Pooling, Voting, and Bayesian Updating

Jan-Willem Romeijn
University of Groningen

Abstract

This paper explores the fact that linear pooling can be represented as a Bayesian update on the opinions of others. It uses this fact to relate pooling to voting theory, especially Condorcet's jury theorem. Relative to certain modelling assumptions the trust parameter from pooling can be equated with the so-called truth-conduciveness of the jurors featuring in Condorcet's result.

1 Introduction

Say that Raquel and Quassim are both pondering over the proposition S. Raquel's belief is $P_R(S) = r$, Quassim's is $P_Q(S) = q$. How can Raquel respond to the opinion of Quassim? One influential model for updating opinions in the light of disagreement, due to Stone (1961), is linear opinion pooling.[1] Linear pooling determines that the posterior opinion of Raquel $P'_R(S)$ is given by

$$r' = wq + (1 - w)r. \tag{1}$$

[1]The model dates back at least to French (1956) and has been developed by DeGroot (1974) and numerous others. For many philosophers the classical treatment is Lehrer and Wagner (1981). Genest and Zidek (1986) offer an extensive mathematical review. More recent philosophical contributions are Dietrich (2010) and Steele (2012).

The parameter $w \in [0, 1]$ specifies to what extent the updated opinion of Raquel will move towards that of Quassim. By way of interpretation: it measures the trust that Raquel has in Quassim.

An entirely different model for Raquel's accommodating Quassim's opinion is based on Bayesian updating. Quassim's opinion is in that case taken to be evidence, and Raquel accommodates this evidence by a Bayesian update. A rigourous theory of how Bayesians can treat the opinions of others was first given in the context of game theory by Harsanyi (1966/67).[2] In what follows it is taken for granted that opinions of Raquel and Quassim can be captured coherently in an algebra, even while those opinions are also expressed in probability assignments over the algebra.

In the Bayesian model we represent the fact that Raquel believes S to degree r, or $P_R(S) = r$ for short, by an element from this algebra, $\ulcorner r \urcorner$. To indicate that Raquel knows her own opinion we write $P_R(S|\ulcorner r \urcorner) = r$. Similarly, we represent $P_Q(S|\ulcorner r \urcorner) = q$, i.e., that Quassim believes S to degree q while knowing that Raquel's degree of belief is r, by the element $\ulcorner q \urcorner$. We assume that Quassim already knows Raquel's opinion for reasons that will become apparent below. We can express Raquel's belief upon learning Quassim's by Bayes' rule:

$$P'_R(S) = P_R(S|\ulcorner q \urcorner \cap \ulcorner r \urcorner) = P_R(S|\ulcorner r \urcorner) \frac{P_R(\ulcorner q \urcorner | S \cap \ulcorner r \urcorner)}{P_R(\ulcorner q \urcorner | \ulcorner r \urcorner)} \tag{2}$$

The second equality is of course Bayes' theorem. If Bayes' rule is adhered to, Raquel's belief after learning Quassim's opinion is fully determined by her prior belief $P_R(S|\ulcorner r \urcorner)$ and the likelihoods $P_R(\ulcorner q \urcorner | S \cap \ulcorner r \urcorner)$.

This paper develops the relation between these two models of updating. It is known that we can provide a Bayesian model of pooling and so make equations (1) and (2) match. Genest and Schervish (1985) have shown how to orchestrate the Bayesian model such that

$$P_R(S|\ulcorner q \urcorner \cap \ulcorner r \urcorner) = r'.$$

[2]Brandenburger (2007) offers an accessible exposition of so-called interactive epistemology, in particular of the basic idea of Harsanyi type spaces.

For mathematicians the formal equivalence of pooling and Bayesian updating may therefore be old hat. However, the result has not enjoyed much attention from a more conceptually engaged audience. Among philosophers more attention has been devoted to how Bayesian updating and DeGroot pooling may be combined, and to the question when the order of these operations matter (cf. Genest et al, 1986; Dietrich, 2010; Leitgeb, 2014). As this paper argues, the relation between pooling and Bayesian updating brings conceptual and interpretative insights that are worth exploring.

