



# Goodness, availability, and argument structure

Anna-Sara Malmgren<sup>1</sup>

Received: 11 December 2019 / Accepted: 29 May 2020 / Published online: 30 June 2020  
© Springer Nature B.V. 2020

## Abstract

According to a widely shared generic conception of inferential justification—‘the standard conception’—an agent is inferentially justified in believing that  $p$  only if she has *antecedently justified beliefs* in all the *non-redundant premises* of a *good argument* for  $p$ . This conception tends to serve as the starting-point in contemporary debates about the nature and scope of inferential justification: as neutral common ground between various competing, more specific, conceptions. But it’s a deeply problematic starting-point. This paper explores three questions that haven’t been given the attention they deserve, that complicate the application of the standard conception to cases, and that reveal it to be underspecified at the core—in ways that aren’t resolved but inherited by more specific (extant) versions of it. The goal isn’t to answer the questions, but to articulate them, explain what turns on them, and invite a critical re-examination of the standard conception.

**Keywords** Epistemology · Inferential justification · Logical form · Argument · Inference

[K]nowledge and justification are inferentially transmissible only if the underlying argument is good. If we start with false or unjustified premises *or* we unreasonably infer a conclusion from them (i.e. infer it invalidly or in an inductively inadmissible way) it’s not to be expected that a belief based on the argument in question constitutes knowledge or is even justified.

(Audi 2011, 199.)

On a widely shared generic conception of inferential justification—henceforth ‘the standard conception’—an agent is inferentially justified in believing that  $p$  only if she has *antecedently justified beliefs* in all the *non-redundant premises* of a *good argument* for  $p$ . This paper explores three questions that haven’t been given the attention they deserve, that complicate the application of the standard conception to cases, and that reveal it to be underspecified at the core—in ways not resolved, but inherited, by more

---

✉ Anna-Sara Malmgren  
malmgren@stanford.edu

<sup>1</sup> Stanford University, 450 Jane Stanford Way, Stanford, CA 94305, USA

specific (extant) versions of it. The goal isn't to answer the questions, but to articulate them, explain what turns on them, and to invite a critical re-examination of the standard conception.

First, what counts as a *good* argument, in the intended sense? (Which argument forms, or schemata or patterns, are fit to constitute inferential justifications?) Second, how should we *initially regiment* or classify ostensibly enthymematic arguments, when theorizing about undecided problem cases? (On what grounds should an apparent enthymeme be treated as a genuine enthymeme, and as instantiating one rather than another familiar, non-enthymematic, form?) Third, what counts as a *relevantly complete* representation of a given argument form—complete vis-à-vis the epistemic and psychological demands that the standard conception places on agents? I'll call these questions the question of 'argument goodness', 'initial regimentation', and 'relevant completeness', respectively.

The questions interact in interesting ways, and they take on particular urgency in relation to certain hard-to-classify cases—*prima facie* instances of justified belief that defy straightforward classification as being either inferentially or non-inferentially justified. Those cases are discussed elsewhere.<sup>1</sup> Here I focus on a would-be easier case: a simple little *prima facie* example of doxastic inferential justification. The discussion of that case brings out the (independent) need for principled answers to the three questions, how difficult it is to answer any one of them without holding answers to the others fixed, and that the answer to the third remains elusive even granting answers to the other two.

The paper is organized as follows. In Sect. 1, I articulate the standard conception and, along the way, the question of argument goodness. In Sects. 2 and 3, I introduce the simple example, and the other two questions. I discuss why, and to whom, they matter. I also isolate the question of relevant completeness, and consider an initial response to it. In Sect. 4, I discuss other responses, but find them wanting. I end on a tentatively skeptical note.

## 1 The standard conception

### 1.1 Preliminaries

I use 'inferential/non-inferential', where some prefer 'derived/foundational', 'non-basic/basic', 'mediate/immediate', or 'indirect/direct'. I use 'rational' and 'justified' interchangeably, as broad umbrella terms; likewise for '(good) reason(s)' and 'evidence'. I write as though reasons are propositions—rather than, say, *true* propositions, or states of affairs—but nothing crucial turns on this: the key points are recoverable on views that treat reasons as factive. And I leave open whether all justification involves anything recognizable as reasons (or evidence).<sup>2</sup> Throughout the paper, it's *epistemic* justification that's at issue—roughly: justification that's neither moral or prudential,

---

<sup>1</sup> See Malmgren (2018).

<sup>2</sup> Thus the ensuing discussion is neutral on the 'content requirement' and the 'premise principle' (Pryor 2005).

and that in the first instance attaches to beliefs (rather than to intentions or actions). Parallel questions arise in the practical domain, but they won't be discussed here. And I often write in terms of *doxastic* rather than *propositional* justification because that simplifies the presentation. (The difference is marked where it matters.) I assume that justification comes in degrees, and that most if not all justification is defeasible. But the bulk of the discussion is conducted in terms of (justification for/to have) all-out, rather than graded, belief since the standard conception trades in that. Whether anything turns on this partly depends on how these attitude-types are related (cf. Sect. 4.3). Last, I take it that the subject-matter of the theory of justification includes the practices of ordinary human agents in the actual world: that at least one of its aims is to correctly describe what *we* do, if and when we do things right, epistemically speaking—what it takes for agents like us to be justified, to be justified in certain ways, and so on. This becomes important later.

## 1.2 What's inferential justification?

I'll start with a minimal partial gloss—a characterization as uncontroversial as they come. That gloss is then elaborated, step-by-step, until we reach the standard conception. An initial negative characterization of non-inferential justification falls out of it (since the distinction is supposed to be strictly dichotomous): non-inferential justification is just justification that's not governed by the constraints that comprise the standard conception.

The standard conception is highly schematic—it's more of a blueprint for a theory than a theory. And it only consists in a set of necessary conditions: a set often supplemented with further requirements—e.g. cognitive or epistemic meta-requirements (e.g. Boghossian 2003, 2014; Dogramaci 2013; Fumerton 1995, 2006; Tucker 2012; Wright 2004). But the viability of any such further requirement is largely orthogonal to the questions I want to raise. Your typical contemporary epistemologist—who's given the inferential/non-inferential justification distinction any thought at all—is *at least* committed to the standard conception (or to something very much like it).

Here's the minimal gloss (cf. Pryor 2005, pp. 182–183):

INF<sub>PROP</sub> S has ('propositional') inferential justification to believe (that) *p* only if her justification to believe *p* at least partly *rests* on her justification to believe at least one other proposition, *q*

INF<sub>DOX</sub> S is inferentially justified in believing *p* (has 'doxastic' inferential justification) only if her justification to believe *p* at least partly rests on her justification to believe *q*—moreover, S *believes p*, and her belief that *p* is *based* on her justified belief that *q*

As per usual assumptions, propositional justification doesn't require belief in the target proposition (*p*), nor suitable basing, whereas doxastic justification requires both. But INF<sub>DOX</sub> goes beyond those assumptions, with the claim that for S's belief that *p* to be

suitably based (on  $q$ )—suitably for inferential justification—is for it to be based on  $S$ 's *justified belief* that  $q$ .<sup>3</sup>

It's hard to say much more about inferential justification without entering controversial territory. But the following elaborations—conditions (a)–(f)—are sufficiently abstract, and sufficiently widely endorsed, to also be included in the standard picture.

First, two non-circularity requirements, pertaining to resting and basing respectively:

- (a)  $S$ 's justification to believe  $p$  rests on her justification to believe  $q$  only if her justification to believe  $q$  doesn't in turn rest on—only if it's *antecedent* to—her justification to believe  $p$ .
- (b)  $S$ 's (justified) belief that  $p$  is based on her (justified) belief that  $q$  only if her justified belief that  $q$  isn't in turn based on her justified belief that  $p$ .

If the basing-relation is or requires a causal relation, (b) is arguably redundant: it's already implicit in  $\text{INF}_{\text{DOX}}$ . Perhaps a similar case can be made for (a): perhaps the proper understanding of the resting-relation entails a ban on circles or loops. Nevertheless, (a) and (b) are worth spelling out, since they help make perspicuous that resting and basing are both asymmetric, when things go well.<sup>4</sup>

A little more can be said about basing:

- (c)  $S$ 's belief that  $p$  is (at least partly) based on her belief that  $q$  only if  $S$ 's believing  $q$  (partly) *explains* her believing  $p$ .

Many take the explanation-type involved in basing to be causal—perhaps causal explanation of a special kind—but we can proceed without committing to this.<sup>5</sup>

It's also part of the standard conception that not *any* set of propositions can constitute an inferential justification—there's a structural or formal constraint on  $p$  and  $q$ :

- (d)  $S$ 's justification to believe  $p$  rests on her justification to believe  $q$  only if  $q$  (inferentially) *evidentially supports*  $p$ —alone or together with other propositions,  $r \dots r_n$ , that are *available* to  $S$

I'll refer to (d) as 'the structural requirement'. It has two components: first, that  $q$  must evidentially support  $p$ , either by itself or in conjunction with auxiliary propositions; second, that any such auxiliaries must be 'available' to the agent. (I here abstract away from routine ways of cashing out that constraint—'available' is just a placeholder—but more on this soon.)

The first component is effectively an instance of the more general demand that, for  $q$  to be a (good) reason to believe  $p$ ,  $q$  must stand in an appropriate logical—or more

<sup>3</sup> With the possible exception of cases where  $q$  is the content of an appropriate *conditional* belief; cf. Sect. 1.4.

<sup>4</sup> The standard conception is supposed to be compatible with coherentism about justification (e.g. Bonjour 1985; Ewing 1934). On one version of that view, resting and basing can indeed be symmetric (when things go well). But we can accommodate this by denying that basing and resting are transitive, or by adding provisos to (a)–(b), permitting circles that are wide enough. (See also DeRose 2005; Thagard 2000.)

<sup>5</sup> On 'doxastic' and 'causal-doxastic' accounts of basing (e.g. Korcz 2000; Tolliver 1982), the debate over possible meta-requirements on inferential justification takes a somewhat different shape (cf. Rhoda 2008). But this difference won't matter here.

broadly: implication—or confirmation-relation—to  $p$  (e.g. Davidson 1986; McDowell 1996; Pryor 2005; Sellars 1953). What relation is that? Which such relation, or which such relations, count(s) as appropriate? This question may well have different answers, depending on the type of reason or justification involved. (What counts as appropriate for *testimonial* reasons may or may not count for *perceptual* reasons, and conversely; what qualifies for *introspective* justification may or may not qualify for *mathematical* justification—etc.) But this point by itself doesn't take us very far. The (umbrella) type under consideration here is that of inferential justification, and it's quite unclear what to say about it—at any rate, what to say of a principled sort.

### 1.3 Goodness

By way of rough overview, the orthodoxy used to be that all and only deductive or monotonic consequence-relations qualify (e.g. Descartes 1641/2008; Hume 1748/1999; Popper 1959). In stark opposition, Harman and others have argued that *no* such relations qualify, and that the job is exclusively reserved for abduction, or 'inference-to-the-best explanation' (e.g. Harman 1965, 1986; Lycan 1988; Poston 2014). Others argue that all ostensibly abductive support reduces to inductive and/or deductive support (e.g. Bird 2005; Fumerton 1980, 1992). Yet others maintain that all and only Bayesian confirmation-relations qualify—roughly: classical deductive consequence, supplemented by probabilistic laws and constraints (e.g. Carnap 1950; Jeffrey 1965; van Fraassen 1989; Williamson 2000).<sup>6</sup> And there are further views in this vein—further attempts at capturing evidential support as part of a unified formal system: either an extension of classical logic (e.g. Horty 2012; Pollock 1970, 1987; Reiter 1980), or a replacement of it (e.g. Alchourrón et al. 1985; Levi 1997).

There's also a cautiously pluralist picture out there: one that many of us, 'informal' or traditional epistemologists, tend to work with. (e.g. Audi 1986, 2011; Boghossian 2008, 2014; Broome 2013; Fumerton 2006; Goldman 2009; Pryor 2005; Schechter 2017; Wedgwood 2012) It recognizes a range of implication-relations, both deductive and non-deductive, as being appropriate. The range is typically introduced by means of a short list of *prima facie* paradigms—which list, typically in turn, includes some basic logical entailment-relations or the corresponding rules (e.g. modus ponens and modus tollens, existential and universal generalization, conjunction introduction and elimination), perhaps some analytic or lexical entailments (e.g. the 'bachelor-rule'), enumerative-statistical inductive implication (to singular and general conclusions), and abductive implication. (Since abductive and inductive implication come in degrees, it's understood that only sufficiently strong relations of these types qualify.) Paradigms from the opposing category—sample rules/relations that *don't* qualify—are sometimes supplied as well. Those may include some familiar formal fallacies (e.g. affirmation of the consequent, denial of the antecedent), extremely weak inductive consequence

<sup>6</sup> But the applicability is complicated by the fact that Bayesian models trade in graded belief (among other things; cf. Sect. 4.3).

