

Justified inference

Ralph Wedgwood

Synthese

An International Journal for
Epistemology, Methodology and
Philosophy of Science


ISSN 0039-7857

Volume 189

Number 2

Synthese (2012) 189:273-295

DOI 10.1007/s11229-011-0012-8



Your article is protected by copyright and all rights are held exclusively by Springer Science+Business Media B.V.. This e-offprint is for personal use only and shall not be self-archived in electronic repositories. If you wish to self-archive your work, please use the accepted author's version for posting to your own website or your institution's repository. You may further deposit the accepted author's version on a funder's repository at a funder's request, provided it is not made publicly available until 12 months after publication.

Justified inference

Ralph Wedgwood

Received: 27 April 2011 / Accepted: 8 September 2011 / Published online: 27 September 2011
© Springer Science+Business Media B.V. 2011

Abstract This paper proposes a general account of the epistemological significance of inference. This account rests on the assumption that the concept of a “justified” belief or inference is a normative concept. It also rests on a conception of belief that distinguishes both (a) between conditional and unconditional beliefs and (b) between enduring belief states and mental events of forming or reaffirming a belief, and interprets all of these different kinds of belief as coming in degrees. Conceptions of “rational coherence” and “competent inference” are then formulated, in terms of the undefeated instances of certain rules of inference. It is proposed that (non-accidental) rational coherence is a necessary and sufficient condition of justified enduring belief states, while competent inference always results in a justified mental event of some kind. This proposal turns out to tell against the view that there are any non-trivial cases of “warrant transmission failure”. Finally, it is explained how these proposals can answer the objections that philosophers have raised against the idea that justified belief is “closed” under competent inference.

Keywords Justification · Inference · Rationality · Normativity · Internalism · Belief · Coherence · Coherentism · Foundationalism · Warrant transmission

In this essay, I shall propose a general conception of the kind of inference that counts as justified or rational. This conception involves a version of the idea that justified belief is “closed” under competent inference—that is, roughly, whenever one competently infers a conclusion from a set of premises each of which one justifiably believes, this inference always puts one in a position to have a justified belief in the conclusion.

R. Wedgwood (✉)
School of Philosophy, University of Southern California, 3709 Trousdale Parkway,
Los Angeles, CA 90089-0451, USA
e-mail: wedgwood@usc.edu

Although many philosophers have advocated versions of this idea of “closure”, the idea has also been widely criticized. The conception proposed here is designed to circumvent these criticisms.

The main focus in this essay will be on articulating this conception of justified inference in detail, and on defending it against objections. I shall not be able in the space available to develop any positive arguments in favour of this conception—let alone to attempt any deep explanatory account of why this conception is true.¹ Nonetheless, in defending this conception against objections, I shall show that it can provide an intuitively acceptable account of a wide range of cases; and this will contribute at least to some extent towards the confirmation of this conception.

1 Justification as a normative concept

I shall rely here on some assumptions about the concept of “justification”—in particular, on the assumption that it is a *normative* concept. To say that ‘justified’ is a normative concept is to say that it is closely related to the concepts that can be expressed by terms like ‘ought’ and the like.

The very word ‘justified’ seems to suggest that the concept that it expresses must be a normative concept, since the word is etymologically related to the Latin words ‘*iustus*’ (‘just’ or ‘righteous’) and ‘*ius*’ (‘law’ or ‘right’). Moreover, it is possible to specify more exactly how the concept of a “justified belief” is related to some of the concepts that can be expressed by paradigmatically normative uses of ‘ought’ and the like.

I shall not distinguish here between the concepts that I am expressing by the terms ‘justified belief’ and ‘rational belief’; and the term ‘rational belief’, as I am using it, expresses the concept of a *rationaly permissible* belief.² So the concept of a “justified belief” that I am using here is the concept of a belief that is in a certain way *permissible*.

Some readers may be surprised by the suggestion that to say that a belief is justified is to say no more than that it is permissible. But there will be many cases in which one simply cannot avoid having some broadly doxastic attitude towards a proposition *p*, while at the same time, the *only* justified or permissible doxastic attitude that one can have towards *p* is to believe *p*. In such cases, it is true, not only that it is justified or permissible for one to believe *p*; it is also true that in a certain sense one *ought* to believe *p*.

This notion of the doxastic attitude that one in this sense *ought* to have at a given time *t* seems to be the same as the notion of the attitude that one is *rationaly required* to have at *t*. Since it is a kind of ‘ought’, the notion of what one is rationally required to think at *t* seems to be a paradigmatically normative concept.

¹ I have started to grapple with these other questions elsewhere: see especially Wedgwood (2011); see also note 10 below.

² Moreover, the currently much-discussed distinction between *propositional* justification and *doxastic* justification (see e.g. Turri 2010) is precisely parallel to the distinction between attitudes that it is “*ex ante* rational” for the relevant thinker to have and those that are actually held in an “*ex post* rational” manner (for the latter distinction, see e.g. Wedgwood 2002a).

To fix ideas, I shall be assuming a particular conception of these normative concepts here—specifically, a “possibilist” and “optimizing” conception of these concepts.³ Thus, you are rationally required to φ at t if and only if you φ in every possible world w such that (i) w is in the relevant way *available* to you at t , and (ii) there is no available possible world $w+$ such that at t you think *more rationally* in $w+$ than in w .

Thus, according to this conception of these concepts, what is rationally required of you at t is determined by: (i) the worlds that are in the relevant way “available” to you at t ; and (ii) the facts about which ways of thinking count as “more rational” than others. In effect, something is rationally required of you if it is a necessary consequence of your thinking in one of the available ways that count as maximally rational.

To say that a world w is “available” to you at a time t is to say something about the opportunities and cognitive capacities that you have at t : in effect, it is to say that it is *possible* for you to realize that world w through the way in which you exercise your capacities at t . But it need not imply that it will be *easy* for you to realize this world w , or that there is a significant *chance* that you will do so. Even if it would be tremendously difficult for you to φ , it does not follow that worlds in which you φ are not available to you.

