

Explosion and the Normativity of Logic

FLORIAN STEINBERGER
Birkbeck, University of London
f.steinberger@bbk.ac.uk

Logic has traditionally been construed as a normative discipline; it sets forth standards of correct reasoning. Explosion is a valid principle of classical logic. It states that an inconsistent set of propositions entails any proposition whatsoever. However, ordinary agents presumably do — occasionally, at least — have inconsistent belief sets. Yet it is false that such agents may, let alone ought to, believe any proposition they please. Therefore, our logic should not recognize explosion as a logical law. Call this the ‘normative argument against explosion’. Arguments of this type play — implicitly or explicitly — a central role in motivating paraconsistent logics. Branden Fitelson (2008), in a throwaway remark, has conjectured that there is no plausible ‘bridge principle’ articulating the normative link between logic and reasoning capable of supporting such arguments. This paper offers a critical evaluation of Fitelson’s conjecture, and hence of normative arguments for paraconsistency and the conceptions of logic’s normative status on which they repose. It is argued that Fitelson’s conjecture turns out to be correct: normative arguments for paraconsistency probably fail.

1. Introduction

The hallmark of paraconsistent logics is their rejection of the rule of explosion (henceforth EXP), or *ex contradictione sequitur quodlibet*, which states that an inconsistent set of premisses entails any proposition whatsoever.¹ EXP is standardly motivated with reference to the classical conception of validity as necessary truth-preservation in virtue of logical form: an argument departing from an inconsistent set of premisses, we tell our bewildered students, can never fail to carry the truth of the premisses over to the conclusion, because the premisses, by dint of their inconsistency, can never be jointly true in the first place. The paraconsistent logician rejects this line of argument

¹ Paraconsistent logics so understood are a broad church. Our chief focus here is on the philosophically better-known representatives of this family of logics: relevant logics in the Anderson-Belnap tradition (Anderson and Belnap 1975) as well as Neil Tennant’s version (Tennant 1997) and dialetheism (Priest 2006). On the other hand, many non-classical logics, perhaps most notably intuitionistic logic, endorse EXP. For simplicity’s sake, I will nevertheless describe the dispute over the validity of EXP as a debate between two (not so mellifluously named) parties: ‘paraconsistentists’ and ‘classicalists’.

along with the classical conception of validity that undergirds it. Her misgivings about the classical view of validity are nicely summarized by Graham Priest in the following:

[T]he notion of validity that comes out of the orthodox account is a strangely perverse one according to which any rule whose conclusion is a logical truth is valid and, conversely, any rule whose premisses contain a contradiction is valid. By a process that does not fall short of indoctrination most logicians have now had their sensibilities dulled to these glaring anomalies. However, this is possible only because logicians have also forgotten that logic is a normative subject: it is supposed to provide an account of correct reasoning. When seen in this light the full force of these absurdities can be appreciated. (Priest 1979, p. 297)

What particularly interests me here is the appeal to the normative status of logic. Priest's core criticism is that the classical conception of consequence pays insufficient heed to logic's essential role as a standard for correct reasoning. Once we remind ourselves of logic's normative role in our cognitive economy, the 'anomalous' aspects of classical logic—EXP, in particular—can be fully appreciated. In adopting this normative perspective, we thus understand why, even after years of classical indoctrination, many of us cannot shake the feeling that there is something fishy about EXP.

But what does Priest mean when he describes logic as a 'normative subject'? As a first stab we might advance the following interpretation: the normative connection between logic and thought consists in an agent's being committed to the logical consequences of her beliefs. Following this suggestion, we can reconstruct what appears to be the argument underpinning Priest's criticism. In Robert Meyer's polemical formulation the argument comes down to this:

[I]t is an evident fact that (1) some people sometimes are committed to some contradictory beliefs. And again, what else is logic for if it is not the case that (2) a man committed to certain beliefs is committed as well to their logical consequences? Friends, it is downright odd, silly and ridiculous that on classical logical terrain (1) and (2) cannot be held together, except on pain of maintaining that some people sometimes are committed to absolutely everything. (Meyer 1971)

In short, the validity of EXP is irreconcilable with logic's essential normativity, provided we assume, as seems eminently reasonable, that ordinary thinkers often (if not always) harbour inconsistent beliefs. Let us refer to arguments that seek to establish the untenability of EXP by showing its validity to be incompatible with the proper normative role of logic, *normative arguments against explosion* (or *normative*

arguments for short). Arguments of this type, I think, have an undeniable appeal. The aim of this paper is to assess their prospects of success.

The plan is as follows. In the following section, I offer a more careful formulation of the paradigmatic normative argument and show why it must fail when formulated in this way. I then articulate and refine Fitelson's conjecture to the effect that there is no 'bridge principle' — no way of precisifying the normative connection between logic and our modes of belief formation, retention and revision — that could salvage the normative argument. §3 addresses the assumption that ordinary thinkers are (and, in certain circumstances, perhaps should be) inconsistent believers. Drawing on and extending upon MacFarlane (2004), §4 introduces a taxonomy of bridge principles that will enable us to provide a *systematic* evaluation of Fitelson's conjecture. The evaluation is then carried out in §§5–7, where the framework is put to good philosophical use: it is argued that there is no successful way of reformulating the normative argument.

2. Normative arguments against EXP

The paraconsistent logician's beef with EXP is not standardly presented as turning on questions of logic's normativity. In the case of relevant logic, for instance, the dispute usually takes the form of an all-out disagreement about what the correct notion of validity should be. The relevantist claims that it is plainly absurd to maintain that the proposition that Mark Spitz is the current president of the United States can be validly inferred from the contradictory propositions that aardvarks are and are not indigenous to Africa. By contrast, the classicist will insist that the inference *is* valid, because necessarily truth-preserving. We thus find ourselves in a deadlock between two competing intuitions about what follows from what, with no clear way of adjudicating between them.²

² There is another option for the paraconsistentist. She can agree with the classical logician that validity is necessary truth-preservation, but maintain that propositions can be both true and false. EXP would indeed turn out to be invalid on this assumption. For suppose P is both true and false. That means that in the argument from $P \wedge \neg P$ to Q the premisses could be true, while the conclusion might not be. But on what grounds are we to believe that propositions can be simultaneously true and false? There are two ways of arguing for this thesis. The first is straight metaphysical dialetheism as advocated by Priest; the second is Nuel Belnap's interpretation of the truth-values in a four-valued logic as 'told-true' and 'told-false' (Belnap 1977). See also Lewis (1982). Here is not the place to discuss these proposals. I suspect, however, that neither of these options is particularly attractive to many paraconsistentists.

However, what arguably makes for the appeal of many paraconsistent approaches are not certain intuitions about the concept of validity but the ‘intuitions about good reasoning’ (MacFarlane 2004, p. 3) that underwrite them — intuitions that EXP-validating conceptions violate. Some advocates of paraconsistent logics like Meyer and Priest explicitly found their arguments on these intuitions.³ But even in the case of those arguments for paraconsistency based on the unintuitive consequences of the classical notion of validity — arguments from the ‘fallacies’ or ‘paradoxes’ of material implication — that do not explicitly appeal to the normative status of logic, it seems that it is the intuitions about correct reasoning that do the heavy lifting.⁴ For it is only when we assume that logic is in the business of providing a standard for correct reasoning that the arguments from the fallacies have the drawing power they do. To see this, suppose with Gilbert Harman that ‘there is no clearly significant way in which logic is specially relevant to reasoning’ (Harman 1986, p. 20) and assume that logic’s aim resides wholly in making an inventory of the argument schemata we deem valid. Viewed from this perspective, it is hard to see how the classical conception of validity could be faulted. Absent our intuitions about correct reasoning, what is the standard that classical validity is supposed to deviate from?

It is for this reason that Harman suggests that the case for relevant logic dries up once we realize that there is no normative link between logic and reasoning of the sort the relevant logician imagines. And it is for this same reason that MacFarlane believes that our only hope of transcending the stale, intuition-mongering debates between paraconsistentists and their adversaries over the correct notion of validity is by transposing ‘questions about logical validity into questions how we ought to think’ (MacFarlane 2004, p. 3).⁵

Does it follow from this that the case for paraconsistent logic stands and falls with the normative argument against explosion? Of course not. I have already pointed out alternative arguments for paraconsistency in note 2 above, and there are others besides. Nevertheless, if the

³ Very similar points are made by Anderson and Belnap (1975, p. 13).

⁴ See Mares (2004, p. 3) and Read (1988, p. 24) for two examples of arguments of this type.

⁵ The claim here is not that there are *no* arguments for paraconsistency that do not explicitly rely on substantive assumptions about the normativity of logic, but rather that such assumptions are implicit in, and in fact central to, such arguments. Paraconsistent logics would not have the following they do could they not rely on the underlying assumptions concerning logic and its role in capturing our intuitions about good reasoning.

normative argument fails, as I believe it does, this would be a result of some moment; it would block the most direct—and arguably the most natural and well-trodden—route to paraconsistency.⁶ Moreover, our investigations offer a case study, which, I hope, will shed light on alternative, normativity-based arguments for paraconsistency, and on the role that considerations pertaining to the normativity of logic can play in arguments for non-classical logical revisions more generally.

To get started, let us consider one way—a rather flat-footed way, as we will see—of spelling out the normative argument more carefully:⁷

- (1) EXP is valid.
- (2) S believes each member of an inconsistent set of propositions Φ .
- (3) If $P_1, \dots, P_n \models Q$, then if S believes the P_i , S ought to believe Q .
- (4) Even if S 's set of beliefs is inconsistent and any proposition Q whatsoever is entailed by it (courtesy of EXP), there are Q s such that S ought not to believe Q .
- (5) $\Phi \models Q$ for some patently unacceptable Q that S ought not to believe (from 1 and 2).
- (6) S ought not to believe Q (from 4).
- (7) S ought to believe Q (from 2, 3 and 5 via *modus ponens*).
- (8) Contradiction (from 6 and 7).
- (9) EXP is invalid (from 1 by *reductio*).