(e.g. generalization from one instance), and truly degenerate counterparts of induction and abduction (e.g. counter-induction, inference-to-the-worst-explanation).<sup>7</sup>

I call the picture ‘cautiously’ pluralist because I take it to be silent on the possibility of successful reductions—of some of the qualifying relations to others—and likewise on the possible existence of a more abstract principle which covers some, or all, of them. As I understand the picture, it’s not committed either way. (Perhaps at least some forms of induction are correctly analyzed in Bayesian terms, perhaps not. Perhaps evidential support can’t be formalized in anything like that way. Perhaps what unifies the qualifying candidates is their being sufficiently conditionally truth-conducive in appropriate domains. Perhaps truth-conduciveness is only part of the story, or it’s not part of it at all. The cautiously pluralist picture is itself neutral on questions like these.)

I’ll be working with this picture, and the paradigms usually associated with it, for much of what follows. But some questions will be raised about it later on (Sect. 4.1).

For now, just another note on usage: in my terminology, any relation that qualifies, relative to inferential justification, instantiates the ‘inferential evidential support-relation’. A ‘good (inference) rule’ is a rule that encodes some such relation. A ‘good argument form’ is a relatively abstract template or schema for particular arguments, at finer levels of grain, that exploit or incorporate a (certain) good rule. And a ‘good argument’ is just an argument that has a good form; that is, fits some such template.<sup>8</sup> (What Wedgwood 2012 calls a ‘rationally accessible argument.’) The first component of the structural requirement can now be restated as: only good arguments are able to constitute inferential justifications.

To say that an argument is good, in this sense, is not to say that it guarantees inferential justification—not even for agents who satisfy all the other requirements on such justification listed here, or some natural expansion of that set. This is important. Complicating factors include the role of defeaters, higher-order epistemic norms (see e.g. Christensen 2010; Harman 1984, 1986; MacFarlane 2004; Pryor 2018; Schechter 2011), and the mechanisms responsible for so-called ‘warrant-transmission failure’, or ‘easy knowledge’ (e.g. Cohen 2002; Davies 2004; Silins 2005; Wright 2002, 2004).<sup>9</sup>

<sup>7</sup> Two other approaches deserve mention: first, the explicitly non-reductive pluralist approach exemplified by so-called ‘argumentation theory’ (e.g. Walton et al. 2008; Walton 2013). Second, the thoroughly ‘psychologistic’ approach on which evidential support is directly reducible to certain non-normative, high-level psychological facts (e.g. Pelletier and Elio 2005; Quine 1969). It seems clear to me that both approaches—at least as elaborated to date—end up misclassifying *prima facie* paradigms; in particular: classifying as good some paradigmatically bad rules/arguments. But further discussion of this will have to be deferred to another occasion.

<sup>8</sup> Note that I’m (still) operating with a relatively loose notion of form. Whether all (or any) good argument forms are fully expressible in some formal language is treated as an open question. (More on this in Sects. 4.2–4.3.) Note also that token arguments are conceived as partly individuated by their forms/rules they exploit. More specifically: I think of an argument as a set, or sequence, of propositions—at least one of which is a conclusion, the rest of which are premises—together with one or more inference-rules. (This conception of arguments has presentational advantages, but nothing of substance turns on it.)

<sup>9</sup> If all (nontrivial) warrant-transmission failure can be analyzed as straightforward premise-circularity (e.g. Tucker 2010), then the non-circularity conditions on resting and basing, properly fleshed out, should ensure immunity to this problem. (Ditto on at least some accounts that treat it as a broader kind of epistemic circularity.) Thanks to an anonymous referee for pointing this out. But it’s controversial whether warrant-transmission failure can be analyzed in some such way (e.g. Neta 2013a; Pryor 2013). And, however that may be, the non-circularity conditions don’t address the problems posed by defeaters, and by higher-order (e.g. ‘clutter avoidance’) norms.

But it's widely thought—hence part of the standard conception—that inferential justifications can't embody arguments of any-and-all forms.<sup>10</sup> A 'good' argument is simply one that makes the cut in this regard; a 'bad' argument is one that doesn't.

#### 1.4 Availability

The structural requirement, (d), also features a notion of availability; how should that be understood? What is it for a proposition  $r$  to be available to  $S$ —available to 'combine' with  $q$ , to evidentially support  $p$ , for *her*? Presumably not just any auxiliary proposition, with the right formal properties (relative to  $p$ ), has what it takes. But what *does* it take?

The standard answer, in effect, is that *what goes for  $q$  goes for  $r$* : the same constraints that govern (the relation between  $S$  and) the original supporting proposition,  $q$ , govern any other proposition  $r$  that figures in  $S$ 's justification to believe—or for believing— $p$ .

For propositional (inferential) justification, this becomes:

(e)  $R$  is available<sub>PROP</sub> to  $S$  only if  $S$  has antecedent justification to believe  $r$ .

And for doxastic:

(f)  $R$  is available<sub>DOX</sub> to  $S$  only if  $S$  is antecedently justified in believing  $r$ .

I'll refer to the conjunction of (e) and (f) as 'the standard availability-constraint', or 'the availability-constraint' for short.<sup>11</sup>

A related, slightly weaker, understanding of availability can also be found in the literature (e.g. mentioned but not endorsed in Pryor 2012). On this view, the demand specified in (e)—antecedent propositional justification vis-à-vis  $r$ —governs the availability of  $r$  to  $S$  for propositional *and* doxastic inferential justification. The weaker constraint has some problems particular to it, but I won't discuss those here. (See Malmgren, *ms.*) And the central issues raised below also arise on this understanding of availability.

As stated, the availability-constraint isn't able to accommodate (what we might call) 'justification by suppositional reasoning'.<sup>12</sup> If, say, conditional proof and reductio ad absurdum are good argument forms—and not because they reduce to other, more congenial, forms—the constraint at least requires modification to include a substitute demand, in lieu of the demand for (antecedent) justification, that's applicable to suppositions or conditional beliefs. But I won't discuss this further, either.

It's however worth repeating that the constraint concerns the availability of a proposition to play a certain contributory role in  $S$ 's inferential justification—the role specified in (d). It doesn't amount to or alone entail the requirement that, to play

<sup>10</sup> The point is sometimes put in terms of neither arguments, rules nor relations—see e.g. Boghossian (2006) on 'transmission-principles'; McHugh and Way (2018) on 'reasoning-patterns'.

<sup>11</sup> For representative endorsements, see e.g. Audi (2011, pp. 183–184, 199), Boghossian (2003, pp. 225–226), Goldman (2008, p. 64), Huemer (2002, pp. 329–330), Fumerton (2006, pp. 38–39, 100), Markie (2005, p. 348), and Schroeder (2011, §1). (And it's arguably entailed by the claim that availability requires knowledge; see Williamson 2000, p. 186.)

<sup>12</sup> On this, see e.g. Dogramaci (2013) and Wright (2014).

that role, a proposition or propositional content must be reflectively accessible to S, or that she must believe—or in some other way ‘take’—that proposition to support (her) believing the conclusion. The availability-constraint just says that she must have antecedent justification to believe that proposition, or (for doxastic justification) be antecedently justified in believing it.

What I’m calling ‘the standard conception’ comprises  $INF_{PROP}$ ,  $INF_{DOX}$ , and (a) through (f).

We can summarize it, somewhat crudely, as follows:

An inferential justification (at least partly) consists in a set of propositions, and one or more inference rules, that together make up a *good argument* for a given conclusion. An agent *has* that justification, *to believe* the conclusion, only if she has *antecedent justification to believe each of those propositions*. She’s justified *in* believing the conclusion, on the basis of that justification, only if she’s *antecedently justified in believing each of the propositions* that constitute it, and her belief in the conclusion is suitably *based* (on her beliefs in those propositions).

This conception tends to serve as the *starting-point*—the smallest common factor—in a range of contemporary debates about inferential justification: e.g. debates that focus on what else such justification may require, what the scope of it is, and/or how to cash out one or more of the controversial theoretical notions that give substance to the conception (e.g. that of antecedence, basing, or evidential support).<sup>13</sup> The starting-point is our topic.

## 2 Sorting cases

### 2.1 Dual asymmetric dependence

Suppose we wonder whether some *prima facie* justified belief (token or type, actual or hypothetical) is inferentially justified. How do we—how should we—settle that question? This is a methodological question: a quest for operational criteria, if you like. But, as we’ll see, part of why it’s hard to pinpoint satisfying operational criteria is that the metaphysical criteria the standard conception provides are underspecified at certain key points.

The two non-circularity conditions, (a) and (b), suggest a certain initial diagnostic: a fallible first-round test for doxastic inferential justification, which I call the ‘dual asymmetric dependence test.’<sup>14</sup> It’s a test that I think we already routinely use, in the provisory classification of cases, and it integrates nicely with the standard conception.

<sup>13</sup> See e.g. Audi (1986, 2011), Bonjour (1978, 1985), Boghossian (2003, 2014), Corbi (2000), Dogramaci (2013), Feldman (2003), Fumerton (1995, 2006, Ch. 3, 6), Gallois (1996), Ginet (2005), Goldman (2008), Greco (1999), Hookway (2000), Huemer (2002, 2016), Leite (2008, 2011), Markie (2005), Neta (2013b), Pollock and Cruz (1999, Ch. 2.3), Pryor (2005), Rhoda (2008), Richard (2003), Tucker (2012), Wedgwood (2012), Wright (2004), Schechter (2017, 2019), and Zalabardo (2011).

<sup>14</sup> For an interesting application of a similar test, see McGrath (2017).

The thought is just this: if the basing- and resting- relations are both asymmetric, an inferentially justified belief should *ceteris paribus* give rise to two asymmetric counterfactual dependences—of belief on belief, and justification on justification. More specifically, (a)–(b) underwrite roughly the following principle: (if S is inferentially justified in believing  $p$ , there's at least one other proposition  $q$  such that) *ceteris paribus*, if S didn't have justification for believing  $q$ , S wouldn't have justification for believing  $p$ , but not vice versa, and if S didn't believe  $q$ , S wouldn't believe  $p$ , but not vice versa. Simplifying: *ceteris paribus*, if S weren't justified in believing  $q$ , S wouldn't be justified in believing  $p$ , but not vice versa. (Whenever the simplification is used below, it's shorthand for the more precise conjunctive claim.)

If that's right then, conversely, a dual dependence of the specified type defeasibly indicates that S is inferentially justified in believing  $p$ —more precisely: that S's justification for believing  $p$  (non-circularly) rests on her justification for believing  $q$ , and that S's belief that  $p$  is (non-circularly) based on her belief that  $q$ . I'll sometimes express this by saying that S's justified belief that  $p$  'inferentially depends' on her justified belief that  $q$ , or that  $q$  is in the 'inferential ground' of her belief that  $p$ . (The dual asymmetric dependence test is in the first instance a test for inferential dependence.)

It seems clear that doxastic inferential justification (or at least inferential dependence) can occur without the asymmetric counterfactuals—most obviously: because both belief and justification can be over-determined (as well as 'pseudo-over-determined'; Swain 1979). Less obviously: because a mischievous demon—or other spurious agent or mechanism—can interfere (and void one or both dependences, or turn them symmetric). Certain patterns of rebutting defeat can also cancel out the asymmetries. (E.g. suppose that my belief that Sven is Scandinavian inferentially depends on my belief that Sven is Swedish, but that in the closest worlds where I'm not justified in believing that he's Scandinavian that's because I believe, and/or have justification to believe, that he's North African.) Conversely, it seems clear that the asymmetric counterfactuals can hold absent inferential justification (or dependence)—e.g., once again, because a demon interferes. They can also be sustained by certain hierarchical relations between belief-contents (e.g. between contents with the forms  $p$  and  $p$  and  $q$ , or the forms  $a$  is  $F$  and  $a$  is  $G$ —where  $F$  is a superordinate-level concept and  $G$  a subordinate-level concept whose extension is included in  $F$ 's). And they can arguably be sustained by relations of 'enabled-on-enabler dependence'—the kind of dependence that a fact or event has on the facts or events that merely enable, as opposed to explain, it. (E.g. the dependence that—in normal circumstances, and against a normal contrast-class—my writing this paper has on my having been born, on my being literate, and on there being oxygen on Earth.)

However, there doesn't seem to be an *open-ended* list of (interestingly different) sources of false positives: of plausible alternate explanations of the asymmetric counterfactuals. That's why the dual asymmetric dependence test isn't a *useless* diagnostic. Indeed, if a subject and a set of (otherwise suitable) propositions pass the test—and no familiar competing explanation suggests itself—it seems reasonable to at least take

the hypothesis of inferential justification seriously, and to proceed to investigate it further.<sup>15</sup>

The issue that we'll turn to now is how to do *that*—how to gauge the overall viability of that hypothesis. A tempting suggestion is to consider whether the other requirements that comprise the standard conception are satisfied in the case under consideration. (If the case is non-actual: whether it could easily be realized in such a way that they're met.) This is where things get more complicated.

## 2.2 The simple example

To illustrate the problematic, let's take a very mundane case.

