

No One Can Serve Two Epistemic Masters

J. Dmitri Gallow [†]

ABSTRACT

Consider two epistemic experts—for concreteness, let them be two weather forecasters. Suppose that you aren't certain that they will issue identical forecasts, and you would like to proportion your degrees of belief to theirs in the following way: first, conditional on either's forecast of rain being x , you'd like your own degree of belief in rain to be x . Secondly, conditional on them issuing different forecasts of rain, you'd like your own degree of belief in rain to be some weighted average of the forecast of each (perhaps with weights determined by their prior reliability). Finally, you'd like your degrees of belief to satisfy the axioms of probability. Moderate ambitions, all. But you can't always get what you want.

AL and Bert are two local weather forecasters. At least when it comes to the question of whether it will rain, they are both quite reliable. And their forecasts tend, for the most part, to agree. When they disagree, it is sometimes Al and sometimes Bert that gets closer to the truth. One reason they agree so often is that they discuss their provisional forecasts with one another before deciding upon a final forecast.

Today you are wondering whether it will rain; and you have not yet checked the forecasts. Should you be certain that Al and Bert will issue identical forecasts? The answer, it seems to me, is “no”. Al and Bert have in the past disagreed with each other, so, it seems to me, you should not be absolutely certain that they will issue identical forecasts today. Let it be so. If your own degrees of confidence are given by the function C , \mathcal{A} is a random variable whose value is Al's forecast of rain, and \mathcal{B} is a random variable whose value is Bert's forecast of rain, then let's stipulate that C should be such that (I).

$$(I) \quad C(\mathcal{A} = \mathcal{B}) < 1$$

Given that Al forecasts that the chance of rain is a , how confident should you be that it will rain? A plausible answer, it seems to me, is “ a ”. Not only is this answer plausible in its own right; it is endorsed (not for Al, but for various other epistemic experts like objective chance and your future self) by GAIFMAN (1988), VAN FRAASSEN (1984, 1995),

Final Draft. Forthcoming in *Philosophical Studies*.

✉: jdmitrigallow@pitt.edu

[†] Thanks to Michael Caie, Daniel Drucker, Harvey Lederman, and an anonymous reviewer for helpful conversations and feedback.

LEWIS (1980), and ELGA (2007), among many others.¹ Let it be so. Where ‘ r ’ is the proposition that it will rain, we will stipulate that C should satisfy (2).

$$(2) \quad \forall a \quad C(r \mid \mathcal{A} = a) = a$$

Given that Bert forecasts that the chance of rain is b , how confident should you be that it will rain? Again, a quite plausible answer is “ b ”. Let it be so. We will stipulate that your degrees of confidence satisfy (3).

$$(3) \quad \forall b \quad C(r \mid \mathcal{B} = b) = b$$

What if Al forecasts that the chance of rain is a , and Bert forecasts that the chance of rain is b , where $a \neq b$? How confident should you be that it will rain, then? A plausible answer, it seems to me, is that your confidence in rain should be some weighted average of a and b (with weights determined by Al and Bert’s prior reliability, perhaps). Let’s use ‘ α ’ for the weight you place in Al’s opinion, conditional on a disagreement between Al and Bert, and let’s use ‘ β ’ for the weight you place in Bert’s opinion, conditional on a disagreement between Al and Bert (where $\alpha, \beta > 0$ and $\alpha + \beta = 1$). Then, let’s stipulate that your degrees of confidence satisfy (4).

$$(4) \quad \forall a, b \quad C(r \mid \mathcal{A} = a, \mathcal{B} = b) = \alpha a + \beta b$$

These are, it seems to me, four plausible constraints to place on your degrees of confidence in an epistemic situation like this. But so too is the following: that C be an orthodox probability function.² And if we add to our endorsement of (1), (2), (3), and (4) the commendation that C be such a probability, we will have contradicted ourselves. For there is no orthodox probability function which satisfies each of (1), (2), (3), and (4).