In particular, this paper shows that an explicit reconstruction of pooling in terms of a Bayesian update illuminates the vexed issue of interpreting the trust parameter w (see also Genest and McConway, 1990). As things stand, the trust parameter is an exogenous component of the model of opinion pooling: it must be set by hand. But this is odd, because the model itself concerns beliefs and it might be expected that trust can be spelled out in terms of beliefs as well. As we will see, the trust parameter of pooling is naturally related to the competence parameters that we find in the context of Condorcet's jury theorem, in particular to beliefs about the truth-conduciveness of others.

2 Linear pooling as Bayesian updating

This section presents a simplified version of the result of Genest and Schervish (1985), which establishes a general equivalence of pooling and Bayesian updating. The result is not well-known among philosophers, which justifies a brief presentation here.[3] The primary concern in the paper by Genest and Schervish is to specify conditions under which some form of pooling may serve as a stand-in for a Bayesian update, in case we do not have a full probability assignment over all the opinions available. Among other things, Genest and Schervish il-

[3] In fact the theorem below was at first proved independently from Genest and Schervish. I am indebted to Carl Wagner for pointing me to the earlier result.

luminate how both pooling and updating might deal with dependencies among multiple expert opinions.[4]

The theorem below is restricted to two agents and to linear pooling. This will help to focus attention on the conceptual links between pooling and voting. As above, we represent the opinions of Quassim and Raquel with the events $\ulcorner q \urcorner$ and $\ulcorner r \urcorner$. It is given in the setup that Quassim is going to reveal his belief regarding S, so the events $\ulcorner q \urcorner$ form a partition. The theorem below establishes that we can always find likelihoods $P_R(\ulcorner q \urcorner | S \cap \ulcorner r \urcorner)$ such that, after updating on $\ulcorner q \urcorner$, Raquel's belief in S equals the result of pooling. In other words, linearly pooling r and q with weight w can always be represented by a Bayesian update.

THEOREM (AFTER GENEST AND SCHERVISH 1985)
Let $P_R(S) = P_R(S | \ulcorner r \urcorner) = r$, $P_Q(S | \ulcorner r \urcorner) = q$, and let $P'_R(S) = r' = wq + (1-w)r$ be the result of linearly pooling these opinions. If we have

$$P_R(\ulcorner q \urcorner | S \cap \ulcorner r \urcorner) = g(q,r) \left(1 - w + \frac{w}{r} q\right) \quad (3)$$

$$P_R(\ulcorner q \urcorner | \neg S \cap \ulcorner r \urcorner) = g(q,r) \left(1 + \frac{r}{1-r} w - \frac{w}{1-r} q\right) \quad (4)$$

where $g(q,r)$ and its integral $G(q,r)$ are such that

$$\int_0^1 g(q,r) dq = 1 \quad (5)$$

$$\int_0^1 G(q,r) dq = 1 - r \quad (6)$$

then the Bayesian update on $\ulcorner q \urcorner$ is identical to the update by linear pooling, $P_R(S | \ulcorner q \urcorner \cap \ulcorner r \urcorner) = r' = P'_R(S)$.

PROOF
To derive equations (3) and (4) we write Bayes' theorem in the following way:

$$P_R(\ulcorner q \urcorner | S \cap \ulcorner r \urcorner) = P_R(\ulcorner q \urcorner | \ulcorner r \urcorner) \frac{P_R(S | \ulcorner q \urcorner \cap \ulcorner r \urcorner)}{P_R(S | \ulcorner r \urcorner)} \quad (7)$$

Notice that in the update of equation (2), we derive the posterior of S from its prior and its likelihood for the event $\ulcorner q \urcorner$. Here we are deriving the likelihood

[4]For an illuminating conceptual take on these issues, see Bradley (2006, section 4).

from the prior and the posterior, $P_R(S|\ulcorner r\urcorner) = r$ and $P_R(S|\ulcorner q\urcorner \cap \ulcorner r\urcorner) = wq + (1-w)r$ respectively. Substituting these formulas in equation (7) we obtain equation (3):

$$P_R(\ulcorner q\urcorner | S \cap \ulcorner r\urcorner) = g(q,r)\left(1 - w + \frac{w}{r}q\right),$$

where we have abbreviated $P_R(\ulcorner q\urcorner|\ulcorner r\urcorner) = g(q,r)$. Equation (4) can be derived in a similar manner.