Suppose Sven comes in the door. He's soaking wet, and he's angry. Aya looks at him, she comes to believe that Sven is angry and that he's soaking wet, and—ostensibly as a result of that—comes to believe that someone stole his umbrella.

I take it that Aya's belief that someone stole Sven's umbrella could be *justified*, and that it's as good a candidate as any for being *inferentially* justified (if it's justified). For one thing, it's hard to see what model of non-inferential justification would be applicable here.<sup>16</sup> For another, there are a couple of propositions that seem fit to figure in an inferential ground of the target belief: viz., *that Sven is angry* and *that he's soaking wet*. To put matters with deliberate imprecision: those propositions aren't wildly off-topic—they look like the kinds of content that could satisfy the structural requirement, (d), in a relatively normal realization of the case. Moreover, by hypothesis Aya *believes* them—and she could easily be *justified* in doing so (perhaps in some or other way that involves her visual experience as Sven enters). In fact—unless the case is fleshed out so as to provide her with independent resources, or with (helpful or disruptive) interference mechanisms—it looks like Aya *wouldn't* be justified in believing that someone stole Sven's umbrella unless she were justified in believing that he's angry

<sup>15</sup> The ubiquity of certain false negatives might however distract—in particular: of commonplace, actual and counterfactual, cases of inferential justification or dependence where the asymmetric counterfactuals fail to hold due to over-determination (of belief and/or justification), or easily accessible fallback routes (psychological/epistemic). These cases are supposed to be precluded by the *cp*-clause, along with more remote possibilities. But their prevalence might make them hard to screen off, when assessing the counterfactuals. So it's worth noting that there's a more discriminating test in the same neighborhood. Since it's independently plausible that basing requires an *explanatory* relation (recall (c))—and that that's why the one belief *ceteris paribus* asymmetrically counterfactually depends on the other—and independently plausible that the asymmetry of resting yields an asymmetry in the transmission of *undercutting defeat* (Pryor 2005, p. 183)—and that that's (partly) why the one justification *ceteris paribus* asymmetrically depends on the other—we can also 'test' for this set of marks, instead or in addition. This test too is fallible, of course, but the distraction the above-mentioned cases might pose has been eliminated. We've moreover pre-empted the appeal to enabled-on-enabler dependences. (For brevity I'll just proceed in terms of the dual asymmetric dependence test, but this makes no important difference.)

<sup>16</sup> E.g. it's implausible that Aya's belief is *perceptually* non-inferentially justified, even supposing that perceptual states can have very rich contents (e.g. McDowell 1994; Siegel 2010). Surely Aya doesn't see (or seem to see) *that someone stole Sven's umbrella*—at *most* she sees that he's angry and soaking wet. Nor is the target belief naturally assimilated to *prima facie* paradigms of non-perceptually non-inferentially justified belief—e.g. my belief that squares have four sides, or that I'm now thinking about you.

and soaking wet, but that the converse is false. That is: the example is naturally read in such a way that the dual asymmetric dependence test returns a positive result.<sup>17</sup>

We can bolster the *prima facie* case for taking the identified contents to play this role by adding the stipulation that, if asked why she believes that someone stole Sven's umbrella, or what justifies her in doing so, Aya would reply: 'he's soaking wet, and he's angry!' (That stipulation doesn't beg any important questions here. The agent's testimony is just treated as another piece of defeasible evidence.)

I've tried to outline a simple realistic example in which there's *prima facie* reason to take an agent to be (doxastically) inferentially justified—no more and no less. What's next? How appraise that hypothesis further?

A natural suggestion is to check the hypothesis against our best theory of inferential justification, or—absent a best theory, or an agreed-on best theory—against an agreed-on generic model, or proto-theory, of it. And the obvious candidate here is the standard conception. (It's the obvious candidate because it's so widely shared, and it's comparatively minimal in its commitments.) The next step, then, would be to consider whether all the requirements making up that conception are plausibly met in the case.

We've already granted the dual asymmetric dependence that (fallibly) indicates that (a) and (b) are satisfied.<sup>18</sup> The conditions that invite the real difficulties, it seems, are the structural requirement, (d), and the availability-constraint, (e)–(f). Simplifying, we can take on both at once by asking: 'is the agent (here, Aya) plausibly antecedently justified in believing all the non-redundant premises of a good argument, from the identified premise set (*that Sven is soaking wet* and *that Sven is angry*) to the conclusion (*that someone stole Sven's umbrella*)?' Let's call this the 'crucial' question.

The crucial question takes us to an interesting dialectical juncture. In essence, it looks like certain other facts about the case must be settled, before we can answer it, but it's not at all clear how to settle those facts—there are just too many moving parts. The task resembles that of solving a set of equations with too many unknowns: barring luck or further constraints, what we have to work with isn't enough to solve for all the variables.

First, *which* argument are we to evaluate for goodness/badness? What form does it have? (How should the apparent enthymeme be initially regimented or categorized, for the purpose of this task?) Next, what does it *take* for an argument form to be good/bad? (What goes on the list? What's the unifying principle, if any?) Last, what counts as a *relevantly complete* representation or specification of a given argument form?<sup>19</sup> (In particular: how fine-grained is that representation, and how is the informational content distributed over premises and rules—what's the 'division of labor'?)

Let me explain.

<sup>17</sup> Likewise for the more discriminating test (fn. 15): the case is naturally read in such a way that Aya believes that someone stole Sven's umbrella partly because she believes that he's angry and soaking wet—but not conversely—and Aya's justification for believing that someone stole Sven's umbrella would be undercut if her justification for believing that he's angry/soaking wet were undercut—but not conversely.

<sup>18</sup> False positives must also be eliminated, but in mundane cases like this the only serious contender is an appeal to enabled-on-enabler dependences. (Cf. Malmgren 2018, §3.2.)

<sup>19</sup> On a certain view of form, this question collapses into that of argument goodness (see Sect. 4.1). On this view, some of what follows needs to be rephrased.

Either we've already identified all the (non-redundant) premises of the argument to be evaluated in the case, or we haven't. If we have, the crucial question turns on whether the identified premise set instantiates a good argument form—*as it stands*—and on whether the agent, Aya, is antecedently justified in believing each member of that set. (But the latter has been granted, at least provisionally.) If not—if we haven't identified all of them—the question also turns on how things are with the 'missing' premises: on their identity, and on Aya's epistemic and psychological relationship to them.

Suppose the former: that we've identified all the premises. Is the resultant little argument *good*—the argument from *Sven is angry* and *Sven is soaking wet*, to *someone stole Sven's umbrella*? Is it able to constitute an inferential justification, as it stands? As already mentioned, it's controversial what it takes for an argument to qualify. I take it that we don't yet have a fully general criterion, for sorting good from bad arguments. (At any rate we don't have one on which there's widespread agreement.) But we arguably have some clear cases: paradigmatic examples of each type. Perhaps the little argument at hand maps on to some familiar argument form—paradigmatically good, or paradigmatically bad—on the list of clear cases? (For now it doesn't matter what exactly the lists include.)

The obvious hitch is that it maps on to *too many*. The given set of propositions clearly fails to determine a unique form, good or bad—by itself, it doesn't even narrow down the options. Then again, this would seem to be true of *any* such set, as long as the identity of the inference-rule (or rules) is left open. Another way of putting the point is to say that, to evaluate the target argument—by checking it against a general criterion or theory, if we have one, or against paradigm cases—we first need to know which argument it *is*, and to know which argument it is, we need to know which inference-rule is at play.

If the present supposition is mistaken—and there are in fact auxiliary premises—we may also want to know what those premises are. In fact, that seems *required* to carry out the next task: to settle whether the availability-constraint is met. This brings us to the question, or problem, of relevant completeness. The problem has a somewhat artificial air, but it seems forced on us by the constraint. And it arises however the target argument is initially regimented, and however the question of argument goodness is answered. (Unless argument forms are so finely individuated that it precludes all representational variation—but then the problem arises in another guise; Sect. 4.1.)

### 3 Relevant completeness

#### 3.1 The problem

By the standard availability-constraint, an argument *A* (from premise *q* to conclusion *p*) constitutes a subject *S*'s doxastic inferential justification for believing *p* only if *S* is antecedently justified in believing each of *A*'s non-redundant premises. It matters, then, *what those premises are*. And the present point—the point I'm now about to make—seems independent of the claim that *S*'s standing depends on *A*'s *form*, in the sense of 'form' that's at work above (and in the earlier discussion of argument goodness; Sect. 1.3).

To see this, let's simply stipulate that A exemplifies some familiar, and *prima facie* good, argument form F. To defer a certain complication, let F be a non-deductive form for now—inductive or abductive—pulled from our paradigm list, and individuated in relatively loose and informal terms, as is customary. (In the contexts where it's being used as a paradigm.<sup>20</sup>) Thus F may be an inductive argument of one of two broad kinds—to a general conclusion about a class of things, from a premise or premises about (perhaps: sufficiently many, representative) members of that class, or to a particular conclusion about some such member, from a premise or premises about the class as a whole. Alternatively—and equally schematically put—F may be an abductive argument: an argument to the conclusion that the hypothesis obtains that best explains, or that if true would best explain, the evidence contained in the premise (or premises).

Whether inductive or abductive, F can be further explicated—or represented or precisified—in multiple different ways, along several different dimensions of variation: ways that feature one or more inference rules, one or more premises, meagre premises paired with relatively complex rules, substantial premises paired with relatively simple rules, and so on. It can also be represented at different levels of grain or specificity.

Some of the representational variation looks substantive—as involving competing descriptions or analyses of F, aspiring to something like correctness.<sup>21</sup> This includes comparatively formal models: attempted systematizations of F in broadly mathematical terms—as well as comparatively informal models: non-trivial explications, cast in natural language, that ascribe some theoretically interesting content or structure to F, sometimes as the first step towards the articulation of a formal model. (Proposed reductions of F, in formal or informal terms, to *prima facie* distinct argument forms also fall in this category.) Other variation doesn't—rather, it seems to involve distinct but equally correct (or incorrect) ways of representing or describing F. The variation with respect to grain is a case in point: a very fine-grained (perhaps formal or semi-formal) rendition of F isn't a genuine alternative to a very coarse-grained (perhaps informal and/or manifestly enthymematic) rendition of it, simply in virtue of *that* difference. Likewise for at least some variation in the division of labor between premises and rules—how information-rich they are, respectively—their order and number, as well as the fine details of the contents (of both premises and rules).

Perhaps all the apparent contrast between substantive and non-substantive variation comes to is that, for some but not all differences in representation of F, there exists a

<sup>20</sup> E.g. here's Wedgwood (2012, p. 6), introducing one type of (good) inductive argument, as “[an argument] from a premise of the form ‘All observed F's have been G’ to the corresponding conclusion ‘The next F to be observed will be G.’” Here's Boghossian (2008, p. 472), introducing another (in rule format): “‘For appropriate Fs and Gs, if you have observed *n* (for some sufficiently large *p*. *n*) Fs and they have all been Gs, you are *prima facie* rationally permitted to believe that all Fs are Gs.’” Weintraub (2013, p. 204) captures both as follows: “inference from a sample to the entire population, or to the next case.” (See also Pollock and Cruz 1999, p. 19.) Next, here's Audi (2011, p. 187), introducing abductive reasoning as “reasoning [...] from a premise stating the likeliest explanation of a presumed fact, to the conclusion that the proposition expressing that explanation is true.” Lipton (2004, p. 184), in turn, glosses abduction as a rule by which “scientists infer the hypothesis that would, if correct, provide the best explanation of the available evidence.” (See also Harman 1965, p. 89; Lycan 2002, p. 408; Pierce 1935, 5.189.)

<sup>21</sup> I say ‘something like’ correctness since some models—some formal models in particular—are arguably best regarded as idealizations aspiring to maximum coherence or utility relative to some external aim (rather than truth).

reasonably well-articulated set of questions and concerns to which they matter. That may be right. Nothing here turns on it.

But some of these differences in representation—differences of both types—make a difference to the precise demands that the standard availability-constraint makes on *S*. Correspondingly: to the explanatory weight that that constraint carries. (And note that the less work the availability-constraint does, in explaining why *A* in particular constitutes *S*'s inferential justification, the more work some other—perhaps not yet specified—requirement must be doing: e.g. a constraint on *S*'s relationship to the operative inference rule/rules, a meta-requirement of the sort mentioned in Sect. 1.2, or something entirely different. More on this in Sect. 3.3.) At one extreme, a representation of *F* that involves *no* further premises *guarantees* that *S* meets the availability-constraint, with respect to *A* (as long as her antecedently justified belief in *q* is granted). At the other, a representation that features further premises some of which are epistemically or psychologically—e.g. conceptually—out-of-reach for *S* guarantees that she violates it. And there would seem to be a spectrum of in-between options: options that issue variably taxing claims on *S*.

Which representation is the (or a) 'right' one—right, relative to the question whether *S* satisfies the availability-constraint? (Or, if rightness is a context-sensitive matter: what are the variables? What kinds of factors fix which representation is the right one, in context?) I'm not aware of a good answer to this question.<sup>22</sup> It raises a difficult and, I think, underappreciated challenge: one that complicates the application of the standard conception to cases—not just controversial cases, but in general.

