

6-1-2013

Group level interpretations of probability : new directions

Darrell Patrick ROWBOTTOM
Lingnan University, Hong Kong

Follow this and additional works at: http://commons.ln.edu.hk/sw_master



Part of the [Philosophy Commons](#)

Recommended Citation

Rowbottom, D. P. (2013). Group level interpretations of probability: New directions. *Pacific Philosophical Quarterly*, 94(2), 188-203.
doi: 10.1111/j.1468-0114.2012.01451.x

This Journal article is brought to you for free and open access by the Lingnan Staff Publication at Digital Commons @ Lingnan University. It has been accepted for inclusion in Staff Publications by an authorized administrator of Digital Commons @ Lingnan University.

Group Level Interpretations of Probability: New Directions

Darrell P. Rowbottom

In this paper, I present some new group level interpretations of probability, and champion one in particular: a consensus-based variant where group degrees of belief are construed as agreed upon betting quotients rather than shared personal degrees of belief. One notable feature of the account is that it allows us to treat consensus between experts on some matter as being on the union of their relevant background information. In the course of the discussion, I also introduce a novel distinction between intersubjective and interobjective interpretations of probability.

1. Subjective, Logical/Objective, and Intersubjective Probabilities

Probabilities are at the heart of our best contemporary accounts of reasoning in science and elsewhere, e.g. Jaynes (2003), Bovens and Hartmann (2004), Jeffrey (2004), Howson and Urbach (2005), and Williamson (2005, 2010). But nowadays they are almost always understood as individual degrees of belief which obey particular (rationality) constraints, although this runs the risk of making ‘rationality’, and hence confirmation, far too subjective a matter.

Almost everyone who thinks of probabilities as individual degrees of belief agrees that no-one should allow a Dutch Book to be made against them; that is, exhibit betting behaviour which means that they can lose, by force, no matter what actually happens. And it turns out that having one’s degrees of belief obey the axioms of

probability ensures that this will not occur. But obeying this constraint, alone, allows for an expansive range of rational difference of opinion.

Here's a simple example. Imagine I think that if I am awarded a Chair in philosophy, then this is a strong indication that God exists. Let p be 'God exists', and q be 'I am awarded a Chair in philosophy'; and let D denote a degree of belief and P a (subjective) probability. For me, $D(p|q) = 0.999$, which means that if I come to believe q then I'll be pretty sure about p . But for this also to be a probability, it need only obey the constraints of probability theory. So provided $D(\sim p|q) = 0.001$, and so forth, my degree of belief is a rational one; $D(p|q) = P(p|q) = 0.999$. For a jealous peer, however, q might strongly indicate that p is false; and for her, $P(p|q)$ would be 0.001. What's more, and this is the crucial point, we would both be rational *even if we shared precisely the same background information*.¹

Earlier philosophers, notably Keynes (1921) and Carnap (1962), avoided this difficulty by instead taking probabilities to be objective relations between propositions. As is now generally accepted, however, the logical approach suffers from a serious fundamental problem of its own. How do we, or can we, recognise and/or calculate these alleged logical relations? Application of the principle of indifference is notoriously problematic, as shown by the paradoxes of Bertrand (1960).² Van Fraassen (1989, pp. 314–316) and Gillies (2000, pp. 48–49) both defend the current majority view, namely that the principle of indifference is little more than a heuristic.

Recall the overarching problem. Even if one's degrees of belief obey the axioms of probability, this doesn't appear to entail that they are rational; on the contrary, they may seem arbitrary despite these constraints. One potential way to solve this problem is to say that they ought to map onto objective logical relations, but this path appears to be doomed to failure. So what are the remaining options?

The first, which is the more popular, is to 'top up' the subjective view of probability by introducing further constraints, which are required of *personal degrees of belief*, for rationality. So-called 'Objective Bayesian' approaches, such as those of Jaynes (2003) and Williamson (2005, 2010), work in this way. The latter, for instance, introduces two extra constraints for degrees of belief *qua* mental entities (Williamson 2006, §13): (i) they should obey any constraints imposed by background information (e.g. reflect knowledge of frequencies), yet (ii) be otherwise maximally non-committal (i.e. have maximum entropy). This way of doing things means that there is only one rational degree of belief in finite cases, although less strict 'top up' accounts, e.g. which only impose empirical constraints, are also possible. The worry with the Objective Bayesian approach is that it has similar problems to those exhibited by the logical interpretation; indeed, Rowbottom (2008b) argues that the main difference between the two interpretative strategies lies in the philosophy of logic (and concerns Platonism vs. nominalism within the realm of weak psychologism).