Proposition 1. *If r is any arbitrary proposition and \mathcal{A} and \mathcal{B} are random variables taking on values in the unit interval, then there is no orthodox probability function C such that (1), (2), (3), and (4) all hold.*

The appendix provides a general proof of **Proposition 1**, but the following special case will prove illustrative.

¹ There are a wide variety of putative experts and a wide variety of principles of expert deference on offer in the current epistemology literature. However, the reader will note that everything we will say here about Al and Bert goes just as well for chance, your rational future self, and so on and so forth. The reader will also note that if Al and Bert have all your evidence and more, and if they are additionally certain of their own forecasts, then every extant principle of expert deference will entail both (2) and (3). For instance, given an expert \mathcal{E} who is certain of their own credences and has all your evidence and more, the following ways of treating \mathcal{E} as an expert are all equivalent: for all p, x, E i) $C(p \mid \mathcal{E}(p) = x) = x$, ii) $C(p \mid \mathcal{E} = E) = E(p)$, iii) $C(p \mid \mathcal{E} = E) = E(p \mid \mathcal{E} = E)$, and $C(p) = \sum_x x \cdot C(\mathcal{E}(p) = x)$. See GALLOW (msb) for a proof of this claim and more on the relationship between various principles of expert deference.

² By an ‘orthodox probability function’, I will mean that C is non-negative, normalized, countably additive, and conglomerable. These assumptions go beyond probabilism in the case where we are considering infinitely many possible values for \mathcal{A} and \mathcal{B} . However, if we suppose that there are at most finitely many potential values for \mathcal{A} and \mathcal{B} , then an ‘orthodox probability function’ is just any finitely additive probability. Thanks to an anonymous reviewer for their clarifying comments on this point.

Example 1. A reliable friend tells you that they saw the forecasts this morning, and they know that both Al and Bert gave a forecast of $1/3$, but they can't recall, in either case, whether those were their forecasts for rain or their forecasts for shine (nor can they recall whether both were forecasting for rain/shine or whether one was forecasting for rain and the other for shine).

In **Example 1**, your credences about \mathcal{A} , \mathcal{B} , and r can be represented by filling in the values in the following 4×2 grid.

	r	$\neg r$
$\mathcal{A} = 1/3 \wedge \mathcal{B} = 1/3$	x	$2x$
$\mathcal{A} = 1/3 \wedge \mathcal{B} = 2/3$	w	u
$\mathcal{A} = 2/3 \wedge \mathcal{B} = 1/3$	v	z
$\mathcal{A} = 2/3 \wedge \mathcal{B} = 2/3$	$2y$	y

To understand the values assigned to the first and fourth rows, note first that, since (4) requires $C(r \mid \mathcal{A} = \mathcal{B} = 1/3)$ to be $1/3$, your credence that r must be one half of your credence that $\neg r$ in the first row. Similarly, since (4) requires $C(r \mid \mathcal{A} = \mathcal{B} = 2/3)$ to be $2/3$, your credence that r must be twice your credence that $\neg r$ in the fourth row. Then, (2) tells us that we must have

$$C(r \mid \mathcal{A} = 1/3) = 1/3$$

$$\frac{x + w}{3x + w + u} = 1/3$$

which implies that $u = 2w$. Similarly, another application of (2) tells us that we must have

$$C(r \mid \mathcal{A} = 2/3) = 2/3$$

$$\frac{2y + v}{3y + v + z} = 2/3$$

which implies that $v = 2z$. So your credences must have the following values, for some choice of x, y, z, w such that $x + y + z + w = 1/3$:

	r	$\neg r$
$\mathcal{A} = 1/3 \wedge \mathcal{B} = 1/3$	x	$2x$
$\mathcal{A} = 1/3 \wedge \mathcal{B} = 2/3$	w	$2w$
$\mathcal{A} = 2/3 \wedge \mathcal{B} = 1/3$	$2z$	z
$\mathcal{A} = 2/3 \wedge \mathcal{B} = 2/3$	$2y$	y

But then, by (3), we must have

$$C(r \mid \mathcal{B} = 1/3) = 1/3$$

$$\frac{x + 2z}{3x + 3z} = 1/3$$

which implies that $z = 0$. Similarly, another application of (3) tells us that we must have

$$C(r \mid \mathcal{B} = 2/3) = 2/3$$

$$\frac{2y + w}{3y + 3w} = 2/3$$

which implies that $w = 0$. And if $z = w = 0$, then $C(\mathcal{A} = \mathcal{B}) = 1$, contradicting (4).