For instance, it complicates its application to mundane—and seemingly straightforward—cases like that of Aya and Sven.

Suppose that, in this example, there are indeed grounds for assimilating the identified little two-premise structure to some familiar, recognizably good, argument form (at the level of generality that, in line with standard practice, we've operated with so far). That's to say: suppose we're able to rationally narrow down the range of inference-rules that might be involved here to rules of one broad familiar type, and a good one at that. How? Perhaps—to a rough first approximation—by relying on some combination of additional testimony, interpretative charity, and various pieces of background knowledge about the agent(s) and the specifics of the situation.<sup>23</sup> Suppose, furthermore, that the broad type in question is abduction. (E.g. if pressed to elaborate, Aya might volunteer: 'why *else* would he be angry and soaking wet?', and 'Sven would never forget his umbrella when the forecast says 'rain''. Background information might include that she lacks directly relevant statistical data—e.g. Sven never turned up in this state before—that she's known him for years, and that her general folk-psychological skills

<sup>22</sup> I'm not even aware of any explicit discussion of it in print. But there are points of contact with the discussion in Kyburg (2008, esp. §2). And the epistemic significance of some seemingly substantive variation, in particular in the explication of induction, has received considerable attention historically—e.g. Carnap (1950, 1968), Hempel (1945), Hume (1748/1975), and Quine (1977).

<sup>23</sup> It's an intriguing question—for another occasion—why, and to what extent, these kinds of clues are in fact reliable, and/or deserve to be treated as evidence (that such-and-such apparent enthymeme instantiates such-and-such familiar form) in the context of ascribing inferential justification to agents, real or imaginary. Here I'll proceed on the assumption that, under some suitably sharpened description, the practice of relying on such clues has at least roughly the merits we pre-theoretically assign to it, in seemingly straightforward cases like this one.

are in shape. And charity might supply that she didn't use, say, a blatantly invalid deductive rule, no rule at all, or inference-to-the-worst-explanation.)

In sum: suppose that there are *prima facie* grounds—furnished by the dual asymmetric dependence test, perhaps together with the agent's testimony—for taking Aya to be inferentially justified in believing that someone stole Sven's umbrella, partly on the basis (of antecedently justifiably believing) that he's angry and soaking wet; and, moreover, that there are *prima facie* grounds—of the local ordinary sort just mentioned—for taking the main inference rule at play in the case to be abductive, rather than inductive or deductive (or of some other contrasting type on our paradigm list).

Of course it might be indeterminate, and in many real-life cases probably *is* indeterminate, what rule is involved in a given case—even at this, very high, level of generality—although it's unclear if there are any such cases where (we'd still want to say that) the target belief is inferentially justified. All I'm suggesting here is that the case at hand is *not* like that; or better: that it's a case in which there's strong evidence, of the defeasible everyday sort, in favor of a determinate option—specifically, evidence that some kind of abductive rule is involved. (It's supposed to be an easy case, in this respect too.)

We've now helped ourselves to an initial regimentation of the argument Aya uses. We've also (by assuming the cautiously pluralist picture) granted that this form—the form the argument is regimented as having—is good. And we've granted that Aya believes the two manifest premises, *Sven is angry* and *Sven is soaking wet*, with antecedent justification. To satisfy the availability-constraint and the structural requirement, however, she must stand in that relation to *all* of the argument's non-redundant premises. And unless there's an answer to the question of relevant completeness, there would seem to be no fact of the matter about what those premises are. (Hence no fact of the matter about whether Aya has what it takes, by the lights of the standard conception, to be inferentially justified after all.) The initial regimentation, of the argument as abductive, points to a certain range of options—of possible would-be complete specifications—but that's all it does.

### 3.2 Illustration: abduction

There are several live—apparently substantive—controversies over how to informally explicate abduction, over its correct formalization, and over the possibility of reducing it to other inference rules (see e.g. Bird 2005; Harman 1965; Fumerton 1980, 1992; Lipton 2004; Psillos 2004; van Fraassen 1989; Weintraub 2013; Weisberg 2009). Crosscutting those controversies, there's a variety of—mutually compatible—ways in which to represent or specify the target rule, and corresponding argument form: ways that differ with respect to grain, division of labor, order and number of premises and rules, etc. Which of those is relevantly complete? Which representation of the (or an) abductive argument, from *Sven is angry* and *Sven is soaking wet* to *someone stole Sven's umbrella*, is the right one—right, relative to the question whether all of its essential premises are available to Aya?

It seems clear upfront that some representational differences don't make an important difference here—e.g. whether we treat the initially identified content/s as one conjunctive premise or as two separate premises—but that others do.

Consider first the question of grain. Arguably abduction breaks down in interestingly different sub-types, corresponding to differences in the kinds of explanation they exploit (e.g. Psillos 2007). Does a relevantly complete representation of the abductive argument that supposedly constitutes Aya's inferential justification make perspicuous which such sub-type is involved? Does it mark that the implicated explanation is, say, *causal* rather than *constitutive* or *nomological*? That it's *causal-rationalizing*? That it's *contrastive* rather than *absolute*, and that it's a *why*- not a *how*-explanation? Or is a more abstract representation of the argument form the 'right' one? How abstract is abstract enough? Does the relevantly complete representation only feature the determinable of which these are determinates: the umbrella notion of an explanation? Does it make perspicuous the criteria by which some such explanation is (deemed) the best? If so: in how much detail? Does the relevantly complete representation feature difficult-to-get-right covering notions like that of explanatory 'loveliness' (e.g. Lipton 2004), or—just barely finer in grain—of simplicity, precision, conservativeness, and so on? Even finer in grain, and assuming the success of a full or partial Bayesian reduction, does it feature a notion of conditional probability, or some more generic notion of likelihood (or neither)?

And *how* do theoretical covering notions like these feature in the right representation of the argument—supposing that they feature there at all? In (the contents of) auxiliary premises? That answer, if generalized, would seem to have the consequence that the availability-constraint is violated in numerous actual and counterfactual cases where—contrary to the verdict such violations suggest—the agents are *prima facie* inferentially justified. The proposed auxiliary premises are arguably both psychologically and epistemically out-of-reach for most ordinary agents (since their conceptual repertoire doesn't include the theoretical covering notions in question).<sup>24</sup>

Perhaps, then, the theoretical notions only belong in the specification of the inference rule(s). But it's in fact not entirely clear that that's *possible*. At least on the usual way of stating (or unpacking) an inference rule, the rule licenses a transition from a *belief* with a certain content—or a set of beliefs with certain contents—to other beliefs.<sup>25</sup> (Again: with the exception of rules that govern suppositional reasoning, and of contexts where the rules are understood to manipulate partial beliefs. But, insofar as the problem is that the conceptual demands are too steep, these differences—between all-out, partial, and conditional belief—don't seem to matter.) E.g. a standard specification of abduction that expressly captures the criterion of explanatory loveliness states a rule that, roughly put, sanctions the move from the belief that *hypothesis H is the loveliest* explanation of evidence E, and the belief that E obtains, to the belief that H is true.

At least it does, when unpacked in the usual way. Perhaps there's a way of stating or unpacking the rule on which it doesn't: one that expressly captures the criterion

<sup>24</sup> I assume that we can't have justification to believe things we can't believe—i.e. that even propositional justification is out-of-reach, with respect to psychologically unavailable contents.

<sup>25</sup> I say 'stating or unpacking' since standard specifications of rules needn't make explicit mention of belief. But it's generally understood that this is how to read them.

without requiring, for the application of the rule, that the agent have beliefs about (and/or justification to believe claims about) loveliness. But that remains to be seen—in particular, it remains to be seen whether there’s a way of doing so that’s well-motivated in context: that of applying the standard conception to cases. (Recall that the question isn’t whether abduction, or abduction done right, involves something like Lipton’s loveliness criterion: it’s whether, *on the assumption that it does*, a relevantly complete representation of abductive arguments makes that perspicuous—and, if so, in what way.)

There are other candidate auxiliaries, of course. The elimination of salient alternative hypotheses seems integral to abduction. Where, in a relevantly complete representation of the target argument, does the exclusion of such candidates appear? (The same question arises about the exclusion of non-salient options or ‘non-starters’.) Does the argument in our example have auxiliary premises to the effect that, say, *Sven would never forget his umbrella*, or that *if he’d done so he wouldn’t be angry*? That *Sven didn’t just take a shower fully dressed*? That *normally, when humans frown like that, they’re angry*? That *cats having four legs doesn’t explain the data at all*? Presumably not every proposition of intuitive relevance to the explanatory force of Aya’s verdict—nor every such proposition that she could believe, or could believe with antecedent justification—goes into the argument that constitutes her justification (on the working hypothesis that some argument does). But it’s unclear what’s ‘in’ and what’s ‘out’—and, loosely put, *how* what’s in, is in. It’s also unclear what a plausible principle that settles these questions would look like.

How about every intuitively relevant proposition that Aya *does* believe—or, even better, believes with antecedent justification? Assuming the standard availability-constraint, which surely is permissible here, those are the only propositions of interest in any case.

Yes. But is it permissible to assume that the constraint is *met*? Our question was: which representation of an argument form F is relevantly complete—with respect to the demands of that constraint—for a given agent and context? The suggestion just aired is that it’s the (or a) representation that features all and only those propositions the agent believes with antecedent justification (and that bear on the viability of the conclusion, in some yet-to-be-precisified way). The question ‘featured *how*?’ matters less, if this is right—we might as well say ‘as premises’. But the suggestion closes off the possibility that the availability-constraint is violated in the case under consideration. The agent could still lack inferential justification, if she violates some condition we haven’t yet articulated (or if F is after all a bad form). But the availability-constraint is satisfied, trivially.

To repeat: what we’re after is a principled criterion for the (or a) representation of argument form F that’s pertinent to the demands of the availability-constraint. What exactly that constraint comes to, in a given case—which propositions the agent must have the requisite relationships to—partly depends on how F is represented or specified: how its internal structure is conceived. So, I’m asking: which representation

is the right one? Proposed answer: (the) one that has her come out satisfying the constraint.<sup>26</sup>

Other proposals are addressed in Sect. 4. (There we also return to other argument forms.) But first, to fix ideas, here's a schematic illustration of just a few of the representational options, for the case of abduction. A sample 'two-stage' explication of the inference rule(s)—adapted from Douven (2017) and Lipton (1993, 2004)—serves as the starting-point. It's immediately followed by two ostensibly compatible variations, each finer in grain than the next, and then a range of alternate renditions of the corresponding argument(s).

Both the starting-point, (IBE1), and the refinements, (IBE2)–(IBE3), are controversial—e.g. the inclusion of a separate stage where non-starters are filtered out, and the role assigned to probability considerations. But to simplify I here ignore manifestly substantive variation. A different starting-point just yields a different range of options downstream.

- (IBE1) Given evidence  $E$  and candidate explanations  $H1... Hn$  of  $E$ , infer the truth (alternatively: the probable/approximate truth<sup>27</sup>) of that  $Hi$  which best explains  $E$ , provided  $Hi$  is good enough *qua* explanation to start with
- (IBE2) Given  $E$  and  $H1... Hn$  of  $E$ , infer the truth of that  $Hi$  which provides the likeliest explanation of  $E$ , provided  $Hi$  is lovely enough to start with
- (IBE3) Given  $E$  and  $H... Hn$  of  $E$ , infer the truth of that  $Hi$  which raises the conditional probability of  $E$  more than any of the other candidates, provided  $Hi$  is simple, unified, conservative, precise, (etc..) enough to start with

In simple argument format:

- (ARG1)  $E$   
 $Hi$  is the best available explanation of  $E$   
 ( $Hi$  is good enough to start with)  
 ----- *ibe*  
 $Hi$
- (ARG2)  $E$   
 $Hi$  is likelier than any of the other candidates  
 ( $Hi$  is lovely enough to start with)  
 ----- *ibe*  
 $Hi$

<sup>26</sup> The variant proposal—on which every intuitively relevant proposition the agent believes is an essential premise in the argument—doesn't completely trivialize the constraint. But both variants have other problematic consequences. For one thing, they threaten to obliterate any distinction between background evidence and auxiliary inferential grounds, and between discrete inferential justifications for the same belief. They also threaten to relativize the requirements that comprise the standard conception to contingent psychological facts about agents in an unacceptable way. (E.g. provided Aya has no further beliefs, or further relevant beliefs, she would now satisfy the availability-constraint.) Depending on how the range of 'intuitively relevant' propositions is demarcated, in more precise statements of the proposal/s, some of these consequences might be avoided. But note that a version of the problem of relevant completeness surfaces at this point—as a demarcation-problem about that range.

<sup>27</sup> For brevity, this alternative is omitted below.