The second, which is relatively unexplored, is to take an intersubjective approach, intended to serve 'as an intermediate between the logical interpretation of the early Keynes and the subjective interpretation of his critic Ramsey' (Gillies 2000, p.174). The basic idea behind this theory, first proposed by Gillies (1991), is that a group can

have a Dutch Book made against it if its members do not have the same degree of belief assignments, and those assignments do not satisfy the probability calculus. Imagine a married couple who are avid gamblers with pooled financial resources. Romeo bets £100 that it will rain tomorrow in Oxford, at even odds. But unbeknownst to Romeo, Juliet has already bet £150, at three to one on, that it will not rain tomorrow in Oxford. The couple are in trouble! They are ‘out’ £250. But no matter whether it rains in Oxford or not tomorrow, they will only get £200 back. Naturally, more would have to be true for us to think they were behaving irrationally. At the bare minimum, they would need some means to communicate in order to co-ordinate their betting strategy, as well as shared interests (in not losing any of the pooled financial resources for no gain in utility). But I will come back to this later.

I have covered this ground to explain why it is worth considering how we can understand group level probabilities; because doing so may hold the key to solving an intractable problem which bears on confirmation theory and scientific method, among other things. A key worry is about how to interpret the probabilities that appear in confirmation functions, such as those preferred by Milne (1996) and Huber (2008). If we go subjective, then confirmation seems merely psychological. If we go intersubjective—or better still, as I will later explain, interobjective—then we may be able to put confirmation back on a less arbitrary footing. In Rowbottom (2008a), I began to champion this move.

Moreover, I might point to developments in philosophy of science in the past few decades—for example, the focus on the division of cognitive labour by Kitcher (1990, 2003), Strevens (2003), Sarkar (2007), and Rowbottom (2011a)³—and to the

blossoming of social epistemology. There is considerable movement toward focusing on group properties, rather than individual properties, when considering the rationality of science. And taking an interest in group degrees of belief and group probabilities—in order to consider how the views expressed by scientific research teams affect one another, and interact with the beliefs of their members, for instance—is a natural development. At the very least, the route deserves more careful investigation. Is it a blind alley, or might it lead us into a pleasant new land?

In order to be in a position to answer this question, we first need to have a clear idea of how we can understand group level probabilities. That is the primary purpose of this paper. The secondary purpose is to advocate one understanding in particular.

In what follows, I will first outline some variants of the intersubjective view with reference to possible understandings of the notion of group degrees of belief. (I will warn you ahead of time, to prevent confusion, that you must not assume that ‘degree of belief’ refers to a psychological or mental state. It is, rather, a term of art in formal epistemology that may be interpreted, and has been interpreted, in a variety of fashions.) I will champion one of these variants in particular, a novel betting variant. I will also introduce a new distinction between intersubjective and interobjective interpretations of probability.

2. Group Degrees of Belief and Variants of the Intersubjective Interpretation

Gillies (1991) takes intersubjective probabilities to require consensus, but there are different ways of construing consensus, and also ways to conceive of group degrees of

belief (and hence group probabilities) such that they do not require consensus. To put it simply, just as the subjective interpretation of probability is based on individual degrees of belief, so the intersubjective interpretation is based on group degrees of belief. And since there are several different possible understandings of the notion of a degree of belief in each case, so there are different variants of each interpretation.⁴

I propose that there are two main ways to understand a group degree of belief:

- (i) As the value of an aggregation function when applied to the individual degrees of belief of the group members.
- (ii) As the end result of a process of the group members reaching consensus.

In so far as Gillies (1991) demands consensus, he rules out the aggregation view. Presumably, the primary reason is that group degrees of belief *qua* aggregates might satisfy the probability calculus as a matter of sheer accident, and in such a way as to be consistent with the group being susceptible to a Dutch Book despite having shared interests. Think back to the example of Romeo and Juliet in the introduction. Even if the aggregates (e.g. averages) of their personal degrees of belief about rain tomorrow in Oxford satisfy the probability calculus, this won't stop them being Dutch Booked as a group.⁵

Having degrees of belief which satisfy the probability calculus when aggregated might nevertheless be a necessary (but not sufficient) condition for group rationality. Imagine Romeo and Juliet reach perfect consensus on whether it will rain tomorrow

in Oxford. If this consensus requires that they have identical degrees of belief concerning that proposition, then the group degrees of belief *qua* aggregates ought, presumably, to take the same values as they do for each individual. (Think of taking an average, for example.) Hence the group degrees of belief *qua* aggregates will satisfy the axioms of probability in the presence of consensus.