Note that the phrase ‘ S ought not to believe Q ’ in 4 is to be understood as ‘ S ought to refrain from believing Q ’ (which is equivalent to the claim that it is not permissible for S to believe Q), as opposed to ‘It is

⁶ Notice that standard arguments from the utility of inconsistent but non-trivial theories also fall into the category of normative arguments. The reason we regard such theories as non-trivial is because the practitioners, mindful of the normative authority of logic over their reasoning, are thought tacitly to operate on the basis of paraconsistent principles. See, for instance, Priest, Tanaka and Weber (2013).

⁷ The formulation is inspired by Fitelson (2008).

not the case that S ought to believe Q (which is compatible with it being permissible for S to believe Q).⁸

What can be said about the normative argument thus spelled out? Well, it is undoubtedly valid. But the question is how effective it is. Not very effective at all, it would seem. The form of the argument is that of an instance of the rule of negation-introduction or *reductio*.⁹ As Fitelson observes, its ineffectiveness stems from the fact that the normative argument can easily be deflected by shifting the blame away from EXP and instead identifying premiss 3 as the culprit in the inconsistent set formed by premisses 1–4. Premiss 3 acts, in MacFarlane's apt terminology (2004), as a 'bridge principle' linking the *logical* concept of entailment and the *epistemological* concepts of inference and belief. As such, it encapsulates the paraconsistentist's assumption regarding the existence of a tight normative connection between instances of the relation of logical consequence and our agent's attitudes vis-à-vis the propositions that stand in the specified logical relations. Fitelson's claim is that *it*, rather than our principle EXP, is at fault: in light of the inconsistency of premisses 1–4, we ought to jettison the simplistic bridge principle along with the erroneous conception of logic's normative status that underlies it, not EXP and our time-honoured classical notion of logical consequence.

Fitelson's attack on premiss 3 is justified, I believe. The claim can be substantiated by two simple objections. The first is familiar from Gilbert Harman's work (e.g. Harman 1984, p. 107). Hence:

Harman's objection: Suppose I believe both P and $P \supset Q$ (and that I am aware of the entailment $P, P \supset Q \models Q$). It simply does not follow that I *may* believe Q , let alone that I *ought* to believe Q as 3 requires. Q may be absurd, or at least discounted by my evidence, in which case the rational course of action for me is not to comply blindly with *modus ponens* and so form the belief Q , but rather to abandon at least one of my antecedent beliefs, P and $P \supset Q$, in light of their unpalatable consequences.

The second argument is, I think, equally simple and equally effective. It is due to John Broome (2000, p. 85). Hence:

Broome's objection: Suppose I find myself believing P . Since $P \models P$ (for any P), premiss 3 entails that I ought to believe P . But that seems patently false. After all, I find myself believing all sorts of things; P may have been

⁸ The deontic modals invoked here will be clarified further in §3.

⁹ Not to be confused with the intuitionistically impermissible rule of *classical reductio*: $\Gamma, \neg P \vdash \perp; \Gamma \vdash P$.

acquired in a doxastically irresponsible way. Surely, then, the mere fact that by happenstance I believe *P* does not in general imply that I ought to, or even that I may, believe *P*.

I take these objections to tell decisively against the proposed bridge principle encapsulated in premiss 3.

At this stage the paraconsistentist will rightly complain that our reconstruction of the normative argument has been uncharitable. Little wonder that the argument (as we have presented it) holds no water: we saddled the paraconsistent logician with a patently false bridge principle! To give the normative argument a fair shake we must replace premiss 3 with a viable one. However, it is not enough for a substitute of premiss 3 to be viable; it must also be logically *strong enough* to support the normative argument. That is to say, in order for the *reductio* to go through it must generate an inconsistent set with the remaining premisses 1, 2 and 4. To emphasize: any candidate bridge principle the paraconsistentist might propose in lieu of 3 must satisfy the following two desiderata.

Plausibility: Any candidate bridge principle must be philosophically defensible.

Strength: Any candidate bridge principle must be sufficiently strong to ensure the argument's validity.

Now of course evaluating the plausibility of bridge principles is no simple feat. Conformity with the criterion of plausibility will presumably be a matter of degree, which will, in turn, be determined on a cost-benefit basis by assessing bridge principles against a range of desiderata.¹⁰ For most purposes, though, a weaker notion will serve fine. To this end, let us simply say that a candidate bridge principle qualifies as *minimally plausible* just in case it is immune to both Harman's and Broome's objections. Minimal plausibility is enough, it turns out, to dismiss a number of bridge principles out of hand.

Thus, the question we now face is whether a replacement bridge principle is to be had that satisfies both Plausibility and Strength. Premiss 3 was strong enough, but failed to meet even the requirement of minimal plausibility. Fitelson is of course fully aware that premiss 3 is but a 'straw man' bridge principle. The implausibility of premiss 3

¹⁰ The five criteria suggested by MacFarlane (2004, p. 11) would at least seem to offer a starting point.

notwithstanding, he advances the following intriguing conjecture in a footnote:

I should also note that in contexts where *S*'s beliefs are inconsistent, I doubt that *any* bridge principle (no matter how sophisticated) will serve the [paraconsistent logician's] purposes. Specifically, I suspect that the [paraconsistent logician] faces a dilemma: any bridge principle will either be *false* (while perhaps being strong enough to make their *reductio* classically valid), or it will be *too weak* to make their *reductio* classically valid (while perhaps being true). The naive bridge principle (3) stated above falls under the first horn of this dilemma. More plausible bridge principles will (I bet) not yield a valid *reductio*. (Fitelson 2008, p. 6 fn. 10)

In the following I propose to put Fitelson's conjecture to the test. That is, I propose to probe whether it is in fact true that any candidate principle will fall foul of at least one of the aforementioned criteria: it will either be philosophically untenable (violating Plausibility) or too weak (violating Strength) to close the gap in the normative argument.¹¹

3. On being inconsistent

For the normative argument to be at all plausible, it must be granted that there are instances in which premiss 2 comes out true. It behooves us, therefore, before we inquire into the prospects of finding a satisfactory replacement for premiss 3, to say a few words about the paraconsistentist's assumption that believers are inconsistent — let us refer to this as the *inconsistency assumption*.

There are several familiar reasons for thinking that ordinary reasoners are inconsistent. An agent may have inconsistent beliefs without being aware of it because she is inattentive or because discovering the inconsistency in a vast network of dispositional and implicit beliefs would simply be too difficult, perhaps even humanly impossible. Moreover, not only individuals are prone to inconsistency. Presumably examples of inconsistency also abound in belief sets shared by groups or societies. For example, the statutes of a club, or legal codes more generally, may harbour inconsistencies. (See Priest 1987, ch. 13.) Similarly, as we noted above, scientific theories may be inconsistent.¹² Paradoxes like the liar — whether or not one wishes to

¹¹ MacFarlane (2004, p. 16) raises the very same question with respect to relevant logic.

¹² Bohr's atomic theory (which is inconsistent with Maxwell's equations) is often mentioned as an example of a theory which, though generally recognized to be inconsistent, was nevertheless heavily relied upon by researchers. It should be noted, however, that some

treat them as dialetheia — may also be thought to be sources of doxastic inconsistency. Crucially for the paraconsistentist, in all of these cases inconsistency need not trivialize the body of propositions in question. Individuals or groups may decide to continue to embrace an inconsistent statute or theory (at least provisionally) because it turns out to be too difficult or costly to restore consistency. They may also continue to engage in their practices because they are simply unaware of the inconsistency of their ways.

All of these considerations rest on descriptive claims about the inconsistency-proneness of ordinary reasoners: *as a matter of fact* agents like us tend to have inconsistent beliefs or endorse inconsistent bodies of propositions. As such they are empirical claims, in principle verifiable by experimental psychology. And this is, of course, one way to substantiate the paraconsistent logician's assumption undergirding the normative argument.

At least in some of the cases we described we would not be inclined to consider the agents or groups in question to be rational, even though the agent may not be blameworthy, as, for example, in the case of inconsistencies that only an agent with superhuman cognitive resources could detect. However, consider now the following stronger claim: not only is the observation that agents typically have inconsistent belief sets descriptively adequate, but there are circumstances in which it is *rational* for agents to have inconsistent beliefs. In other words, there are situations in which agents not only do but *should* have inconsistent beliefs.

According to Harman (1986, p. 15), even upon recognizing her beliefs to be inconsistent, a rational agent may be within her rational rights in failing to take measures to resolve the inconsistency (at least in the short term), for instance, if the cost in terms of time, computational power, etc., of straightening out her belief set would simply be prohibitive. Instead of attempting to restore consistency, the reasonable thing for her to do might be simply to quarantine her inconsistent beliefs, seeing to it that she does not exploit them in inference.

Examples of this type are characterized by their refusal to abstract away from our limitations of time, focus, cognitive resources, and so on. But it is often held that a certain amount of idealization is

philosophers will balk at the idea that scientific theories are proper objects of belief. Some, most notably constructive empiricists, maintain that our attitude towards the propositions that constitute our scientific theories is one of acceptance rather than belief, where acceptance is a cognitive mental state characterized by greater susceptibility to our voluntary control than belief.

necessary in laying down norms of rationality. Norms of rationality, the thought goes, just are *regulative ideals* which, though they may be unattainable for ordinary agents, nevertheless set the standard to which we hold ourselves, and which we seek to approximate as much as possible. The notion of rationality appealed to in the case described by Harman, then, might not be taken to be a purely epistemic one. Rather, it might be taken to represent a case in which considerations of practical and/or prudential rationality illicitly encroach upon our standards of epistemic evaluation. A properly epistemic notion of rationality demands that we abstract away from our limitations, cognitive and otherwise.