(ARG3) *E*  
*Hi* raises the conditional probability of *E* more than any of the other candidates  
(*Hi* is simple, unified, conservative, precise, etc., enough to start with)  
----- *ibe*  
*Hi*

‘----- *ibe*’ is introduced to designate abductive implication. Alternatively, the overall rule could be cast as a conditional bridging premise; e.g. as in ARG3\*. We can then use familiar deductive rules—perhaps most plausibly, modus ponens (with material implication).

(ARG3\*) *E*  
*Hi* is the best available explanation of *E*  
(*Hi* is good enough to start with)  
If *Hi* is the best available explanation of *E* (and good enough to start with), *Hi*  
-----  
*Hi*

Another option is to operate with more specific rules (and/or more specific bridging premises)—rules particular to different sub-types of abduction; e.g. as in ARG3\*\*.

(ARG3\*\*) *E*  
*Hi* is the best available causal explanation of *E*  
(*Hi* is good enough to start with)  
----- *ibe-causal*  
*Hi*

Further variations:

(ARG4) *E*  
*Hi* is *F*  
(*Hi* is *G*)  
----- *ibe*  
*Hi*

----- where *F* is what *makes* (or perhaps: what *seems to the agent to make/is rational for her to believe makes*) *Hi* the best available explanation of *E*, and *G* is what *makes* (or: what seems to the agent to make/is rational for her to believe make) *Hi* good enough to start with.

(ARG5) *E*  
----- *ibe*  
*Hi*

(ARG6) *E*  
*E* wouldn’t be the case unless *Hi*  
----- *ibe*  
*Hi*

(ARG7) *E*  
*E* because *Hi*  
----- *ibe-causal*  
*Hi*

(ARG8) *E*  
If *Hi* were the case, *E* would be the case  
----- *ibe*  
*Hi*

Like (ARG3\*\*), (ARG6)–(ARG8) are perhaps best seen as representations of distinct sub-types of abduction (that admit of further breakdown in turn). But each is still compatible with all the remaining variations, considered as representations of the umbrella category. (ARG5) might seem absurdly bare, but it's a natural candidate for relevant completeness on a 'material-inference' approach to abduction (e.g. Brandom 2000, p. 87; Brigandt 2010; Norton *forthcoming*). I'll say a little more about this soon.

And it's clear that the above examples don't exhaust the options. Even finer-grained representations (albeit almost all of them partial) can be found in the formal literature.<sup>28</sup> Another dimension of variation is buried in the clunky use of 'E' to stand for all the relevant evidence.<sup>29</sup> As discussed in connection with the mundane case, it's far from obvious what that includes.

### 3.3 Interim discussion

The problem of relevant completeness is perplexing. It suggests that it's indeterminate whether Aya meets the requirements that comprise the standard conception (or whether the case could easily be realized in such a way that she does), because those requirements leave open which if any further propositions must be available to an agent in her situation. And it's not clear how to solve the problem, or that it even has a solution. Relatedly, it does have the ring of a pseudo-problem. Perhaps it's just an artifact of a poorly formulated principle—presumably, the availability-constraint (and/or the structural requirement that that constraint feeds into). But can the constraint be reformulated so as to avoid the problem? If not, is it an option to get rid of it?

Not without loss. Consider the job the availability-constraint does, in the relevant (proto-)theory: the standard conception. It supplies a principled criterion for when—and a partial explanation of how—an otherwise suitable proposition can contribute to an agent's justification (specifically: her *inferential* justification, but the core achievement transcends that category). In other words: the constraint answers a certain selection problem—'why  $x$  rather than  $y$  or  $z$  (or nothing at all)?'—and it does so in a principled way, while giving a partial explanation of how that answer could indeed be the right one: of how it's possible for an auxiliary proposition  $r$  to play a certain justificatory and explanatory role (for  $S$ , with respect to the belief that  $p$ ). Presumably any adequate reformulation—or substantive amendment or replacement—of the constraint would have to do this job at least equally well. And it should either solve, or avoid, the problem of relevant completeness.

Meanwhile: for *whom* is it a problem? Since the standard conception is the best generic account of inferential justification we have (and specific extant versions of it don't resolve the issue), the fact that the conception has this 'blank' in it does in principle bear on *any* undecided case—in particular, any *prima facie* case of inferential justification or knowledge, into which it's not simply stipulated what the requisite

<sup>28</sup> See e.g. Aliseda (2006), Climenhaga (2017a, b), Douven (1999, 2017), Schupbach (2017).

<sup>29</sup> Cf. Kyburg (2008, 294): "What is 'the evidence' that the agent has? It is all very well to symbolize the evidence with the single letter 'E', but if the point is to rationalize or even represent the relation between evidence and hypothesis this is not much help." (Kyburg is discussing statistical-inductive arguments here, but my concern applies equally to those.)

premises are. The crucial question about Aya, articulated in Sect. 2.2—whether she’s antecedently justified in believing all the essential premises of a good argument, from the initially identified premises to the conclusion—presupposes that there’s a fact of the matter about what counts as the full set of essential premises, of an argument with a given (good) form; it has no answer if there isn’t. The point generalizes to all structurally similar examples.

And the problem takes on *weight* with respect to controversial cases. For instance, cases that feature justification of a ‘source-individuated’ kind whose status in precisely this regard—as inferential or non-inferential—is in dispute; such as the *testimonial*, *memory-based*, or *introspective* kind. Whether justifications of these kinds are inferential is generally taken to depend on whether all the non-redundant premises of suitable good arguments are available to agents in canonical cases. But what these premises are, or might be—which propositions qualify—turns on how the questions of goodness and completeness are answered.<sup>30</sup> Furthermore, if argument representations that are *very* light on premises (or otherwise austere with respect to premise-content) are or can be relevantly complete, then *something else* must distinguish inferential justifications that embody (say) abductive arguments from justifications that embody deductive arguments—and from things that aren’t justifications at all, because the arguments involved are bad. Whatever principle picks up the explanatory slack, it introduces another variable: another factor to contend with when interpreting canonical cases (and beyond).

To clarify, suppose that (ARG5) is indeed a relevantly complete representation of abduction, and ditto for the corresponding representations of contrasting argument forms, and that *Sven is angry* and *Sven is soaking wet* exhausts Aya’s evidence in the mundane example. Then something else—something other than her meeting the availability-constraint with regard to further premises—must do all the work, in explaining *why the initial regimentation of the argument as abductive is correct* (if it is): in virtue of what Aya is using an abductive rather than, say, a deductive or inductive or counter-inductive rule, using a certain abductive rule rather than another, or using no rule at all. The present point isn’t (just) that we’re owed a story; it’s that the story will inevitably introduce substantial new demands. (E.g. that the agent ‘register’ or represent certain auxiliary propositions in some other way, that she manifest some distinctive cognitive disposition, and/or that she register a suitably differentiating bridge-proposition—e.g. *that the premises F-ly imply the conclusion*. See e.g. Boghossian 2014; Rosa 2019.)<sup>31</sup> Whether those demands are satisfied, in ways that

<sup>30</sup> There’s some discussion of goodness, but none of completeness, in the relevant literature. (For the case of introspection, see e.g. Aydede 2003; Barnett 2016; Byrne 2005, 2011; Gallois 1996. For memory, e.g. Audi 1995; Bonjour 2002; Senor 1993. For testimony, e.g. Fricker 1994, 1995; Lackey 2006; Schiffer 2001). The reductionism/non-reductionism debate in the epistemology of testimony is a clear example where the question of completeness matters—see e.g. Burge 1993, 1997; Coady 1973, 1992; Fricker *ibid.*; Graham 2006; Malmgren 2006; Reid 1764/1997. (Both questions also affect the foundationalism/coherentism debate; in particular, the question of ‘epistemic ascent’—see e.g. Bonjour and Sosa 2003; Pollock and Cruz 1999, Ch. 2.)

<sup>31</sup> On Norton’s version of the material-inference approach, the agent’s (antecedent) knowledge of/justified belief in certain domain-specific ‘material postulates’ distinguishes applications of inductive and abductive rules that can generate knowledge/justification from applications that can’t (e.g. Norton 2003, 2014). Since the postulates vary in content, depending in part on which rule is involved—which of those broad types,

procure good argument forms, is then another major consideration that should inform how cases are sorted—including cases that feature justification of source-individuated kinds whose status as inferential or non-inferential is already a matter of controversy.

The problem also takes on weight with respect to controversial cases that feature justification for believing certain challenging claims, or claims on certain challenging topics—e.g. that *there's an external world*, that *modus ponens is valid*, or that *that [salient object of perceptual experience] is a vice-consul* (assuming that perception can't represent vice-consuls as such). And it takes on weight with respect to the hard-to-classify cases mentioned at the outset (and further discussed in Malmgren 2018, and *ms.*)

Let's take stock. The question of relevant completeness is part of a nexus of questions that complicate the application of the standard conception to cases. To settle if the availability-constraint is met, in a given undecided case, we need to know how the target argument should be represented, for the purpose of that task; to settle if that argument is good, we need to know what it takes to be good—and, to do either, we need to know which argument is (or arguments are) at issue in the first place: what form it has. The last question—of initial regimentation—is especially pressing when we face apparent enthymemes, and the kinds of clues we added to the mundane case in Sect. 3.1 are absent. (The hard-to-classify cases are cases in point.) But what the discussion of the mundane case makes clear is that such clues don't—they can't—settle the question of relevant completeness. The picture of argument goodness we've been working with so far doesn't settle that question either.

One of the issues to be investigated next is whether that picture of goodness is at fault. Another is whether relevant completeness can be understood with appeal to *logical* form. This suggestion is particularly problematic in application to ampliative arguments—partly for familiar reasons to do with such arguments' resistance to exhaustive formalization—but, as we'll see, it's unclear whether it even works for simple structural entailments.

## 4 Responses

### 4.1 Goodness revisited

It might be suggested that the relevantly complete representation of an argument form F is one that reveals whatever feature makes F good (or, presumably, bad). But, again, what feature is that? In line with the cautiously pluralist picture, I've been assuming that we don't have a unified answer to this question—at any rate not yet. Being on the list of *prima facie* paradigms, or being relevantly related to those paradigms, isn't a unifying feature of a helpful sort (even if '*prima facie*' is dropped). Not only is it

---

Footnote 31 continued

and which specific rule within each category (2003, §3–4)—they could also help distinguish between applications of different rules. However, given that the agent's *epistemic and psychological relationship* to a suitable postulate is integral to the explanation (cf. Kelly 2010, §3), these postulates effectively just function as further premises. That is: on Norton's view, very austere ('premise-minimalist') argument representations are not after all relevantly complete. Rather, the view vindicates a solution to the problem of completeness that marks off rather fine-grained sub-types of ampliative argument forms from one another, in part by virtue of including different material postulates among the respective non-redundant premises.

not explanatory—presumably the paradigms instantiate goodness/badness, if they do, in virtue of some other feature: one that also explains why they're on the list—but it doesn't help at all with the current problem: that some argument form is among the paradigms doesn't, by itself, tell us what the relevantly complete representation of that form is.

A more substantive suggestion would be that the way—whatever way—an argument is specified on standard extant versions of that list is relevantly complete. But this can't be right either; not because of too much variation on this score, between extant versions, but because the paradigms—especially the ampliative ones—are so coarsely carved, and loosely described, on those lists (cf. Sects. 1.3, 3.1). That's arguably what prompted the quest for the relevantly complete representations of these argument forms in the first place: the proffered descriptions don't settle all the details needed for the availability-constraint to issue determinate demands—in particular, they don't fix the arguments' essential premises. (The deductive paradigms seem better off in this regard; more on this shortly.)

Perhaps that's the root of the problem: perhaps the paradigms are less coarsely carved, and less loosely specified, on the *correct* list. (But—to anticipate—then what is that list?) Perhaps it was a mistake to try to differentiate good from bad argument forms, and paradigms of each type, in such rough-and-ready terms, and at such high levels of generality—e.g. count (sufficiently strong) inductive arguments as good, and counter-inductive arguments as bad, while saying close to nothing about the internal structure of these arguments. The correct list (or theory) might feature argument forms so precisely specified that no problematic representational variation is possible—at the limit, it may not even *make sense* to suppose that they can be represented in different ways: a different representation of an argument is simply a different argument.<sup>32</sup>

This diagnosis shifts the weight of the problem a bit, but doesn't make it go away. With exactly *how* much precision—and with precision in *what respects*—are the arguments or rules on the (correct) list specified? We were asking for a principled way of singling out the relevantly complete representation(s) of an argument form, from other representations of it—complete, relative to the application of the availability-constraint (for a given agent and context). We now face the need for a principled way of singling out the representation that's apt to the question of argument goodness—to the determination of an argument as good or bad—or, as per the limit-case suggestion, a principled way of individuating argument forms (that's apt to that question). It doesn't look like we've made progress yet.

Let's return to the opening thought: that a relevantly complete representation of F reveals whatever feature makes F good (or bad). We may lack a fully general account of argument goodness, but we can surely say a *little* more—notably, that all good deductive forms are *valid* (or valid-in-a-system<sup>33</sup>), and all good inductive and abductive forms are (inductively/abductively) *strong*. Perhaps a relevantly complete representation of F contains all and only those premises and rules that are needed to reveal F's validity, and inductive or abductive strength, respectively? (Alternatively: perhaps that's the kind of representation that's pertinent to the question of argument goodness.