Moreover, as we shall see towards the end of this paper, it seems that there is not in fact a single univocal concept of “availability” here, but a spectrum of such concepts, ranging from a rather weak concept (on which a great many worlds count as available) to a much stronger concept (on which many fewer worlds count as available). It may be that the context in which we use terms like ‘ought’ or ‘rationally required’ determines which concept of “availability” is relevant in that context. The relatively weak concepts of availability correspond to relatively *idealized* notions of what is “rationally required”—whereas the stronger concepts of availability correspond to much more “realistic” (or less “demanding”) notions of what is rationally required.

For most of the remainder of this essay, I shall simplify matters by using an extremely weak concept of “availability”. Specifically, I shall assume that in this sense, there is always a possible world “available” to every thinker in which that thinker thinks in a *perfectly rational* way. The notion of a rational requirement corresponding to this weak concept of “availability” is an extremely idealized notion. In this sense, you are “rationally required” to φ at t if and only if you φ in every available possible world in which you think in a *perfectly rational* way at t . So, with this assumption in place, we can give a characterization of justified inference (in this idealized sense) by answering the question: *What is it to think in a perfectly rational way?*

There is one further assumption about justified belief that I shall rely on here. Specifically, I shall rely on the assumption that an *internalist* view of justified or rational thinking is correct. In effect, this is the view that the way in which you are rationally required to think at t supervenes on the internal mental states—that is, the

³ This “possibilist” view of the concepts expressed by ‘ought’ is opposed by Jackson (1985). I have defended this possibilist and optimizing view of the logic of normative concepts elsewhere; see especially Wedgwood (2007, Chaps. 4–5).

“non-factive” mental states⁴—that you have at t and the internal mental events that are happening to you at or shortly before t .⁵

In short, according to the assumptions that I shall rely on here, the notion of a justified or rational belief is a normative notion—but it also has a special “internalist” character, so that all the facts about the beliefs or doxastic states that it is justified or rational for a thinker to have supervene on facts about the thinker’s internal mental events and states.

2 The nature of belief

To give a good account of justified inference, we shall have to work with a conception of the nature of belief. It is important to note that there are many different sorts of belief.

For example, one basic distinction that needs to be drawn here is between (i) *enduring* mental states (such as the background beliefs that one holds for many years, such as my belief that Dublin is the capital of Ireland), and (ii) mental *events*, which occur at a particular point in time (such as *judgments*—mental events in which the thinker consciously *forms* or *reaffirms* a particular belief).

In general, for every type of enduring belief-state, there is a corresponding type of mental event, in which the thinker consciously forms or reaffirms that enduring belief-state. Of course, it sometimes happens that when one forms an enduring belief-state, this belief-state does not in fact endure for more than a moment, but is forgotten almost immediately afterwards. Still, the nature of the mental event in question is to be the event of forming that enduring belief-state.

A second crucial feature of belief is that beliefs come in *degrees*: thinkers believe some propositions with greater confidence than others. For example, you might believe that $1 + 1 = 2$ more strongly, and with greater confidence, than that Dushanbe is the capital of Tajikistan. In effect, then, your belief-system creates an (at least partial) *ranking* of propositions, with respect to how much confidence you have in them. One way to visualize this ranking of propositions is as a partially ordered collection of boxes. Your belief-system effectively sorts propositions into these boxes: the proposition that $1 + 1 = 2$ goes into the box at the very top, the proposition that Dushanbe is the capital of Tajikistan goes into a box that comes somewhat lower in the ranking, while the proposition that $1 + 1 = 5$ goes into the box at the very bottom, along with the other propositions in which you have the strongest possible degree of disbelief.

It may be that strictly speaking, a thinker’s belief system involves more than one ranking of propositions—where the thinker ranks the propositions in different ways for different purposes. For example, perhaps for certain practical purposes the thinker treats it as effectively as certain as anything can be that she will still be alive in a week’s time—while for more theoretical purposes, she might treat it as somewhat less certain that she will still be alive in a week’s time than that $1 + 1 = 2$. I shall ignore

⁴ For the notion of a “factive” mental state, see especially Williamson (2000, Chap. 1).

⁵ I have defended this sort of “internalism” elsewhere; see Wedgwood (2002a) and (2006).

this complication here, and suppose that the thinker's belief-system involves only one ranking of propositions that we need to be concerned with.

Your confidence ranking of these propositions can be in a sense *modelled* by means of a set of real-valued credence functions. For a set of credence functions to model this confidence ranking, every credence function Cr in this set must meet the following four conditions:

- a. If you have *more* confidence in p than in q , then $Cr(p) > Cr(q)$
- b. If you have *equal* confidence in p and in q , then $Cr(p) = Cr(q)$
- c. If you have the *lowest possible* confidence in p , then $Cr(p) = 0$
- d. If you have the *highest possible* confidence in p , then $Cr(p) = 1$

It may be that your belief-system is in a sense *indeterminate*: there may be a pair of propositions p and q such that although p and q both appear in your belief-system, your belief-system does not rank p and q in relation to each other. In this case, there clearly cannot be a *unique* credence function that models your belief-system. Instead, your belief-system can only be modelled by a *set* of credence functions, including some credence functions that rank p above q , and also some that rank q above p . Each of these credence functions contains much more detailed structure than your belief-system itself. It is only the whole set—which captures the structure that is *common* to all these credence functions—that models the belief-system itself.

Moreover, there may also be many propositions towards which the believer has *no doxastic attitudes at all*. As Harman (1986, p. 15) puts it, it seems positively desirable to avoid “cluttering” one's mind with unnecessary beliefs. So there seems to be nothing irrational in the slightest about having a belief-system that has “gaps” in this way—that is, being such that for many propositions, one's belief-system includes absolutely no attitude towards those propositions at all.

On the face of it, to be completely “attitudeless” towards a proposition p is not the same as having a maximally indeterminate attitude towards p .⁶ So the property of being completely attitudeless towards p cannot be modelled simply by means of an infinite set of credence functions that for every single real number n , includes a credence function Cr such that $Cr(p) = n$. This infinite set of credence functions would model the state of having a maximally indeterminate attitude towards p , not the property of being completely attitudeless towards p .