Even so, there are arguably situations in which agents count as *epistemically* rational in spite of holding inconsistent beliefs. The thought is simple. A good epistemic agent seeks to have informative true beliefs about the world. In order to further that aim, she must collect and evaluate evidence. But clearly, even if her evidential situation is such that it strongly supports each of her beliefs individually, the evidence may still be misleading in that the total set of propositions believed is inconsistent. Scenarios such as these are dramatized in the familiar lottery and preface paradoxes. Here we focus on a version of the preface paradox.¹³ Let S be our agent and let P_1, \dots, P_n be a large, non-trivial set of her justified beliefs.¹⁴ Assuming that S is an interesting inquirer, it is likely and eminently reasonable for her to form the belief, Q , that she is mistaken about at least one of her original beliefs. But now the belief set composed of all the P_i and Q is inconsistent. For either at least one of the P_i is false or — if, miraculously, all of them turn out to be true — Q is false. Pending further evidence, the rational thing for S to do seems to be to tolerate the inconsistency in her belief system. Indeed, this is presumably the predicament of any agent without the benefit of having worldly truths directly *revealed* to her.¹⁵

¹³ Harman (1984, p. 109) hints at this ‘global’ version of the preface paradox.

¹⁴ ‘Non-trivial’ and ‘interesting’ simply serve to rule out silly ‘agents’ whose entire set of beliefs consists of banal propositions such as ‘0 is a natural number’, ‘1 is a natural number’, ‘2 is a natural number’, ...

¹⁵ David Christensen (2004, ch. 3) makes a convincing case for the significance and inevitability of preface paradox scenarios. Incidentally, Fitelson himself rejects global norms of consistency. Along with Kenny Easwaran (Easwaran and Fitelson 2015), he has developed a weaker coherence norm for full belief on the basis of Joyce-style accuracy-dominance considerations (Joyce 1998). Very roughly, a belief set is coherent if there is no alternative belief set

However, even authors who are unswayed by preface paradox–like considerations can hardly deny the weak descriptive version of the inconsistency assumption—that there are ordinary agents who at times harbour inconsistent belief sets. And this is all that is needed to make the normative argument valid.¹⁶

4. Parameters and principles

Let us return to our central question, ‘Is there a bridge principle that meets the requirements of Plausibility and of Strength that can successfully take the place of premiss 3?’ Suppose the verdict is positive. All that is required of us to establish this fact is that we produce a bridge principle that fits the bill. Granted, there is the further difficulty of settling the question of Plausibility (as opposed to mere minimal plausibility), but let us set that problem to one side for now. But what could possibly warrant a negative conclusion? In order to refute the normative argument definitively, we would need to examine *all* possible bridge principles. But how could we possibly be confident that we have exhaustively examined all eligible candidates? The problem is that we have no clear conception of what is to count as an eligible candidate bridge principle. It would seem that, short of an oracle that provides us with an exhaustive list of bridge principles (along with a certificate of its completeness), we can never justifiably arrive at a negative conclusion. What to do?

In the absence of an oracle, I suggest we engage in good old-fashioned conceptual analysis—an analysis of the very notion of a bridge principle. We analyse bridge principles into their elementary constituents or parameters. Once we have identified all of the parameters and explored all the various ways in which these can be varied—all the possible ‘parameter settings’, as it were—we can map out the logical space of possible bridge principles—all the possible ways bridge principles can be generated on the basis of the initial stock of parameters and the (discrete) range of parameter settings. In this way, we arrive at a complete taxonomy of bridge principles (complete relative to the adequacy of our analysis), which we can then investigate in a *systematic* fashion. Based on such an analysis, we can say that the

that outperforms it in terms of its lower measure of inaccuracy across all possible worlds, that is, just in case it is not weakly dominated with respect to accuracy.

¹⁶ The stronger claim to the effect that it may be rational at times to hold inconsistent belief sets will occupy us again in §§6 and 7.

normative argument fails if a convincing case can be made that *none* of the bridge principles that can be generated within the scheme (i.e. none of the eligible candidates) are up to the job.

Thankfully, John MacFarlane (2004) has already done the bulk of the taxonomic work. In the following, I will show how MacFarlane's classification of bridge principles (along with some extensions of my own) can be used to make the evaluation of the normative argument philosophically tractable. But first let me reconstruct the analysis of bridge principles and the taxonomy it induces.

Let us begin by delineating the general shape — the blueprint, if you like — of a bridge principle. A bridge principle is a material conditional of the following form:¹⁷

$$\text{If } P_1, P_2, \dots, P_n \models Q, \text{ then } \Phi(P_1, P_2, \dots, P_n, Q)$$

where the antecedent states a 'fact' about logical consequence and the consequent takes the form of a normative claim featuring the agent's attitudes towards the propositions in question. Our premiss 3 can be seen to fit this mould:

$$\text{If } P_1, \dots, P_n \models Q, \text{ then if } S \text{ believes all the } P_i, S \text{ ought to believe } Q.$$

Moreover, it illustrates that the consequent *itself* often takes the form of a material conditional (though not always — this, as we will see, is one of the ways in which my taxonomy goes beyond MacFarlane's). I will often refer to the conditional in the consequent as the *embedded conditional* (as opposed to the *main conditional*). So much for the basic form of bridge principles. What now are the parameters, and what is their range of variability? Well, normative claims require deontic vocabulary.¹⁸ Following MacFarlane, I distinguish three deontic operators: *ought* (o), *may/permission* (p) and (*defeasible*) *reason* (r). And each of these operators can be given three distinct types of scope with respect to the embedded conditional:

- (C) Narrow scope with respect to the consequent: $(A \supset O(B))$.
- (W) Wide scope across the entire embedded conditional: $O(A \supset B)$.

¹⁷ I take the simplifying assumption that the conditionals in question be material conditionals to be harmless in the present context, even from the paraconsistentist's point of view. No potentially objectionable features of the material conditional are exploited.

¹⁸ Unless, of course, the norm is expressed via the imperatival mood of the sentences by means of which it is expressed. Such normative claims need not concern us here.

- (B) Binding both the antecedent and the consequent of the embedded conditional: $(O(A) \supset O(B))$.

(‘*O*’ here functions as a placeholder for deontic operators.) Furthermore, we can distinguish two types of polarities for normative claims: our claim could either be a *positive* obligation/permission/reason to *believe* (+); or it could be a *negative* obligation/permission/reason *not to disbelieve* (–).¹⁹ Note that disbelieving is a mental state ‘that stands in the same relation to believing as denying stands to asserting’ (MacFarlane 2004, p. 8). Hence, ‘disbelieving’ is to be distinguished from ‘not believing’. There are many propositions I neither believe nor disbelieve, either because I have not considered the proposition in question or because, upon consideration, I choose to suspend judgement about it. By contrast, if a proposition I have considered is discredited by the evidence, *disbelief* is the prima facie appropriate mental attitude. For simplicity, I will identify disbelieving *P* with believing not-*P*, but nothing much hinges on this here.

To illustrate the workings of the framework, notice that premiss 3 corresponds to

(Co+) If $P_1, \dots, P_n \models Q$, then if *S* believes all the P_i , *S*
ought to believe *Q*

Here ‘*C*’ designates the scope of the operator (that the operator is given narrow scope with respect to the consequent); ‘*o*’ marks the type of operator (*ought*); and ‘+’ tracks the polarity (that it is an ‘ought-to-believe’ rather than an ‘ought-not-to-disbelieve’). In this way, MacFarlane’s nomenclature (which I will adopt and extend) enables us to designate bridge principles uniquely.

There is more. Some will find premiss 3, and indeed all of the bridge principles that can be generated by means of the machinery introduced thus far, to be excessively demanding. The antecedent of the main (as opposed to the embedded) conditional states a fact about logical consequence that is in no way sensitive to the agent’s logical knowledge or her capacities for recognizing logical consequences of her beliefs. (Co+), i.e. premiss 3, states that I ought to endorse the

¹⁹ As has often been noted, a problem with ‘ought’ in English is that it does not admit of nominalization. ‘Obligation’ is an imperfect surrogate because not every true ought-claim entails a corresponding obligation. It may be, for instance, that I ought to get new shoes, but I am under no obligation to do so. My talk of obligation, here and throughout, should therefore be taken with a grain of salt.

consequences of all of my beliefs, even when the complexity and/or length of the shortest (non-trivial) deductive proof of these consequences far surpasses my cognitive and temporal resources. In particular, I ought to believe each and every theorem of Peano arithmetic provided I believe the axioms. On the principle that *ought* implies *can*, and given that it would not be humanly possible literally to conform to (Co+), some will push for the following weaker variant of premiss 3:

(Co + k) If S knows that $P_1, \dots, P_n \models Q$, then if S believes the P_i , S ought to believe Q .

Let us call this the *epistemically constrained* variant of (Co+) — MacFarlane refers to such principles as the ‘k-variants’ of the corresponding unconstrained principles.²⁰ While the epistemically constrained principle avoids the potential objection of being excessively demanding, it invites objections of its own. According to it, our obligations or permissions extend only as far as our logical knowledge. But as MacFarlane points out, one might worry that this in fact creates a disincentive to extend one’s logical knowledge; after all, the more (logically) ignorant I am, the freer I am to believe as I please.