<sup>32</sup> Thanks to R.A. Briggs for pressing this point.

<sup>33</sup> Or, strictly speaking, in the meta-theory of a system.

The suggestion could also be adapted into an individuation criterion for argument forms. For all we've seen so far, there's no obvious advantage to one of these ways of talking over the others.<sup>34</sup>)

## 4.2 Logical form

What, then, does it take to 'reveal' those features? For deductive validity, it's tempting to appeal to *logical form* at this point—at least for structural or formal (and explicitly logical) validity, in contrast to analytic or lexical validity. (To illustrate: the little argument from *Sven is angry and soaking wet* to *Sven is angry* is ostensibly structurally valid. The argument from *Sven is Aya's brother* to *Sven is Aya's sibling* isn't, but it's ostensibly analytically valid.) On a traditional picture, logical form *explains* structural validity: an argument is structurally valid, when it is, in virtue of instantiating a more abstract form that's logically valid (e.g. Davidson 1984; Montague 1973). In straightforward cases, the validity of the more abstract form simply falls out of the axioms that govern its fixed (or logical) content-components, and the ways these and other content-parts are arranged in premises and conclusion. Any argument of that form—with the same fixed parts, and the same overall combinatorics—is guaranteed to preserve truth, and this can be proved (within the given formal system).

If structural validity is, at least in part, what makes arguments that have it good arguments, and if structural validity reduces to logical validity in roughly this way, then specifications of such arguments that make explicit their logical form might serve present purposes too—they might count as revealing (part of) these arguments' goodness-making features. (Conversely, presumably, for structurally *invalid* arguments with a claim to validity—however exactly that's to be understood.) The proposal, then, is that a relevantly complete representation of a structurally valid argument form *F* makes explicit *F*'s logical form (alternatively: the logical form of a canonical statement of *F*).<sup>35</sup>

There's arguably more than one notion of logical form at work in the classical debate.<sup>36</sup> In addition to accounting for structural validity, logical form has traditionally been expected to play a central role in the theory of linguistic meaning, and it's unclear whether there's a unique notion of form that can play both roles (see Iacona 2016). But we can set this aside. The present proposal is only committed to there being a notion of logical form that explains structural validity (and to some non-arbitrary way of mapping arguments, identified in other terms, onto such forms). The claim is that

<sup>34</sup> But the last way of talking is highly revisionary, since it obliterates the distinction between substantive and non-substantive variation. (It would be quite surprising if, e.g., none of the apparent disagreements over how to explicate induction or abduction turned out to be genuine.)

<sup>35</sup> Logical form is often attributed to sentences—in natural or formal languages—and to arguments conceived as sets of such sentences. But for defense of the view that the primary bearers are propositions, see e.g. Cargile (2010) and Russell (1914).

<sup>36</sup> Even bracketing (the notion of) covert syntactic structure that some linguists call 'Logical Form'/'LF' (e.g. Chomsky 1995). LF corresponds many-one to (and perhaps partly determines) logical form, in the sense that's relevant to compositional semantics—let's call that 'semantic form.' If Iacona 2016 is right, semantic form corresponds many-one to (and partly determines) logical form, in the sense that's needed to explain structural validity—'truth-conditional form.' (See also Jackson 2006; Peregrin and Svoboda 2013.)

it—that notion—can also help resolve the problem of relevant completeness, at least for structurally valid arguments (and, again, for structurally invalid arguments with a claim to validity).

To see how this plays out, let's start by supposing that the language of standard (classical, propositional and predicate) logic is at least on the right track, with respect to that notion of form—that it has the resources to represent a core range of structurally valid arguments such that their validity is shown to be a function of their logical form(s). Using those resources, we can cast the little argument from *Sven is angry and soaking wet* to *Sven is angry* as a straightforward instance of conjunction elimination—e.g. as ' $Fa \wedge Ga \therefore Fa$ ' or, even simpler, as ' $p \wedge q \therefore p$ '. For another example, we can cast the argument from *Sven is angry and soaking wet* and *if Sven is angry and soaking wet then someone stole his umbrella* to *someone stole his umbrella*, as a straightforward instance of modus ponens—e.g. as ' $((Fa \wedge Ga) \wedge (Fa \wedge Ga \rightarrow Ha)) \therefore Ha$ ' or, simpler, as ' $p \wedge (p \rightarrow q) \therefore q$ '.

But the proposal can hardly be that these schematic descriptions per se are relevantly complete representations of the arguments. (It's not even clear how to understand this suggestion.) Nor can it be that any consistent interpretation of each schema is relevantly complete. The availability-constraint comes out issuing an absurd demand on that reading: that the agent have antecedently justified beliefs in *some or other* premises of the same logical form—relative to her belief in *some or other* conclusion of the same logical form—as the target premises and conclusion. It's extremely hard to see how to motivate a demand like that, and it's much too weak to do the work the constraint is supposed to be doing (cf. Sects. 1.4, 3.3).

The proposal must be, or amended to be, that a relevantly complete representation of a structurally valid argument reveals its logical form *and its non-logical content* (if any). Bracketing hard questions about how precisely to demarcate logical from non-logical content (e.g. MacFarlane 2015; Jackson 2007), we can think of the 'logical' content as what's expressed—or at least approximated—by the arrangement of logical constants, letters and variables, that make up the schematic descriptions of the argument that are designed for use in proofs: descriptions of the sort exemplified above. Whatever is left—the content expressed by ordinary singular terms, predicates and proper names, in natural-language renditions of the argument—is its non-logical content. The amended criterion yields semi-formal descriptions of the two sample arguments along the following lines: '*Sven is angry*  $\wedge$  *Sven is soaking wet*  $\therefore$  *Sven is angry*', and ' $(Sven\ is\ angry \wedge Sven\ is\ soaking\ wet) \wedge ((Sven\ is\ angry \wedge Sven\ is\ soaking\ wet) \rightarrow \exists(x) x(stole\ Sven's\ umbrella)) \therefore \exists(x) x(stole\ Sven's\ umbrella)$ '. The proposal, finally, is that *those* are the kinds of representation that are pertinent to the application of the availability-constraint.

A natural worry is that some of the logical expressions in these hybrid representations—still taking them to be expressions in the language of standard logic—differ in meaning from their natural language correlates: the material conditional arguably doesn't capture any of the implication-relations denoted in ordinary English by 'if... then' (nor does the strict conditional), logical conjunction doesn't discriminate between coordinating and contrasting conjunction—etc. (See e.g. Edgington 1995; Kratzer 1986; Grice 1989.) Could it really be that the relevantly complete representation of the argument that goes to constitute the agent's putative justification,

in a mundane case like that of Aya and Sven, features premises containing those concepts—concepts she doesn't plausibly employ in the context described (or even possesses)? The agent would then violate the availability-constraint, with respect to those premises.<sup>37</sup> But that seems very implausible.

Perhaps standard logic is the culprit. After all, some non-classical systems—notably relevance logic (e.g. Anderson and Belnap 1975; Priest 1979, 2008)—were developed in large part to avoid the apparent paradoxes and fallacies that indicate a mismatch between the classical logical concepts and their informal counterparts. With respect to deductive arguments, and everyday cases, perhaps we should defer to the best (possibly future) version of non-classical logic—and say that the relevantly complete representations reveal logical form, as captured in the language of *that* logic (and, again, non-logical content)?

This doesn't yet commit us to saying that some non-classical system does a better job overall, in sorting good from bad (deductive) arguments, than does standard logic. The proposal is just that, for agreed-on paradigms, the problem of relevant completeness can be solved with appeal to the forms assigned to these arguments by the best non-classical logic. Depending on what further desiderata there are on the system—by what total combination of virtues it counts as the best—it may not matter whether it also sanctions some arguments that aren't good (in our sense) or conversely.<sup>38</sup>

But before we know what the proffered system looks like, and what its wider theoretical ambitions are, it's very difficult to properly evaluate the suggestion.<sup>39</sup> And this by itself is some cause for concern. One might have thought that the solution to, or dissolution of, the problem of relevant completeness shouldn't require that we adjudicate between competing formal systems—and then work our way back from the winner, as it were.

It's also hard to know how much of a departure to expect, from the previous (formal and hybrid) representations within the scope of the proposal—representations that display 'classical' logical form. In the first instance there's just the difference we wanted: in the interpretation of the constants. (E.g. ' $\rightarrow$ ' may now express a non-classical implication-relation.) For some systems, and some argument forms, this may be the only difference with consequences for availability. But for others it may not be. It depends, again, on which non-classical system is at issue. It also depends on how certain disagreements are interpreted.<sup>40</sup>

<sup>37</sup> E.g. on the current interpretation, '(Sven is angry  $\wedge$  Sven is soaking wet)  $\rightarrow$   $\exists(x) x$ (stole Sven's umbrella)' expresses a different proposition from 'if Sven is angry and soaking wet, someone stole his umbrella' (used in ordinary contexts). Perhaps the justificatory demands are the same here, but they needn't be—and the psychological demands aren't, since the contents differ.

<sup>38</sup> Perhaps some non-classical system does a better job than standard logic of capturing a fragment of the total set of (structurally valid) arguments fit to constitute justifications, although it fails to capture the whole set. Perhaps different such systems win out for different fragments, and perhaps there's no single consistent system, classical or other, which captures all.

<sup>39</sup> It may exist already, but there are many contenders, and the terms of the debate are quite unclear. (See e.g. Beall and Restall 2000; Russell 2008, 2018.)

<sup>40</sup> To illustrate: on one way of understanding key disputes between classical and free logic (e.g. Lambert 1958, 1981; Leonard 1956) they imply disputes about logical form—e.g. of arguments by universal elimination and existential introduction. In a free logic, a further premise (expressing the assumption that the domain of quantification is non-empty) is required to get from the claim that everything is *G* to the claim

### 4.3 Scope

Finally, suppose that some version of the logical form criterion does in fact work, for the structurally valid arguments under consideration. This still leaves the problem hanging for analytically valid and ampliative arguments. It's not even clear that these argument-types *have* logical forms, in the sense we've granted to structural entailments—that they fall under more abstract schemata whose fixed content-parts and combinatorics guarantee the corresponding goodness-making feature(s) of their instances: analytic validity and inductive/abductive strength, respectively. Ambitious attempts at identifying such forms, and providing a germane logic, have been made—in particular for enumerative-statistical induction (e.g. Carnap 1950; Hempel 1945; Keynes 1921). But there are familiar problems with these attempts (e.g. Goodman 1955; Titelbaum 2010). The non-logical content of an inductive argument can make all the difference to its strength. Correspondingly for analytic entailments (*pace* Carnap 1952) and, it would seem, abductive arguments (*pace* e.g. Aliseda 2006).

The non-logical content can also make a difference to the epistemic and psychological demands issued by the availability-constraint—even for structurally valid arguments. That's why we settled for hybrid representations of such arguments earlier: representations that reveal logical form *and non-logical content*. Perhaps we could extend that suggestion to analytic and ampliative arguments; thus bypass the last mentioned complication?

There are several problems with this idea. For one thing, it's highly unclear what logical forms—or *ersatz* logical forms—to assign to these argument-types (even when the hope that they're fully tractable has been abandoned). Once again there's non-substantive variation—e.g. between (here, formal) representations that differ with respect to the division of labor, and between representations pitched at different levels of grain.<sup>41</sup> There's also substantive variation, some of it consequent on disputes over how to informally explicate the target arguments or rules.<sup>42</sup> The choice of formal model is moreover affected by whether, and how, abduction is treated in full generality or divided into sub-types upfront. Likewise, of course, for induction.

---

Footnote 40 continued

that some particular object  $o$  is  $G$ , or from the claim that  $o$  is  $G$  to the claim that there's some particular  $o$  that is  $G$ . This is sometimes expressed by saying that universal elimination and existential introduction are invalid rules, in this system, but that they have (weaker) valid analogues (see e.g. Nolt 2014). Alternatively, the rules and arguments sanctioned by the free logic could be seen as competing specifications of form—specifications that depart from classical options in ways that notably affect the demands of the availability-constraint. Indeed, it's only under that description that the free logic schemata even make sense, as candidate formalizations of these and other arguments containing quantifiers! Thus the former description shifts the weight of our problem (again) but doesn't make it go away. If we favor that description, but think that free logic should at least be in the running here, we may need to expand the paradigm list—to include arguments by the corresponding free logic rules—or restate the affected entries in terms that leave this issue open. (*Mutatis mutandis* for relevance logic and arguments by modus ponens, by conjunction introduction, etc.)

<sup>41</sup> E.g. the formalism in Horty (2012) is abstract enough that descriptions of arguments in terms of it are compatible with finer-grained descriptions that, say, trade in probabilities.

<sup>42</sup> E.g. if the correct informal explication of abduction makes no mention of likelihoods—*pace* (IBE2) and (IBE3) above—then presumably an even partly Bayesian, or otherwise probabilistic, formal explication is a non-starter.