To model such “gaps” in the believer's belief-system, we should not require that the credence functions in the set that models the thinker's belief-system must be *complete*, in the sense of assigning a real-valued credence to every proposition in a whole propositional algebra. Instead, the credence functions in this set may be *partial* functions—that is, there may be many propositions in the algebra for which these functions are undefined, and do not assign any credence at all.

Since these credence functions will often be incomplete in this way, they need not be probability functions, since probability functions are never incomplete in this way. Moreover, if the believer is not perfectly rational, the believer's belief-system may be *probabilistically incoherent*. For example, perhaps there are two logically equivalent propositions p and q , such that the believer is more confident of p than of q ; then the

⁶ For a compelling argument for this distinction, see Friedman (2011).

credence functions in the set that models this believer's belief-system will themselves be probabilistically incoherent. These incoherent credence functions could not even be *extended* into a probability function.

There is a further complication that we need to introduce into our conception of belief. So far, every kind of belief that we have considered is an attitude towards a *single* proposition. But in addition to these simple unconditional beliefs, there are also *conditional* beliefs. A conditional belief of this sort is not an attitude towards a *single* proposition, but towards a *pair* of propositions. It involves *conditionally* believing the second proposition, conditionally on the *assumption* (or *supposition*) of the first proposition.⁷

Moreover, we can generalize this idea of a conditional belief, so that it includes mental states that involve conditionally accepting a proposition, conditionally on a *set* of assumptions or propositions. In general, we can say that an *argument* is a structure built up out of propositions. The simplest arguments consist of a *pair*—where the first item in this pair is a set of propositions (the *premises*), and the second item is also a proposition (the *conclusion*). (With some arguments, the set of premises is simply the empty set.) There are also more complex arguments that contain *sub-arguments* as well as premises; but ultimately these more complex arguments are structured out of arguments of this simpler sort.

In the case of the simple arguments that contain no sub-arguments, the state of *accepting an argument* is also not an attitude towards a single proposition, but towards the whole argument—the attitude of *conditionally* believing the conclusion, on the assumption(s) of the premise(s). A simple conditional belief is simply the state of accepting a single-premise argument; but there is no reason why we should not recognize the state of accepting other arguments as well. (If the argument's set of premises is empty, then accepting the argument is effectively just an *unconditional* belief in the argument's conclusion.)

As I understand it, the state of accepting an argument is an *enduring* mental state. But there is also a mental event that corresponds to this enduring state. As I shall use the term, the mental event of *drawing an inference* is the mental event that involves *forming* the enduring mental state of accepting the corresponding argument.

To avoid misunderstanding, I should emphasize that I am using the terms 'accepting an argument' and 'drawing an inference' as technical terms here. In ordinary English, it might seem odd to say that a thinker "accepted an argument" from certain premises to a certain conclusion if the thinker does not in fact regard the premises as supporting the conclusion in any way. But whenever the thinker conditionally believes the conclusion, given the assumption of the premises, it will be true in my sense to say that the thinker "accepts" this argument.

In sum, there are four kinds of broadly doxastic or belief-like phenomena that we need to take account of:

⁷ For an account of this sort of conditional belief, see especially Edgington (1995). Edgington believes that conditional sentences, such as English sentences of the form 'If *p*, then *q*', express such conditional beliefs rather than beliefs in conditional propositions; I am not assuming this view of conditional sentences here—just the view that these conditional beliefs exist.

	Mental events	Enduring states
Attitudes towards <i>single propositions</i>	Events of <i>forming</i> or <i>reaffirming</i> a belief	States of <i>enduring belief</i>
Attitudes towards <i>whole arguments</i>	Events of <i>drawing an inference</i>	States of <i>accepting an argument</i>

According to the conception of belief that I shall be assuming here, *all four* kinds of doxastic attitudes come in *degrees*—degrees that can be modelled by sets of (incomplete) credence functions of the sort that I have described above.

3 Rationally accessible arguments

If every structure that is built up out of propositions in this way counts as an argument, then there will certainly be bad arguments as well as good ones. A rational agent will accept good arguments with greater confidence than bad ones. In this section, I shall make some tentative suggestions about what this difference between good and bad arguments comes to.

Since the term ‘a good argument’ can be used in many ways, I shall use a technical term here: I shall speak of “rationally accessible arguments” instead—intuitively, the arguments that can be exploited in rational or justified inference.

First, it seems clear that if it is rational for a thinker to believe, with a reasonably high degree of confidence, that a certain argument is truth-preserving in the actual world, then the thinker can exploit this argument in rational inferences. However, this point cannot give us a complete account of which arguments are rationally accessible, since we would still have to explain why it is rational to believe these arguments to be truth-preserving in the actual world in the first place.

It is evidently a huge task to give a full account of exactly which arguments are rationally accessible in this way. However, it seems plausible that for each thinker at each time, there are certain rules of *logical* or *deductive* inference which count as *basic*. To put it another way, there are certain rules of deductive inference that it is “primitively rational” for the thinker to follow at this time. Many different theories could be given of what makes a rule of deductive inference basic in this way; elsewhere (Wedgwood 2011), I have suggested that they include all the rules that are built into the thinker’s possession of the relevant concepts. But to fix ideas, let us assume that these basic rules of deductive inference include *modus ponens*, disjunction introduction, existential generalization, and the like, and also all instances of certain *analytically valid* rules of inference—such as the argument from ‘*x* is an uncle’ to ‘*x* is male’, the argument from ‘*x* knows that *p*’ to ‘*p*’, and so on. In general, then, we can say that besides all arguments that the thinker rationally believes (with a reasonably high of confidence) to be truth-preserving in the actual world, these rationally accessible arguments include all instances of these basic rules of deductive inference.

Evidently, however, a complete theory of inference will have to include non-deductive or “ampliative” inferences as well. I tentatively propose that the rationally accessible non-deductive (or ampliative) arguments include all of what we could call the *non-defeated* instances of certain special rules of non-deductive inference. (Since these

arguments are non-deductive or ampliative, they must be instances of rules that are not *analytically valid*, but in some sense *synthetic* instead.⁸)

Plausibly, these special rules will have to include the rule of *taking experience at face value* (a rule that we might perhaps call “external-world introduction”). For example, instances of this rule include arguments from a premise of the form ‘It looks to me as though x is red’ to the corresponding conclusion ‘ x is red’. There may also be similar rules for taking one’s *memories* and one’s *moral intuitions* at face value, and so on.