But this looks backwards. We seek logical knowledge so that we will know how we ought to revise our beliefs: not just how we will be obligated to revise them when we acquire this logical knowledge, but how we are obligated to revise them even now, in our state of ignorance. (MacFarlane 2004, p. 12)

I will make no attempt to settle the question of whether epistemically restricted or unrestricted bridge principles are ultimately to be preferred. For expositional simplicity, I will continue formulating all bridge principles in their epistemically unconstrained form. But it is important to note that nothing hinges on this; our discussion would proceed in just the same way were we to consider k-variants of all the principles we are investigating.²¹

²⁰ The literature on epistemic closure principles knows a variety of different types of constraints: ‘If S is justified in believing that ...’, ‘If S ought to know that ...’, etc. Moreover, many internalists take epistemic rationality to supervene on non-factive mental states. They would consider only bridge principles whose antecedents are restricted to non-factive attitudes, that is, typically, to the agent’s beliefs about what follows from what. I will pass over these issues here.

²¹ In fact, a related question arises with respect to premisses 2 (the paraconsistent logician’s inconsistency assumption) and 4. Premiss 4 affirms that there are propositions that S ought not to believe in light of the inconsistency of her belief set. However, premiss 2 admits of both an epistemically constrained and an epistemically unconstrained reading. On our formulation of it, the normative argument rests on the epistemically *unconstrained* assumption that S has

So much for our analysis of the bridge principles. However, before we proceed further, a number of comments concerning the deontic modals figuring in our bridge principles are in order. Our three deontic operators — *ought*, *may* and *reasons* — express norms of *theoretical rationality*. Also, I will treat all deontic modals as propositional operators. This is not uncontroversial. Peter Geach (1982) and, recently, Mark Schroeder (2011) have argued that so-called deliberative *oughts* — roughly, the kind of *oughts* that guide us in first-personal deliberation — are best analysed, not as propositional operators, but as relations between agents and actions.²² Nevertheless, I assume here without further argument that the operator-reading can be made to work. (For defences of this position see, for example, Broome 2000, Broome 2013, and Wedgwood 2006.)

While *ought* and *may* are understood to be strict or ‘all-things-considered’ notions, having *reasons* is a defeasible and *pro tanto* or contributory notion. Having a reason to ϕ is compatible with simultaneously having reasons not to ϕ and even with it being the case that I ought not to ϕ . Reasons, unlike *oughts*, may be weighed against each other; the side that wins out determines what ought to be done. It may be perfectly proper for me to have reasons to ϕ and yet not to ϕ because my reasons are overridden, whereas if ϕ -ing is what I ought to do and I fail to do so, I am, in Broome’s words, ‘not entirely as I should be’.

Next, we must be clear what kind of *ought* is being invoked. I shall assume that we are dealing with agential ‘practical’ or ‘deliberative *oughts*’. Roughly, I take *oughts* of this sort to allocate responsibility for an action to an agent (e.g. ‘Noa really ought to call her mother’). ‘Evaluative *oughts*’, by contrast, present a certain state of affairs as generally desirable (e.g. ‘Pasta ought to be cooked *al dente*’) without imputing responsibility to anyone in any obvious way.²³ This assumption too is not wholly uncontroversial. It might be thought that the tension between the demands of epistemic (including logical) norms and doxastic involuntarism is best finessed by treating doxastic *oughts*

(possibly unbeknownst to her) an inconsistent belief set. However, we might have opted for an epistemically constrained version, whereupon premiss 2 concerns itself only with recognized inconsistency. As we have seen in the foregoing section, agents plausibly find themselves in both types of situation. Again, the normative argument can be run on the basis of both the epistemically constrained and the epistemically unconstrained reading of premiss 2.

²² See (Schroeder 2011, §2.1) for a detailed characterization of deliberative or practical *oughts*.

²³ For the underlying distinction, see Sellars (1969) and Humberstone (1971).

as evaluative *oughts*, or ‘ought-to-be’s in the Sellarsian phrase, and hence not as *can*-implying *oughts*. Matthew Chrisman (2008) has recently proposed an account of doxastic *oughts* along evaluative lines. Richard Feldman’s approach (2000), which appeals to ‘role oughts’ falls into the same general category, but is criticized by Chrisman. Chrisman’s approach has merit, but suffers, as he concedes, from the lack of a clear connection between the rules of criticism that are the ‘ought-to-be’s and the corresponding rules of conduct that are the ‘ought-to-do’s.

A further choice point is whether the *oughts* in question are ‘subjective’ or ‘objective’ (the distinction straightforwardly carries over to our remaining deontic concepts).²⁴ The underlying thought is that an agent’s conduct can be appraised from the standpoint of the informational state of the agent or from the standpoint of a superior or ideal state of information. Imagine our friend George wandering through a maze.²⁵ We are perched on a tree above the maze overseeing the event. From our privileged vantage point, we are able to say, ‘George has no way of knowing it, but he ought to take a left there.’ The *ought* in question is the objective *ought*. But we can also imagine a situation in which we might employ a subjective *ought*. Suppose George has information (from an otherwise trustworthy source) about the layout of the maze which, unbeknownst to him, is erroneous. We might then find ourselves saying, ‘By George’s lights it makes most sense to turn right there. So he ought to turn right.’

We are now in a position to see that there are natural pairings between subjective deontic operators and epistemically constrained readings of our bridge principles, on the one hand, and objective deontic operators and epistemically unconstrained operators, on the other. For instance, it makes sense to read the following epistemically unconstrained version of (Co+),

If $P_1, \dots, P_n \models Q$, then if S believes the P_i , S ought to believe Q

²⁴ I largely follow Wedgwood (forthcoming) here. But nothing much hinges on this particular account. All that matters is that our account is able to track the sensitivity of deontic modals to the agent’s informational state (and possibly those of other agents advising or evaluating her). How this sensitivity is accommodated in a semantic account of deontic modals need not concern us. For instance, we might, for our purposes, opt equally well for a contextualist or for a relativist treatment like those discussed by Kolodny and MacFarlane (2010, p. 29).

²⁵ The example is borrowed from Wedgwood (2012).

as a claim, where the *ought* is to be understood in the objective sense. By contrast,

If S recognizes that $P_1, \dots, P_n \models Q$, then if S believes the P_i , S ought to believe Q

is a claim where *ought* receives a subjective reading that is sensitive to the speaker's informational context. Accordingly, let us call the first type of bridge principle *objective*, the second type *subjective*. Most of what follows is compatible with both types of bridge principles.

5. The Cs and Ps

So much by way of general comments about our bridge principles. Let us turn to the possible replacements for premiss 3. We have already seen that (Co+) will not do the trick. A moment's reflection reveals that the remaining members of the C-family with strict deontic operators fail, and fail for exactly the same reason: they are sufficiently strong to support the argument, but violate minimal plausibility.

(Co+) If $P_1, \dots, P_n \models Q$, then if S believes all the P_i , S ought to believe Q .

(Co-) If $P_1, \dots, P_n \models Q$, then if S believes all the P_i , S ought not to disbelieve Q .

(Cp+) If $P_1, \dots, P_n \models Q$, then if S believes all the P_i , S may believe Q .

(Cp-) If $P_1, \dots, P_n \models Q$, then if S believes all the P_i , S is permitted to not disbelieve Q .

What goes wrong in each of these cases is that the combination of narrow scope and strict operators makes the normative force of logic implausibly demanding. According to these principles, the normative force of logic is too strong relative to that of our non-logical doxastic norms; for example, according to (Co+) I ought to believe Q because it follows from my beliefs, no matter how good my reasons are for disbelieving Q . A natural reaction in light of these failings is to retreat to the weaker *reasons* operator, thus giving rise to

(Cr+) If $P_1, \dots, P_n \models Q$, then if S believes all the P_i , S has (defeasible) reason to believe Q .

(Cr-) If $P_1, \dots, P_n \models Q$, then if S believes all the P_i , S has (defeasible) reason not to disbelieve Q .

And this move does have some initial promise. Unlike its close cousins the (Co)s and the (Cp)s, the (Cr)s appear to be immune against

Harman's objections. For simplicity, let us focus on (Cr+). My having a defeasible reason for believing *Q* as a result of *Q*'s being entailed by my antecedent beliefs is perfectly compatible with my simultaneously having good, perhaps better, reasons for not believing *Q* (e.g. the available evidence supports not-*Q*). This seems initially plausible: my recognizing that a certain proposition *Q* is entailed by my beliefs gives me (defeasible) reasons to believe *Q*, but these reasons can be trumped by epistemic reasons for disbelieving *Q*.

What about Broome's objection? Things are somewhat less straightforward here. Still, a case can arguably be made that (Cr+) can parry Broome's objection, provided one is willing to accept certain epistemological background assumptions. The epistemological accounts germane to these assumptions are ones characterized by the fact that they promote what we might call an innocent-until-proven-guilty policy concerning belief maintenance. The idea behind such approaches—Harman (2002, p. 10) has dubbed them *general foundations theories*—is a kind of conservatism about belief: an agent's belief set enjoys a kind of default justification until she encounters sufficiently strong countervailing evidence. On such views, then, it seems proper to say, at first blush at least, that I *do* have reason to believe any proposition *I* in fact believe. I have reason to stick to my beliefs unless and until I am presented with sufficiently strong grounds for abandoning them. Of course, adopting (Cr+) might therefore mean tying the fate of the normative argument to that of a general foundations theory of belief maintenance, but perhaps that is a lot the paraconsistent logician should embrace.