The Gospels teach that “No one can serve two masters. Either you will hate the one and love the other, or you will be devoted to the one and despise the other.”³ Serving two masters is impossible when the masters issue contrary commands. Moreover, even being *disposed* to serve two masters is impossible if it is *possible* that the masters issue contrary commands. Similarly, serving two *epistemic* masters—deferring to two experts, Al and Bert, in the manner of (2), (3), and (4)—is impossible if you are less than certain that Al and Bert won’t issue contrary recommendations, *i.e.*, if (1).

Perhaps (2) and (3) are not so plausible, after all. Suppose that Al’s forecast of rain is 100%. How confident should you be of rain, conditional on that supposition? Perhaps, on that supposition, you should have some degree of belief shy of 100%—perhaps certainty is never justified by the testimony of experts. Similar concerns arise for (4). Suppose that Al’s forecast is 100% and Bert’s forecast is 100%− ϵ . Perhaps, on that supposition (for especially small values of ϵ), your own credence in rain should be less than 100%− ϵ . For Al and Bert, such suppositional credences seem sensible. However, we are not only interested in showing epistemic deference to experts like Al and Bert. We are additionally interested in showing epistemic deference to experts like chance. And, conditional on the chance of r being 1, surely your credence in r should be 1, too. Moreover, note that, even if you think that you should only epistemically defer in Al and Bert within some limited range of values—like, for instance, the interval $[1/3, 2/3]$ —proposition 1 shows that this, too, will be impossible if you come to know that both Al and Bert’s forecasts lie in this range. For proposition 1 does not assume anything about what information you do or do not have about Al and Bert’s forecasts. Suppose you learn, from a reliable source, that both Al and Bert made a forecast of somewhere between a 1/3 and a 2/3 chance of rain, but you’ve yet to learn what precisely those forecasts are. Then, we may suppose that $C(\mathcal{A}, \mathcal{B} \in [1/3, 2/3]) = 1$. Even so, you will be incapable of satisfying (1*), (2*), (3*), and (4*) while your degrees of belief are given by an orthodox probability.

(1*)		$C(\mathcal{A} = \mathcal{B}) < 1$
(2*)	$\forall a \in [1/3, 2/3]$	$C(r \mid \mathcal{A} = a) = a$
(3*)	$\forall b \in [1/3, 2/3]$	$C(r \mid \mathcal{B} = b) = b$
(4*)	$\forall a, b \in [1/3, 2/3]$	$C(r \mid \mathcal{A} = a, \mathcal{B} = b) = \alpha a + \beta b$

And the same goes for any other interval of values within which you may wish to treat Al and Bert as epistemic experts.

In fact, there is a deeper issue with the plausible principles (2), (3), and (4). Whether or not you decide to *defer* to Al and Bert in the manner described by those principles, you cannot regard both Al and Bert as unbiased indicators of each other’s opinions (at least, not if you are to retain an orthodox probability function for your credences.) It is because (2), (3), and (4) require you to regard Al’s forecast as an unbiased indicator

³ Matthew 6:24.

of Bert's, and Bert's forecast as an unbiased indicator of Al's, that they end up being incompatible with (1).