These special non-deductive rules would presumably also have to include some rule of *induction*. For example, instances of this rule might include arguments 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 ’. These special rules may probably also have to include some rule of *inference to the best explanation*, and so on.

All of these special non-deductive rules of inference have instances that are not truth-preserving even in the actual world; they may even have instances that are known not to be truth-preserving in the actual world. Nonetheless, it may be that it is in a sense the *default* position that instances of these special rules are rationally accessible. That is, the instances of these special rules are normally rationally accessible, unless exceptional *defeating* conditions arise, preventing instances of these rules from being rationally accessible. In general, when an instance of a rule of inference is *defeated*, this is because the believer is somehow committed, by his or her other beliefs, to having only a *low* degree of conditional confidence in the conclusion of that instance, conditional on the relevant premises.⁹

In some way, cases where instances of these rules are defeated in this manner count as *exceptional* rather than as *typical*; when such special defeating conditions are absent, instances of these special ampliative rules of inference seem to be rationally accessible arguments—that is, they are arguments that can be exploited in rational inferences.

4 Rational coherence

These rationally accessible arguments have implications for when beliefs count as justified. In this section, I shall focus on our *enduring* belief-states, starting with enduring states that take a whole *argument* (as opposed to a single proposition) as their object.

I propose that it is a necessary condition on the rationality or justification of your current enduring belief-states that they should meet certain conditions of *coherence*.

⁸ If—as I argue elsewhere (Wedgwood 2011)—all basic rules of inference are a priori, the basic rules of non-deductive inference would have to consist of *synthetic a priori* rules.

⁹ This characterization of defeat would allow for an account of the distinction between undercutting and rebutting defeaters. One has a “rebutting defeater” for an inference when one’s other beliefs commit one, not only to having a low degree of conditional confidence in the conclusion, conditionally on the relevant premises, but also to having a low degree of unconditional confidence in the conclusion as well—whereas an “undercutting defeater” would be a defeater that is not rebutting in this way. Some philosophers would object to this account of the distinction, since it implies that there can only be an undercutting defeater for an inference if the relevant premises are not certainly true. But since it is unclear what theoretical work the distinction between undercutting and rebutting defeaters is supposed to do, it is not obvious that this is a defect in this account of the distinction.

These conditions of coherence correspond, in a systematic way, to these rationally accessible arguments.

First, let us start with the arguments that count as rationally accessible because they are instances of *deductive* or *analytically valid* rules of inference. Suppose that A_1 is an argument of this sort. Then, it seems, the only *coherent* way to accept A_1 must involve having the *maximum* degree of *conditional* confidence in the conclusion of A_1 , conditionally on the assumption(s) of the premise(s).

For example, consider the following analytically valid argument:

Ralph is an uncle
Ralph is male

The only coherent way to accept this argument is by having the *highest possible* degree of conditional confidence in the conclusion that Ralph is male, conditionally on the premise that Ralph is an uncle.

Secondly, consider the arguments that count as rationally accessible because they are non-defeated instances of the special non-deductive rules of inference that I described in the previous section (such as induction and the like). Suppose that the argument A_2 is a rationally accessible argument of this sort. Then, I propose, the *coherent* way to accept A_2 must involve having some *reasonably high* degree of conditional confidence in the conclusion of A_2 , conditionally on the assumption(s) of the premise(s); however, it may still be coherent if one's degree of conditional confidence in the conclusion, on the assumption(s) of the premise(s), is somewhat *less* than the highest possible level.

For example, consider the following inductive argument:

The sun rose in the East on all observed days in the past
The sun will rise in the East tomorrow

The coherent way to accept this argument is by having a *high* but possibly *not maximal* degree of conditional confidence in the conclusion that the sun will rise in the East tomorrow, conditionally on the premise that the sun rose in the East on all observed days in the past. (A full treatment of this topic would have to investigate whether we could give a more precise specification of exactly *how high* one's conditional confidence in the conclusion should be, conditionally on the assumption of the premise; but for present purposes, this relatively imprecise specification will have to suffice.)

If there are these coherence requirements that our *conditional* beliefs—and more generally, our *acceptance of arguments*—must meet, it is plausible that there will also be corresponding coherence requirements on *unconditional* beliefs. The reason for this is that it seems that coherence requires that there should be a certain fundamental connection between one's conditional beliefs and one's unconditional beliefs. Specifically, if the set of credence functions that can model one's unconditional beliefs includes Cr , then one's conditional belief in q given p (at least when one is not absolutely certain that p is false) can be modelled by the quotient $Cr(p \& q)/Cr(p)$.

More generally, each of these credence functions must in effect constitute an assignment of propositions to regions in a kind of *space*—effectively, a “logical space”—such

that one's acceptance of an argument from the premises p_1, \dots, p_n to the conclusion q can be modelled by the *proportion* of the part of this space where the regions assigned to the premises p_1, \dots, p_n all overlap that also overlaps with the region that is assigned to the conclusion q .¹⁰

If coherence does indeed require this connection between one's unconditional beliefs and one's acceptance of arguments, then we can derive many further consequences about what coherence requires of one's unconditional beliefs. For example, if the argument from p to q is a rationally accessible *deductive* argument, and one has some degree of confidence both in p and in q , then coherence requires that one's degree of confidence in q should be *no less* than one's degree of confidence in p . Similarly, if the argument from p to q is a rationally accessible *non-deductive* (ampliative) argument, then coherence requires that one's degree of confidence in q should not be *too much* less than one's degree of confidence in p .

In general, so long as all arguments that are truth-functionally valid in classical logic count as rationally accessible, then any fully coherent belief-system must be *probabilistically coherent*. For one's belief-system to be probabilistically coherent, it must be possible to *extend* each of the credence functions in the set that models this belief-system into a complete *probability* function.