It is not obvious, however, that one should construe consensus in the same way as Gillies. Instead, consensus might require agreement only on a shared betting strategy. So in short, reaching consensus on degrees of belief concerning p might be understood either as:

(iii) Arriving at shared personal degrees of belief concerning p (following Gillies 1991).

Or as:

(iv) Agreeing to adopt identical betting quotients concerning p (i.e. to act in accordance with a specific betting strategy).

If one follows De Finetti (1937) in understanding personal degrees of belief in an operational fashion, where degrees of belief *just are* dispositions to select particular betting quotients, then (iii) entails (iv). (As we will later see, however, (iv) needn't entail (iii) under such circumstances; one may have one's own preferred betting strategy but elect to implement someone else's.) If one instead understands personal degrees of belief as mental or psychological states, however, then (iii) entails (iv) only

in the event that these map on to betting quotients, i.e. can be accurately measured by betting behaviour.

It may help to think of these different views of personal degrees of belief by reference to the different available accounts of ‘on-off’ belief in the philosophy of mind. De Finetti’s account has a dispositionalist (if not behaviourist) flavour, after Price (1969) or Marcus (1990); roughly, different dispositional profiles are equated with different degrees of belief (rather than merely different beliefs).⁶ But it is possible, instead, to take a representational approach; to suggest that having a personal degree of belief centrally involves possessing a mental representation, in line with what Fodor (1975) or Cummins (1996), for instance, say about beliefs *simpliciter*. Indeed, Ramsey (1926, p. 125) hints at this when he writes that: ‘it is... conceivable that degrees of belief could be measured by a psychogalvanometer or some such instrument.’ One option would be to take different degrees of belief in the same proposition to involve different, although often highly similar, representations. Another would be to take any two degrees of belief in a given proposition to involve the same representation, albeit playing a different (functional or causal) role. This is the path that Ramsey (1926, p. 169) appears to take when he settles on the view that: ‘the degree of a belief is a causal property of it, which we can express vaguely as the extent to which we are prepared to act on it.’ In what follows, I will use ‘credence’ to refer to a degree of belief understood in any such representational sense.

Now given some of the objections to the Dutch Book argument at the level of the individual—see, for instance, Seidenfeld et al. (1990), Hájek (2005) and Rowbottom (2007a)—we may conclude that credences (if they exist) infrequently map on to

betting quotients. As such, if Romeo and Juliet arrive at shared personal credences concerning p that satisfy the axioms of probability, this doesn't guarantee that they will avoid being Dutch Booked in their bets on p in a real life betting scenario. In short, sharing credences does not entail sharing betting strategies. And it is also dubious that groups possess credences in anything other than a figurative sense.⁷

Another worry with the view that consensus requires shared personal degrees of belief is that the intersubjective view then seems to amount to little more than the subjective view plus a rationality constraint. Perhaps it is true that there are some situations in which individuals ought to consult with members of the groups of which they are a part, in order to arrive at rational *personal* degrees of belief. (Perhaps it's important also to keep channels of communication open, to allow critical discourse, and so forth.) But why not simply consider the personal degree of belief of any member of the group after the consensus generating process has taken place? In short, the consensus group degree of belief doesn't appear to be interestingly distinct from the personal degrees of belief of the members of the group on (iii).

My own view is therefore that we should take group degrees of belief to involve consensus (ii), but that we should understand consensus as agreement to employ identical betting quotients (iv)—or more generally, implement a shared betting strategy—in *circumstances in which group interests are relevant*.

This brings us to an outstanding question, which should be answered before we proceed. When, precisely, do rationality constraints apply to a group? An answer is provided by Gillies (1991), who explains, first, that his group-level Dutch Book

argument only concerns—i.e., is only effective for—groups with a shared interest (or interests). More precisely, he writes:

The members of the group must be linked by a common purpose; whether the common purpose leads to solidarity or rivalry within the group does not matter much; the important point is that the members have an interest in acting together and reaching consensus; love or fear would create, in this case, similar bonds. The common purpose might be financial, but need not be; for example, a group of soldiers might have the common purpose of taking an enemy position with the minimum injury and loss of life to the group.

So in the earlier case of Romeo and Juliet, there is a shared interest in (not wasting any of) the jointly owned pool of money. The existence of common interest in a group is not sufficient, however, for a rationality constraint to be in place. In addition, the members of the group must *be able* to achieve consensus. Imagine Romeo and Juliet are in different parts of the world, and not able to communicate, despite both having access to a shared bank account. (Maybe Romeo is on the run, and can't use a phone for fear of his call being traced, although he is confident that the police don't know about his shared bank account with Juliet and therefore can't trace his transactions.) Gillies (1991, p. 518) deals with this by introducing a *flow of information* requirement:

There must be a flow of information between the members though it does not matter whether the communication is organized centrally or peripherally or

whether it is direct (between any two members) or indirect (through the intervention of third parties).