To illustrate: suppose that Al's forecast of rain is a . Then, what should you expect Bert's forecast of rain to be? A plausible answer, it seems to me, is " a ". And suppose that Bert's forecast of rain is b . Then, what should you expect Al's forecast of rain to be? A sensible answer, it seems to me, is " b ". So it seems natural to suppose that your degrees of confidence should satisfy (5) and (6) (where ' $\mathbb{E}[\mathcal{A} \mid \mathcal{B} = b]$ ' is your conditional expectation of \mathcal{A} 's value, given that \mathcal{B} 's value is b , and likewise for ' $\mathbb{E}[\mathcal{B} \mid \mathcal{A} = a]$ ').

$$(5) \quad \forall b \quad \mathbb{E}[\mathcal{A} \mid \mathcal{B} = b] = b$$

$$(6) \quad \forall a \quad \mathbb{E}[\mathcal{B} \mid \mathcal{A} = a] = a$$

This much appears unobjectionable, and is implied by (2), (3), and (4) above. But, once again, appearances are misleading. If we recommend (1), (5), and (6), and we follow these recommendations up with the advice that C be an orthodox probability, we will have contradicted ourselves.

Proposition 2. *If \mathcal{A} and \mathcal{B} are random variables taking on values in the unit interval, then there is no orthodox probability function C such that (1), (5), and (6) hold.*

The case we have been considering is interesting in its own right, but it additionally has implications for broader issues in epistemology. In the first place, contemporary epistemology is filled with principles of expert deference which implore you to treat various experts—*e.g.*, your future rational self and the future chances—as epistemic masters.^{4,5} Such principles are often suggested and endorsed independently, and little is said about how those principles should relate to one another. But, of course, your Monday rational self and the Sunday chances might disagree about the probability of rain on Thursday.⁶ Supposing that they do, how confident should you be in rain on Thursday? A plausible initial thought is: your confidence in rain on Thursday should be some weighted average of your rational Monday self's credence and the chances on Sunday, with weights determined by your beliefs about the reliability of your rational Monday self and the Sunday chances with respect to the question of whether it will rain on Thursday. What we've seen is that such an attitude would be inconsistent with the claims that a) you should not be certain that the Sunday chances and your rational Monday self will agree; b) you should treat your rational Monday self as an expert; c) you should treat the Sunday chances as an expert; and d) your degrees of confidence should be given by an orthodox probability function. So more needs to be said about how these principles of epistemic deference are related; how, conditional

⁴ See, for starters, GAIFMAN (1988), LEWIS (1980, 1994), HALL (1994), VAN FRAASSEN (1984, 1995), CHRISTENSEN (2010), and ELGA (2013).

⁵ Many of these principles of expert deference look different from (2) and (3); fortunately, in most cases, we can present them in the form of (2) and (3) by simply shifting our attention to a different expert. For instance, LEWIS (1994) and HALL (1994) both say that we should not defer to the judgments of chance, but rather the judgments of chance, conditionalized on the proposition that it *is* chance. In such a case, we could consider the relevant expert to be, not chance itself, but rather chance conditionalized on chance, and we will get back a principle looking like (2) and (3). *Cf.* HALL & ARNTZENIUS (2003) and SCHAFFER (2003). (See also footnote 1.)

⁶ *Cf.* GALLOW (*msa*).

on disagreements, we should adjudicate between the verdicts of the various experts to which we are enjoined to defer.

In the second place, contemporary epistemology has focused quite a bit on the special case of our original scenario in which *you yourself* are one of the epistemic masters—for instance, the case in which you are *A*.⁷ Suppose that you have yet to form your weather forecast or hear news of Bert's forecast. You pore over the evidence and you settle on a forecast of *a* that it will rain. You then hear that Bert, who has all the same evidence as you, has settled on a forecast of *b* that it will rain, where $b \neq a$. How confident should you be of rain, then? The *conciliacionist* answers: 'some intermediate degree of belief between *a* and *b*'. One method for implementing the conciliacionist approach says that you should take a weighted average of your forecast and Bert's forecast, with the weights α and β given by your prior degrees of belief that, conditional on a disagreement, either you or Bert would make the more rational (or perhaps, the more accurate) judgment. This method of conciliating, it has been pointed out,⁸ fails to commute with conditionalizing on your evidence if the weights remain fixed over time. In response, some have considered using *variable* weights—weights that change as you acquire new evidence. If this is done in just the right way, then linear averaging and conditionalization can commute.⁹

