_k \|q_k - \varphi^*\|_2^2 = \min_{\varphi \in \mathcal{P}} \sum_{k=1}^K w_k \|q_k - \varphi\|_2^2$$

i.e., $q = \varphi^*$ solves (163).

APPENDIX J
PROOF OF THEOREM 17
(UNCONSTRAINED MINIMIZATION
OF THE WEIGHTED AVERAGE
OF GENERAL DISTANCES)

Let $\chi_\varphi(\theta) \triangleq \chi(\varphi(\theta))$ and $\chi_{q_k}(\theta) \triangleq \chi(q_k(\theta))$. We want to find

$$\tilde{q} = \arg \min_{\varphi} \sum_{k=1}^K w_k \|\chi_{q_k} - \chi_\varphi\|_2^2. \quad (166)$$

To this end, we first derive

$$\chi^* = \arg \min_{\chi} \sum_{k=1}^K w_k \|\chi_{q_k} - \chi\|_2^2. \quad (167)$$

Following the same steps as in Appendix I with q_k replaced by χ_{q_k} and φ replaced by χ , it is easy to see that

$$\chi^*(\theta) = \sum_{k=1}^K w_k \chi(q_k(\theta)). \quad (168)$$

Because $\chi(q_k(\theta)) \in (a, b)$, the convex combination $\sum_{k=1}^K w_k \chi(q_k(\theta))$ is again in (a, b) . Thus, $\chi^*(\theta)$ is in the range of χ , and we can define

$$\varphi^*(\theta) \triangleq \chi^{-1}(\chi^*(\theta)). \quad (169)$$

This implies that

$$\chi_{\varphi^*}(\theta) = \chi(\varphi^*(\theta)) = \chi^*(\theta). \quad (170)$$

We claim that \tilde{q} defined in (166) equals φ^* . Indeed, we have for any φ

$$\begin{aligned} \sum_{k=1}^K w_k \|\chi_{q_k} - \chi_\varphi\|_2^2 &\geq \min_{\chi} \sum_{k=1}^K w_k \|\chi_{q_k} - \chi\|_2^2 \\ &\stackrel{(167)}{=} \sum_{k=1}^K w_k \|\chi_{q_k} - \chi^*\|_2^2 \\ &\stackrel{(170)}{=} \sum_{k=1}^K w_k \|\chi_{q_k} - \chi_{\varphi^*}\|_2^2 \end{aligned}$$

from which we conclude that φ^* achieves the minimum in (166) and, thus, equals \tilde{q} . We then obtain the optimal nonnormalized pooling function as

$$\begin{aligned} \tilde{q}(\theta) &= \varphi^*(\theta) \\ &\stackrel{(169)}{=} \chi^{-1}(\chi^*(\theta)) \\ &\stackrel{(168)}{=} \chi^{-1}\left(\sum_{k=1}^K w_k \chi(q_k(\theta))\right). \end{aligned}$$

APPENDIX K
PROOFS OF THE FUSION RULE FOR
A SCALAR PARAMETER

A. Proof of Theorem 21

For $d_\theta = 1$, the local observation likelihood functions from (71) are given by

$$\ell_k(\theta) \propto \exp\left(-\frac{\theta^2 \mathbf{h}_k^\top \Sigma_{kk}^{-1} \mathbf{h}_k}{2} + \theta \mathbf{h}_k^\top \Sigma_{kk}^{-1} \mathbf{h}_k t_k\right) \quad (171)$$

where $t_k = \mathbf{v}_k^\top \mathbf{y}_k = \mathbf{h}_k^\top \Sigma_{kk}^{-1} \mathbf{y}_k / (\mathbf{h}_k^\top \Sigma_{kk}^{-1} \mathbf{h}_k)$ according to (82) and (83). Furthermore, the global likelihood function (81) can be rewritten as

$$\begin{aligned} \lambda(\theta) &\propto \exp\left(-\frac{(\mathbf{t} - \mathbf{1}_K \theta)^\top \tilde{\Sigma}^{-1} (\mathbf{t} - \mathbf{1}_K \theta)}{2}\right) \\ &\propto \exp\left(-\frac{\theta^2}{2\hat{\sigma}^2} + \theta \mathbf{1}_K^\top \tilde{\Sigma}^{-1} \mathbf{t}\right) \end{aligned} \quad (172)$$

where $\hat{\sigma}^2 = 1/(\mathbf{1}_K^\top \tilde{\Sigma}^{-1} \mathbf{1}_K)$. The relation (84) follows from

$$\begin{aligned} &\prod_{k=1}^K (\ell_k(\theta))^{w_k} \\ &\stackrel{(171)}{\propto} \exp\left(\sum_{k=1}^K w_k \left(-\frac{\theta^2 \mathbf{h}_k^\top \Sigma_{kk}^{-1} \mathbf{h}_k}{2} + \theta \mathbf{h}_k^\top \Sigma_{kk}^{-1} \mathbf{h}_k t_k\right)\right) \\ &\stackrel{(85)}{=} \exp\left(-\frac{\sum_{k=1}^K \theta^2 \mathbf{1}_K^\top \tilde{\Sigma}^{-1} \mathbf{e}_k}{2} + \sum_{k=1}^K \theta \mathbf{1}_K^\top \tilde{\Sigma}^{-1} \mathbf{e}_k t_k\right) \\ &\stackrel{(a)}{=} \exp\left(-\frac{\theta^2}{2\hat{\sigma}^2} + \theta \mathbf{1}_K^\top \tilde{\Sigma}^{-1} \mathbf{t}\right) \\ &\stackrel{(172)}{\propto} \lambda(\theta) \end{aligned}$$

where we used in (a) that $\sum_{k=1}^K \mathbf{e}_k = \mathbf{1}_K$ and $\sum_{k=1}^K \mathbf{e}_k t_k = \mathbf{t}$. Finally, the fusion rule for the posteriors in (86) easily follows from (84), i.e.,

$$\begin{aligned} p(\theta | \mathbf{t}) &\propto p(\theta) \lambda(\theta) \\ &\stackrel{(84)}{\propto} p(\theta) \prod_{k=1}^K (\ell_k(\theta))^{w_k} \\ &\propto p(\theta) \prod_{k=1}^K \left(\frac{\pi_k(\theta)}{p(\theta)}\right)^{w_k} \\ &= (p(\theta))^{1 - \sum_{k=1}^K w_k} \prod_{k=1}^K (\pi_k(\theta))^{w_k}. \end{aligned}$$

B. Calculation of the Weights in Example 3

We will show expression (92) for w_k . The vectors \mathbf{v}_k^\top in (83) are given as

$$\mathbf{v}_k^\top = \frac{1}{\mathbf{1}_{r_0+r_k}^\top \mathbf{1}_{r_0+r_k}} \mathbf{1}_{r_0+r_k}^\top = \frac{1}{r_0+r_k} \mathbf{1}_{r_0+r_k}^\top$$

and, in turn, the matrix $\tilde{\Sigma}$ in (79) is given by the entries

$$\begin{aligned}\tilde{\Sigma}_{kk'} &= \mathbf{v}_k^\top \Sigma_{kk'} \mathbf{v}_{k'} \\ &= \frac{1}{r_0 + r_k} \mathbf{1}_{r_0+r_k}^\top \Sigma_{kk'} \frac{1}{r_0 + r_{k'}} \mathbf{1}_{r_0+r_{k'}} \\ &= \frac{r_0}{(r_0 + r_k)(r_0 + r_{k'})}\end{aligned}$$

for $k \neq k'$ and

$$\tilde{\Sigma}_{kk} = \frac{1}{r_0 + r_k}.$$

It is easily verified that we can rewrite $\tilde{\Sigma}$ as the following sum of a diagonal matrix and a rank one matrix:

$$\begin{aligned}\tilde{\Sigma} &= \begin{pmatrix} \frac{r_1}{(r_0+r_1)^2} & & \\ & \ddots & \\ & & \frac{r_K}{(r_0+r_K)^2} \end{pmatrix} \\ &+ \begin{pmatrix} \frac{1}{r_0+r_1} \\ \vdots \\ \frac{1}{r_0+r_K} \end{pmatrix} r_0 \begin{pmatrix} \frac{1}{r_0+r_1} & \cdots & \frac{1}{r_0+r_K} \end{pmatrix}.\end{aligned}$$