This is somewhat imprecise; it is important not only that there is flow of information, but also that this flow can be employed to enable consensus to be reached (in principle, at the very least). But the basic idea is clear. A tricky question not covered by Gillies, however, is whether one-way communication links may suffice. Even if Juliet can only pass on to Romeo how she is betting with their shared funds (i.e. Romeo cannot reply), for example, Romeo can avoid making *some* bets that will lose shared funds whatever happens. And so he should. But is that a requirement of *group* rationality, or just *personal* rationality (on Romeo's part)? For present purposes, to be on safer ground, let's assume the latter. Let us stipulate that each member of a group must be able to transmit to, and receive from, every other member (whether directly or indirectly) in order for group rationality to be a consideration.

It is therefore evident that when one construes group degrees of belief in a betting-based way, they are non-trivially distinct from individual degrees of belief. Take a group of two for illustrative purposes. Neither member of the group may have an individual degree of belief that matches the group degree of belief *qua* agreed betting quotient. And this holds independently of which of the available accounts of personal degrees of belief, outlined by Eriksson and Hájek (2007), one prefers. You may interpret personal degrees of belief as either representational or dispositional in character, for example.

Furthermore, the betting-based view does not succumb to an objection to the notion of group rationality raised by Christensen (1991, p. 239–240 & 244–246). As part of his argument against the view that reflection is a diachronic rationality constraint—roughly, that agents should have confidence in their future degrees of belief (and act on these if they know them)—Christensen considers a ‘Double Agent Dutch Book’, similar to the one involving Romeo and Juliet discussed above. This involves the same bookie betting with each of the agents (if he needs to in order to win).⁸ However, Christensen (1991, p. 240) draws a rather different moral from the one that Gillies and I do:

The reason the Double Agent Dutch Book does not show anyone to be irrational, I think, is this: although my beliefs are in a clear sense inconsistent with my wife’s, that is a perfectly reasonable state of affairs.

But the advocate of the betting view of group probability may wholeheartedly agree that it is reasonable for Romeo’s beliefs—or more precisely, his degrees of belief—to be inconsistent with Juliet’s. And he may turn matters around. Surely Christensen cannot think that it *is* rational—in an ordinary language sense of ‘rational’⁹—for a husband and wife to make individual bets with a bookie such that they lose joint funds whatever happens, and then just continue to make such bets, come what may, despite having communication channels by which they could easily prevent this occurring (and joint interest in not losing money for no utility)? But it would appear that this is a consequence of denying the notion that avoiding Dutch Books at the group level is ever a (practical) rationality requirement. Before I ‘bet’ with a significant sum of money shared with my wife, e.g. in making an investment, I discuss this her. And the

same is true, *mutatis mutandis*, when she wishes to ‘bet’. I trust we are not the only working couple with a joint bank account to take this precaution!

I should also emphasise that talk of ‘betting strategies’ should be understood to apply to decision-making scenarios in general, where utilities are at stake.¹⁰ Think of the senior members of a political party, who make a joint assessment of the chance that they will improve their popularity by advocating a peculiar policy. In the interests of the group, they will ‘bet’ in line with the joint assessment in their dealings with the media, and so forth. (Their opponents would make political capital out of any perceived divisions.) But no member of the group may have a personal assessment that precisely matches the joint assessment. Moreover, if any member of the political party were to leave the group, and quit politics, she would ‘bet’ differently. She would feel free to voice her personal view on the policy.

This is perfectly compatible with the notion that there is only one best way for groups to decide upon intersubjective probabilities, in any particular situation. One possibility is that groups should sometimes employ judgement aggregation strategies—as discussed by List and Pettit (2002) and Dietrich (2007), for example—although I do not wish to be drawn into that issue in the present paper.

Of course, the question also arises as to whether agreed betting quotients should always be coherent, on pain of irrationality, and thus satisfy the axioms of probability. Some objections to the Dutch Book argument at the individual level, such as Rowbottom (2007a), work precisely by showing that it is sometimes best to select incoherent betting quotients. But this is not to question the Dutch Book theorem that

an individual is *susceptible* to having a Dutch Book drawn up against her if her betting quotients fail to satisfy the axioms. And such a theorem holds in the case of groups too. All one need add, in order to link coherence with rationality *in some contexts*, is that there is a significant class of circumstances where it is best to avoid susceptibility to a Dutch Book.