What we've just seen, however, is that this *diachronic* difficulty with linear averaging as a method of conciliating (*vis.*, that it does not commute with conditionalization) is no more than a symptom of an underlying *synchronic* difficulty with treating both yourself and Bert as experts, in the sense of (2) and (3), and taking a weighted average of your respective opinions—according to *any* system of weights—conditional on a disagreement. Suppose that, before you've looked at any of the evidence or heard news of Bert's forecast, you have the following conditional degrees of belief: given that your forecast of rain is *a*, your credence in rain is *a* (2); given that Bert's forecast of rain is *b*, your credence of rain is *b* (3); and, given that you and Bert disagree in your forecasts, your credence in rain is some weighted average of your own forecast and Bert's forecast (4). Then, you will either be certain that you and Bert *won't* disagree, or else your degrees of belief will not be an orthodox probability measure.

What should we say about principles of epistemic deference and peer disagreement in light of proposition 1? We might note that (4) presupposes that the weights α and β remain fixed for all potential values of *a* and *b*. Noticing this, we may think that we should relax this assumption, and allow that the weights given to *A*'s and *B*'s opinions, conditional on a disagreement, be a function of *A*'s and *B*'s forecasts.¹⁰ This would get around proposition 1, but would not, I think, be a satisfactory solution of the problem in general. Note that, in *Example 1*, we only needed to apply (4) in the case where *A* and *B* both agreed on their forecast. So no amount of futzing with the weights α and β will help in cases like *Example 1*. Denying that $C(r \mid \mathcal{A} = \mathcal{B} = x)$ should be *x* does not look very plausible. Nor does the evidence you've received in *Example 1* make it seem any more rational to be certain that $\mathcal{A} = \mathcal{B}$. Others will have to make up their

⁷ See, for starters, KELLY (2005), ELGA (2007), CHRISTENSEN (2007, 2010) and CHRISTENSEN (2011).

⁸ See, e.g., SHOGENJI (ms) and FITELSON & JEHL (2009)

⁹ See, e.g., WAGNER (1985) and STAFFEL (2015, §6).

¹⁰ Cf. LEVINSTEIN (2015).

own minds, but, for my money, the culprit is the conjunction of (2) and (3). Perhaps, when we wish to use both Al's and Bert's forecasts to calibrate our own views about rain, there can be evidence we receive (like the evidence in *Example 1*) which makes it permissible to defer to neither Al nor Bert individually. More generally, when there are multiple experts whose opinions we wish to use to calibrate our own, there can be evidence we receive which makes it permissible to defer to none of them individually. If we take this tack, we might end up thinking that the following is a rational credence distribution in *Example 1*:

	r	$\neg r$
$\mathcal{A} = 1/3 \wedge \mathcal{B} = 1/3$	2/24	4/24
$\mathcal{A} = 1/3 \wedge \mathcal{B} = 2/3$	3/24	3/24
$\mathcal{A} = 2/3 \wedge \mathcal{B} = 1/3$	3/24	3/24
$\mathcal{A} = 2/3 \wedge \mathcal{B} = 2/3$	4/24	2/24

Even though this distribution violates both (2) and (3), it satisfies (4), with $\alpha = \beta = 1/2$.

Or perhaps, when we wish to use both Al's and Bert's forecasts to calibrate our own views about rain, we will need to have some priority ranking between them. Perhaps, given a conflict of opinion between Al and Bert, you should no longer defer to Bert, but continue deferring to Al. More generally, conditional on a collection of experts disagreeing, you should have some expert whose opinion trumps the opinions of the others. If we take this tack, we might end up thinking that the following is a rational credence distribution in *Example 1*:

	r	$\neg r$
$\mathcal{A} = 1/3 \wedge \mathcal{B} = 1/3$	2/24	4/24
$\mathcal{A} = 1/3 \wedge \mathcal{B} = 2/3$	2/24	4/24
$\mathcal{A} = 2/3 \wedge \mathcal{B} = 1/3$	4/24	2/24
$\mathcal{A} = 2/3 \wedge \mathcal{B} = 2/3$	4/24	2/24