Our next task is to derive the constraints on the function $g(q,r)$, as given in equations (5) and (6). Because the events $\ulcorner q\urcorner$ form a partition, we must have that

$$\int_0^1 P_R(\ulcorner q\urcorner|\ulcorner r\urcorner)dq = \int_0^1 g(q,r)dq = 1.$$

This is constraint (5). Furthermore, by the law of total probability we have

$$P_R(S|\ulcorner r\urcorner) = \int_0^1 P_R(\ulcorner q\urcorner|\ulcorner r\urcorner)P_R(S|\ulcorner q\urcorner \cap \ulcorner r\urcorner)\, dq \tag{8}$$

or equivalently,

$$r = \int_0^1 g(q,r)\,(wq + (1-w)r)\; dq.$$

Notice that a companion constraint, that can be based on a similar expression for $P_R(\neg S|\ulcorner r\urcorner) = 1 - r$, is entailed by the constraint (5) in conjunction with (8).

We can employ a so-called partial integration to transform equation (8) into a constraint on $g(q,r)$ alone. Using $G(q,r)$ for the integral of $g(q,r)$, and setting $G(1,r) = 1$ and $G(0,r) = 0$ to satisfy constraint (5), we can write

$$\begin{aligned} r &= \int_0^1 g(q,r)\,(wq + (1-w)r)\; dq \\ &= (w + (1-w)r)\,G(1,r) - (1-w)rG(0,r)\; - \int_0^1 w\,G(q,r)dq \\ &= (w + r - wr)\; - \; w\int_0^1 G(q,r)dq. \end{aligned}$$

With some algebra we obtain

$$\int_0^1 G(q,r)dq = 1 - r.$$

And this is precisely constraint (6). □

The constraints (5) and (6) still leave a considerable amount of freedom for the likelihood functions of equations (3) and (4). A uniform distribution cannot be made to fit but many simple non-negative functions $g(q, r)$ can. Following Genest and Schervish (1985) and Bonnay and Cozic (2014), the constraint (6) can also be written as a constraint on the expectation value for q:

$$\begin{aligned} r &= \int_0^1 g(q,r)\,(wq + (1-w)r)\ dq \\ &= (1-w)r\,(G(1,r) - G(0,r))\ +\ w\int_0^1 q\,g(q,r)dq \\ &= (1-w)r\ +\ w\int_0^1 q\,g(q,r)dq. \end{aligned}$$

so that

$$\int_0^1 q\,g(q,r)dq = r.$$

The interpretation is rather natural: in a Bayesian rendering of pooling, Raquel's distribution for Quassim's opinion is centred on her own. And this seems exactly right. Because she puts some trust in Quassim, any deviation in expectations would force Raquel to adapt her opinion even before she learns his opinion.

To get a sense of what $g(q, r)$ may look like, notice that Raquel may use a Beta-distribution centred on r:

$$P_R(\ulcorner q \urcorner | \ulcorner r \urcorner) = \frac{\Gamma(n)}{\Gamma(rn)\Gamma((1-r)n)} q^{rn-1}(1-q)^{(1-r)n-1},$$

where n is a free parameter and the Γ function is an extension of the factorials. If $rn > 1$ and $(1-r)n > 1$, the distribution $P_R(\ulcorner q \urcorner | \ulcorner r \urcorner)$ shows a peak at r, expressing that Raquel thinks it more probable that Quassim's opinion sits close to her own, while assigning decreasing probabilities to Quassim's opinion being further removed from hers. Other members from the versatile family of Beta-distributions can model yet other expectations but of course nothing hinges on using specifically this family of functions.

If these constraints seem too restrictive, we might consider dropping the assumption that the events $\ulcorner q \urcorner$ form a partition. Quassim could for instance return a blank when Raquel asks him for his opinion on S. This allows us to lift the constraints (5) and (6) from the Bayesian reconstruction of pooling. Rather than satisfying those, we can tweak the probability assignments for returning a blank to make updating match pooling. For example, we may choose $g(q, r) = 1$ and obtain linear likelihood functions that are more easy to interpret. However, the standard story on pooling is that agents do not return blanks and reveal real-valued degrees of belief.