This is not to say that coherence requires no more than probabilistic coherence: on the contrary, I have already implied that it requires many other conditions as well. (For example, if the argument from 'x looks red' to 'x is red' is an undefeated instance of the rule of taking experience at face value, then coherence requires that one should not simultaneously have an extremely high degree of confidence in 'x looks red' and an extremely low degree of confidence in 'x is red'.) There may be many other coherence requirements as well: for example, perhaps another such requirement is the so-called "Principal Principle" (which concerns the conditional degrees of confidence that one should have in certain propositions conditionally on various assumptions about these propositions' chances).¹¹ It will clearly be a challenging task to give a complete account of all these coherence requirements, but perhaps enough has been said to convey an idea of what these requirements of coherence are like.

I am now in a position to make my first main proposal about the epistemological significance of such rationally accessible arguments:

Proposal 1: A thinker's overall set of enduring belief-states is rational or justified *only if* it is (non-accidentally) meets all of these requirements of coherence.

So far, I have articulated this principle only as a *necessary* condition of the rationality of enduring beliefs. It is a controversial issue whether it is also a sufficient condition. Some philosophers—specifically, those who adhere to a more *foundationalist* conception of justified belief—would deny that it is also a sufficient condition.¹²

¹⁰ I suspect that it would be possible to explain why coherence requires accepting rationally accessible arguments in the way that I have outlined, and why it requires this connection between one's acceptance of arguments and one's unconditional beliefs, by developing some of the seminal ideas of Joyce (2009). But unfortunately I cannot take the time to explore this suspicion here.

¹¹ For discussions of the Principal Principle, see Lewis (1986) and Hall (1994).

¹² For an example of such a foundationalist conception of justification, see Audi (2001).

However, there are reasons to be doubtful of any such foundationalist conception. Most of my enduring beliefs—such as my belief that Hume was born in 1711, or that Tashkent is the capital of Uzbekistan—no longer seem to be based on any “foundations”. I no longer hold these beliefs because of the bases on which they were originally acquired (indeed, I have long since forgotten what those original bases were). I continue to hold these beliefs purely because I held them in the past, and they cohere with the rest of my current beliefs.

According to the version of internalism that I am assuming here, one's enduring beliefs count as rational or justified at a given time t in virtue of the facts about the internal mental states and events that are present in one's mind at and shortly before t ; their rationality cannot be due to facts about the remote past—not even to facts about the bases on which those beliefs were originally acquired. Given this assumption, the apparent fact that enduring beliefs of this kind can be rational seems to support a *coherentist* picture, on which (non-accidental) coherence is *necessary and sufficient* for the rationality of such enduring beliefs.

Traditional versions of coherentism about justified belief have often been criticized—especially on the grounds that a thinker's beliefs could be perfectly coherent even if they are not in any way supported by the thinker's sensory experiences or perceptions.¹³ However, traditional versions of coherentism attempted to account *both* for the justification of enduring belief states *and* for the justification of events of belief revision. The coherentist view that I am proposing here is not designed to account for the justification of events of belief revision: it is solely concerned with the justification of enduring belief states. A coherentist view of enduring belief states is quite compatible with a foundationalist view of events of belief revision—that is, with the view that the events in which we form or revise our beliefs are justified in virtue of the conscious mental events (such as the experiences or processes of reasoning) on which they are based.

So this form of coherentism is compatible with the claim that under normal circumstances, rational thinkers must respond to their sensory experiences by forming new beliefs. The normal result of an event of one's forming a new belief is that the new belief becomes one of one's enduring beliefs. So this coherentist view is compatible with the idea that a normal rational thinker's system of enduring beliefs will constantly be inundated with new beliefs as they pour in through sensory experience. So, if one's enduring beliefs are non-accidentally coherent at a given time t , one will normally have had to adjust those beliefs to take account of the new beliefs that flowed into one's belief-system in this way just before t . So the coherentist view that one's enduring beliefs are justified at a given time t if and only if they are non-accidentally coherent at t need not be vulnerable to the objection that coherent beliefs can be wholly cut off from one's sensory experience.

So far, I have only proposed an account of what it is for a thinker's whole system of enduring beliefs to be justified: such a system is justified if and only if it non-accidentally meets all the requirements of coherence. But what if the thinker's belief-system is *not* perfectly coherent? Couldn't *some* of her enduring beliefs nonetheless be entirely

¹³ For a classic statement of this objection, see Plantinga's (1993, p. 82) “case of the epistemically inflexible climber”.

justified? Roughly, it is plausible to suggest that in cases of this sort, a particular belief in a proposition p is justified if and only if none of the optimal ways of revising the belief-system so as to make it perfectly coherent involve radically revising the believer's belief in p . Still, to flesh out this suggestion in more detail, we would have to explain when one way of revising a belief-system to make it coherent counts as in the relevant way “better” than another; and that task would take us too far afield at this point.¹⁴ So for the purposes of this essay, I shall confine myself to the proposal that I have made about the justification of a whole belief-system; I cannot explore the justification of particular enduring beliefs in any further detail.

Still, this conception of optimal ways of revising a belief-system may illuminate other questions as well. In particular, it may help to explain what it is for one to be “committed” to having a certain doxastic attitude towards a proposition p : one is committed to having this attitude if and only if every optimal coherent extension of one's current belief-system that includes any doxastic attitude at all towards p includes this attitude towards p . This understanding of what it is to be “committed” to a certain doxastic attitude may help to explain the conditions under which an instance of a non-deductive rule of inference can be “defeated”, at least if I was right to suggest that such a non-deductive argument will be defeated if and only if the thinker is committed, by his or her other beliefs, to having only a low degree of conditional confidence in the conclusion of that argument, conditionally on the argument's premises.

5 Rational inferences

The proposal that I made at the end of the previous section concerned the conditions under which *enduring belief states* are rational or justified. To give a full account of justified inference, however, we shall also have to consider *mental events*, such as the event of drawing an inference, and the event of forming a new belief as a result of such an inference.

I have assumed that the rational believer's belief-system will contain many “gaps”—that is, there will be many propositions towards which the believer has no doxastic attitude at all. The primary function of these mental events of drawing inferences is to *fill* some of these gaps. That is, through drawing an inference one can come to have a doxastic attitude towards a proposition or an argument that one had never had any such attitude towards before.