The crucial question now is whether (Cr+) also fulfils Strength. That is, is (Cr+) strong enough to support the *reductio*? For the *reductio* to succeed, premisses 1, 2, 4 and (Cr+) must form an inconsistent set, thereby laying the ground for rejecting EXP. It is here, it turns out, that (Cr+) founders. (Since (Cr+) is the strongest of the (Cr)s, the argument from insufficient strength will apply *a fortiori* to the remaining (Cr)s.) In fact, *the very reason* that (Cr+) is able to evade Harman's and Broome's objections is also the reason for its demise, the fact, namely, that having an obligation *not* to believe *P* (e.g. as dictated by certain doxastic norms) is compatible with having *pro tanto* reasons for believing *P* on account of its being entailed by one's extant beliefs. In more detail: suppose *S* has inconsistent beliefs. EXP and (Cr+) together imply that *S* has, for any proposition *P*, defeasible reasons for believing *P*. But merely having defeasible reasons for believing *P* is consistent with *P* being a proposition of the sort

premiss 4 mandates us not to believe. (Cr+) can be superseded by the epistemic norms that underlie premiss 4. Hence, (Cr+) is compatible with 1, 2 and 4, and so fails to satisfy Strength.²⁶

This is a strong blow against the (Cr)s. But the paraconsistentist might try the following tack. It is true that (Cr+) does not support the *original* normative argument. But we can easily tweak premiss 4 to make it work. Simply replace premiss 4 with the following strengthened version:

- 4'. Even if *S*'s set of beliefs is inconsistent and any proposition *Q* whatsoever is entailed by it (courtesy of EXP), there are *Q*s such that *S* has *no reason* to believe *Q*.

The combination of (Cr+) and 4' will indeed support the *reductio*. But should we buy into 4', or does its extra strength also render it less plausible? After all, why should there be propositions that I have no reason to believe? Before addressing the question directly, we should note that there is a scope ambiguity in the statement of 4' (and 4). Is 4' saying that there is some proposition *P* such that, for every agent *S*, *S* has no reason to believe *P*? Not so. At least not if we focus on factual propositions (as I will for present purposes). Arguably, there could be agents in bizarre enough epistemic situations that could produce deranged enough belief systems such that an agent with such a belief system *does* have a (subjective) reason to believe that Mark Spitz is the current president of the United States, as well as any other inane proposition you might come up with. The more plausible reading, therefore, is that for every agent *S*, there exists some proposition *P* such that *S* has no reason to believe *P*. But even this reading raises some doubt. For if (Cr+) is philosophically tenable, as the advocate of 4' maintains, then I *do* have a reason (albeit a defeasible one) for believing any logical consequence of my beliefs. And so, saying that I have no reason whatsoever to believe *P* given my background beliefs and my evidence is to say that I *could not possibly believe* (not merely that I do not have any reason to believe) any set of propositions that entails *P* (otherwise, by (Cr+), I *would* have a reason to believe *P*). But that seems implausibly strong.²⁷

²⁶ MacFarlane (2004, pp. 9–10) presents an interesting alternative argument to the effect that anyone who adopts (Cr) must in fact be committed to a further principle, (Br), which he deems problematic. While I will not discuss MacFarlane's claim to the effect that (Br) is the only plausible motivation one may have for adhering to (Cr), I will discuss the Bs in §5.

²⁷ Notice that this is importantly different from the claim that for every agent there are propositions such that if they are entailed by the agent's beliefs, the agent has reason to revise

I think this is a major strike against such attempts at salvaging the (Cr)s. More importantly, though, even if the normative argument could be modified so as to restore its validity by means that do not undermine the tenability of its premisses, the paraconsistentist would still have to vindicate (Cr+). Saying that (Cr+) is minimally plausible is one thing; saying that it satisfies Plausibility is quite another. I submit that (Cr+) is indefensibly weak, and therefore not philosophically tenable. To see this, consider again the ‘Harman scenario’: I believe $P \supset Q$ and P . (Cr+) gives me a reason to believe Q . But suppose I have decisive reasons for disbelieving Q . In that case, my logic-induced reason for believing Q will be trumped by my epistemic reasons for disbelieving Q . In a way, that is as it should be, as we have seen: it avoids Harman’s objection. However, according to (Cr+), that is *all* that logic requires of me. So long as I am appropriately sensitive to the reasons for believing Q stemming from its being entailed by my antecedent beliefs, I have discharged my logical duties. In particular, (Cr+) in no way requires of me that I should revise my beliefs in light of their consequences. Surely, though, logical coherence *does* demand that I modify my belief set so as to avoid blatant inconsistencies. True, we have seen that there may be reasons for thinking that there are cases of rational doxastic inconsistency. However, the adoption of (Cr+) would license *pandemic* inconsistencies that go well beyond such rarefied cases of ‘reasonable inconsistency’. The solution, therefore, cannot lie with the (Cr)s (at least not taken on their own).

So far, then, Fitelson’s conjecture has held up. But then again, none of the bridge principles considered up to this point — with the exception, perhaps, of the (Cr)s — held any philosophical promise. At this stage I want to consider a further family of bridge principles, not treated in MacFarlane (2004), which deserve our consideration. The key idea is that the embedded conditional in the bridge principles we have been considering should be replaced by a primitive, indecomposable, dyadic conditional deontic operator. In the case of *ought*, this amounts to ‘ $Ought_S(Q|P)$ ’, read as ‘ S ought to Q conditional on P ’,²⁸ which is true just in case Q holds at all the most ideal worlds at which P is true. Conditional operators of this type have been proposed to deal with well-known paradoxes in deontic logic arising from

her initial beliefs. (Cr+), being a narrow-scope principle, makes no provision for belief revisions.

²⁸ I will henceforth suppress the index.

‘contrary-to-duty obligations’ (for example, *if I kill my neighbour, I ought to do it humanely*).²⁹

On the basis of such operators we can create a novel family of principles (call it ‘P’ for ‘primitive’):

(Po+) If $P_1, \dots, P_n \models Q$, then Ought (S believes $Q|S$ believes the P_i).

(Pp+) If $P_1, \dots, P_n \models Q$, then May (S believes $Q|S$ believes the P_i).

(Pr+) If $P_1, \dots, P_n \models Q$, then Reasons (S believes $Q|S$ believes the P_i).³⁰

The trouble with conditional operators from our perspective is that they fail to satisfy Strength, on account of the fact that the ‘consequent’ in our conditional operators—the operator-involving claim—does not detach. In other words, the pattern of inference

- (1) S ought to Q conditional on P
- (2) P
- (3) Therefore, S ought to Q

is invalid.³¹ (If it were valid, our dyadic operators would simply reduce to the Cs, and so would be inadequate in dealing with the aforementioned paradoxes.) However, the conclusion of the argument would need to detach in order for us to be able to generate the desired contradiction. Accordingly, the Ps do not generate a contradiction with 4, and therefore turn out to be too weak to support the normative argument. This type of weakness equally afflicts the Ws, as we will see.³²

6. The Ws

In the face of the inadequacies of strict narrow-scope bridge principles and of dyadic conditional operators, the choice one is presented with is either to opt for a slack reasons operator or to become a ‘wide-scope’. Offered this choice, many authors have plumped for the latter option.³³ And indeed, in the light of our discussion of the (Cr)s, this

²⁹ See von Wright (1956).

³⁰ It is not hard to imagine the corresponding negative variants.

³¹ For simplicity I focus on the case of *ought*. The cases of the other two operators are analogous.

³² For further objections against dyadic conditional deontic operators, see Kolodny and MacFarlane (2010).

³³ The question arises, not only in the case of logical coherence norms, but also for other principles of theoretical and practical rationality (for instance, that one ought to bring about

may seem to be the paraconsistent logician's best bet. In the present context this amounts to espousing a member of what we might dub the 'W-family'. Its distinctive feature, recall, is that the deontic operators are given *wide* scope over the embedded conditional: $O(A \supset B)$. We arrive at the following three wide-scope principles and their negative counterparts:

(Wo+) If $P_1, \dots, P_n \models Q$, then S ought to see to it that (if S believes all the P_i , S believes Q).

(Wp+) If $P_1, \dots, P_n \models Q$, then S may see to it that (if S believes all the P_i , S believes Q).

(Wr+) If $P_1, \dots, P_n \models Q$, then S has reason to see to it that (if S believes all the P_i , S believes Q).

To facilitate the exposition I will focus on (Wo+). Everything I say in the discussion to follow carries over straightforwardly to the remaining Ws.³⁴

(Wo+) has considerable intuitive upside. It elegantly dodges both of our criteria for minimal plausibility: neither Harman's objection nor Broome's point get any purchase on (Wo+). As for the former, the wide-scope reading provides just the wiggle room needed to neutralize Harman's objection. Suppose I believe P and $P \supset Q$. According to (Wo+), I am given the choice of either retaining my beliefs and also coming to believe Q or ditching at least one of my antecedent beliefs in P and $P \supset Q$ so as to absolve me from the obligation to believe Q . Should Q turn out to be untenable, the latter course of action recommends itself: we revise our beliefs in the light of their unpalatable consequences. This not only meets Harman's challenge but also seems to get the normative link between logic and reasoning exactly right: our processes of belief maintenance are constrained by facts about what follows from what, but they are not so narrowly constrained as to reduce us to mere theorem provers. We do not and should not merely churn out and endorse every last consequence of some initial belief set. Put another way, we should not look to logic to tell us what to believe. Rather, reasoning consists in negotiating global

the means deemed necessary to promote one's ends, that if one ought to ϕ one ought to intend to ϕ , and so on). See Way (2010) for a survey and further references.

³⁴ We may note immediately, though, that the (Wp)s are out of contention. The (Wp)s imply that I have permission to see to it that if I believe P and Q , I also believe $P \wedge Q$, but this does not even so much as provide me with a reason to believe the latter when I believe the former. As MacFarlane puts it, 'the difference between the (Wp)s and the position that there are no logical norms for belief seems slim indeed' (2004, p. 10).

logical demands of coherence with local epistemic norms that guide us in our doxastic deliberation. This element of negotiation is well captured by the wide-scope reading.