With the Bayesian underpinning of pooling firmly in place, we now turn to its possible uses. First off, the result may serve to justify pooling as a method of belief change. There are well-known defenses of the Bayesian model of belief change, and because pooling fits this model, it may fall back on these defenses. Furthermore, we may motivate applications of opinion pooling by pointing to the probabilistic assumptions that the Bayesian model brings out. If those assumptions are met, then arguably pooling is warranted. And this is a useful conclusion for agents with limited computational resources: under specific constraints pooling can be taken as an epistemic shortcut.

Perhaps it is not very surprising that there is a representation of pooling in terms of Bayesian updates. There are several results showing the wide range of belief update operations that can be subsumed under a Bayesian header (cf. van Fraassen, 1989). Specifically pooling may be linked to well-known results on the reflection principle, in virtue of their formal similarity: setting one's opinion to a linear combination of opinions of others is much like setting one's prior opinion to a linear combination of possible posterior opinions.[5] However, our concern here is not with the formal link between opinion pooling and Bayesian updating itself. This paper investigates how this link may be employed to link pooling and voting, and thus help to interpret aspects of these models of social deliberation.

More in particular, the following sections will offer an interpretation of the notion of trust, which is central to pooling, in terms of beliefs. For a general impression, consider the role of the trust w in the likelihood functions (3) and (4), which determine how Quassim's opinion impacts on Raquel's. If Raquel

[5] I owe this suggestion to Richard Bradley.

puts some trust in Quassim and hence chooses $w > 0$, then she finds it more probable that Quassim will report a degree of belief in S higher than her own if indeed S is true, than if S is false. Similarly, she finds it more probable that Quassim reports a degree of belief in S lower than her own if indeed S is false, than if S is true. The trust parameter thus relates to the ability that Raquel takes Quassim to have in determining the truth or falsity of S.

3 Truth-conduciveness in voting

In what follows we briefly review Condorcet's result and a Bayesian reformulation of it. This will bring out a particular aspect, namely the truth-conduciveness of jurors, that is crucial for the adequacy of a jury verdict. In the following section it is then shown that the trust parameter w in pooling can be equated with a parameter that measures truth-conduciveness.

Consider the setting of the original jury theorem of Condorcet ([omitted for blind reviewing]). Jurors are asked for a categorical vote on a proposition S. It is assumed that jurors are competent, meaning that if S is true, they are more likely to vote in support of it while if S is false they are more likely to vote against. We denote the event of voting for or against S by V and $\neg V$ respectively. This can be expressed by the values of two competence parameters, c_S and $c_{\neg S}$:

$$
\begin{aligned}
c_S &= P(V|S) > \tfrac{1}{2}, \\
c_{\neg S} &= P(\neg V|\neg S) > \tfrac{1}{2}.
\end{aligned}
$$

In words, the assumption of juror competence is that they are better than a fair coin in determining the truth or falsity of S. The original result of Condorcet is that for competent jurors, the majority vote of the jury will tend to the truth with increasing jury size. More precisely, if S is true then for ever larger juries the probability that more jurors vote for S than against will tend to 1, and similarly for S being false. The theorem is essentially a version of the law of large numbers.

There is also a Bayesian version of Condorcet's result, to the effect that the posterior probability of the true proposition, either S or $\neg S$, will tend to 1 when updating on ever more juror votes. Importantly, to arrive at this result we need not assume that the jurors are competent, i.e., that they perform better than a fair coin. The assumption needed here is that the jurors are truth-conducive: if S is true it is more probable that jurors cast a vote for S than that they cast a vote against it, and conversely if S is false. Formally, we may express this requirement as follows:

$$\Delta \;=\; P(V|S) - P(V|\neg S) = c_S + c_{\neg S} - 1 > 0 \qquad (9)$$

where Δ is defined as a measure of truth-conduciveness. Notice that this requirement entails the truth-conduciveness for $\neg S$ as well because

$$P(\neg V|\neg S) - P(\neg V|S) = P(V|S) - P(V|\neg S).$$

The point of truth-conduciveness is that a vote V will boost the probability of S, while a vote $\neg V$ will boost the probability of $\neg S$. It may still be that a vote V boosts the probability of S much less than that a vote $\neg V$ boosts the probability of $\neg S$. One of the two competences c_S and $c_{\neg S}$ might thereby drop below half.[6]