By the matrix inversion lemma [199, eq. (0.7.4.2)], we can, hence, calculate $\tilde{\Sigma}^{-1}$ as

$$\begin{aligned}\tilde{\Sigma}^{-1} &= \begin{pmatrix} \frac{(r_0+r_1)^2}{r_1} & & \\ & \ddots & \\ & & \frac{(r_0+r_K)^2}{r_K} \end{pmatrix} \\ &- \left(\sum_{k=0}^K \frac{1}{r_k} \right)^{-1} \begin{pmatrix} \frac{r_0+r_1}{r_1} \\ \vdots \\ \frac{r_0+r_K}{r_K} \end{pmatrix} \begin{pmatrix} \frac{r_0+r_1}{r_1} & \cdots & \frac{r_0+r_K}{r_K} \end{pmatrix}.\end{aligned}$$

To calculate the weights w_k in (85), we have to sum over the k th column of $\tilde{\Sigma}^{-1}$ and divide by $\mathbf{1}_{r_0+r_k}^\top \mathbf{1}_{r_0+r_k} = r_0 + r_k$, i.e.,

$$\begin{aligned}w_k &= \frac{1}{r_0 + r_k} \left(\frac{(r_0 + r_k)^2}{r_k} - \frac{\sum_{k'=1}^K \frac{(r_0 + r_k)(r_0 + r_{k'})}{r_k r_{k'}}}{\sum_{k'=0}^K \frac{1}{r_{k'}}} \right) \\ &= \frac{r_0 + r_k}{r_k} - \frac{\sum_{k'=1}^K \frac{r_0 + r_{k'}}{r_k r_{k'}}}{\sum_{k'=0}^K \frac{1}{r_{k'}}} \\ &= \frac{r_0 + r_k}{r_k} - \frac{\frac{K}{r_k} + \frac{r_0}{r_k} \sum_{k'=1}^K \frac{1}{r_{k'}}}{\sum_{k'=0}^K \frac{1}{r_{k'}}} \\ &= \frac{r_0 + r_k}{r_k} - \frac{\frac{K-1}{r_k} + \frac{r_0}{r_k} \sum_{k'=0}^K \frac{1}{r_{k'}}}{\sum_{k'=0}^K \frac{1}{r_{k'}}} \\ &= \frac{r_0}{r_k} + 1 - \frac{\frac{K-1}{r_k}}{\sum_{k'=0}^K \frac{1}{r_{k'}}} - \frac{r_0}{r_k} \\ &= 1 - \frac{K-1}{r_k} \left(\sum_{k'=0}^K \frac{1}{r_{k'}} \right)^{-1}.\end{aligned}$$

APPENDIX L PROOFS OF THE FUSION RULE FOR A VECTOR PARAMETER

A. Proof of Theorem 23

We can rewrite (81) as

$$\lambda(\boldsymbol{\theta})$$

$$\begin{aligned}&\propto \exp \left(-\frac{((\mathbf{1}_K \otimes \mathbf{I}_{d_\theta}) \boldsymbol{\theta} - \mathbf{t})^\top \tilde{\Sigma}^{-1} ((\mathbf{1}_K \otimes \mathbf{I}_{d_\theta}) \boldsymbol{\theta} - \mathbf{t})}{2} \right) \\ &\propto \exp \left(-\frac{\boldsymbol{\theta}^\top (\mathbf{1}_K \otimes \mathbf{I}_{d_\theta})^\top \tilde{\Sigma}^{-1} (\mathbf{1}_K \otimes \mathbf{I}_{d_\theta}) \boldsymbol{\theta}}{2} \right. \\ &\quad \left. + \boldsymbol{\theta}^\top (\mathbf{1}_K \otimes \mathbf{I}_{d_\theta})^\top \tilde{\Sigma}^{-1} \mathbf{t} \right) \\ &\stackrel{(99)}{=} \exp \left(-\frac{\boldsymbol{\theta}^\top \tilde{\Sigma}^{-1} \boldsymbol{\theta}}{2} + \boldsymbol{\theta}^\top (\mathbf{1}_K \otimes \mathbf{I}_{d_\theta})^\top \tilde{\Sigma}^{-1} \mathbf{t} \right). \quad (173)\end{aligned}$$

Furthermore, from (71), we see that

$$\ell_k(\boldsymbol{\theta}) \propto \exp \left(-\frac{\boldsymbol{\theta}^\top \mathbf{H}_k^\top \Sigma_{kk}^{-1} \mathbf{H}_k \boldsymbol{\theta}}{2} + \boldsymbol{\theta}^\top \mathbf{H}_k^\top \Sigma_{kk}^{-1} \mathbf{H}_k \mathbf{t}_k \right)$$

where $\mathbf{t}_k = \mathbf{V}_k \mathbf{y}_k$. Thus, we have

$$\begin{aligned}&\prod_{k=1}^K \ell_k(\mathbf{W}_k \boldsymbol{\theta}) \\ &\propto \exp \left(-\frac{\boldsymbol{\theta}^\top \left(\sum_{k=1}^K \mathbf{W}_k^\top \mathbf{H}_k^\top \Sigma_{kk}^{-1} \mathbf{H}_k \mathbf{W}_k \right) \boldsymbol{\theta}}{2} \right. \\ &\quad \left. + \boldsymbol{\theta}^\top \sum_{k=1}^K \mathbf{W}_k^\top \mathbf{H}_k^\top \Sigma_{kk}^{-1} \mathbf{H}_k \mathbf{t}_k \right) \\ &\stackrel{(98)}{=} \exp \left(-\frac{\boldsymbol{\theta}^\top (\hat{\Sigma}^{-1} - \mathbf{G}) \boldsymbol{\theta}}{2} + \boldsymbol{\theta}^\top \sum_{k=1}^K \mathbf{W}_k^\top \mathbf{H}_k^\top \Sigma_{kk}^{-1} \mathbf{H}_k \mathbf{t}_k \right) \\ &= \frac{1}{\xi_0(\boldsymbol{\theta})} \exp \left(-\frac{\boldsymbol{\theta}^\top \hat{\Sigma}^{-1} \boldsymbol{\theta}}{2} + \boldsymbol{\theta}^\top \sum_{k=1}^K \mathbf{W}_k^\top \mathbf{H}_k^\top \Sigma_{kk}^{-1} \mathbf{H}_k \mathbf{t}_k \right) \quad (174)\end{aligned}$$

with $\xi_0(\boldsymbol{\theta})$ as defined in (97). By comparing (173) and (174), we see that (95) holds, provided that

$$(\mathbf{1}_K \otimes \mathbf{I}_{d_\theta})^\top \tilde{\Sigma}^{-1} \mathbf{t} = \sum_{k=1}^K \mathbf{W}_k^\top \mathbf{H}_k^\top \Sigma_{kk}^{-1} \mathbf{H}_k \mathbf{t}_k. \quad (175)$$

Inserting (96) into the right-hand side of (175), we obtain

$$\begin{aligned}\sum_{k=1}^K \mathbf{W}_k^\top \mathbf{H}_k^\top \Sigma_{kk}^{-1} \mathbf{H}_k \mathbf{t}_k &= (\mathbf{1}_K \otimes \mathbf{I}_{d_\theta})^\top \tilde{\Sigma}^{-1} \sum_{k=1}^K (\mathbf{e}_k \otimes \mathbf{I}_{d_\theta}) \mathbf{t}_k \\ &= (\mathbf{1}_K \otimes \mathbf{I}_{d_\theta})^\top \tilde{\Sigma}^{-1} \mathbf{t}\end{aligned}$$

concluding the proof of (95).