In this section, I shall make two proposals about when it is rational to undergo such mental events. First, I shall make a proposal about the mental event of *drawing an inference*—where drawing an inference is a mental event that involves forming (or reaffirming) an attitude towards a whole *argument* (rather than an unconditional belief in a single proposition). Then I shall make a proposal about when it is rational to undergo the event of forming (or reaffirming) an unconditional belief in a *single proposition*.

¹⁴ One *prima facie* attractive way of understanding this notion of better and worse ways of revising a belief-system towards coherence would of course embrace an idea of “conservatism” along the lines suggested by Harman (1986, p. 32).

At this point, we shall have to explore the notion of a *competent inference*. I propose to understand this notion in terms of the manifestation of a suitable sort of *disposition*—namely, a disposition that is specifically a disposition to accept rationally accessible arguments within a certain range. So every disposition of this sort would concern a certain range R of rationally accessible arguments. In effect, it would be a disposition to respond to the event of one's *considering an argument* within that range R , with the mental event of *drawing the corresponding inference*—that is, with the event of forming (or reaffirming) the attitude of coherently accepting that argument. Whenever one manifests such a disposition, one's drawing the inference in question can be called a case of *competent inference*.

On this conception, a competent inference consists in the manifestation of a disposition that can *only* be manifested in the coherent acceptance of a rationally accessible argument. The acceptance of an argument that is not rationally accessible (say, because it is a *defeated* instance of a non-deductive rule of inference) cannot count as the manifestation of this disposition. In this way, the disposition that one manifests in making a competent inference is a necessarily rational disposition.

At the same time, I have not specified exactly what the “range” of rationally accessible arguments must be. Exactly which dispositions the individual thinker actually has, and is manifesting on a particular occasion, must surely be an empirical question, to be studied by the cognitive psychologists who investigate the individual's reasoning competence. For example, some reasoners might have a disposition that is specifically keyed to instances of *modus ponens* as such, while other reasoners might have a disposition that responds to a different category of rationally accessible arguments—including both some (but perhaps not all) instances of *modus ponens* and also some other rationally accessible arguments as well. These differences between different reasoners would not matter for my purposes. So long as they have some dispositions that can be manifested *only* in the coherent acceptance of rationally accessible arguments, these reasoners will be capable of competent inference.

My proposal about when mental events of drawing inferences are rational can be stated using this notion of competent inferences:

Proposal 2: Such competent inferences are, without exception, always rational.

This proposal should be especially plausible within the context of a coherentist conception of rational belief, of the sort that was sketched at the end of the previous section. This is because every such competent inference is an event that involves non-accidentally forming (or reaffirming) an attitude towards an argument of the sort that is required by coherence. If, as the coherentist conception implies, non-accidental coherence is sufficient for the rationality of enduring belief-states, then manifesting this sort of competence should be sufficient for the rationality of the event of drawing an inference.

This proposal concerns the rationality of *mental events of drawing inferences* (that is, events of forming or reaffirming the attitude of *accepting certain arguments*). But we also need to understand how inferences can be used to form (or to reaffirm) *unconditional* beliefs in particular propositions. This includes cases where we use inference to form new beliefs—beliefs in propositions towards which we previously had no belief at all.

Suppose that you also have the following disposition: in every normal case in which you rationally draw an inference from a set of premises each of which you have some kind of rational unconditional belief in, you respond to your drawing this inference by incorporating the conclusion into your belief-system, in such a way that the new belief-system is just like the old one, except for the fact that (i) the new belief-system incorporates the conclusion into its overall ranking of propositions, and (ii) the new belief-system also meets all conditions of rational coherence.

My final proposal about the epistemological significance of inference is the following:

Proposal 3: The mental event of incorporating a proposition into your belief-system by manifesting this sort of disposition is always a rational mental event.¹⁵

This proposal should also seem particularly plausible from the perspective of the sort of coherentist conception that I have sketched above, since when one adds a proposition to one's belief-system in this way, this is again a way of ensuring that one's belief-system incorporates that proposition at the same time as non-accidentally meeting all conditions of rational coherence.

Although this conception of rational inference is particularly well suited to a coherentist conception of the rationality of enduring belief-states, it is in another way akin to a foundationalist conception of what makes such mental events of inference rational. I shall try to explain this point by considering the idea that there are cases of “warrant transmission failure”—that is, cases where a competent inference fails to “transmit” warrant from the premises to the conclusion of the relevant argument.

6 Warrant transmission failure?

The picture that I have been developing here could be summed up by the slogan: “Coherentism for enduring beliefs, foundationalism for mental events.”

According to this picture, if an enduring belief is rational, this is always purely in virtue of (i) its being held, and (ii) its forming part of a non-accidentally coherent belief-system. Its rationality no longer owes anything to any particular inference by means of which it was originally acquired, or is currently being sustained. On this picture, a rational enduring belief is not sustained by any particular inference: it is simply sustained by its belonging to a belief-system that is non-accidentally coherent.

However, this picture involves a quite different account of the rationality of mental events. On this picture, a mental event is always rational in virtue of the *process* that results in that event, and this process always involves that mental event's occurring

¹⁵ An anonymous referee for this journal wondered whether this proposal might “underdescribe” the situation. If a thinker competently infers a conclusion from certain premises, and also has a rational unconditional credence in each of those premises, would it not be positively *irrational* for the thinker *not* to incorporate the conclusion into his belief-system by manifesting this disposition? Strictly speaking, however, I do not believe that it would be irrational for the thinker suddenly to *forget* all about the inference immediately after performing it, so that the thinker never gets to the point of incorporating the conclusion into his belief-system at all. However, it does seem true that in this case, every alternative way of forming a doxastic attitude towards the conclusion would be irrational: that is, to think in a rational way, the thinker must either somehow remain totally “attitudeless” towards the conclusion, or else manifest this disposition.

in *response to* some other mental event—which could for that reason be called the “*basis*” of the mental event in question.

For example, according to Proposal 2, if the mental event of drawing a rational inference is rational, it is rational in virtue of its being a case of *competent inference*. As I have characterized it, a competent inference is a certain sort of *process* (involving the manifestation of a disposition of the appropriate sort); and this process involves drawing the relevant inference in *response to* the mental event of *considering* the argument in question. So, in this way, it is a process in which the event of drawing the inference is a response to, and in that sense *based on*, the event of considering the relevant argument.