With the notion of truth-conduciveness in place, we can now turn to the central result of this paper, concerning the relation between voting and opinion pooling. We approach this relation by representing voting in terms of the more fine-grained Bayesian model which was set up to represent pooling. This will lead to an expression for the truth-conduciveness of voters in which the trust parameter of pooling is the dominant term.

[6]See footnote 1 of [omitted for blind reviewing] for a general expression of what supermajority is needed to get the required limiting behaviour, relative to the values of c_S and $c_{\neg S}$.

4 Trust as truth-conduciveness

To capture voting in the more fine-grained Bayesian model, we view Quassim as the sole member of a jury that advises Raquel. But rather than taking his categorical vote at face value, we now frame Quassim's vote in terms of the probabilistic opinions that he might have. That is, we make a more fine-grained model of the opinion that Quassim is expressing, in order to seek a connection with the Bayesian model of pooling.

Specifically, we identify the votes V and $\neg V$ with distinct events in the Bayesian model of pooling. Recall that we assumed that Quassim knows Raquel's opinion, $P_Q(S|\ulcorner r \urcorner) = q$. Accordingly we may suppose that Quassim votes in favour or against S with an eye on the effect that this has on Raquel's degree of belief. More precisely, he votes in favour if his opinion is higher than hers, and against if his degree of belief is lower. We denote these two events by $\ulcorner q > r \urcorner$ and $\ulcorner q < r \urcorner$ respectively, and we stipulate that the events V and $\neg V$ run parallel to $\ulcorner q > r \urcorner$ and $\ulcorner q < r \urcorner$.

The central result of this paper employs this parallel. The idea is that Raquel conceives of Quassim's vote as a shorthand for a range of probabilistic opinions that he might have. Raquel sees the event V as a shorthand for Quassim expressing the opinion $\ulcorner q > r \urcorner$, and similarly for the event $\neg V$ expressing $\ulcorner q < r \urcorner$. Raquel then accommodates Quassim's opinion much like she accommodates a categorical vote, namely by a Bayesian update. In the remainder of this section we investigate what this leads to, if we fill in the details of this update by means of the Bayesian model of pooling developed above, i.e., by means of the likelihoods (3) and (4) and the constraints (5) and (6).

Before we work out the details, let us briefly pause and consider the philosophical motivation for representing voting in the more fine-grained Bayesian model. The parallel between the categorical vote and the more fine-grained expression of a range of probabilities may be considered insufficiently motivated, or even somewhat contrived. But remember what the goal of this representation is: to interpret the trust parameter that is used in pooling in terms of beliefs.

Now the Bayesian representation of pooling already offers an interpretation of sorts: it shows up as the slope of the likelihood function. But in what follows, we will see that the fine-grained Bayesian representation of voting occasions a far more natural interpretation of the trust parameter. This motivates the representation of voting just given. We do not need to claim that that the representation of voting is unique or forced on us in any sense. It suffices that the representation follows from reasoned choices, and that it offers a particular voting-related perspective on the Bayesian model of pooling.

We now provide the formal details of the relation between voting and pooling. To fully specify the update that Raquel performs, we need to determine the distribution $P_R(\ulcorner q \urcorner | \ulcorner r \urcorner)$, expressing Raquel's expectations about Quassim's opinions. Raquel might choose the following step function:

$$P_R(\ulcorner q \urcorner | \ulcorner r \urcorner) = \begin{cases} l & \text{if } q < \epsilon r \\ h & \text{if } q > 1 - \epsilon(1-r) \\ 0 & \text{else.} \end{cases} \tag{10}$$

where $0 < \epsilon \leq 1$. If we take $\epsilon = 1$, Raquel distinguishes between Quassim offering a lower or a higher degree of belief in S, but within these two ranges the distribution is uniform. For small ϵ, Raquel takes Quassim to offer a probability below ϵr and hence close to zero, or above $1 - \epsilon(1-r)$ and hence close to one. A distribution with small ϵ expresses that Quassim's opinions, though probabilistic, resemble categorical votes.