There are also disputes—generating further variation—which turn on how competing formal frameworks fare with respect to certain overall theoretical desiderata (some of which are themselves contested). This isn't the place for a general discussion of these desiderata, or of the wider aim(s) and scope of formal models of ampliative arguments (see e.g. Christensen 2004). Let me just note that it's hard to see why systems that exhibit them—whatever else they might be good for—should be expected to yield argument representations that are *relevantly complete*, or from which such representations are easily constructed (by the addition of case-specific non-logical content).

Extant systems in this family certainly don't. Most recent attempts at formalizing both induction and abduction trade in graded rather than categorical belief, and it's unclear how to think of the relation between the two frameworks (and the two attitude-types). But we need a good grip on that—or at least: a plausible conversion scheme—to get argument specifications that suit present purposes out of what these systems provide.<sup>43</sup> Relatedly, some of the heavy idealizations that are built into leading degree-theoretic systems are troubling.<sup>44</sup> On the assumption that actual human agents are sometimes inferentially justified, and that the standard conception is supposed to get those cases right (see Sect. 1.1)—we'd also need to know how to convert norms that govern (the behavior of) relevantly idealized agents to norms that govern *us*.

To clarify, the question of relevant completeness concerns how to specify a given argument form, relative to the demands of the availability-constraint—the requirement that agents be related in the right way to all the premises that are essential to arguments with that form. For an agent to be inferentially justified, by the lights of the standard conception, the form must also be *good*—its instances must be fit to constitute inferential justifications. Good *for whom*? Presumably: the very agent(s) under consideration. But the extent and nature of the idealizing assumptions that are built into leading degree-theoretic systems suggest that the arguments these systems sanction—and offer formal descriptions of—at best approximate arguments that are good for (even the best of) us, ordinary human agents. And it's unclear what remains, when the idealizations have been removed, and the approximations accordingly sharpened. It's correspondingly unclear how the arguments of interest would be specified—what (*ersatz*) logical forms they would have—in a system that dispenses with the idealizations. So it's simply an open question whether the representations to be found in such a system are relevantly complete (or can be turned into representations that are). As before, this by itself is some cause for concern.

#### 4.4 Conclusion

The preceding discussion exposes a deep unclarity at the core of the standard conception of inferential justification. Perhaps there's in fact a plausible head-on answer to

<sup>43</sup> For a notable attempt, see Foley (1992, 2009). (But cf. Christensen 2004, Ch. 4; Horgan 2017.)

<sup>44</sup> The attribution to the agent of numerically precise degrees of belief is one such idealization; others include her having the intellectual and computational resources to be logically omniscient. (The extent to which the idealizations are problematic depends on what the system presupposing them is used to do. Here I'm just concerned with the suggestion that it can help with the problem of relevant completeness.)

the question of relevant completeness forthcoming—perhaps it falls out of the right view of argument goodness—but it’s not at all clear what that answer, or that view, would be.<sup>45</sup> Perhaps this question has no answer save by stipulation. That, by itself, wouldn’t suffice to show that there’s no fact of the matter about whether a case like that of Aya and Sven exemplifies inferential justification. But it would at least establish that the standard conception is remarkably unhelpful, for understanding the intended distinction.

**Acknowledgements** I’m very grateful to Robert Audi, Brendan Balcerak-Jackson, David Barnett, Sinan Dogramaci, Patrick Greenough, Hilary Kornblith, Ram Neta, Michael Raven, Joshua Schechter, and two anonymous referees for this journal, for feedback on earlier drafts. Thanks also to Endre Begby, R.A. Briggs, Mark Crimmins, John Horty, Thomas Icard, Krista Lawlor, Helen Longino, Øysten Linnebo, and Kenneth Taylor for helpful discussion, and to audiences at Arizona State University, Cologne University, Konstanz University, Lund University, Miami University, Notre Dame University, Oberlin College, Oslo University, Simon Fraser University, UC Irvine, UNC Chapel Hill, UCLA, and UQAM, where previous versions of this material were presented in 2015–2016 and 2018–2019.

## References

- Alchourrón, C., Gärdenfors, P., & Makinson, D. (1985). On the logic of theory change. *Journal of Symbolic Logic*, 50, 510–530.
- Aliseda, A. (2006). *Abductive reasoning: Logical investigations into discovery and explanation*. Berlin: Springer.
- Anderson, A. R., & Belnap, N. P. (1975). *Entailment: The logic of relevance and necessity*. Princeton: PUP.
- Audi, R. (1986). Belief, reason and inference. *Philosophical Topics*, 14, 27–65.
- Audi, R. (1995). Memorial justification. *Philosophical Topics*, 23(1), 31–45.
- Audi, R. (2011). *Epistemology: A contemporary introduction* (3rd ed.). New York: Routledge.
- Aydede, M. (2003). Is introspection inferential? In B. Gertler (Ed.), *Privileged access: Philosophical accounts of self-knowledge*. Farnham: Ashgate Publishing.
- Barnett, D. (2016). Inferential justification and the transparency of belief. *Nous*, 50(1), 184–212.
- Beall, J. C., & Restall, G. (2000). Logical pluralism. *Australasian Journal of Philosophy*, 78(4), 475–493.
- Besson, C. (2018). Norms, reasons and reasoning: a guide through Lewis Carroll’s Regress Argument. In D. Star (Ed.), *The Oxford handbook of reasons and normativity*. Oxford: OUP.
- Bird, A. (2005). Abductive knowledge and holmesian inference. In J. Hawthorne & T. Gendler (Eds.), *Oxford studies in epistemology* (Vol. 1, pp. 1–31). Oxford: OUP.
- Boghossian, P. (2003). Blind reasoning. *Supplement to Proceedings of the Aristotelian Society*, 77, 225–248.
- Boghossian, P. (2006). *Fear of knowledge*. Oxford: OUP.

<sup>45</sup> It might be suggested that the question of completeness disappears on a certain view of goodness: the view that the inclusion of a premise to the effect that an argument (or rule) is good is enough for it to be good. If that view is correct, perhaps it doesn’t matter how the argument is specified—as long as the crucial bridging premise is in? (The view isn’t automatically too permissive: antecedently justified belief in the bridging premise might be hard to achieve.) As mentioned in Sect. 1.2, some endorse a meta-requirement on inferential justification—popularly: that the agent be suitably related to (something like) the claim that the argument/rule involved is good. The nature of the requisite relation is hotly debated (e.g. Boghossian 2003, 2014; Brewer 1999; Hlobil 2019; Tucker 2012; Valaris 2014; Wright 2004, 2011). But a straightforward thought—which brings us close to the current suggestion—is that it’s the very same relation that the availability-constraint stipulates: antecedently justified belief. On this reading, however, the meta-requirement faces severe familiar problems—notably, the threat of a vicious regress (Carroll 1895; cf. Besson 2018; Boghossian 2003; Rosa 2019). The present suggestion was just that (the inclusion of) the premise that an argument is good suffices for argument goodness—but, to get a general criterion for relevant completeness out of this, we’d better say that it’s necessary (too). And once we do that, the suggestion inherits all the familiar problems facing versions of the meta-requirement that incorporate the straightforward thought.

- Boghossian, P. (2008). Epistemic rules. *The Journal of Philosophy*, 105, 475–500.
- Boghossian, P. (2014). What is inference?'. *Philosophical Studies*, 169, 1–18.
- Bonjour, L. (1978). Can empirical knowledge have a foundation? *American Philosophical Quarterly*, 15(1), 1–14.
- Bonjour, L. (1985). *The structure of empirical knowledge*. Cambridge: HUP.
- Bonjour, L. (2002). Internalism and externalism. In P. K. Moser (Ed.), *The Oxford handbook of epistemology* (pp. 234–264). Oxford: OUP.
- Bonjour, L., & Sosa, E. (2003). *Epistemic justification: Internalism vs. externalism, foundations vs. virtues*. Malden, MA: Wiley-Blackwell.
- Brandom, R. (2000). *Articulating reasons*. Cambridge: HUP.
- Brewer, B. (1999). Mental causation: Compulsion by reason. *Proceedings of the Aristotelian Society*, 69, 237–253.
- Brigandt, I. (2010). Scientific reasoning is material inference: Combining confirmation, discovery, and explanation. *International Studies in the Philosophy of Science*, 24(1), 31–43.
- Broome, J. (2013). *Rationality through reasoning*. Malden, MA: Wiley-Blackwell.
- Burge, T. (1993). Content preservation. *Philosophical Review*, 102(4), 457–488.
- Burge, T. (1997). Interlocution, perception, and memory. *Philosophical Studies*, 86(1), 21–47.
- Byrne, A. (2005). Introspection. *Philosophical Topics*, 33(1), 79–104.
- Byrne, A. (2011). Knowing what I want. In J. Liu & J. Perry (Eds.), *Consciousness and the self: New essays*. Cambridge: CUP.
- Cargile, J. (2010). Logical form. In T. Smiley, J. Lear, & A. Oliver (Eds.), *The force of argument*. London: Routledge.
- Carnap, R. (1950). *Logical foundations of probability*. Chicago: Routledge & Kegan Paul.
- Carnap, R. (1952). Meaning postulates. *Philosophical Studies*, 3, 65–73.
- Carnap, R. (1968). Inductive logic and inductive intuition. In I. Lakatos (Ed.), *The problem of inductive logic* (pp. 258–314). Amsterdam: North-Holland.
- Carroll, L. (1895). What the tortoise said to achilles. *Mind*, 4(14), 278–280.
- Chomsky, N. (1995). *The minimalist program*. Cambridge MA: MIT Press.
- Christensen, D. (2004). *Putting logic in its place*. Oxford: OUP.
- Christensen, D. (2010). Higher-order evidence. *Philosophy and Phenomenological Research*, 81, 185–215.
- Climenhaga, N. (2017a). Inference to the best explanation made incoherent. *Journal of Philosophy*, 114, 251–273.
- Climenhaga, N. (2017b). How explanation guides confirmation. *Philosophy of Science*, 84, 359–368.
- Coady, C. A. (1973). Testimony and observation. *American Philosophical Quarterly*, 10(2), 149–155.
- Coady, C. A. (1992). *Testimony: A philosophical study*. Oxford: Clarendon Press.
- Cohen, S. (2002). Basic knowledge and the problem of easy knowledge. *Philosophy and Phenomenological Research*, 65, 309–329.
- Corbi, J. E. (2000). The principle of inferential justification, skepticism, and causal beliefs. *Nous*, 34, 377–385.
- Davidson, D. (1984). *Essays on truth and interpretation*. Oxford: OUP.
- Davidson, D. (1986). A coherence theory of truth and knowledge. In E. LePore (Ed.), *Truth and interpretation* (pp. 307–319). Oxford: Blackwell.
- Davies, M. (2004). Epistemic entitlement, warrant transmission, and easy knowledge. *Supplement to Proceedings of the Aristotelian Society*, 78, 213–245.
- DeRose, K. (2005). Direct warrant realism. In A. Dole & A. Chignell (Eds.), *God and the ethics of belief: New essays in the philosophy of religion*. Cambridge: CUP.
- Descartes, R. (1641/2008). *Meditations on first philosophy* (Trans. Moriarty). Oxford: OUP.
- Dogramaci, S. (2013). Intuitions for inferences. *Philosophical Studies*, 165, 371–399.
- Douven, I. (1999). Inference to the best explanation made coherent. *Philosophy of Science*, 66, 424–435.
- Douven, I. (2017). Abduction. In E. N. Zalta (Ed.), *Stanford encyclopedia of philosophy*. Stanford: Stanford University.
- Edgington, D. (1995). On conditionals. *Mind*, 104, 235–329.
- Ewing, A. C. (1934). *Idealism: A critical survey*. London: Methuen.
- Feldman, R. (2003). *Epistemology*. Prentice Hall.
- Foley, R. (1992). The epistemology of belief and the epistemology of degrees of belief. *American Philosophical Quarterly*, 29(2), 111–124.