3. Intersubjective Versus Interobjective Probabilities

I also wish to introduce another novel distinction, between intersubjective and interobjective interpretations of probability. Both involve consensus group degrees of belief. But on a simple intersubjective account, it does not matter how consensus is reached; it is enough that it is present. On an interobjective account, by contrast, particular procedures are also required in order to form consensus, e.g. critical discussion with input from all members of the group who have relevant degrees of belief or relevant expertise. In short, the idea is that consensus reached in an inappropriate way will fail to generate *rational* group degrees of belief, and hence group probabilities.

Compare the subjective and objective Bayesian interpretations, which operate at the level of the individual. On a pure subjective view, having degrees of belief which satisfy the axioms of probability is sufficient for rationality. But objective Bayesians introduce further rationality requirements, e.g. that degrees of belief should reflect observed frequencies where appropriate, as ‘top ups’. Ultimately, it is perhaps best to understand this in terms of a spectrum; strong objective Bayesians hold that a personal probability should always have a special unique value, whereas pure

subjective Bayesians require only that it lies within a particular range. Those in the middle of the spectrum think that personal probabilities should have unique values in some contexts, but not in others.¹¹ Such a spectrum also exists when we discuss group, rather than individual, degrees of belief.

I take it for granted that there are good and bad ways to reach consensus. This should be uncontroversial; one way of reaching consensus is by brainwashing, but clearly that is not of much interest in so far as rational decision-making is concerned. Nonetheless, many possible interobjective accounts remain. Rather than develop a particular interobjective interpretation here, I will just highlight several kinds of conditions and processes that might feature in such an account. A caveat to bear in mind, in what follows, is that how it is best for a group to reach a *considered* consensus may be highly context sensitive.

One important consideration is expertise. For example, one might specify:

Strong expertise condition: The group members participating in any particular process relevant to the ultimate assignment of group probabilities should have expertise in that process.

Or one might opt for something rather weaker:

Weak expertise condition: The results/findings/contributions of group members participating in a process relevant to the ultimate assignment of

group probabilities should be weighted according to the relevant degrees of expertise of those members.

Notice that such conditions are expressed in an external fashion—i.e. it is not assumed that the group has any evidence about the expertise of its members, and/or any ability to determine relative expertise—although internal equivalents are possible. These would involve appeal to *perceived* expertise instead.

Another possible condition, which may clash with *strong expertise*, is:

Freedom of input: Any member of the group should be free to share any information that he/she perceives to be relevant to the decision-making process (or a subsidiary process).¹²

Alternatively, it may be possible to preserve *strong expertise* by instead requiring only that an information pooling process should occur (subject to the expertise condition):

Information pooling process: Information relevant to the decision making process should be pooled and made available to relevant members of the group.

On a related note, it may also be wise to allow for new data to be gathered in order to further inform a decision:

Information gathering process: Information relevant to the decision making process should be gathered (subject to appropriate expected utility and/or opportunity cost considerations).

Note that *solicitation* of opinions from experts who are not part of the group, i.e. gathering of extra-group testimony, may constitute a part of *information gathering* in appropriate circumstances; this will not be a general rule, however, in so far as solicitation will not be generally possible or advisable (e.g. if the experts outside the group are hostile and liable to offer false testimony).

Maybe the group should also be open to expansion or contraction before the decision making process is completed. (Groups may be defined with reference to their histories, and not just their actual membership at a point in time; as such, the identity of a group may be preserved despite a change in its active membership.) So *recruitment* and *exclusion* processes—constrained by appropriate considerations of expertise and intent—may also be required in suitable contexts.

In closing this section, I shall simply mention two further processes, the potential benefits of which I take to be reasonably self-evident:

Brainstorming/hypothesis generation process: An attempt should be made to identify relevant possibilities that have not yet been made explicit and/or considered by the group (or any members thereof).

Error correction process: Arguments and calculations should be checked for errors, and observation statements should be intersubjectively tested (e.g. by the repetition of experiments).

It should go without saying that this list is incomplete, and that there may be some overlap between processes as defined here. But for illustrative purposes, it suffices. In closing the paper, I will emphasise one special advantage of the kind of betting-based group level interpretation that I have argued for above.

4. The Scope of Background Information

It is usually claimed that consensus probabilities have to be based on *shared* background information. One might therefore imagine that a group deciding on $P(h|b)$, say, must be composed of members who each understand b in its entirety (in addition to their other personal information). According to this view, a group of scientists trying to solve a problem should pool all the information that they think is relevant, make sure that each and every member of the group understands it, and only then try to arrive at consensus.