If we want the Bayesian update of Raquel to emulate pooling, the distribution $P_R(\ulcorner q \urcorner | \ulcorner r \urcorner)$ must comply to the constraints (5) and (6). Solving for these constraints yields:

$$l = \frac{1-r}{\epsilon r}, \qquad h = \frac{r}{\epsilon(1-r)} \tag{11}$$

Notice that this entails

$$P_R(\ulcorner q < r \urcorner | \ulcorner r \urcorner) dq \ = \ 1 - r. \tag{12}$$

The constraints entail that Raquel finds it more probable that Quassim will report a value q that is lower than her opinion when she is herself assigning a low probability r to S. This is not specific for the distribution above. To comply to pooling, Raquel must always expect that Quassim's opinions will reinforce what she is already thinking about S.

Together with the likelihoods of equations (3) and (4), the step function (11) offers a Bayesian representation of Raquel's pooling operation, if she were given a sharp probability by Quassim. However, Quassim merely offers his opinion in the form of a categorical vote, which is then interpreted as either of the events $\ulcorner q > r \urcorner$ or $\ulcorner q < r \urcorner$. Fortunately the Bayesian representation still allows us to compute an updated probability for Raquel, using the marginal likelihoods of S and $\neg S$ for the event $\ulcorner q > r \urcorner$, to wit, $P(\ulcorner q > r \urcorner | S \cap \ulcorner r \urcorner)$ and $P(\ulcorner q > r \urcorner | \neg S \cap \ulcorner r \urcorner)$, as well as the marginal likelihoods for the event $\ulcorner q < r \urcorner$.

On the basis of the foregoing we now determine these marginal likelihoods. For the event $\ulcorner q > r \urcorner$ we find

$$\begin{aligned} P(\ulcorner q > r \urcorner | S \cap \ulcorner r \urcorner) &= (1-w)r + w - (1-r)w\frac{\epsilon}{2}, \\ P(\ulcorner q > r \urcorner | \neg S \cap \ulcorner r \urcorner) &= (1-w)r + rw\frac{\epsilon}{2}. \end{aligned}$$

Similar equations can of course be derived for the event $\ulcorner q < r \urcorner$. Now observe that

$$P(\ulcorner q > r \urcorner | S \cap \ulcorner r \urcorner) - P(\ulcorner q > r \urcorner | \neg S \cap \ulcorner r \urcorner) = w\left(1 - \frac{\epsilon}{2}\right) \tag{13}$$

and equivalently for $\ulcorner q < r \urcorner$. In words, we find that a certain measure for the impact that the event $\ulcorner q > r \urcorner$ has on Raquel's opinion of S and $\neg S$, namely the difference of the likelihoods for these events, is proportional to the trust parameter. And if we take the limit of ϵ to zero, this proportionality is replaced by an equality. We have thus arrived at a clean expression of the trust parameter of pooling in terms of beliefs.

Recall that in the context of voting, the truth-conduciveness of equation (9) showed up as a crucial quantity. It will be clear that equation (13), pertaining to trust in the context of a Bayesian reconstruction of pooling, shows similarity to equation (9) concerning truth-conduciveness. Drawing on the parallel between votes V and $\neg V$ on the one hand, and opinions $\ulcorner q > r \urcorner$ and $\ulcorner q < r \urcorner$ on the other, we can talk about truth-conduciveness in the context of pooling as well:

$$\Delta = P(\ulcorner q > r \urcorner | S \cap \ulcorner r \urcorner) - P(\ulcorner q > r \urcorner | \neg S \cap \ulcorner r \urcorner).$$

In words, we label the difference between the likelihoods of S and $\neg S$ for the event $\ulcorner q > r \urcorner$ as the truth-conduciveness that Raquel attributes to Quassim. Hence we have $\Delta = w\big(1 - \frac{\epsilon}{2}\big)$. For diminishing values of ϵ, the match between truth-conduciveness and trust becomes exact:

$$\Delta = w \tag{14}$$

That is to say, we can think of the trust parameter w as a measure of the truth-conduciveness that Raquel attributes to Quassim. When viewing a vote as a coarse-grained version of pooling, the trust parameter from pooling shows up as the truth-conduciveness from voting.