Finally, the fusion rule (100) easily follows from (95), i.e.,

$$\begin{aligned} p(\boldsymbol{\theta} | \mathbf{t}) &\propto p(\boldsymbol{\theta})\lambda(\boldsymbol{\theta}) \\ &\propto p(\boldsymbol{\theta})\xi_0(\boldsymbol{\theta}) \prod_{k=1}^K \ell_k(\mathbf{W}_k \boldsymbol{\theta}) \\ &\propto p(\boldsymbol{\theta})\xi_0(\boldsymbol{\theta}) \prod_{k=1}^K \frac{\pi_k(\mathbf{W}_k \boldsymbol{\theta})}{p(\mathbf{W}_k \boldsymbol{\theta})}. \end{aligned}$$

B. Proof of Corollary 24

We start directly from $p(\boldsymbol{\theta} | \mathbf{t}) \propto p(\boldsymbol{\theta})\lambda(\boldsymbol{\theta})$. By (173) and our choice of prior $p(\boldsymbol{\theta}) = \mathcal{N}(\boldsymbol{\theta}; \boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0) \propto \exp(-\boldsymbol{\theta}^\top \boldsymbol{\Sigma}_0^{-1} \boldsymbol{\theta} / 2 + \boldsymbol{\theta}^\top \boldsymbol{\Sigma}_0^{-1} \boldsymbol{\mu}_0)$, we have that

$$\begin{aligned} p(\boldsymbol{\theta})\lambda(\boldsymbol{\theta}) &\propto \exp\left(-\frac{\boldsymbol{\theta}^\top \hat{\boldsymbol{\Sigma}}^{-1} \boldsymbol{\theta} + \boldsymbol{\theta}^\top \boldsymbol{\Sigma}_0^{-1} \boldsymbol{\theta}}{2}\right. \\ &\quad \left. + \boldsymbol{\theta}^\top (\mathbf{1}_K \otimes \mathbf{I}_{d_\theta})^\top \tilde{\boldsymbol{\Sigma}}^{-1} \mathbf{t} + \boldsymbol{\theta}^\top \boldsymbol{\Sigma}_0^{-1} \boldsymbol{\mu}_0\right) \end{aligned}$$

REFERENCES

- [1] S. Bandyopadhyay and S.-J. Chung, "Distributed Bayesian filtering using logarithmic opinion pool for dynamic sensor networks," *Automatica*, vol. 97, pp. 7–17, Nov. 2018.
- [2] D. Clark, S. Julier, R. Mahler, and B. Ristic, "Robust multi-object sensor fusion with unknown correlations," in *Proc. Sens. Signal Process. Defence (SSPD)*, London, U.K., 2010, p. 14.
- [3] J. Hu, L. Xie, and C. Zhang, "Diffusion Kalman filtering based on covariance intersection," *IEEE Trans. Signal Process.*, vol. 60, no. 2, pp. 891–902, Feb. 2012.
- [4] W. Maddern and P. Newman, "Real-time probabilistic fusion of sparse 3D LiDAR and dense stereo," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Daejeon, South Korea, Oct. 2016, pp. 2181–2188.
- [5] K. Da, T. Li, Y. Zhu, H. Fan, and Q. Fu, "Recent advances in multisensor multitarget tracking using random finite set," *Frontiers Inf. Technol. Electron. Eng.*, vol. 22, no. 1, pp. 5–24, Jan. 2021.
- [6] C. Fantacci, B.-N. Vo, B.-T. Vo, G. Battistelli, and L. Chisci, "Robust fusion for multisensor multiobject tracking," *IEEE Signal Process. Lett.*, vol. 25, no. 5, pp. 640–644, May 2018.
- [7] F. Meyer, O. Hlinka, H. Wymeersch, E. Riegler, and F. Hlawatsch, "Distributed localization and tracking of mobile networks including noncooperative objects," *IEEE Trans. Signal Inf. Process. Netw.*, vol. 2, no. 1, pp. 57–71, Mar. 2016.
- [8] M. Üney, D. E. Clark, and S. J. Julier, "Distributed fusion of PHD filters via exponential mixture densities," *IEEE J. Sel. Topics Signal Process.*, vol. 7, no. 3, pp. 521–531, Jun. 2013.
- [9] B. Lakshminarayanan, A. Pritzel, and C. Blundell, "Simple and scalable predictive uncertainty estimation using deep ensembles," in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, vol. 30, 2017, pp. 1–12.
- [10] H. Lee, S. Hong, and E. Kim, "Neural network ensemble with probabilistic fusion and its application to gait recognition," *Neurocomputing*, vol. 72, nos. 7–9, pp. 1557–1564, Mar. 2009.
- [11] Q. Lu, G. Karanikolas, Y. Shen, and G. B. Giannakis, "Ensemble Gaussian processes with spectral features for online interactive learning with scalability," in *Proc. Int. Conf. Artif. Intell. Statist.*, Palermo, Italy, Jun. 2020, pp. 1910–1920.
- [12] A. T. Thorgeirsson and F. Gautierin, "Probabilistic predictions with federated learning," *Entropy*, vol. 23, no. 1, p. 41, Dec. 2020.
- [13] E. Alam, R. Mehmood, I. Katib, N. N. Albogami, and A. Albeshri, "Data fusion and IoT for smart ubiquitous environments: A survey," *IEEE Access*, vol. 5, pp. 9533–9554, 2017.
- [14] E. Kats, J. Goldberger, and H. Greenspan, "A soft STAPLE algorithm combined with anatomical knowledge," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent. (MICCAI)*, Shenzhen, China: Springer, 2019, pp. 510–517.
- [15] B. Kolosz, S. Grant-Müller, and K. Djemame, "Modelling uncertainty in the sustainability of intelligent transport systems for highways using probabilistic data fusion," *Environ. Model. Softw.*, vol. 49, pp. 78–97, Nov. 2013.
- [16] A. Chlingaryan, S. Sukkarieh, and B. Whelan, "Machine learning approaches for crop yield prediction and nitrogen status estimation in precision agriculture: A review," *Comput. Electron. Agricul.*, vol. 151, pp. 61–69, Aug. 2018.
- [17] G.-B. Li, L.-L. Yang, Y. Xu, W.-J. Wang, L.-L. Li, and S.-Y. Yang, "A combined molecular docking-based and pharmacophore-based target prediction strategy with a probabilistic fusion method for target ranking," *J. Mol. Graph. Model.*, vol. 44, pp. 278–285, Jul. 2013.
- [18] A. H. Murphy and R. L. Winkler, "Probability forecasting in meteorology," *J. Amer. Stat. Assoc.*, vol. 79, no. 387, pp. 489–500, Sep. 1984.
- [19] R. Marty, V. Fortin, H. Kuswanto, A.-C. Favre, and E. Parent, "Combining the Bayesian processor of output with Bayesian model averaging for reliable ensemble forecasting," *J. Roy. Stat. Soc., C, Appl. Statist.*, vol. 64, no. 1, pp. 75–92, Jan. 2015.
- [20] J. Mitchell and S. G. Hall, "Evaluating, comparing and combining density forecasts using the KLIC with an application to the bank of England and NIESR 'Fan' charts of inflation," *Oxford Bull. Econ. Statist.*, vol. 67, no. s1, pp. 995–1033, Dec. 2005.
- [21] E. Moral-Benito, "Model averaging in economics: An overview," *J. Econ. Surv.*, vol. 29, no. 1, pp. 46–75, Feb. 2015.
- [22] S. Barak, A. Arjmand, and S. Ortobelli, "Fusion of multiple diverse predictors in stock market," *Inf. Fusion*, vol. 36, pp. 90–102, Jul. 2017.
- [23] C. Genest and J. V. Zidek, "Combining probability distributions: A critique and an annotated bibliography," *Stat. Sci.*, vol. 1, no. 1, pp. 114–135, Feb. 1986.
- [24] S. G. Hall and J. Mitchell, "Combining density forecasts," *Int. J. Forecasting*, vol. 23, no. 1, pp. 1–13, Jan. 2007.
- [25] T. Gneiting, "Editorial: Probabilistic forecasting," *J. Roy. Stat. Soc., A, Statist. Soc.*, vol. 171, no. 2, pp. 319–321, Apr. 2008.
- [26] M. Stone, "The opinion pool," *Ann. Math. Statist.*, vol. 32, no. 4, pp. 1339–1342, Dec. 1961.
- [27] T. Bailey, S. Julier, and G. Agamennoni, "On conservative fusion of information with unknown non-Gaussian dependence," in *Proc. Int. Conf. Inf. Fusion (FUSION)*, Singapore, 2012, pp. 1876–1883.
- [28] C. Genest, "A characterization theorem for externally Bayesian groups," *Ann. Statist.*, vol. 12, no. 3, pp. 1100–1105, Sep. 1984.
- [29] M. B. Hurley, "An information theoretic justification for covariance intersection and its generalization," in *Proc. Int. Conf. Inf. Fusion (FUSION)*, Annapolis, MD, USA, 2002, pp. 505–511.
- [30] N. Lehrer, O. Tslil, and A. Carmi, "Log-linear Chernoff fusion for distributed particle filtering," in *Proc. Int. Conf. Inf. Fusion (FUSION)*, Ottawa, ON, Canada, 2019, pp. 1–9.
- [31] M. Gunay, U. Orguner, and M. Demirekler, "Chernoff fusion of Gaussian mixtures based on sigma-point approximation," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 52, no. 6, pp. 2732–2746, Dec. 2016.
- [32] S. J. Julier and J. K. Uhlmann, "A non-divergent estimation algorithm in the presence of unknown correlations," in *Proc. Amer. Control Conf.*, Albuquerque, NM, USA, 1997, pp. 2369–2373.
- [33] K.-C. Chang, C.-Y. Chong, and S. Mori, "Analytical and computational evaluation of scalable distributed fusion algorithms," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 46, no. 4, pp. 2022–2034, Oct. 2010.
- [34] I. Urteaga, M. F. Bugallo, and P. M. Djurić, "Sequential Monte Carlo methods under model uncertainty," in *Proc. IEEE Stat. Signal Process. Workshop (SSP)*, Palma de Mallorca, Spain, Jun. 2016, pp. 1–5.
- [35] O. Hlinka, F. Hlawatsch, and P. M. Djurić, "Consensus-based distributed particle filtering with distributed proposal adaptation," *IEEE Trans.*