However, intersubjective probabilities need not be limited in such a way. Gillies (2000, p.173) seems to be on the verge of spotting this when he writes:

[Background information] may be more extensive than the knowledge possessed by any individual members of the group. Since there is flow of information and exchange of ideas within the group, if one member has a

piece of relevant knowledge which the others lack, he or she can communicate it to the others.

What Gillies misses is that it may often *not* be possible for individuals to communicate all the relevant information that they have. It might take years of work for Dr Black, an engineer specialising in turbines, to fully understand Prof. Plum's work in nuclear physics. (It may take more years, indeed, that the ageing Dr Black has left.) Yet Black and Plum might nevertheless find themselves working together in order to design a nuclear submarine, and required to provide politicians with an accurate estimate of when the work will be completed.

Imagine that Black and Plum meet and discuss the matter in some depth, several times. Although each realises that some of the other's specialist knowledge, which cannot be shared, is relevant to answering the question, they nevertheless manage to agree about when the submarine will be completed.¹³ They make a joint estimate that a prototype of the submarine will be completed in twenty years, with which they are both entirely satisfied.

The crucial point to note, however, is that they are not agreeing on (the value of) the same conditional probability when we consider only their personal degrees of belief. Let p be 'A prototype of the submarine will be completed in twenty years', and B be the shared background information of Dr Black and Prof. Plum which is relevant to p —e.g. about the funding available for the project and the number of staff available to work on it. Let T_1 be the testimony given to Dr Black by Prof. Plum, and T_2 be the testimony given by Dr Black to Prof. Plum. Finally, let S_1 be the specialist knowledge

of Dr Black *which cannot be communicated* to Prof. Plum and S_2 be the specialist knowledge of Prof. Plum *which cannot be communicated* to Dr Black. We find:

$$(\alpha) \quad P_{Black}(p|B \& S_1 \& T_1) = P_{Plum}(p|B \& S_2 \& T_2)$$

The question is: what significance, if any, does this have?

Imagine T_1 is entailed by $S_2 \& B$ and T_2 is entailed by $S_1 \& B$. In that event, each party will have taken account of some of the products of the other's specialist knowledge without needing to possess such knowledge. Moreover, entailment may not always be required. For example, T_1 may only be rendered highly probable (in an aleatory sense) by S_2 .

Now let K be the union of B and S_1 and S_2 . It may even be suggested that under the circumstances described, on a consensus account of group degrees of belief:

$$(\beta) \quad P_{Black}(p|B \& S_1 \& T_1) = P_{Plum}(p|B \& S_2 \& T_2) \cong P_{Black \& Plum}(p|K)$$

This is too strong if we read ‘ \cong ’ in its standard mathematical sense (i.e. as ‘is approximately equal to’); clearly, S_1 and S_2 (in conjunction with B) may together entail (or render probable) propositions unforeseen by either Black or Plum. But if we read ‘ \cong ’ instead as ‘is our best estimate of’, then β appears considerably more plausible. Neither Plum nor Black do, or *could*, have K as their personal background information. But they have made up for that as best they can, and our progress relies on such efforts.

Now there may be serious doubts about this if consensus requires arriving at shared personal degrees of belief. ‘ \cong ’ is used in place of ‘=’, after all, precisely because neither Plum nor Black have degrees of belief involving K . But on the understanding of intersubjective probabilities which I advocated in section 2—where reaching consensus involves only ‘agreeing to adopt identical betting quotients’—the following may hold irrespective of the inability of either party to access the entirety of K :

$$(\chi) \quad P_{Black}(p|B \& S_1 \& T_1) = P_{Plum}(p|B \& S_2 \& T_2) = P_{Black \& Plum}(p|K)$$

In short, Black and Plum may agree to bet in a particular way on the basis of the sum of their background knowledge despite their inability to completely share it. Thus their report to their government may be understood, probabilistically, *in the way that it typically is in the real world*: as their joint recommendation on the basis of the sum of their relevant expert knowledge.

Note that this discussion does not preclude the notion that $P_{Black \& Plum}(p|K)$ should be calculated on the basis of rationality constraints that go beyond those required by (inter)subjectivists. On the contrary, one is free to adopt an interobjective view as outlined in section three.

Department of Philosophy,

Lingnan University

Acknowledgements

I am grateful to an audience at Oxford University, and especially Antony Eagle and Simon Saunders, for comments on an early version of this paper. I should also like to thank Matthew Kopec and an anonymous referee for comments on more recent drafts. Finally, I should like to thank the British Academy and the Sydney Centre for the Foundations of Science for supporting much of my work on the paper.

References

Bangu, S. 2010. On Bertrand's Paradox. *Analysis* 70, 30–35.

Bertrand, J. 1960 [1889]. *Calcul des Probabilités*. New York: Chelsea, 3rd Ed.