Of course we can contest the result of equation (14), by challenging specific choices that were made in forging the link between the models of voting and pooling. For starters, the choice of a step function as distribution over $\ulcorner q \urcorner$ seems rather arbitrary. But note that the exact shape of the distribution looses import if we take ever smaller values of ϵ. In fact in the limit, when taking more and more probability mass towards the extremes, nothing hinges on the exact shape of the distribution.

Another main criticism was already touched on in the foregoing: the representation of votes in terms of the Bayesian model of pooling is insufficiently motivated. In addition to what was argued above, we can now point to characteristics of the model for small ϵ. There are well-known ties between probabilistic and categorical belief that employ the so-called Lockean thesis: above a certain threshold for the probability, categorical belief is warranted. By choosing ϵ small enough Quassim's opinion will pass the thresholds, so that categorical votes V and $\neg V$ are indeed adequate representations. For small ϵ, it makes sense to draw a parallel between V and the event $\ulcorner q > 1 - \epsilon(1 - r) \urcorner$, and between $\neg V$ and $\ulcorner q < \epsilon r \urcorner$. It seems natural to represent a categorical voter in terms of the distribution of equation (10) if we choose small values of ϵ.

There may well be other criticism of how we forged the link between voting and pooling. But the general message can be appreciated even if some of the details of the link are deemed problematic. In a Bayesian model of how an agent accommodates the trustworthy opinion of someone else, as for example in the setting of Condorcet, the likelihoods express a notion of truth-conduciveness. If we construct a Bayesian model of pooling, the trust parameter shows up as a characterstic of the likelihood functions in, by and large, the same way. Certain expressions involving the trust parameter can be interpreted as expressing this truth-conduciveness, and this helps us to think about the trust parameter in a way that relates to beliefs.

5 Conclusions

What precisely is gained now that we have a formal relation between voting, pooling, and Bayesian updating? The main result is that the trust parameter from pooling can be given a natural interpretation by means of the Bayesian model, and by means of locating categorical voting in it. This seems a gain from the point of view of representational economy. As intimated earlier, it is odd that a model of belief change would employ a trust parameter that is not itself representable in terms of the beliefs that are being modelled. The trust should somehow be located among those beliefs. This paper shows that the Bayesian representation of pooling allows us to identify trust with intuitive epistemic notions.

Research that exploits the Bayesian representation of pooling does not stop here. There is a vast body of literature that can fruitfully be related to pooling because pooling has been given a Bayesian reconstruction. Much of this literature goes under the headers of interactive epistemology and epistemic game theory (Brandenburger, 2007; Pacuit and Roy, 2014). The present paper has made only a modest start with that. Interactive epistemology was developed mostly in the context of game theory, to model players of a game who reason about each others opinions. It seems that the relevance of this literature has not been fully appreciated by all philosophers working on social epistemology. I suggest that the current result may stimulate research at the intersection of these disciplines.

One natural application of the foregoing concerns a debate in traditional epistemology. Since a number of years epistemologists have been interested in disagreement among peers (Christensen and Lackey, 2010). With some exceptions (e.g., Russell et al, 2015), this debate has relatively few points of contact with the literature in statistics and epistemic game theory. The current paper may serve a constructive role in facilitating more extensive contact. By way of example, an influential view on how to resolve disagreement employs the strategy of splitting-the-difference: when two agents have different degrees of belief about some proposition S, then they can resolve their disagreement by both adopting a, possibly weighted, average of the two degrees, i.e., by a pooling operation. Many have wondered what may justify this intuitive response (Jehle and Fitelson, 2009). The Bayesian reconstruction of pooling may provide such a justification, by stating precisely what beliefs the agents must have for splitting-the-difference to be probabilistically coherent.