Similarly, consider a rational mental event of using such an inference to form or reaffirm a belief in the conclusion of the relevant argument (and so to “incorporate” that conclusion into one’s belief-system, as I put above). According to Proposal 3, this mental event is rational in virtue of its being based on (i) a rational inference and (ii) rational beliefs in the premises of the argument. In such cases, there is in a sense *always* “warrant transmission”. If a thinker rationally forms or reaffirms a belief in the conclusion of the argument in this way, this mental event derives its rationality at least in part from the rationality of the thinker’s beliefs in the argument’s premises. In that sense, the rational inference “transmits” rationality or “warrant” from the premises to the conclusion.

On the other hand, there is in effect *never* “warrant transmission” for rational enduring beliefs. If rational enduring beliefs are never really based on (or sustained by) an inference anyway, there is no question of their having their rationality “transmitted” to them from the premises of any argument: on the contrary, they derive their rationality from their being held as part of a non-accidentally coherent belief-system.

This picture, then, is opposed to the claim that has been made by Wright (1985) among others that there are non-trivial cases of “warrant transmission failure”. On this picture, there are no non-trivial cases of warrant transmission failure: there is always warrant transmission for the event of forming or reaffirming a belief through inference, and there is never warrant transmission for enduring belief-states.

For example, consider a thinker who runs through the “proof of an external world” that was given by Moore (1939). That is, consider a thinker who draws the inference from the premises ‘This is a hand’ and ‘If this is a hand, an external world exists’ to the conclusion ‘An external world exists’. Drawing this inference is the mental event of forming the enduring state of accepting this argument—which is in turn a matter of *conditionally* believing the conclusion of the argument, conditionally on the assumptions of the argument’s premises. On my picture, the mental event of drawing this inference is rational in virtue of its being a competent inference, responding to the event of considering this argument.

Now suppose that the thinker responds to the event of her having drawn this inference with the event of forming—or more likely, *reaffirming*—a belief in the proposition that an external exists. This particular event of forming or reaffirming the belief that an external world exists in this way on this particular occasion is rational in virtue of its being based on this rational inference, and on rational beliefs in the premises; in that sense this mental event has its rationality in part “transmitted” to it from the rational beliefs in the premises.

However, we should not conclude from this that it is only because of this inference that it is rational for the thinker to hold the *enduring* belief that an external world exists. On the contrary, the enduring state of having a reasonably high degree of confidence that an external world exists is rational because it coheres in the appropriate way with the believer's overall belief-system. So, even Moore's proof of an external world does not seem to be a non-trivial example of "warrant transmission failure".

Someone might object to this conclusion: surely there is a phenomenon of "begging the question"; and how are we to account for begging the question if not in terms of warrant transmission failure? However, some forms of reasoning that might be called "begging the question" are clearly irrational according to the proposals that have been made here: for example, it would clearly be irrational for you to respond to your acceptance of an argument that has a proposition p both as one of its premises and as its conclusion by *raising* your level of confidence in p ! Other phenomena that are sometimes called "begging the question" may not be fundamentally epistemological phenomena, but *dialectical* phenomena instead, concerned with what is a useful contribution to the kind of *debate* that takes place between conscientious debaters, who cannot be assumed to be *perfectly* rational. At all events, it is not clear that we need to appeal to the idea of "warrant transmission failure" to explain the phenomenon of begging the question.

7 Objections and replies

7.1 The lottery paradox

My account involves a cautious kind of "closure" principle: at least, it implies that competent inferences will always put the rational thinker in a position to form new beliefs (or to reaffirm old beliefs) in the logical consequences of propositions that the thinker already believes. This raises the question of how I propose to respond to familiar worries about the *lottery paradox* that was originally raised by Kyburg (1961).

In fact, there is a straightforward answer to this. Since I interpret beliefs as consisting fundamentally in degrees of confidence, my account can endorse the familiar *probabilistic* solution to the lottery paradox. When a rational thinker infers the conjunctive proposition 'Ticket m won't win & ticket n won't win' from the conjuncts 'Ticket m won't win' and 'Ticket n won't win', the thinker can easily have a *lower* degree of confidence in the conclusion than in *either* of the premises (at least so long as the thinker does not have a *maximal* degree of confidence in these premises).

The reason for this is that on the account that I have offered above, a rational thinker's degree of confidence in the conclusion of such an argument is not determined purely by the inference and by the degree of confidence that the thinker has in the premises. It is also determined, in a holistic way, by the thinker's belief-system as a whole. According to Proposal 3, as I formulated it above, the rational way for the thinker to form (or to reaffirm) a belief in the conclusion of this argument involves incorporating this conclusion into his or her belief-system in such a way that the new belief-system is as much as possible like the old one, except that it incorporates a doxastic attitude towards this conclusion, and the new belief-system as a whole meets all constraints of rational coherence.

In the case of the lottery paradox, it is assumed that the thinker has the background belief that the lottery is fair, and that as a result every ticket has an equal chance of winning, and there certainly will be exactly one winning ticket. Given these background beliefs, the only coherent attitude to have towards the conjunction ‘Ticket m won’t win & ticket n won’t win’ is to be *less* confident in this conjunctive proposition than in either of its conjuncts.¹⁶ So the kind of closure principle that I have defended here does not have the implausible result that competent inferences can lead a rational thinker to have a high degree of confidence that no ticket will win the lottery.

7.2 Computational intractability

Many philosophers—most notably Harman (1986, pp. 25–27)—will object to the account that I have given here, along the following lines: “It is asking too much of creatures like us to demand that we should root out all incoherencies in our beliefs. So rationality cannot require that our belief-system should be (non-accidentally) perfectly coherent!”

This objection is certainly right about something. For one’s belief-system to be non-accidentally perfectly coherent, it would have to be sensitive to *every* rationally accessible argument that can be built up out of the set of propositions that one has doxastic attitudes towards. That is, one would have to be sensitive to every rationally accessible argument from *any* subset of this set of propositions to *any* member of that set. With n propositions, this could be up to $n \times 2^n$ arguments.