Even though this distribution violates (3) and (4), it still satisfies (2). On this approach, Al's opinion completely trumps Bert's, and we should only align our opinions with Bert's insofar as they give us some indication of the opinions of Al. Thus, we may end up thinking that, just as the Monday objective chances completely trump the Sunday objective chances—just as we should disregard the Sunday objective chances conditional on the Monday objective chances—so too must your Monday rational self trump the Sunday objective chances—so too must you completely disregard the Sunday objective chances conditional on the opinion of your Monday rational self.

A PROOF OF PROPOSITIONS 1 AND 2

Proof. We establish three lemmas, from which the theorems follow immediately. (Note: throughout, I will use ' C ' indiscriminately for i) a joint probability density function over the values of \mathcal{A} and \mathcal{B} , as well as ii) the corresponding marginal densities, and iii) the corresponding probability function. In the event that there are at most finitely many possible values of \mathcal{A} and \mathcal{B} , ' C ' will everywhere denote a probability function and integrals may be exchanged for sums throughout.)

Lemma 1. *If (2), (3), and (4) hold, then so do (5) and (6).*

Proof. Since C is a countably additive, conglomerable probability, for all a ,

$$\begin{aligned} C(r \mid \mathcal{A} = a) &= \int_0^1 C(r \mid \mathcal{A} = a, \mathcal{B} = b) \cdot C(\mathcal{B} = b \mid \mathcal{A} = a) \cdot db \\ &= \int_0^1 (\alpha a + \beta b) \cdot C(\mathcal{B} = b \mid \mathcal{A} = a) \cdot db \\ &= \alpha a \cdot \int_0^1 C(\mathcal{B} = b \mid \mathcal{A} = a) \cdot db + \beta \int_0^1 b \cdot C(\mathcal{B} = b \mid \mathcal{A} = a) \cdot db \\ &= \alpha a + \beta \mathbb{E}[\mathcal{B} \mid \mathcal{A} = a] \end{aligned}$$

Then, because $C(r \mid \mathcal{A} = a) = a$ and $\beta = 1 - \alpha$, we have (6). Following the same procedure, with ‘ \mathcal{A} ’ and ‘ \mathcal{B} ’ exchanged throughout, establishes (5). \square

Lemma 2. *If (5) and (6) hold, then so does (7).*

$$(7) \quad \mathbb{E}[\mathcal{A}\mathcal{B}] = \mathbb{E}[\mathcal{A}^2] = \mathbb{E}[\mathcal{B}^2]$$

Proof.

$$\begin{aligned} \mathbb{E}[\mathcal{A}\mathcal{B}] &= \int_0^1 \int_0^1 ab \cdot C(\mathcal{A} = a, \mathcal{B} = b) \cdot da \cdot db \\ &= \int_0^1 a \cdot C(\mathcal{A} = a) \cdot \left[\int_0^1 b \cdot C(\mathcal{B} = b \mid \mathcal{A} = a) \cdot db \right] \cdot da \\ &= \int_0^1 a \cdot C(\mathcal{A} = a) \cdot \mathbb{E}[\mathcal{B} \mid \mathcal{A} = a] \cdot da \\ &= \int_0^1 a^2 \cdot C(\mathcal{A} = a) \cdot da \\ &= \mathbb{E}[\mathcal{A}^2] \end{aligned}$$

The same procedure, with ‘ \mathcal{A} ’ exchanged for ‘ \mathcal{B} ’ throughout, establishes that $\mathbb{E}[\mathcal{A}\mathcal{B}] = \mathbb{E}[\mathcal{B}^2]$. \square

Lemma 3. *If (7) holds, then so does (8).*

$$(8) \quad \mathbb{E}[(\mathcal{A} - \mathcal{B})^2] = 0$$

Proof.