$$\begin{aligned} &= \exp\left(-\frac{\boldsymbol{\theta}^\top (\hat{\boldsymbol{\Sigma}}^{-1} + \boldsymbol{\Sigma}_0^{-1}) \boldsymbol{\theta}}{2}\right. \\ &\quad \left. + \boldsymbol{\theta}^\top ((\mathbf{1}_K \otimes \mathbf{I}_{d_\theta})^\top \tilde{\boldsymbol{\Sigma}}^{-1} \mathbf{t} + \boldsymbol{\Sigma}_0^{-1} \boldsymbol{\mu}_0)\right) \\ &= \exp\left(-\frac{\boldsymbol{\theta}^\top \boldsymbol{\Sigma}_1^{-1} \boldsymbol{\theta}}{2} + \boldsymbol{\theta}^\top \boldsymbol{\Sigma}_1^{-1} \boldsymbol{\mu}_1\right) \\ &\propto \exp\left(-\frac{(\boldsymbol{\theta} - \boldsymbol{\mu}_1)^\top \boldsymbol{\Sigma}_1^{-1} (\boldsymbol{\theta} - \boldsymbol{\mu}_1)}{2}\right) \end{aligned} \quad (176)$$

with $\boldsymbol{\mu}_1$ and $\boldsymbol{\Sigma}_1$ given by (101) and (102), respectively. Expression (176) is proportional to the pdf of a Gaussian with mean $\boldsymbol{\mu}_1$ and covariance matrix $\boldsymbol{\Sigma}_1$.

Acknowledgment

The authors would like to thank the anonymous reviewers for insightful comments and constructive criticism, which have resulted in a significant improvement of this article. In particular, they are thankful for a comment on a conceptual relationship between probability density function (pdf) fusion and robust hypothesis testing as mentioned in Section V. They are also grateful to Thomas Kropfleiter for sharing his expertise in target tracking.

Signal Process., vol. 62, no. 12, pp. 3029–3041, Jun. 2014.

[36] V. Savic, H. Wymeersch, and S. Zazo, “Belief consensus algorithms for fast distributed target tracking in wireless sensor networks,” *Signal Process.*, vol. 95, no. 2, pp. 149–160, Feb. 2014.

[37] H. Wymeersch, J. Lien, and M. Z. Win, “Cooperative localization in wireless networks,” *Proc. IEEE*, vol. 97, no. 2, pp. 427–450, Feb. 2009.

[38] C. G. Lopes and A. H. Sayed, “Diffusion least-mean squares over adaptive networks: Formulation and performance analysis,” *IEEE Trans. Signal Process.*, vol. 56, no. 7, pp. 3122–3136, Jul. 2008.

[39] O. Puntsa, “Bayesian approaches to multi-sensor data fusion,” M.S. thesis, Signal Process. Commun. Lab., Univ. Cambridge, Cambridge, U.K., 1999.

[40] G. Battistelli and L. Chisci, “Kullback–Leibler average, consensus on probability densities, and distributed state estimation with guaranteed stability,” *Automatica*, vol. 50, no. 3, pp. 707–718, Mar. 2014.

[41] Z. Deng, P. Zhang, W. Qi, J. Liu, and Y. Gao, “Sequential covariance intersection fusion Kalman filter,” *Inf. Sci.*, vol. 189, pp. 293–309, Apr. 2012.

[42] M. Tang, Y. Rong, J. Zhou, and X. R. Li, “Information geometric approach to multisensor estimation fusion,” *IEEE Trans. Signal Process.*, vol. 67, no. 2, pp. 279–292, Jan. 2019.

[43] G. Li, G. Battistelli, L. Chisci, W. Yi, and L. Kong, “Distributed multi-view multi-target tracking based on CPHD filtering,” *Signal Process.*, vol. 188, Nov. 2021, Art. no. 108210.

[44] T. Li, X. Wang, Y. Liang, and Q. Pan, “On arithmetic average fusion and its application for distributed multi-Bernoulli multitarget tracking,” *IEEE Trans. Signal Process.*, vol. 68, pp. 2883–2896, 2020.

[45] T. Li and F. Hlawatsch, “A distributed particle-PHD filter using arithmetic-average fusion of Gaussian mixture parameters,” *Inf. Fusion*, vol. 73, pp. 111–124, Sep. 2021.

[46] W. Yi, G. Li, and G. Battistelli, “Distributed multi-sensor fusion of PHD filters with different sensor fields of view,” *IEEE Trans. Signal Process.*, vol. 68, pp. 5204–5218, 2020.

[47] L. Gao, G. Battistelli, and L. Chisci, “Fusion of labeled RFS densities with minimum information loss,” *IEEE Trans. Signal Process.*, vol. 68, pp. 5855–5868, 2020.

[48] A. K. Gostar et al., “Cooperative sensor fusion in centralized sensor networks using Cauchy–Schwarz divergence,” *Signal Process.*, vol. 167, Feb. 2020, Art. no. 107278.

[49] S. Li, G. Battistelli, L. Chisci, W. Yi, B. Wang, and L. Kong, “Computationally efficient multi-agent multi-object tracking with labeled random finite sets,” *IEEE Trans. Signal Process.*, vol. 67, no. 1, pp. 260–275, Jan. 2019.

[50] R. P. S. Mahler, “Optimal/robust distributed data fusion: A unified approach,” *Proc. SPIE*, vol. 4052, pp. 128–138, Aug. 2000.

[51] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, “Federated learning: Challenges, methods, and future directions,” *IEEE Signal Process. Mag.*, vol. 37, no. 3, pp. 50–60, May 2020.

[52] M. Yurochkin, M. Agarwal, S. Ghosh, K. Greenewald, N. Hoang, and Y. Khazaeni, “Bayesian nonparametric federated learning of neural networks,” in *Proc. Int. Conf. Mach. Learn. (ICML)*, Long Beach, CA, USA, 2019, pp. 7252–7261.

[53] R. T. Clemen, “Combining forecasts: A review and annotated bibliography,” *Int. J. Forecasting*, vol. 5, no. 4, pp. 559–583, 1989.

[54] J. S. Armstrong, “Combining forecasts,” in *Principles of Forecasting*, J. S. Armstrong, Ed. Boston, MA, USA: Springer, 2001, pp. 417–439.

[55] K. F. Wallis, “Combining density and interval forecasts: A modest proposal,” *Oxford Bull. Econ. Statist.*, vol. 67, no. s1, pp. 983–994, Dec. 2005.