A possibly more exciting application concerns the existing models for describing deliberation and consensus formation. DeGroot (1974) and Lehrer and Wagner (1981) describe conditions under which iterated pooling leads to consensus. It seems rather natural that such results can be related to Aumann's well-known agreement theorem, and Aumann himself indeed thought of his result as providing an underpinning for iterated pooling (Aumann, 1976, p. 1238). The result of Aumann is static, i.e., it does not mention belief change. But Genneakoplos and Polemarchakis (1982) offer a dynamic formulation of the same result, and [omitted for blind reviewing] provide the construction of a prior probability assignment such that repeated Bayesian updates on the revealed opinions of others replicates the consensus formation process effected by iterated pooling. That result relies in part on the results of this paper. In fact the current results were arrived at in the context of work on the relation between Aumann and DeGroot.

Once a relation has been established between the agreement theorem and consensus formation through repeated pooling, this may invite further research into pooling as a form of information sharing, and into failures of consensus as resulting from differences in the priors of different agents. In general, the fact that pooling is a short-hand for a more fine-grained Bayesian update may help to mitigate doubts over the nature and justification of this procedure, and illuminate the presuppositions of iterated opinion pooling as a dynamics for social deliberation.

Acknowledgements

Omitted for the purpose of blind reviewing.

References

Aumann, R. J. (1976), Agreeing to disagree, Annals of Statistics 4(6), 1236–1239.

Bonnay, D. and M. Cozic (2014), Weighted Averaging, Bayesian update, and Jeffrey Conditioning, manuscript available from the authors.

Bradley, R. (2007), Taking Advantage of Difference in Opinion, Episteme: A Journal of Social Epistemology 3(3), 141–15.

Brandenburger, A. (2007), The power of paradox, International Journal of Game Theory 35, 465–492.

Christensen, D. and J. Lackey (2010), The Epistemology of Disagreement, Oxford University Press.

Russell, J. S. , J. Hawthorne and L. Buchak (2015), Groupthink, Philosophical Studies 172(5), 1287–1309.

DeGroot, M. (1974), Reaching a consensus, Journal of the American Statistical Association 69, 118–121.

Dietrich, F. (2010), Bayesian group belief, Social Choice and Welfare 35(4), 595–626.

French, John R. Jr. (1956), A formal theory of social power, Psychological Review 63(3), 181–194.

Geanakoplos, J. D. and H. M. Polemarchakis (1982), We Cannot Disagree Forever, Journal of Economic Theory 28, 192–200.

Genest, C. and M. J. Schervish (1985), Modelling Expert Judgments for Bayesian Updating, Annals of Statistics 13(3), 1198–1212.

Genest, C. , K. J. McConway, and M. J. Schervish (1986), Characterization of Externally Bayesian Pooling Operators, Annals of Statistics 14(2), 487–501.

Genest, C. and K. J. McConway (1990), Allocating the Weights in the Linear Opinion Pool, Journal of Forecasting 9, 53–73.

Genest, C. and J. V. Zidek (1986), ?Combining Probability Distributions: A Critique and Annotated Bibliography?, Statistical Science 1(1), 113–135.

Harsanyi, J. (1966-67), Games of incomplete information playes by Bayesian players, Parts I-III, Management Science 14, 159–182, 320-334, 486–502.

Jehle, D. and B. Fitelson (2009), What is the 'Equal Weight View'?, Episteme, 6, 280–293.

Lehrer, K. and C. Wagner (1981), Rational Consensus in Science and Society: A Philosophical and Mathematical Study, Springer Verlag.

Leitgeb, H. (2014), Imaging all the people, manuscript available from the author.

Pacuit, E. and O. Roy (2014), Interactive Rationality, manuscript available from the authors.

Steele, K. (2012), Testimony as Evidence: More Problems for Linear Pooling, Journal of Philosophical Logic 41, 983–999.

Stone, M. (1961), The linear opinion pool, Annals of Mathematical Statistics 32, 1339–1342.

van Fraassen, B. C. (1989), Laws and Symmtery, Clarendon Press, Oxford.