- Foley, R. (2009). Beliefs, degrees of belief, and the Lockean thesis. In F. Huber & C. Schmidt-Petri (Eds.), *Degrees of belief* (pp. 37–47). Berlin: Springer.
- Fricker, E. (1994). Against gullibility. In A. Chakrabarti & B. K. Matilal (Eds.), *Knowing from words*. New York: Kluwer.
- Fricker, E. (1995). Telling and trusting: Reductionism and anti-reductionism in the epistemology of testimony. *Mind*, 104, 393–411.
- Fumerton, R. A. (1980). Induction and reasoning to the best explanation. *Philosophy of Science*, 47, 589–600.
- Fumerton, R. A. (1992). Skepticism and reasoning to the best explanation. *Philosophical Topics*, 2, 149–169.
- Fumerton, R. A. (1995). *Metaepistemology and skepticism*. Lanham MA: Rowman & Littlefield.
- Fumerton, R. A. (2006). *Epistemology*. Malden, MA: Wiley-Blackwell.
- Gallois, A. (1996). *The world without, the mind within: An essay on first-person authority*. Cambridge: CUP.
- Ginet, C. (2005). Infitism is not the answer to the regress problem. In M. Steup & E. Sosa (Eds.), *Contemporary debates in epistemology* (pp. 140–149). Malden: Blackwell.
- Goldman, A. (2008). Immediate justification and process reliabilism. In Q. Smith (Ed.), *Epistemology: New essays*. Oxford: OUP.
- Goldman, A. (2009). Internalism, externalism, and the architecture of justification. *Journal of Philosophy*, 106, 309–338.
- Goodman, N. (1955). *Fact, fiction, and forecast*. Cambridge MA: HUP.
- Graham, P. (2006). Testimonial justification: Inferential or non-inferential? *Philosophical Quarterly*, 56(222), 84–95.
- Greco, J. (1999). Agent reliabilism. *Philosophical Perspectives*, 13, 273–296.
- Grice, P. (1989). *Studies in the ways of words*. Cambridge MA: HUP.
- Harman, G. (1965). Inference to the best explanation. *Philosophical Review*, 74, 88–95.
- Harman, G. (1984). Logic and reasoning. *Synthese*, 60, 107–127.
- Harman, G. (1986). *Change in view: Principles of reasoning*. Cambridge, MA: MIT Press.
- Hempel, C. G. (1945). Studies in the logic of confirmation. *Mind*, 54, 1–26.
- Hlobil, U. (2019). Inferring by attaching force. *Australasian Journal of Philosophy*, 97(4), 701–714.
- Hookway, C. (2000). Skepticism and the principle of inferential justification. *Nous*, 32(1), 344–365.
- Horgan, T. (2017). Troubles for Bayesian formal epistemology. *Res Philosophica*, 94(2), 233–255.
- Horty, J. F. (2012). *Reasons as defaults*. Oxford: OUP.
- Huemer, M. (2002). Fumerton's principle of inferential justification. *Journal of Philosophical Research*, 27, 329–340.
- Hume, D. (1748/1999). In T. Beauchamp (Ed.), *An enquiry concerning human understanding*. Oxford/New York: OUP.
- Iacona, A. (2016). Two notions of logical form. *Journal of Philosophy*, 113(12), 617–643.
- Jackson, B. (2006). Logical form: Classical conception and recent challenges. *Philosophy Compass*, 1, 303–316.
- Jackson, B. (2007). Beyond logical form. *Philosophical Studies*, 132, 347–380.
- Jeffrey, R. C. (1965). *The logic of decision*. Chicago: UCP.
- Kelly, T. (2010). Hume, Norton, and induction without rules. *Philosophy of Science*, 77(5), 754–764.
- Keynes, J. M. (1921). *A treatise on probability*. London: MacMillan & Co.
- Korcz, K. A. (2000). The causal-doxastic theory of the basing relation. *Canadian Journal of Philosophy*, 30, 525–550.
- Kratzer, A. (1986). Conditionals. *Proceedings of the Chicago Linguistic Society*, 22, 1–15.
- Kyburg, H. E. (2008). Inductive logic and inductive reasoning. In J. E. Adler & L. J. Rips (Eds.), *Reasoning*. Cambridge: CUP.
- Lackey, J. (2006). It takes two to tango: Beyond reductionism and non-reductionism in the epistemology of testimony. In J. Lackey & E. Sosa (Eds.), *The epistemology of testimony* (pp. 160–189). Oxford: OUP.
- Lambert, K. (1958). Notes on E! *Philosophical Studies*, 9, 60–63.
- Lambert, K. (1981). On the philosophical foundations of free logic. *Inquiry*, 24, 147–203.
- Leite, A. (2008). Believing one's reasons are good. *Synthese*, 161, 419–441.
- Leite, A. (2011). Immediate warrant, epistemic responsibility, and Moorean dogmatism. In A. Reisner & A. Steglich-Petersen (Eds.), *Reasons for belief* (pp. 158–179). Cambridge: CUP.
- Leonard, H. S. (1956). The logic of existence. *Philosophical Studies*, 7, 49–64.
- Levi, I. (1997). *The covenant of reason: Rationality and the commitments of thought*. Cambridge: CUP.
- Lipton, P. (1993). Is the best good enough? *Proceedings of the Aristotelian Society*, 93, 89–104.

- Lipton, P. (2004). *Inference to the best explanation* (2nd ed.). London: Routledge.
- Lycan, W. (2002). Explanation and epistemology. In P. Moser (Ed.), *The Oxford handbook of epistemology*. Oxford: Oxford University Press.
- Lycan, W. G. (1988). *Judgment and justification*. New York, NY: CUP.
- MacFarlane, J. (2004). In what sense (if any) is logic normative for thought? ms.
- MacFarlane, J. (2015). Logical constants. In E. N. Zalta (Ed.), *Stanford encyclopedia of philosophy*. Stanford: Stanford University.
- Malmgren, A. S. (ms). *Understanding inference*.
- Malmgren, A. S. (2006). Is there a priori knowledge by testimony? *Philosophical Review*, 115(2), 199–241.
- Malmgren, A. S. (2018). Varieties of Inference. *Philosophical Issues*, 28(1), 221–254.
- Markie, P. (2005). The mystery of direct perceptual justification. *Philosophical Studies*, 126, 347–373.
- McDowell, J. (1994). The content of perceptual experience. *Philosophical Quarterly*, 44(175), 190–205.
- McDowell, J. (1996). *Mind and world*. Cambridge, MA: HUP.
- McGrath, M. (2017). Knowing what things look like. *Philosophical Review*, 126, 1–41.
- McHugh, C., & Way, J. (2018). What is good reasoning? *Philosophy and Phenomenological Research*, 96(1), 153–174.
- Montague, R. (1973). The proper treatment of quantification in ordinary English. In P. Suppes, J. M. E. Moravcsik, & K. J. J. Hintikka (Eds.), *Approaches to natural language* (pp. 221–242). Dordrecht: Reidel.
- Neta, R. (2013a). Easy knowledge, transmission failure, and empiricism. *Oxford Studies in Epistemology*, 4, 166–184.
- Neta, R. (2013b). What is an inference? *Philosophical Issues*, 23, 388–407.
- Nolt, J. (2014). Free logic. In E. N. Zalta (Ed.), *Stanford encyclopedia of philosophy*. Stanford: Stanford University.
- Norton, J. D. (2003). A material theory of induction. *Philosophy of Science*, 70, 647–670.
- Norton, J. D. (2014). A material dissolution of the problem of induction. *Synthese*, 191, 671–690.
- Norton, J. D. (forthcoming). *The material theory of induction*, ms.
- Pelletier, F. J., & Elio, R. (2005). The case for psychologism in default and inheritance reasoning. *Synthese*, 146, 7–35.
- Peregrin, J., & Svoboda, V. (2013). Criteria for logical formalization. *Synthese*, 190, 2897–2924.
- Pierce, C. S. (1935). In C. Hartshorne & P. Weiss (Eds.), *Collected papers of Charles Sanders Peirce* (Vol. VI). Cambridge MA: Harvard University Press.
- Pollock, J. L. (1970). The structure of epistemic justification. *American Philosophical Quarterly*, 4, 62–78.
- Pollock, J. L. (1987). Defeasible reasoning. *Cognitive Science*, 11, 481–518.
- Pollock, J. L., & Cruz, J. (1999). *Contemporary theories of knowledge* (2nd ed.). Lanham: Rowman & Littlefield.
- Popper, K. (1959). *The logic of scientific discovery*. London: Hutchinson & Co.
- Poston, T. (2014). *Reason and explanation*. Basingstoke: Palgrave Macmillan.
- Priest, G. (1979). The logic of paradox. *Journal of Philosophical Logic*, 8, 219–224.
- Priest, G. (2008). *An introduction to non-classical logic: from Ifs to Is* (2nd ed.). Cambridge: CUP.
- Pryor, J. (2005). There is immediate justification. In M. Steup & E. Sosa (Eds.), *Contemporary debates in epistemology* (pp. 181–202). Cambridge MA: Blackwell.
- Pryor, J. (2012). When warrant transmits. In A. Coliva (Ed.), *Mind, meaning, and knowledge: Themes from the philosophy of Crispin Wright* (pp. 269–303). Oxford: OUP.
- Pryor, J. (2018). The merits of incoherence. *Analytic Philosophy*, 59, 112–141.
- Psillos, S. (2004). Inference to the best explanation and Bayesianism. In F. Stadler (Ed.), *Induction and deduction in the sciences* (pp. 83–91). Dordrecht: Kluwer.
- Psillos, S. (2007). The fine structure of inference to the best explanation. *Philosophy and Phenomenological Research*, 74, 441–448.
- Quine, W. V. (1969). *Ontological relativity and other essays*. New York: CUP.
- Quine, W. V. (1977). Intensions revisited. *Midwest Studies in Philosophy*, 2, 5–11.
- Reid, T. (1764/1997). In Brookes (Ed.), *An inquiry into the human mind on the principles of common sense*. University Park: Pennsylvania State University Press.
- Reiter, R. (1980). A logic for default reasoning. *Artificial Intelligence*, 13, 81–137.
- Rhoda, A. R. (2008). Fumerton's principle of inferential justification, skepticism, and the nature of inference. *Journal of Philosophical Research*, 33, 215–234.
- Richard, F. (2003). *Epistemology*. Saddle River: Prentice Hall.

- Rosa, L. (2019). Reasoning without regress. *Synthese*, 196, 2263–2278.
- Russell, B. (1914). *Our knowledge of the external world*. Chicago & London: Open Court.
- Russell, G. (2008). One true logic? *Journal of Philosophical Logic*, 37(6), 593–611.
- Russell, G. (2018). Logical nihilism: Could there be no logic? *Philosophical Issues*, 28, 308–324.
- Schechter, J. (2011). Rational self-doubt and the failure of closure. *Philosophical Studies*, 163, 429–452.
- Schechter, J. (2017). Internalism and entitlement to rules and methods. In N. Jang-Pedersen & P. Graham (Eds.), *Epistemic internalism*. Oxford: OUP.
- Schiffer, S. (2001). Communication. In N. J. Smelser & P. B. Baltes (Eds.), *The international encyclopedia of the social and behavioral sciences*. Amsterdam: Elsevier.
- Schroeder, M. (2011). What does it take to 'have' a reason? In A. Resiner & A. Steglich-Peterson (Eds.), *Reasons for belief*. Cambridge: CUP.
- Schupbach, J. N. (2017). Inference to the best explanation cleaned up and made respectable. In T. Poston & K. McCain (Eds.), *Best explanations: New essays on inference to the best explanation*. Oxford: OUP.
- Sellars, W. (1953). Inference and meaning. *Mind*, 62, 313–338.
- Senor, T. D. (1993). Internalistic foundationalism and the justification of memory belief. *Synthese*, 94(3), 453–476.
- Siegel, S. (2010). *The contents of visual experience*. New York: OUP.
- Silins, N. (2005). Transmission failure failure. *Philosophical Studies*, 126, 71–102.
- Swain, M. (1979). Justification and the basis of belief. In G. S. Pappas (Ed.), *Justification and knowledge* (pp. 25–49). Dordrecht: D. Reidel.
- Thagard, P. (2000). Coherence in thought and action. *Philosophical Books*, 43(2), 123–140.
- Titelbaum, M. (2010). Not enough there there: Evidence, reasons, and language independence. *Philosophical Perspectives*, 24, 477–528.
- Tolliver, J. (1982). Basing beliefs on reasons. *Grazer Philosophische Studien*, 15, 149–161.
- Tucker, C. (2010). When transmission fails. *Philosophical Review*, 119, 497–529.
- Tucker, C. (2012). Movin' on up: Higher-level requirements and inferential justification. *Philosophical Studies*, 157, 323–340.
- Valaris, M. (2014). Reasoning and regress. *Mind*, 123(489), 101–127.
- van Fraassen, B. C. (1989). *Laws and symmetry*. Oxford: OUP.
- Walton, D. (2013). *Methods of argumentation*. Cambridge: CUP.
- Walton, D., Reed, C., & Macagno, F. (2008). *Argumentation schemes*. Cambridge: CUP.
- Wedgwood, R. (2012). Justified inference. *Synthese*, 189, 1–23.
- Weintraub, R. (2013). Induction and inference to the best explanation. *Philosophical Studies*, 166, 203–216.
- Weisberg, J. (2009). Locating IBE in the Bayesian framework. *Synthese*, 167, 125–143.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: OUP.
- Wright, C. (2002). Anti-skeptics simple and subtle: G.E. Moore and John McDowell. *Philosophy and Phenomenological Research*, 65, 330–348.
- Wright, C. (2004). Warrant for nothing (and foundations for free)? *Supplement to Aristotelian Society Proceedings*, 78, 167–212.
- Wright, C. (2014). Comment on Paul Boghossian, 'What Is Inference?'. *Philosophical Studies*, 169(1), 27–37.
- Zalabardo, J. L. (2011). Boghossian on inferential knowledge. *Analytic Philosophy*, 52(2), 124–139.