[56] R. L. Winkler, Y. Grushka-Cockayne, K. C. Lichtendahl, Jr., and V. R. R. Jose, “Probability forecasts and their combination: A research perspective,” *Decis. Anal.*, vol. 16, no. 4, pp. 239–260, Dec. 2019.

[57] J. A. Hoeting, D. Madigan, A. E. Raftery, and C. T. Volinsky, “Bayesian model averaging: A tutorial,” *Statist. Sci.*, vol. 14, no. 4, pp. 382–401, 1999.

[58] T. M. Fragoso, W. Bertoli, and F. Louzada, “Bayesian model averaging: A systematic review and conceptual classification,” *Int. Stat. Rev.*, vol. 86, no. 1, pp. 1–28, Apr. 2018.

[59] D. Posada, “jModelTest: Phylogenetic model averaging,” *Mol. Biol. Evol.*, vol. 25, no. 7, pp. 1253–1256, Apr. 2008.

[60] D. Darriba, G. L. Taboada, R. Doallo, and D. Posada, “jModelTest 2: More models, new heuristics and parallel computing,” *Nature Methods*, vol. 9, no. 8, p. 772, 2012.

[61] M. F. J. Steel, “Model averaging and its use in economics,” *J. Econ. Literature*, vol. 58, no. 3, pp. 644–719, Sep. 2020.

[62] C. F. Dormann et al., “Model averaging in ecology: A review of Bayesian, information-theoretic, and tactical approaches for predictive inference,” *Ecol. Monograph*, vol. 88, no. 4, pp. 485–504, Nov. 2018.

[63] F. E. Turkheimer, R. Hinz, and V. J. Cunningham, “On the undecidability among kinetic models: From model selection to model averaging,” *J. Cerebral Blood Flow Metabolism*, vol. 23, no. 4, pp. 490–498, 2003.

[64] J. M. Montgomery and B. Nyhan, “Bayesian model averaging: Theoretical developments and practical applications,” *Political Anal.*, vol. 18, no. 2, pp. 245–270, 2010.

[65] W. Neiswanger, C. Wang, and E. P. Xing, “Asymptotically exact, embarrassingly parallel MCMC,” in *Proc. Conf. Uncertainty Artif. Intell. (UAI)*, Arlington, VA, USA, 2014, pp. 623–632.

[66] X. Wang and D. B. Dunson, “Parallelizing MCMC via Weierstrass sampler,” 2013, *arXiv:1312.4605*.

[67] R. Barnden, A. Doucet, and C. Holmes, “On Markov chain Monte Carlo methods for tall data,” *J. Mach. Learn. Res.*, vol. 18, no. 1, pp. 1515–1557, 2017.

[68] J. Rabin, G. Peyré, J. Delon, and M. Bernot, “Wasserstein barycenter and its application to texture mixing,” in *Proc. Int. Conf. Scale Space Variational Methods Comput. Vis. (SSVM)*, Ein-Gedi, Israel, 2011, pp. 435–446.

[69] S. Srivastava, C. Cevher, Q. Dinh, and D. Dunson, “WASP: Scalable Bayes via barycenters of subset posteriors,” in *Proc. Int. Conf. Artif. Intell. Stat. (AISTATS)*, San Diego, CA, USA, 2015, pp. 912–920.

[70] D. H. Wolpert, “Stacked generalization,” *Neural Netw.*, vol. 5, no. 2, pp. 241–259, 1992.

[71] L. Breiman, “Stacked regressions,” *Mach. Learn.*, vol. 24, no. 1, pp. 49–64, 1996.

[72] N. Shazeer et al., “Outrageously large neural networks: The sparsely-gated mixture-of-experts layer,” in *Proc. Int. Conf. Learn. Represent. (ICLR)*, Toulon, France, 2017.

[73] M. Hoang, N. Hoang, B. K. H. Low, and C. Kingsford, “Collective model fusion for multiple black-box experts,” in *Proc. Int. Conf. Mach. Learn. (ICML)*, Long Beach, CA, USA, 2019, pp. 2742–2750.

[74] H. Liu, Y.-S. Ong, X. Shen, and J. Cai, “When Gaussian process meets big data: A review of scalable GPs,” *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 31, no. 11, pp. 4405–4423, Nov. 2020.

[75] M. Woźniak, M. Graña, and E. Corchado, “A survey of multiple classifier systems as hybrid systems,” *Inf. Fusion*, vol. 16, pp. 3–17, Mar. 2014.

[76] K. P. Burnham and D. R. Anderson, *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach*, 2nd ed. New York, NY, USA: Springer, 2002.

[77] K. F. Wallis, “Combining forecasts—Forty years later,” *Appl. Financial Econ.*, vol. 21, nos. 1–2, pp. 33–41, Jan. 2011.

[78] R. M. Cooke, *Experts in Uncertainty: Opinion and Subjective Probability in Science*. New York, NY, USA: Oxford Univ. Press, 1991.

[79] R. T. Clemen and R. L. Winkler, “Combining probability distributions from experts in risk analysis,” *Risk Anal.*, vol. 19, no. 2, pp. 187–203, 1999.

[80] F. Dietrich and C. List, “Probabilistic opinion pooling,” in *The Oxford Handbook of Probability and Philosophy*. Oxford, U.K.: Oxford Univ. Press, 2016.

[81] R. T. Stewart and I. O. Quintana, “Probabilistic opinion pooling with imprecise probabilities,” *J. Phil. Log.*, vol. 47, no. 1, pp. 17–45, Feb. 2018.

[82] R. T. Stewart and I. O. Quintana, “Learning and pooling, pooling and learning,” *Erkenntnis*, vol. 83, no. 3, pp. 369–389, Jun. 2018.

[83] A. E. Abbas, “A Kullback–Leibler view of linear and log-linear pools,” *Decis. Anal.*, vol. 6, no. 1, pp. 25–37, Mar. 2009.

[84] A. Garg, T. S. Jayram, S. Vaithyanathan, and H. Zhu, “Generalized opinion pooling,” in *Proc. Int. Symp. Artif. Intell. Math. (ISAIM)*, Fort Lauderdale, FL, USA, 2004.

[85] K. Da, T. Li, Y. Zhu, H. Fan, and Q. Fu, “Kullback–Leibler averaging for multitarget density fusion,” in *Proc. Int. Symp. Distrib. Comput. Artif. Intell. (DCAI)*, Ávila, Spain, 2019, pp. 253–261.

[86] M. Aguech and G. Carlier, “Barycenters in the Wasserstein space,” *SIAM J. Math. Anal.*, vol. 43, no. 2, pp. 904–924, 2011.

[87] R. L. Winkler, “The consensus of subjective probability distributions,” *Manage. Sci.*, vol. 15, no. 2, pp. 61–75, 1968.

[88] P. A. Morris, “Combining expert judgments: A Bayesian approach,” *Manage. Sci.*, vol. 23, no. 7, pp. 679–693, Mar. 1977.

[89] R. L. Winkler, “Combining probability distributions from dependent information sources,” *Manage. Sci.*, vol. 27, no. 4, pp. 479–488, Apr. 1981.

[90] D. Lindley, “Reconciliation of probability distributions,” *Oper. Res.*, vol. 31, no. 5, pp. 866–880, Oct. 1983.

[91] R. T. Clemen and R. L. Winkler, “Limits for the precision and value of information from dependent sources,” *Oper. Res.*, vol. 33, no. 2, pp. 427–442, Apr. 1985.

[92] B. Anderson and J. Moore, *Optimal Filtering*. Englewood Cliffs, NJ, USA: Prentice-Hall, 1979.

[93] B. Ristic, S. Arulampalam, and N. Gordon, *Beyond the Kalman Filter: Particle Filters for Tracking Applications*. Boston, MA, USA: Artech House, 2003.

[94] C. Chong and S. Mori, “Convex combination and covariance intersection algorithms in distributed fusion,” in *Proc. Int. Conf. Inf. Fusion*, Montréal, QC, Canada, 2001.

[95] Y. Bar-Shalom and X. Li, *Multitarget-Multisensor Tracking: Principles and Techniques*. Storrs, CT, USA: YBS Publishing, 1995.

[96] S. Blackman and R. Popoli, *Design and Analysis of Modern Tracking Systems*. Norwood, MA, USA: Artech House, 1999.