Moreover, (Wo+) is untroubled by Broome's reflexivity worries. The flaw with the *ought*- and *may*-based principles considered so far was that because P entails itself, the agent incurs an obligation (or permission) to believe P for any (perhaps unfounded) belief P she happens to be saddled with at a given point in time. (Wo+) imposes no such bogus obligations. The fact that I happen to believe P and that $P \models P$ only entails an obligation to believe P if I believe P , which is no obligation at all. It is because of these features that the W s enjoy considerable popularity.³⁵

The question now is whether the *prima facie* attractive W s are a good fit for the paraconsistentist. But here we immediately face a difficulty. We have seen that the normative argument is premised on what we called the inconsistency assumption—the assumption that reasoners do at times (perhaps even should at times) hold inconsistent beliefs. It is because (even rational) agents are likely to find themselves with logically incoherent beliefs *and* because logic ought to be normative for reasoning that the paraconsistent logician contends we should reject EXP. But now take an agent S who harbours an inconsistent belief set Φ , and suppose we accept EXP. (Wo+) requires S to do one of two things: it requires of S either that she renege on her belief in some of the Φ so as to restore consistency or, if she retains her beliefs in all of the members of Φ , that she must believe all of Φ 's consequences, i.e. *any* Q whatsoever. But this seems to show that (Wo+) is *consistent* with the normative argument's remaining premisses. This will always be so if one regards it to be a correct blanket policy always to avoid inconsistency. In other words, facing the choice between abandoning some of her present beliefs to avoid inconsistency and having to endorse any proposition whatever, the agent ought *always* to avoid inconsistency by revising her beliefs. On this (staunchly objective) reading of epistemic *oughts*, the W s would not produce the necessary contradiction with premiss 4 (in the presence of EXP), and so the naive argument would not go through.

However, I do not think the paraconsistentist should be overly troubled by this. So long as one is willing to grant that there are circumstances (such as preface paradox-type scenarios) in which

³⁵ MacFarlane himself falls into this category of advocates, as do Field (2009), Restall (2005) and Sainsbury (2002).

agents not only have but *ought to have* inconsistent belief sets, for the reasons mentioned in §2, we thereby grant that there are circumstances in which (Wo+) *would* enjoin us *not* to take consistency-restoring measures, but rather to stick to our inconsistent belief set and hence to incur a commitment to believing all of its consequences — and in the presence of EXP, this amounts to a commitment to believing any proposition whatsoever.

A version of the preface paradox will illustrate the point. Suppose I author a meticulously researched non-fiction book on cuttlefish. My book is composed of a large set of non-trivial propositions P_1, \dots, P_n about the extraordinary physiology and ethology of cuttlefish. Seeing that all of my claims are the product of scrupulous research, I have every reason firmly to believe each of the P_i individually. But I also have overwhelming inductive evidence for Q : that at least one of my beliefs is in error. The P_i and Q cannot be jointly true. How ought I react to the inconsistency in my belief set? Surely I cannot simply ditch some of the P_i willy-nilly to ensure consistency.³⁶ It would be irrational for me to abandon any particular proposition in the absence of any specific countervailing evidence. And it would be mad, of course, to ditch the whole lot of my beliefs. Indeed, as MacFarlane (2004, p. 15) rightly emphasizes, even if, irresponsibly, I wanted to sacrifice some or all of the claims in my book, unless one maintains a wildly far-fetched form of belief voluntarism, it is not something I can simply *decide* to do. Hence, consistency-restoring belief revision is not an option in the situation described (in the absence of further relevant evidence). The upshot of this, we said, is that (Wo+) requires of S that she believe all the consequences of her beliefs whenever she has good reason to hold inconsistent beliefs.

Have we thereby rehabilitated (Wo+) as a candidate bridge principle? For this to be the case it would have to meet both Plausibility and Strength. Despite its initial attraction, much more would have to be said about the philosophical viability of the W_s .³⁷ But even if we grant for the time being that a good case can be made for them,

³⁶ Acting in this way would carry the considerable risk of what Niko Kolodny calls satisfying a coherence requirement ‘against reason’ (Kolodny 2007). The idea is this. An agent with an inconsistent belief set can reimpose consistency in a variety of ways that are not sensitive to the evidence she is presented with. Suppose I have good epistemic reasons for believing P and that I also happen to believe not- P on rather flimsy grounds. As far as the consistency requirement is concerned it does not matter which of the two beliefs I abandon. Yet it is clear that given my evidential situation I should ditch not- P .

³⁷ See MacFarlane’s helpful discussion (2004, pp. 11–14).

I believe the *Ws* are not an option for the paraconsistentist. This is because the *Ws* still fail to satisfy Strength. This seems surprising. After all, as we have seen, in the presence of EXP *S* ought to believe any proposition whatsoever in scenarios in which *S* rationally holds inconsistent beliefs. And this seems to be exactly the type of absurd consequence the paraconsistent logician needs in order to get the normative argument off the ground. But this is a mistake. What is needed for the *reductio* to go through is an *explicit deontic claim* to the effect that *S* ought to or may believe an absurd or implausible proposition, thereby contradicting premiss 4. However (Wo+) yields no such claim. In fact (Wo+) is *consistent* with premisses 1, 2 and 4 *even* in cases in which consistency-restoring belief revision is out of the question. The reason is that the principle merely requires the agent to see to it that the conditional in the scope of the *ought*-operator is true. However, even in cases in which the antecedent inconsistent belief set is retained, no *ought*-claim can be detached. Indeed their non-detaching nature is one of the characteristic features of wide-scope principles (see Broome 2003). In other words, from ‘*S* ought to (believe *Q* if *S* believes Φ)’ (where Φ is an instance of a rationally held inconsistent belief set and *Q* an untenable proposition which *S* ought not to believe, according to premiss 4, but which obviously follows from Φ via EXP) and ‘*S* believes Φ ’ it follows that ‘*S* believes *Q*’ (provided that *S* complies with (Wo+)), but it does not follow that ‘*S* ought to believe *Q*’. But it is this latter, explicitly normative, consequence that the paraconsistentist needs to generate a contradiction with premiss 4 so as to make the *reductio* work. To be sure, in the circumstances described *S* will hold any number of beliefs which, according to premiss 4, she ought not to hold, and for this she may be liable to epistemic criticism, but this is not enough to render premisses 1, 2, (Wo+) and 4 inconsistent. It follows that even if we take into account scenarios in which inconsistency is inevitable, (Wo+) is too weak. Seeing that (Wo-) and the (Wr)s are strictly weaker than (Wo+), they suffer the same fate.

Is there a comeback for the paraconsistentist? Well, one might try to argue for an additional principle that *would* enable us to detach deontic consequences from (Wo+), thus closing the gap in the normative argument. Either one of the following two principles would do the trick:

- (A) $O(A \supset B) \supset (A \supset O(B))$
- (B) $O(A \supset B) \supset (O(A) \supset O(B))$

What can we say about them? (A) can immediately be seen to be unacceptable. In its presence the Ws would straightforwardly collapse into the Cs, which we have already discarded. (B) is of course just a version of the principle K that characterizes normal modal logics (modal logics admitting of a Kripke semantics). However, K, it seems fair to say, is rejected by most deontic logicians.³⁸ But even if the paraconsistentist were willing to bite the bullet on (B), adding it to the mix in effect amounts to endorsing a new, distinct family of bridge principles: the Bs. And it is to the Bs that we must turn in the next section.

7. The Bs

So far, then, Fitelson's conjecture appears to be right on the money: the *prima facie* most plausible bridge principles have revealed themselves to be too weak to support the normative argument's paraconsistent conclusions, while the sufficiently strong bridge principles do not stand up to scrutiny. But let us not jump to conclusions; we have yet to examine the Bs — the class of bridge principles whose characteristic feature is that the deontic operator occurs both in the antecedent and in the consequent of the embedded conditional. By varying the deontic operator, we can once again generate three principles along with their negative variants.³⁹

(Bo+) If $P_1, \dots, P_n \models Q$, then if S ought to believe all the P_i , S ought to believe Q .

(Bp+) If $P_1, \dots, P_n \models Q$, then if S may believe all the P_i , S may believe Q .

(Br+) If $P_1, \dots, P_n \models Q$, then if S has reason to believe all the P_i , S has reason to believe Q .

Are the Bs any better suited for the job at hand? Whether this is so depends on the norms underwriting the doxastic obligations and

³⁸ For reasons largely already provided by Chisholm (1963).

³⁹ We could further extend MacFarlane's classificatory scheme by allowing for 'mixed' Bs in which the deontic operators featuring in the antecedent and in the consequent of the embedded conditional could be distinct. For example, in addition to (Bo+), we could consider also

(Bop+): If $P_1, \dots, P_n \models Q$, then if S ought to believe all the P_i , S may believe Q .

(Bor+): If $P_1, \dots, P_n \models Q$, then if S ought to believe all the P_i , S has reason to believe Q .

and so on for all the possible combinations. However, it turns out that for present purposes there is no need to distinguish these additional cases.

permissions in question. Let me explain. The Bs really incorporate references to two distinct kinds of norms.⁴⁰ The *oughts* in the antecedents and those in the consequents of the embedded conditionals seem to stem from different normative sources. In the embedded conditional ‘if *S* ought to believe all the P_i , *S* ought to believe *Q*’, the *ought* in the antecedent refers to whatever doxastic norms (evidential norms, perhaps) make it the case that *S* ought to believe the P_i (for simplicity we may assume that the P_i are not themselves acquired by logical inference, and so the norms in question will not themselves be logical or logic-induced). The obligation to believe *Q*, by contrast, appears to stem from the normative force (if any) induced by logical consequence, along with that inherited from the doxastic norms that oblige us to believe the P_i . *S* ought to believe *Q* on the strength of it being the case that she ought to believe the P_i and it being the case that *Q* is logically entailed by the P_i ; the positive epistemic status of *S*’s beliefs in the P_i are propagated to their logical consequences.