Bovens, L. & Hartmann, S. 2004. *Bayesian Epistemology*. Oxford: Oxford University Press.

Carnap, R. 1962 [1950]. *Logical Foundations of Probability*. Chicago: University of Chicago Press, 2nd Ed.

Christensen, D. 1991. 'Clever Bookies and Coherent Beliefs', *The Philosophical Review* 100, 229–247.

Cummins, R. 1996. *Representations, Targets, and Attitudes*. Cambridge, MA: MIT Press.

De Finetti, B. 1937. 'Foresight, Its Logical Laws, Its Subjective Sources', in H. E. Kyburg & H. E. Smokler (eds), *Studies in Subjective Probability*. (New York: John Wiley & Sons), pp. 93–158.

Dietrich, F. 2007. 'A Generalised Model of Judgment Aggregation', *Social Choice and Welfare* 28, 529–565.

Eriksson, L. & Hájek, A. 2007. 'What are Degrees of Belief?', *Studia Logica* 86, 183–213.

Fodor, J. 1975. *The Language of Thought*. Cambridge, MA: Harvard University Press.

Gillies, D. A. 1991. 'Intersubjective Probability and Confirmation Theory', *British Journal for the Philosophy of Science* 42, 513–533.

Gillies, D. A. 2000. *Philosophical Theories of Probability*. London: Routledge.

Hájek, A. 2005. 'Scotching Dutch Books?', *Philosophical Perspectives* 19, 139–151.

Howson, C. & Urbach, P. 2005 [1989]. *Scientific Reasoning: The Bayesian Approach*. La Salle: Open Court, 3rd Ed.

Huber, F. 2008. 'Milne's Argument for the Log-Ratio Measure', *Philosophy of Science* 75, 413–420.

Jaynes, E. T. 2003. *Probability Theory: The Logic of Science*. Cambridge: Cambridge University Press.

Jeffrey, R. 2004. *Subjective Probability: The Real Thing*. Cambridge: Cambridge University Press.

Keynes, J. M. 1921. *A Treatise on Probability*. London: Macmillan.

Kitcher, P. 1990. 'The Division of Cognitive Labor', *Journal of Philosophy* 87, 5–22.

Kitcher, P. 2003. *The Advancement of Science: Science without Legend, Objectivity without Illusions*. Oxford: Oxford University Press.

List, C. and P. Pettit 2002. 'Aggregating Sets of Judgments: An Impossibility Result', *Economics and Philosophy* 18, 89–110

Marcus, R. B. 1990. 'Some Revisionary Proposals about Belief and Believing', *Philosophy and Phenomenological Research* 50, 132–153.

Milne, P. 1996. 'Log[P(h/eb)/P(h/b)] Is the One True Measure of Confirmation', *Philosophy of Science* 63, 21–26

Price, H. H. 1969. *Belief*. London: Allen & Unwin.

Ramsey, F. P. 1926. ‘Truth and Probability’, in H. E. Kyburg and H. E. Smokler (eds), *Studies in Subjective Probability*. New York: Wiley, pp. 63–92.

Rowbottom, D. P. 2007a. ‘The Insufficiency of the Dutch Book Argument’, *Studia Logica* 87, 65–71.

Rowbottom, D. P. 2007b. “‘In Between Believing’ and Degrees of Belief”, *Teorema* XXVI, 131–137.

Rowbottom, D. P. 2008a. ‘Intersubjective Corroboration’, *Studies in History and Philosophy of Science* 39, 124–132.

Rowbottom, D. P. 2008b. ‘On the Proximity of the Logical and “Objective Bayesian” Interpretations of Probability’, *Erkenntnis* 69, 335–349.

Rowbottom, D. P. 2011a. ‘Kuhn vs. Popper on Criticism and Dogmatism in Science: A Resolution at the Group Level’, *Studies in History and Philosophy of Science* 42, 117–124.

Rowbottom, D. P. 2011b. ‘Objective Bayesianism Defended?’, review of *In Defence of Objective Bayesianism* by Jon Williamson. *Metascience* 21, 193–196.

Rowbottom, D. P. Forthcoming. ‘Information Versus Knowledge in Confirmation Theory’, *Logique et Analyse*.

Rowbottom, D. P. and N. Shackel. 2010. ‘Bangu’s Random Thoughts on Bertrand’s Paradox’, *Analysis* 70, 689–692.

Ryder, J. M. 1981. ‘Consequences of a Simple Extension of the Dutch Book Argument’, *British Journal for the Philosophy of Science* 32, 164–167.

Sarkar, H. 2007. *Group Rationality in Scientific Research*. Cambridge: Cambridge University Press.