[97] R. Mahler, *Statistical Multisource-Multitarget Information Fusion*. Norwood, MA, USA: Artech House, 2007.

[98] Y. Bar-Shalom, P. K. Willett, and X. Tian, *Tracking and Data Fusion*. Storrs, CT, USA: YBS Publishing, 2011.

[99] S. Challa, M. R. Morelande, D. Mušicki, and R. J. Evans, *Fundamentals of Object Tracking*. New York, NY, USA: Cambridge Univ. Press, 2011.

[100] R. P. S. Mahler, *Advances in Statistical Multisource-Multitarget Information Fusion*. Boston, MA, USA: Artech House, 2014.

[101] W. Koch, *Tracking and Sensor Data Fusion*. Berlin, Germany: Springer, 2016.

[102] F. Meyer et al., “Message passing algorithms for scalable multitarget tracking,” *Proc. IEEE*, vol. 106, no. 2, pp. 221–259, Feb. 2018.

[103] D. J. Daley and D. Vere-Jones, *An Introduction to the Theory of Point Processes: Elementary Theory and Methods*, vol. 1. New York, NY, USA: Springer, 2003.

[104] L. M. Kaplan, Y. Bar-Shalom, and W. D. Blair, “Assignment costs for multiple sensor track-to-track association,” *IEEE Trans. Aerosp. Electron. Syst.*, vol. 44, no. 2, pp. 655–677, Apr. 2008.

[105] S. Maresca, P. Braca, J. Horstmann, and R. Grasso, “Maritime surveillance using multiple high-frequency surface-wave radars,” *IEEE Trans. Geosci. Remote Sens.*, vol. 52, no. 8, pp. 5056–5071, Aug. 2014.

[106] M. Üney, S. Julier, D. Clark, and B. Ristić, “Monte

Carlo realisation of a distributed multi-object fusion algorithm," in *Proc. Sensor Signal Process. Defence (SSPD)*, London, U.K., Sep. 2010, pp. 1–5.

[107] R. Mahler, "Toward a theoretical foundation for distributed fusion," in *Distributed Data Fusion for Network-Centric Operations*, D. Hall, C.-Y. Chong, J. Llinas, and M. Liggins, II, Eds. Boca Raton, FL, USA: CRC Press, 2013, ch. 8.

[108] T. Li, H. Fan, J. García, and J. M. Corchado, "Second-order statistics analysis and comparison between arithmetic and geometric average fusion: Application to multi-sensor target tracking," *Inf. Fusion*, vol. 51, pp. 233–243, Nov. 2019.

[109] J. Y. Yu, M. Coates, and M. Rabat, "Distributed multi-sensor CPHD filter using pairwise gossiping," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, Shanghai, China, Mar. 2016, pp. 3176–3180.

[110] T. Li, J. M. Corchado, and S. Sun, "Partial consensus and conservative fusion of Gaussian mixtures for distributed PHD fusion," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 55, no. 5, pp. 2150–2163, Oct. 2019.

[111] L. Gao, G. Battistelli, and L. Chisci, "Multiobject fusion with minimum information loss," *IEEE Signal Process. Lett.*, vol. 27, pp. 201–205, 2020.

[112] C. Fantacci, B.-N. Vo, B.-T. Vo, G. Battistelli, and L. Chisci, "Consensus labeled random finite set filtering for distributed multi-object tracking," 2015, *arXiv:1501.01579*.

[113] S. Li, W. Yi, R. Hoseinnezhad, G. Battistelli, B. Wang, and L. Kong, "Robust distributed fusion with labeled random finite sets," *IEEE Trans. Signal Process.*, vol. 66, no. 2, pp. 278–293, Jan. 2018.

[114] T. Kropfreiter and F. Hlawatsch, "A probabilistic label association algorithm for distributed labeled multi-Bernoulli filtering," in *Proc. Int. Conf. Inf. Fusion (FUSION)*, Rustenburg, South Africa, Jul. 2020, pp. 1–8.

[115] Z. Ghahramani, "Probabilistic machine learning and artificial intelligence," *Nature*, vol. 521, no. 7553, pp. 452–459, May 2015.

[116] Y. Gal and Z. Ghahramani, "Dropout as a Bayesian approximation: Representing model uncertainty in deep learning," in *Proc. Int. Conf. Mach. Learn.*, New York, NY, USA, Jun. 2016, pp. 1050–1059.

[117] R. V. Krem, "Bayesian machine learning for quantum molecular dynamics," *Phys. Chem. Chem. Phys.*, vol. 21, no. 25, pp. 13392–13410, 2019.

[118] C. Leibig, V. Alkén, M. S. Ayhan, P. Berens, and S. Wahl, "Leveraging uncertainty information from deep neural networks for disease detection," *Sci. Rep.*, vol. 7, no. 1, pp. 1–14, Dec. 2017.

[119] E. Begoli, T. Bhattacharya, and D. Kusnezov, "The need for uncertainty quantification in machine-assisted medical decision making," *Nature Mach. Intell.*, vol. 1, no. 1, pp. 20–23, Jan. 2019.

[120] A. Kendall, V. Badrinarayanan, and R. Cipolla, "Bayesian SegNet: Model uncertainty in deep convolutional encoder-decoder architectures for scene understanding," 2015, *arXiv:1511.02680*.

[121] J. Ching and K.-K. Phoon, "Constructing site-specific multivariate probability distribution model using Bayesian machine learning," *J. Eng. Mech.*, vol. 145, no. 1, Jan. 2019, Art. no. 04018126.

[122] Y. Ovadia *et al.*, "Can you trust your model's uncertainty? Evaluating predictive uncertainty under dataset shift," in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, pp. 13991–14002.

[123] M. Abdar *et al.*, "A review of uncertainty quantification in deep learning: Techniques, applications and challenges," *Inf. Fusion*, vol. 76, pp. 243–297, Dec. 2021.

[124] H. Wang and D.-Y. Yeung, "Towards Bayesian deep learning: A framework and some existing methods," *IEEE Trans. Knowl. Data Eng.*, vol. 28, no. 12, pp. 3395–3408, Dec. 2016.

[125] A. G. Wilson and P. Izmailov, "Bayesian deep learning and a probabilistic perspective of generalization," 2020, *arXiv:2002.08791*.

[126] C. E. Rasmussen and C. K. Williams, *Gaussian Processes for Machine Learning*. Cambridge, MA, USA: MIT Press, 2006.

[127] H. Salimbeni and M. Deisenroth, "Doubly stochastic variational inference for deep Gaussian processes," in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, pp. 4588–4599.

[128] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning: Data Mining, Inference and Prediction*. New York, NY, USA: Springer, 2009.

[129] K. P. Murphy, *Machine Learning: A Probabilistic Perspective*. Cambridge, MA, USA: MIT Press, 2012.

[130] O. Sagi and L. Rokach, "Ensemble learning: A survey," *WIREs Data Mining Knowl. Discovery*, vol. 8, no. 4, p. e1249, Jul. 2018.

[131] D. Opitz and R. Maclin, "Popular ensemble methods: An empirical study," *J. Artif. Intell. Res.*, vol. 11, pp. 169–198, Aug. 1999.

[132] R. Polikar, "Ensemble based systems in decision making," *IEEE Circuits Syst. Mag.*, vol. 6, no. 3, pp. 21–45, Sep. 2006.

[133] S. Rothe and D. Söflker, "Comparison of different information fusion methods using ensemble selection considering benchmark data," in *Proc. Int. Conf. Inf. Fusion (FUSION)*, Heidelberg, Germany, 2016, pp. 73–78.

[134] S. Savazzi, M. Nicoli, and V. Rampa, "Federated learning with cooperating devices: A consensus approach for massive IoT networks," *IEEE Internet Things J.*, vol. 7, no. 5, pp. 4641–4654, May 2020.

[135] M. Mohri, G. Sivek, and A. T. Suresh, "Agnostic federated learning," in *Proc. Int. Conf. Mach. Learn. (ICML)*, Long Beach, CA, USA, 2019, pp. 4615–4625.

[136] A. Vehtari *et al.*, "Expectation propagation as a way of life: A framework for Bayesian inference on partitioned data," *J. Mach. Learn. Res.*, vol. 21, pp. 1–17, Dec. 2020.