$$\mathbb{E}[(\mathcal{A} - \mathcal{B})^2] = \mathbb{E}[\mathcal{A}^2] - 2\mathbb{E}[\mathcal{A}\mathcal{B}] + \mathbb{E}[\mathcal{B}^2] = 0$$

\square

If the expectation of $(\mathcal{A} - \mathcal{B})^2$ is 0, then $C(\mathcal{A} = \mathcal{B}) = 1$, and (1) is violated. \square

REFERENCES

- CHRISTENSEN, DAVID. 2007. "Epistemology of Disagreement: The Good News." *Philosophical Review*, vol. 116 (2): 187–217. [6]
- . 2010. "Rational Reflection." *Philosophical Perspectives*, vol. 24: 121–140. [5], [6]
- . 2011. "Disagreement, Question-Begging, and Epistemic Self-Criticism." *Philosopher's Imprint*, vol. 11 (6). [6]
- ELGA, ADAM. 2007. "Reflection and Disagreement." *Noûs*, vol. 41 (3): 478–502. [2], [6]
- . 2013. "The puzzle of the unmarked clock and the new rational reflection principle." *Philosophical Studies*, vol. 164: 127–139. [5]
- FITELSON, BRANDEN & DAVID JEHL. 2009. "What is the 'Equal Weight View'?" *Episteme*, vol. 6 (3): 280–293. [6]
- GAIFMAN, HAIM. 1988. "A Theory of Higher Order Probabilities." In *Causation, Chance, and Credence: Proceedings of the Irvine Conference on Probability and Causation*, BRIAN SKYRMS & WILLIAM L. HARPER, editors, vol. 1, 191–220. Kluwer Academic Publishers, Dordrecht. [1], [5]
- GALLOW, J. DMITRI. msa. "Expert Deference and News from the Future." [5]
- . msb. "Kinds of Experts, Forms of Deference." [2]
- HALL, NED. 1994. "Correcting the Guide to Objective Chance." *Mind*, vol. 103 (412): 505–517. [5]
- HALL, NED & FRANK ARNTZENIUS. 2003. "On What We Know About Chance." *The British Journal for the Philosophy of Science*, vol. 54 (2): 171–179. [5]
- KELLY, THOMAS. 2005. "The Epistemic Significance of Disagreement." In *Oxford Studies in Epistemology*, JOHN HAWTHORNE & TAMAR GENDLER, editors, vol. 1, 167–196. Oxford University Press. [6]
- LEVINSTEIN, BENJAMIN ANDERS. 2015. "With All Due Respect: The Macro-Epistemology of Disagreement." *Philosopher's Imprint*, vol. 15 (13): 1–20. [6]
- LEWIS, DAVID K. 1980. "A Subjectivist's Guide to Objective Chance." In *Studies in Inductive Logic and Probability*, RICHARD C. JEFFREY, editor, vol. II, 263–293. University of California Press, Berkeley. [2], [5]
- . 1994. "Humean Supervenience Debugged." *Mind*, vol. 103 (412): 473–490. [5]
- SCHAFFER, JONATHAN. 2003. "Principled Chances." *The British Journal for the Philosophy of Science*, vol. 54 (1): 27–41. [5]
- SHOGENJI, TOMOJI. ms. "A Conundrum in Bayesian Epistemology of Disagreement." [6]

STAFFEL, JULIA. 2015. "Disagreement and Epistemic Utility-Based Compromise." *Journal of Philosophical Logic*, vol. 44: 273–286. [6]

VAN FRAASSEN, BAS C. 1984. "Belief and the Will." *The Journal of Philosophy*, vol. 81 (5): 235–256. [1], [5]

—. 1995. "Belief and the Problem of Ulysses and the Sirens." *Philosophical Studies*, vol. 77: 7–37. [1], [5]

WAGNER, CARL. 1985. "On the formal properties of weighted averaging as a method of aggregation." *Synthese*, vol. 62 (1): 97–108. [6]