What non-logical doxastic norm might underwrite the *ought* in the antecedent? Perhaps the most obvious contender for the non-logical doxastic norm in question is the truth norm (TN):⁴¹

(TN) For all *S*, for all *P*, if *S* considers or ought to consider *P*, (*S* ought to believe *P*) if and only if *P* is true.

However, a moment’s reflection reveals that the (TN)-based approach, like the blanket policy of inconsistency avoidance we discussed in the previous section, is not an option for the paraconsistent logician. The reason is that (TN) does not countenance scenarios in which an agent ought to believe an inconsistent set of propositions. But it is precisely this property that any doxastic norm underwriting (Bo+) would have to enjoy for (Bo+) to be a genuine contender. This, in turn, is because for (Bo+) to satisfy Strength, the consequent of the embedded conditional (‘*S* ought to believe *Q*’ for some unacceptable *Q*) would have to be detachable in some cases so as to generate an inconsistency with premiss 4. There would thus have to be instances in which the antecedent of the embedded conditional (‘*S* ought to believe all the P_i ’) is true even when the set of P_i is inconsistent. Clearly, though, on the (TN)-based interpretation of the antecedent this is impossible—it is never the case that an agent ought to believe an inconsistent set of

⁴⁰ For expositional convenience, my discussion will focus on (Bo+). But everything I say carries over *mutatis mutandis* to the remaining Bs, unless explicitly noted otherwise.

⁴¹ Anything I go on to say about the truth norm applies equally to Timothy Williamson-style knowledge norms (Williamson 2000).

propositions, because by definition at least one of the propositions in an inconsistent set cannot be true. It follows that the agent must be flouting (TN) with respect to at least one of the P_i . In other words, (TN) entails the norm of logical consistency.⁴²

A general lesson emerges from these considerations: due to the structure of (Bo+), where *ought* acts also on the antecedent of the embedded conditional, the simple inconsistency assumption is not enough. Not only must agents occasionally hold inconsistent belief sets, to satisfy the antecedent it must be possible to do so rationally. What is needed, therefore, is a doxastic norm to underwrite the *ought* in the antecedent that tolerates (indeed, in the case of the (Bo)s, sometimes mandates) inconsistent belief sets. Only then can the antecedent of the embedded conditional ever be satisfied and so the consequent detached. To see this, suppose that the doxastic norm S is subject to inconsistency-mandating and that they can be made out to be plausible. The paraconsistent logician's argument can then be seen to go through as follows. Let Φ be an inconsistent belief set that S ought to believe according to the doxastic norm in question.

- (a) $\Phi \models Q$ for some absurd proposition Q (by premiss 1, i.e. by the supposed validity of EXP).
- (b) If S ought to believe each member of Φ , then S ought to believe Q (by *modus ponens* from (Bo+) and (a)).
- (c) S ought to believe Q (by *modus ponens* from 2 and the assumption that S ought to believe each member of Φ).
- (d) S ought *not* believe Q (from premiss 3 and the fact that Q is absurd).
- (e) Contradiction (from (c) and (d)).

Notice that if, contrary to our assumption, we were to enforce a policy of strict consistency with respect to belief systems (which (TN) entails), the step from (b) to (c) (and hence the argument as a whole) no longer goes through.

It follows from this that (Bo+) satisfies Strength just in case the paraconsistentist can make it plausible that there is a doxastic norm

⁴² There may be other reasons for dismissing (TN). Some of them are quite general (see Bykvist and Hattiangadi 2007 and Glüer and Wikforss 2009 for two recent criticisms). Others may pertain only to particular brands of paraconsistentists (for instance, dialetheists presumably cannot accept (TN)).

that underpins the *ought* in the antecedent and which is such that it mandates inconsistent belief sets at least under certain circumstances. But is there a plausible norm that fits this description?

We saw that the prototypical cases in which an agent may be said to rationally hold inconsistent beliefs are ones in which the agent is highly confident in each member of a set of propositions taken individually but where those propositions cannot be jointly true. Since high confidence does not guarantee truth, an alethic norm like (TN) is ill-suited for the job. What the paraconsistentist is after, rather, is a *sub-truth* norm: a norm that allows an agent to form beliefs on less than conclusive grounds. Consequently, an evidential norm appears to be the most natural non-alethic substitute for (TN). That is, instead of requiring that an agent ought to believe a proposition just in case it is true, the paraconsistent logician might try something along the lines of ‘*S* ought to believe *P* just in case the evidence available to *S* makes it sufficiently likely that *P*’, where the ‘sufficiently-likely threshold’ will be determined by contextual factors (such as what is at stake in the deliberative situation at hand). The question, then, is ‘Can the paraconsistent logician, equipped with a suitable sub-truth norm, make a case for a version of the Bs?’

I believe there is a principled reason for doubting that there can be any such norm. The trouble is that any sub-truth norm that does what the paraconsistentist needs it to do is bound to be incompatible with (Bo+). Now, it is important to be clear about the kind of incompatibility I am after here. Suppose *N* is a sub-truth norm and let Φ be an *N*-tolerated (or mandated) inconsistent set of propositions believed by *S*. In the presence of EXP, Φ entails *P* for some patently absurd, premiss 4-violating proposition *P*. But if *N* is a sensible doxastic norm, as we are assuming, it should tell against *P*. In other words, according to *N*, *S* ought not to believe *P*. But given that *P* follows from a set all of whose members *S* ought to believe according to *N*, *S* also ought to believe *P* by (Bo+). This is just the kind of incompatibility that is needed for the normative argument against explosion to go through. It is *not* the kind of incompatibility I have in mind.

Rather, the tension I wish to highlight stems from the fact that the paraconsistent logician’s desired sub-truth norm is likely not to be closed under conjunction. That is, the norm might prescribe belief in each of a number of propositions, while disallowing belief in the conjunction of all these propositions. The idea is perhaps best illustrated

by the following familiar example. Suppose the paraconsistent logician goes in for the evidential norm of belief mentioned above:

(EN) For all S , for all propositions P : if S considers or ought to consider P , S ought to believe P if and only if P is sufficiently likely in light of S 's evidence.

(EN) is a sub-truth norm that presumably does have the desired property of countenancing instances of rationally held inconsistent belief sets: a set of propositions may be such that each of its members individually exceeds the appropriate likelihood threshold and so satisfies (EN), and yet it may be impossible for all of the propositions contained in the set to be true together. (EN) may give rise to such cases because, unlike the truth norm, it allows for a certain margin for error: it may be that I ought to believe a proposition in a given epistemic situation, even though there is a chance that my evidence is misleading me and so that my belief is false. The trouble now, however, is that errors add up. If we consider a sufficiently large number of interesting propositions, the likelihood of all of the propositions being jointly true will dip below the threshold. Therefore, sub-truth norms like (EN) will typically have the consequence that it is not the case that one ought to believe the conjunction of such a large set of propositions. Indeed, given certain reasonable assumptions, they will prompt us to *disbelieve* the conjunction of such a large number of propositions.

Again, this is just the lesson of the epistemic paradoxes. It can be seen most clearly if we spell out (EN)'s sufficient likelihood condition in probabilistic terms—in terms of subjective probabilities, say.⁴³ Given a threshold expressed in the form of a real number t in the unit interval, we can easily conceive of scenarios in which each member of a set of propositions $\{P_1, \dots, P_n\}$ exceeds the threshold ($Pr(P_i) > t$, where $1 \leq i \leq n$), but where the conjunction of the propositions in question fails to do so ($Pr(P_1 \wedge \dots \wedge P_n) \leq t$).

Where is the problem? Well, we have observed that there will be cases where, according to (EN), S ought to believe each of the P_i , but where (according to the very same norm) it is not the case that S ought to believe their conjunction. The problem, now, resides in the fact that (Bo+) dictates that S ought to believe all of the logical consequences of the P_i , including the conjunction of the P_i . Thus (EN)

⁴³ I invoke subjective probabilities only for the sake of concreteness; the familiar point I am making can be made with respect to any probability function, regardless of how we choose to interpret it.

and (Bo+) are incompatible! Jointly they entail that it is and is not the case that S ought to believe $P_1 \wedge \dots \wedge P_n$. Crucially, though, the conflict between these two norms cannot be chalked up to the presence of EXP. Therefore *this* tension between (EN) and (Bo+), unlike the one discussed above, *undermines* the paraconsistent logician's case.

Moreover, the phenomenon we encounter here is not merely an artefact of (EN) or our probabilistic elaboration of it. Rather, it seems plausible that *any* sub-truth norm will be afflicted by similar problems when paired with (Bo+). This is so by virtue of the fact that any such norm, if it is to support the normative argument against explosion, must jointly satisfy two seemingly irreconcilable conditions: it must (i) mandate (or at least tolerate) inconsistent belief sets, and (ii) tolerate closure under (at least known) entailment. However, as we have observed, a doxastic norm like (TN) that complies with (ii) is bound to flout (i); and conversely, a norm like (EN) that allows for inconsistent belief sets (and so respects (i) will flout (ii). The reason why sub-truth norms violate (ii), we have said, is because they tolerate a certain margin of error. On the principle that error accumulates, any such norm should deem conjunctions of non-trivial propositions to be less likely (and hence less worthy of belief) in proportion to the number of propositions conjoined. Indeed, supposing, as seems reasonable, that a workable sub-truth norm requires that propositions that are sufficiently unlikely ought to be *disbelieved*, it turns out that such sub-truth norms are incompatible with *any* of the bridge principles in the B-family that involve strict deontic operators (including those with negative polarity), not merely (Bo+).