[137] P. J. Brockwell and R. A. Davis, *Introduction to Time Series and Forecasting*, 3rd ed. Basel, Switzerland: Springer, 2016.

[138] R. J. Berndowitz and E. A. Zurnedorfer, "Automated guidance for predicting quantitative precipitation," *Monthly Weather Rev.*, vol. 107, no. 2, pp. 122–128, Feb. 1979.

[139] A. S. Tay and K. F. Wallis, "Density forecasting: A survey," *J. Forecasting*, vol. 19, no. 4, pp. 235–254, 2000.

[140] W. Freiburger and U. Grenander, "On the formulation of statistical meteorology," *Revue l'Institut Int. Statistique*, vol. 33, no. 1, pp. 59–86, 1965.

[141] E. S. Epstein, "Stochastic dynamic prediction," *Tellus*, vol. 21, no. 6, pp. 739–759, Dec. 1969.

[142] T. Gneiting and R. Ranjan, "Combining predictive distributions," *Electron. J. Statist.*, vol. 7, pp. 1747–1782, Jan. 2013.

[143] A. E. Raftery, T. Gneiting, F. Balabdaoui, and M. Polakowski, "Using Bayesian model averaging to calibrate forecast ensembles," *Monthly Weather Rev.*, vol. 133, no. 5, pp. 1155–1174, 2005.

[144] B. Glahn *et al.*, "MOS uncertainty estimates in an ensemble framework," *Monthly Weather Rev.*, vol. 137, no. 1, pp. 246–268, Jan. 2009.

[145] F. Bassetti, R. Casarin, and F. Ravazzolo, "Bayesian nonparametric calibration and combination of predictive distributions," *J. Amer. Stat. Assoc.*, vol. 113, no. 522, pp. 675–685, Apr. 2018.

[146] M. P. Clements and D. I. Harvey, "Combining probability forecasts," *Int. J. Forecasting*, vol. 27, no. 2, pp. 208–223, Apr. 2011.

[147] V. Zarnowitz, "The new ASA-NBER survey of forecasts by economic statisticians," *Amer. Statistician*, vol. 23, no. 1, pp. 12–16, Feb. 1969.

[148] J. Tastu, P. Pinson, P.-J. Trombe, and H. Madsen, "Probabilistic forecasts of wind power generation accounting for geographically dispersed information," *IEEE Trans. Smart Grid*, vol. 5, no. 1, pp. 480–489, Jan. 2014.

[149] S. Baran and S. Lerch, "Mixture EMOS model for calibrating ensemble forecasts of wind speed," *Environmetrics*, vol. 27, no. 2, pp. 116–130, Mar. 2016.

[150] T. Hong and S. Fan, "Probabilistic electric load forecasting: A tutorial review," *Int. J. Forecasting*, vol. 32, no. 3, pp. 914–938, 2016.

[151] J. Nowotarski and R. Weron, "Recent advances in electricity price forecasting: A review of probabilistic forecasting," *Renew. Sustain. Energy Rev.*, vol. 81, no. 1, pp. 1548–1568, 2018.

[152] K. Doubleday, S. Jascourt, W. Kleiber, and B.-M. Hodge, "Probabilistic solar power forecasting using Bayesian model averaging," *IEEE Trans. Sustain. Energy*, vol. 12, no. 1, pp. 325–337, Jan. 2021.

[153] C. Genest, "Pooling operators with the marginalization property," *Can. J. Statist.*, vol. 12, no. 2, pp. 153–163, Jun. 1984.

[154] C. Genest, K. J. McConway, and M. J. Schervish, "Characterization of externally Bayesian pooling operators," *Ann. Statist.*, vol. 14, no. 2, pp. 487–501, Jun. 1986.

[155] P. S. Bullen, *Handbook of Means and Their Inequalities*. Dordrecht, The Netherlands: Springer, 2013.

[156] K. J. McConway, "Marginalization and linear opinion pools," *J. Amer. Stat. Assoc.*, vol. 76, no. 374, pp. 410–414, Jun. 1981.

[157] R. Laddaga, "Lehrer and the consensus proposal," *Synthese*, vol. 36, no. 4, pp. 473–477, Dec. 1977.

[158] A. Madansky, "Externally Bayesian groups," RAND Corp., Santa Monica, CA, USA, Tech. Rep. RM-4141-PR, 1964.

[159] M. J. Rufo, J. Martín, and C. J. Pérez, "Log-linear pool to combine prior distributions: A suggestion for a calibration-based approach," *Bayesian Anal.*, vol. 7, no. 2, pp. 411–438, Jun. 2012.

[160] C. Genest, "A conflict between two axioms for combining subjective distributions," *J. Roy. Stat. Soc., B, Methodol.*, vol. 46, no. 3, pp. 403–405, 1984.

[161] I. Csiszár, "Eine informationstheoretische Ungleichung und ihre Anwendung auf den Beweis der Ergodizität von Markoffschen Ketten," *Magyar Tud. Akad. Mat. Kutató Int. Közl.*, vol. 8, nos. 1–2, pp. 85–108, 1963.

[162] I. Csiszár, "Information-type measures of difference of probability distributions and indirect observation," *Studia Scientiarum Mathematicarum Hungarica*, vol. 2, nos. 3–4, pp. 299–318, 1967.

[163] I. Vajda, "On the f -divergence and singularity of probability measures," *Periodica Math. Hungarica*, vol. 2, nos. 1–4, pp. 223–234, 1972.

[164] I. Sason, "On f -divergences: Integral representations, local behavior, and inequalities," *Entropy*, vol. 20, no. 5, p. 383, May 2018.

[165] V. V. Veeravalli, T. Basar, and H. V. Poor, "Minimax robust decentralized detection," *IEEE Trans. Inf. Theory*, vol. 40, no. 1, pp. 35–40, Jan. 1994.

[166] M. Fauß, A. M. Zoubir, and H. V. Poor, "Minimax robust detection: Classic results and recent advances," *IEEE Trans. Signal Process.*, vol. 69, pp. 2252–2283, 2021.

[167] F. Liese and I. Vajda, "On divergences and informations in statistics and information theory," *IEEE Trans. Inf. Theory*, vol. 52, no. 10, pp. 4394–4412, Oct. 2006.

[168] T. van Erven and P. Harremans, "Rényi divergence and Kullback–Leibler divergence," *IEEE Trans. Inf. Theory*, vol. 60, no. 7, pp. 3797–3820, Jul. 2014.

[169] K. Dedecius and P. M. Djurić, "Sequential estimation and diffusion of information over networks: A Bayesian approach with exponential family of distributions," *IEEE Trans. Signal Process.*, vol. 65, no. 7, pp. 1795–1809, Apr. 2017.

[170] H. Zhu and R. Rohwer, "Information-geometric measurements of generalisation," Aston Univ., Birmingham, U.K., Tech. Rep. NCRG/4350, 1995.

[171] T. Minka, "Divergence measures and message passing," Microsoft Res., Cambridge, U.K., Tech. Rep. MSR-TR-2005-173, 2005.

[172] A. Cichocki, S. Cruces, and S.-I. Amari, "Generalized alpha-beta divergences and their application to robust nonnegative matrix factorization," *Entropy*, vol. 13, no. 1, pp. 134–170, 2011.

[173] K. Pearson, "X. On the criterion that a given system of deviations from the probable in the case of a correlated system of variables is such that it can be reasonably supposed to have arisen from random sampling," *Philos. Mag.*, vol. 50, no. 302, pp. 157–175, 1900.

[174] D. P. Malladi and J. L. Speyer, "A new approach to multiple model adaptive estimation," in *Proc. IEEE Conf. Decis. Control (CDC)*, San Diego, CA, USA, Dec. 1997, pp. 3460–3467.

[175] J. Wang and M. R. Taafe, "Multivariate mixtures

of normal distributions: Properties, random vector generation, fitting, and as models of market daily changes," *INFORMS J. Comput.*, vol. 27, no. 2, pp. 193–203, Apr. 2015.

[176] P. D. Hanlon and P. S. Maybeck, "Multiple-model adaptive estimation using a residual correlation Kalman filter bank," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 36, no. 2, pp. 393–406, Apr. 2000.