Schwitzgebel, E. 2001. ‘In-between Believing’, *Philosophical Quarterly* 51, 76–82.

Schwitzgebel, E. 2010. ‘Acting Contrary to Our Professed Beliefs, or the Gulf between Occurrent Judgment and Dispositional Belief’, *Pacific Philosophical Quarterly* 91, 531–553.

Seidenfeld, T., Schervish, M. J., and Kadane, J. B. 1990. ‘When Fair Betting Odds are not Degrees of Belief’, *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, 517–524.

Strevens, M. 2003. ‘The Role of the Priority Rule in Science’, *Journal of Philosophy* 100, 55–79.

Van Fraassen, B.C. 1989. *Laws and Symmetry* (Oxford: Clarendon Press)

Williamson, J. 2005. *Bayesian Nets and Causality: Philosophical and Computational Foundations*. Oxford: Oxford University Press.

Williamson, J. 2006. ‘Philosophies of Probability: Objective Bayesianism and its Challenges’, in A. Irvine (ed.): *Handbook of the Philosophy of Mathematics, Volume 4 of the Handbook of the Philosophy of Science* (London: Elsevier)

Williamson, J. 2010. *In Defence of Objective Bayesianism*. Oxford: Oxford University Press.

¹ See De Finetti (1937, pp.146–147).

² For a recent defence of the principle of indifference, see Bangu (2010). See Rowbottom and Shackel (2010) for a refutation.

³ To be more specific, one might think that how labour should be distributed within some group is a function of group probabilities.

⁴ See Eriksson and Hájek (2007) for the best contemporary discussion of how we may understand subjective degrees of belief.

⁵ I recognise that averaging may not be a good strategy for judgement aggregation, but I use this merely as an example; nothing rides on this in the present context.

⁶ I say ‘roughly’ because we could allow, like Price (1969), that dispositional profiles will be affected by propositional attitudes other than beliefs. A charitable way of reading De Finetti (1937) is as suggesting that we can create (measurement) scenarios where behaviour will be identical if degrees of belief *qua* dispositions are. For more on dispositionalism and degrees of belief, see Schwitzgebel (2001), Rowbottom (2007b), and Schwitzgebel (2010).

⁷ If personal degrees of beliefs are credences, then this explains why they are so difficult to measure (or render operational). Certainly, actual behavioural patterns (like the dispositions that guide these, if one allows for such) are determined by rather more than mental representations or the properties thereof! See also the previous endnote.

⁸ See also a similar point made earlier by Ryder (1981, p. 165), namely: ‘If we have two (or more) people with different degrees of belief in the same simple event *E*, a Dutch Book can be made against them. This is just as ‘disastrous’ and ‘obviously unreasonable’ [quotation from Carnap] as it is for an individual. It means that Subjectivists never actually make the bets which are envisaged by the Dutch Book argument. If they did someone could come along and find two or more subjectivists with different degrees of belief and make a system of bets which would result in a certain loss to the subjectivists considered as a group.’ Gillies (1991, p. 520–521) addresses Ryder’s point by appeal to the aforementioned common interest and flow of information requirements.

⁹ Christensen (1991) also argues that the force of Dutch Book arguments is not practical. In fact, he concedes (Christensen 1991, p. 240) that: ‘If the force of Dutch Book considerations were practical—if the reason for obeying the probability calculus were to avoid actual monetary loss—then perhaps the Double Agent Dutch Book could be used to support a demand for probabilistic consistency between individuals (or at least between spouses who share their assets).’ However, I believe that the force is practical *when it comes to considering groups*. (And we may also note that ‘probabilistic consistency’ is vague; it could refer to betting quotients or personal degrees of belief.) I also concede that sometimes it may be rational for group degrees

of belief, *qua* betting quotients, not to satisfy the axioms of probability. I address this point below, in closing this section.

¹⁰ In the words of Ramsey (1926, p. 183): ‘[A]ll our lives we are in a sense betting. Whenever we go to the station we are betting that a train will really run, and if we had not a sufficient degree of belief in this we should decline the bet and stay at home.’

¹¹ Williamson (2010) is somewhere in the middle, in so far as he thinks that there are some situations, e.g. Bertrand’s paradox, in which it is reasonable to equivocate in different ways. See also Rowbottom (2011b).

¹² Note that this is a stronger requirement than the ‘flow of information’ one specified previously, following Gillies (1991).

¹³ I use ‘knowledge’ here because it seems natural to do so. Do not read ‘knowledge’ in the sense of traditional epistemology. ‘Information’ would suffice as a replacement. See also Rowbottom (Forthcoming).