In light of the foregoing considerations, the paraconsistentist is thus faced with the following decision. Either she can retreat to one of the (Br)s,⁴⁴ or she can point to a way of understanding our sub-truth norms which does satisfy (ii) after all, and so does not clash with (Bo+) (or the other B-type bridge principles involving strict deontic operators).

However, neither of these options holds much promise. Begin with the former. It is true that a weaker (Br)-type bridge principle *is* compatible with (EN) and sub-truth norms like it: according to the (Br)s, S has a reason to believe $P_1 \wedge \dots \wedge P_n$ whenever her beliefs in the P_i taken individually satisfy the operative non-logical doxastic norm, but as we know, S 's having a reason to believe $P_1 \wedge \dots \wedge P_n$ may be

⁴⁴ I mean here to include also the 'mixed' B-type principles mentioned in fn. 39, the (Bor)s and the (Bpr)s.

trumped by a weightier reason, namely the reason, stemming from the said doxastic norm, that $P_1 \wedge \dots \wedge P_n$ is highly unlikely to be true. The problem with this, however, is that the (Br)s are subject to the very same objections that already disqualified their narrow-scope cousins, the (Cr)s (see §5 above). Like them, the (Br)s violate both Strength and Plausibility (if not minimal plausibility).

This leaves the paraconsistentist with the second option, that of proposing a sub-truth norm that can be understood so as to meet conditions (i) and (ii). We have seen that if there is such a proposal it must reject the principle that error accumulates so as to decrease the likelihood of large conjunctions. Does a norm like (EN) remain intelligible when this extra principle is abandoned? One way of doing so would be to adopt John Pollock's view to the effect that any argument is only 'as good as its weakest link' (1983, p. 248). On this view, even a conjunction composed of a very large number of propositions would be no less likely, and hence according to (EN) no less worthy of belief, than the least likely of its conjuncts. I believe that Christensen (2004, §4.3) successfully demonstrates the 'weakest link principle' to be untenable. I am willing to concede, however, that, for all I have said, there may be alternative principles that the paraconsistentist could avail herself of. But the onus is on the paraconsistent logician to present such a principle.

8. Conclusion

Fitelson's dilemma has proved to be real. Our investigation demonstrates that the normative argument is in serious trouble. Vindicating the argument would require either casting reasonable doubt on our analysis of bridge principles or exploiting a loophole we have missed. It is also worth noting that our discussion is perfectly compatible with all-out Harmanesque scepticism about the normativity of logic. Importantly, though, it does not presuppose it. One can maintain that logic is normative for thinking, while denying that this assumption furthers the paraconsistentist's case. Indeed, we have encountered a range of bridge principles that are *prima facie* attractive and yet incapable of supporting the normative argument against EXP because they fail to meet Strength.⁴⁵

⁴⁵ I am grateful to Arif Ahmed, Mic Detlefsen, Jesper Kallestrup, Peter Milne, Graham Priest, Jeff Speaks, Asbjørn Steglich-Petersen and Fritz Warfield, as well as to audiences at the universities of St Andrews, Notre Dame, Aarhus, Edinburgh and Cambridge for helpful discussions. Special thanks are due to Yoon Choi, Mara-Daria Cojocaru, Branden Fitelson, Hannes Leitgeb, John MacFarlane, Julien Murzi and Diego Tajer for commenting on earlier

References

- Anderson, Alan R., and Nuel Belnap 1975: *Entailment: The Logic of Relevance and Necessity*, Vol. 1. Princeton, NJ: Princeton University Press.
- Belnap, Nuel 1977: 'A Useful Four-Valued Logic'. In Michael Dunn and George Epstein (eds.), *Modern Uses of Multiple-Valued Logic*, pp. 8–37. Dordrecht: Reidel.
- Broome, John 1999: 'Normative Requirements'. *Ratio*, 12, pp. 398–419.
— 2013: *Rationality Through Reasoning*. Oxford: Wiley Blackwell.
- Bykvist, Kirster, and Anandi Hattiangadi 2007: 'Does Thought Imply Ought?' *Analysis*, 67, pp. 277–85.
- Chisholm, Roderick 1963: 'Contrary-to-Duty Imperatives and Deontic Logic'. *Analysis*, 24, pp. 33–6.
- Chrisman, Matthew 2008: 'Ought to Believe'. *Journal of Philosophy*, 105, pp. 346–70.
- Christensen, David 2004: *Putting Logic in its Place: Formal Constraints on Rational Belief*. Oxford: Oxford University Press.
- Easwaran, Kenny, and Branden Fitelson 2015: 'Accuracy, Coherence, and Evidence'. In Tamar Szabó Gendler and John Hawthorne (eds.), *Oxford Studies in Epistemology*, Vol. 5, pp. 61–96. Oxford: Oxford University Press.
- Feldman, Richard 2000: 'The Ethics of Belief'. *Philosophy and Phenomenological Research*, 60, pp. 667–95.
- Field, Hartry 2009: 'What is the Normative Role of Logic?' *Proceedings of the Aristotelian Society*, 83, pp. 252–68.
- Fitelson, Branden 2008: 'Goodman's "New Riddle"'. *Journal of Philosophical Logic*, 37, pp. 613–43.
- Geach, Peter 1982: 'Whatever Happened to Deontic Logic?' *Philosophia*, 11, pp. 1–12.
- Glüer-Pagin, Kathrin, and Åsa Wikforss 2009: 'Against Content Normativity'. *Mind*, 118, pp. 31–70.
- Harman, Gilbert 1984: 'Logic and Reasoning'. *Synthese*, 60, pp. 107–27.
— 1986: *Change in View: Principles of Reasoning*. Cambridge, MA: MIT Press.
— 2002: 'Reflections on *Knowledge and Its Limits*'. *Philosophical Review*, 111, pp. 417–29.

drafts of this paper, which led to significant improvements. I bear the sole responsibility for any remaining errors. Research on this paper was funded by the Alexander von Humboldt Foundation, whose support I gratefully acknowledge.

- Humberstone, I. L. 1971: 'Two Sorts of "Ought"s'. *Analysis*, 32, pp. 8–11.
- Joyce, James M. 1998: 'A Nonpragmatic Vindication of Probabilism'. *Philosophy of Science*, 65, pp. 575–603.
- Kolodny, Niko 2007: 'How Does Coherence Matter?' *Proceedings of the Aristotelian Society*, 107, pp. 229–63.
- Kolodny, Niko and John MacFarlane 2010: 'Ifs and Oughts'. *Journal of Philosophy*, 107, pp. 115–43.
- Lewis, David 1982: 'Logic for Equivocators'. *Noûs*, 16, pp. 431–41.
- MacFarlane, John 2004: 'In What Sense (If Any) is Logic Normative for Thought?' Unpublished.
- Mares, Edwin 2004: *Relevant Logic: A Philosophical Interpretation*. Cambridge: Cambridge University Press.
- Meyer, Robert K. 1971: 'Entailment'. *Journal of Philosophy*, 68, pp. 808–18.
- Pollock, John L. 1983: 'Epistemology and Probability'. *Synthese*, 55, pp. 231–52.
- Priest, Graham 1979: 'Two Dogmas of Quineanism'. *Philosophical Quarterly*, 117, pp. 289–301.
- 2006: *In Contradiction*. 2nd edn. Oxford: Oxford University Press.
- Priest, Graham Koji Tanaka and Zach Weber 2013: 'Paraconsistent Logic'. In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, Fall 2013 edition. <http://plato.stanford.edu/archives/fall2013/entries/logic-paraconsistent/>. Accessed 18 July 2013.
- Read, Stephen 1988: *Relevant Logic: A Philosophical Examination of Inference*. Oxford: Blackwell.
- Restall, Greg 2005: 'Multiple Conclusions'. In Peter Hájek, Luis M. Valdés-Villanueva and Dag Westerståhl (eds.), *Logic, Methodology and Philosophy of Science: Proceedings of the Twelfth International Congress*, pp. 189–205. London: King's College Publications.
- Sainsbury, Mark 2002: 'What Logic Should We Think With?' In Anthony O'Hear (ed.), *Logic, Thought and Language*, pp. 1–18. Cambridge: Cambridge University Press.
- Schroeder, Mark 2011: 'Ought, Agents, and Actions'. *Philosophical Review*, 120, pp. 1–41.
- Sellars, Wilfrid 1969: 'Language as Thought and as Communication'. *Philosophy and Phenomenological Research*, 29, pp. 506–27.
- Tennant, Neil 1997: *The Taming of the True*. Oxford: Oxford University Press.

- von Wright, Georg H. 1956: 'A Note on Deontic Logic and Derived Obligation'. *Mind*, 65, pp. 507–9.
- Way, Jonathan 2010: 'The Normativity of Rationality'. *Philosophy Compass*, 5, 1057–68.
- Wedgwood, Ralph 2006: 'The Meaning of "Ought"'. In Russ Shafer-Landau (ed.), *Oxford Studies in Metaethics*, Vol. 1, pp. 127–60. Oxford: Oxford University Press.
- forthcoming: 'Objective and Subjective "Ought"'. To appear in Nate Charlow and Matthew Chrisman (eds.), *Deontic Modality*. Oxford: Oxford University Press.
- Williamson, Timothy 2000: *Knowledge and Its Limits*. Oxford: Oxford University Press.