[177] C. Robert and G. Casella, *Monte Carlo Statistical Methods*. New York, NY, USA: Springer, 2013.

[178] A. Hinrichs, J. Prochno, and M. Ullrich, "The curse of dimensionality for numerical integration on general domains," *J. Complex.*, vol. 50, pp. 25–42, Feb. 2019.

[179] C. Genest and K. J. McConway, "Allocating the weights in the linear opinion pool," *J. Forecasting*, vol. 9, no. 1, pp. 53–73, Jan. 1990.

[180] M. H. DeGroot and J. Mortera, "Optimal linear opinion pools," *Manage. Sci.*, vol. 37, no. 5, pp. 546–558, May 1991.

[181] R. T. Clemen, "Comment on Cooke's classical method," *Rel. Eng. Syst. Saf.*, vol. 93, no. 5, pp. 760–765, 2008.

[182] T. Heskes, "Selecting weighting factors in logarithmic opinion pools," in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, Denver, CO, USA, 1997, pp. 266–272.

[183] L. M. de Carvalho, D. A. M. Villela, F. C. Coelho, and L. S. Bastos, "Combining probability distributions: Extending the logarithmic pooling approach," 2015, *arXiv:1502.04206*.

[184] D. W. Bunn, "Two methodologies for the linear combination of forecasts," *J. Oper. Res. Soc.*, vol. 32, no. 3, pp. 213–222, Mar. 1981.

[185] D. W. Bunn and E. Kappos, "Synthesis or selection of forecasting models," *Eur. J. Oper. Res.*, vol. 9, no. 2, pp. 173–180, Feb. 1982.

[186] R. E. Barlow, R. W. Mensing, and N. G. Smiriga, "Combination of experts' opinions based on decision theory," in *Reliability and Quality Control*, A. P. Basu, Ed. Amsterdam, The Netherlands: Elsevier, 1986, pp. 9–19.

[187] M. H. DeGroot, "Reaching a consensus," *J. Amer. Statist. Assoc.*, vol. 69, no. 345, pp. 118–121, Mar. 1974.

[188] A. Carvalho and K. Larson, "A consensual linear opinion pool," in *Proc. Int. Joint Conf. Artif. Intell. (IJCAI)*, Beijing, China, 2013, pp. 2518–2524.

[189] M. J. Rufo, C. J. Pérez, and J. Martín, "A Bayesian approach to aggregate experts' initial information," *Electron. J. Statist.*, vol. 6, pp. 2362–2382, Jan. 2012.

[190] G. Casella and R. L. Berger, *Statistical Inference*, 2nd ed. Pacific Grove, CA, USA: Duxbury, 2002.

[191] L. D. Brown, *Fundamentals of Statistical Exponential Families: With Applications in* *Statistical Decision Theory*. Hayward, CA, USA: Institute of Mathematical Statistics, 1986.

[192] L. C. Evans and R. F. Gariepy, *Measure Theory and Fine Properties of Functions*. Boca Raton, FL, USA: CRC Press, 1992.

[193] J. Qiu, Q. Wu, G. Ding, Y. Xu, and S. Feng, "A survey of machine learning for big data processing," *EURASIP J. Adv. Signal Process.*, vol. 2016, Dec. 2016, Art. no. 67.

[194] Y.-H. Kim and B. Pass, "Wasserstein barycenters over Riemannian manifolds," *Adv. Math.*, vol. 307, pp. 640–683, Feb. 2017.

[195] S. Ghosal and A. Van der Vaart, *Fundamentals of Nonparametric Bayesian Inference*. Cambridge, U.K.: Cambridge Univ. Press, 2017.

[196] J. Aczél, *Lectures on Functional Equations and Their Applications*. New York, NY, USA: Academic, 1966.

[197] R. G. Bartle, *The Elements of Integration and Lebesgue Measure*. New York, NY, USA: Wiley, 1995.

[198] W. Rudin, *Real and Complex Analysis*, 3rd ed. New York, NY, USA: McGraw-Hill, 1986.

[199] R. A. Horn and C. R. Johnson, *Matrix Analysis*, 2nd ed. Cambridge, U.K.: Cambridge Univ. Press, 2013.

ABOUT THE AUTHORS

Günther Koliander received the master's degree (honors) in technical mathematics and the Ph.D. degree (honors) in electrical engineering from TU Wien, Vienna, Austria, in 2011 and 2015, respectively.

From 2015 to 2017, he was a Postdoctoral Researcher with the Institute of Telecommunications, TU Wien. He twice held visiting researcher positions at Chalmers University of Technology, Gothenburg, Sweden, and once at Georg-August-Universität Göttingen, Göttingen, Germany. Since 2017, he has been a Postdoctoral Researcher with the Acoustics Research Institute, Austrian Academy of Sciences, Vienna, with a short break in 2021 when he was with the Faculty of Mathematics, University of Vienna, Vienna. His research interests are in the areas of information theory, machine learning, compressed sensing, point processes, and time-frequency and time-scale analysis.



Yousef El-Laham received the B.E. and Ph.D. degrees in electrical engineering from Stony Brook University, Stony Brook, NY, USA.

He is currently with JP Morgan AI Research. His research interests include Bayesian inference for complex systems, Monte Carlo methods, Bayesian filtering, stochastic optimization, synthetic data generation, and probabilistic machine learning.

Dr. El-Laham was a recipient of the Best Student Paper Award at IEEE ICASSP 2019 and the Junior Researcher Award from Stony Brook University's Institute of Advanced Computational Sciences.



Petar M. Djurić (Fellow, IEEE) received the B.S. and M.S. degrees in electrical engineering from the University of Belgrade, Belgrade, Yugoslavia, and the Ph.D. degree in electrical engineering from the University of Rhode Island, Kingston, RI, USA.

He is currently a SUNY Distinguished Professor and the Chair of the Department of Electrical and Computer Engineering, Stony Brook University, Stony Brook, NY, USA.



His research has been in the areas of signal and information processing with primary interests in the theory of machine learning; Monte Carlo-based methods; signal and information processing over networks; signal modeling, detection, and estimation; Radio-Frequency Identification (RFID); and the Internet of Things (IoT).

Dr. Djurić is a Fellow of the European Association for Signal Processing (EURASIP). He was a recipient of the IEEE Signal Processing Magazine Best Paper Award in 2007 and the EURASIP Technical Achievement Award in 2012. He was the Chair of Excellence of Universidad Carlos III de Madrid-Banco de Santander in 2008. He has been on numerous committees of the IEEE Signal Processing Society and many professional conferences and workshops. He was the Editor-in-Chief of IEEE TRANSACTIONS ON SIGNAL AND INFORMATION PROCESSING OVER NETWORKS. From 2008 to 2009, he was a Distinguished Lecturer of the IEEE Signal Processing Society.

Franz Hlawatsch (Fellow, IEEE) received the Diplom-Ingenieur, Dr.Techn., and Univ.-Dozent (Habilitation) degrees in electrical engineering/signal processing from TU Wien, Vienna, Austria, in 1983, 1988, and 1996, respectively.

He is currently an Associate Professor with the Institute of Telecommunications, TU Wien. From 1991 to 1992, as a recipient of an Erwin Schrödinger Fellowship, he spent a sabbatical year with the Department of Electrical Engineering, University of Rhode Island, Kingston, RI, USA. His research interests include statistical and compressive signal processing methods and their application to inference and learning problems.

Dr. Hlawatsch is a Fellow of the European Association for Signal Processing (EURASIP). He was a member of the IEEE SPCOM Technical Committee from 2004 to 2009. He coauthored papers that won the IEEE Signal Processing Society Young Author Best Paper Award and Student Paper Awards at the 36th IEEE International Conference on Acoustics, Speech and Signal Processing and at the 23rd International Conference on Information Fusion. He was the Technical Program Co-Chair of the European Signal Processing Conference 2004. He has served on the technical committees of numerous IEEE conferences. He was an Associate Editor of IEEE TRANSACTIONS ON SIGNAL PROCESSING, IEEE TRANSACTIONS ON INFORMATION THEORY, and IEEE TRANSACTIONS ON SIGNAL AND INFORMATION PROCESSING OVER NETWORKS.