So it is unlikely that any of us will ever have a (non-accidentally) perfectly coherent belief-system. But it is at least *possible* for a thinker to be perfectly coherent. So it is not obviously wrong to claim that every *perfectly rational* thinker would be perfectly coherent. There is no reason to expect that being perfectly rational has to be easy, only at most that it is *possible*. Even if the notion of what is “rationally required” of a thinker at a given time is a kind of ‘ought’, and ‘ought’ entails some sort of possibility, it may be that this kind of ‘ought’ only entails a relatively weak sort of possibility.

However, in some contexts in which we talk about what is “rationally required” of a given thinker at a time t , it may *not* be presupposed that there is a relevantly available possible world in which that thinker thinks in a perfectly rational way at t . Effectively, this would be a context in which the kind of ‘ought’ that corresponds to this “rational requirement” entails a *stronger* kind of ‘can’. In such contexts, it would not be true to say that the thinker is rationally required to achieve such perfect coherence: it would only be true to say that the thinker must only achieve such coherence as is a necessary part of the *least irrational* way of thinking that is available to the thinker at that time.

In this way, it seems to me, the truth conditions of statements about what is “rationally required” of a thinker at a given time are not completely context-independent: what it takes for such statements to be true depends on which possible worlds are viewed as available in the context. The fundamental context-independent truth here

¹⁶ A similar solution can also be given to the preface paradox: it can be perfectly rational to have a high degree of confidence in the preface proposition ‘Some of the beliefs that I have expressed in this book are false’, so long as one does not have a *maximal* degree of confidence in each of the beliefs in question.

does not concern what is rationally required of the thinker at the relevant time: it concerns which ways of thinking are more irrational and which are less irrational—where the perfectly rational ways of thinking are the ways of thinking that are less irrational than any others. It does not seem implausible to claim that these perfectly rational ways of thinking would require the sort of perfect coherence that I have described above.

7.3 Arguments with implausible conclusions

Another objection that philosophers like Harman (1986, pp. 11–12) might make to the account that I have given here is that it denies the seemingly plausible view that there are many cases where the rational way to respond to drawing an inference is not by coming to believe the conclusion, but by abandoning or revising one's beliefs in some of the premises.

For example, suppose that you suddenly consider a rationally accessible argument that has an implausible conclusion. Couldn't this argument give you reason to *doubt the premises* rather than to *come to believe* the conclusion?

My reply to this objection is similar to my reply to the previous objection. I am trying to characterize what counts as *perfectly rational* thinking. If all your beliefs are *perfectly rational*, then the degrees of confidence that you attached to the premises, prior to your drawing the inference in question, must already be utterly unimpeachable.

However, it now seems much less clear that you could really *learn* anything that casts doubt on those unimpeachable beliefs by means of an *inference from those very beliefs*. The crucial point here is that inference is not a process through which one gains any *new* information from the *external world*. It is simply a process whereby one teases out what is *implicit* in the belief-system that one *already* has.

If one were perfectly rational, one would *already* have perfectly adjusted all of one's beliefs to the information that one received from the external world. So just by teasing out some of what is implicit in one's belief-system, one will surely not discover anything new about the external world. So, it is not clear how a perfectly rational believer could ever be led to revise her degrees of confidence in the premises of an inference just in response to drawing the inference.

I have claimed that if one were perfectly rational, one would already have perfectly adjusted all of one's beliefs to the information that one received from the external world. It is natural to wonder how exactly one could have done that. Presumably, if one is perfectly to adjust one's beliefs to information from the external world in this way, then one's response to that information must have been determined, not just by the intrinsic content of that information, but also, in a holistic fashion, by the whole of one's belief-system as well. This is, as we saw in considering the previous objection, a highly demanding requirement, which it will certainly not be easy for any actual thinker to meet. But it seems not to raise any further considerations over and above those that we have already addressed.

7.4 Reasons to doubt one's own competence

According to the second proposal that I made above, competent inferences are, without exception, always rational. This might seem far too sweeping. Aren't there some

cases where what is in fact a competent inference can fail to be rational? For example, suppose that, although in fact you are inferring perfectly competently, you have highly persuasive (but misleading) evidence that you are in fact reasoning fallaciously. For example, we could imagine a case that has been explored by Christensen (2010), in which you know that you have taken a drug which has an 80% chance of leading you to reason in a fallacious way. Wouldn't a competent inference fail to be rational in this case?¹⁷

This case is not covered by the general conception of inferential *defeat* that I suggested earlier, in Sect. 2 and 3. According to that conception, special cases can arise in which non-deductive (ampliative) arguments are defeated; and when a non-deductive argument is defeated, it ceases to be rationally accessible—that is, it can no longer be exploited in rational inferences. In general, I suggested that such defeat occurs if and only if, exceptionally, the thinker is committed to having only a *low* conditional degree of confidence in the conclusion of the argument, conditionally on the argument's premises; and for a rational thinker to be “committed” to having such an attitude towards this argument is for it to be the case that every optimally coherent extension of the thinker's belief-system includes that attitude towards the argument.

However, in the case that has just been described, it clearly need not be true that the optimally coherent extensions of your belief-system will include this sort of attitude towards this argument. If this case counts as a counterexample to my account at all, then it could surely count as such a counterexample even if the inference in question involved a *deductive* or *analytically valid* argument; but no coherent extension of your belief-system would involve having a low degree of conditional confidence in the conclusion of such a deductive or analytically valid argument, conditionally on the relevant premises. On the contrary, the only fully coherent attitude to take towards such a deductive or analytically valid argument is to have a *maximal* degree of conditional confidence in the conclusion, conditionally on the assumptions of the argument's premises. So, at least according to the conception of inferential defeat that I sketched above, this case does not involve any such inferential defeat.

Some philosophers might think that the sort of “internalism” that I have been assuming here supports the thought that your inference cannot be rational unless it is simultaneously rational for you to *believe* it to be rational. But in fact this thought is not supported by the kind of internalism that is being assumed here. According to my kind of internalism, whether or not a mental process (like drawing an inference) is rational depends on the *real* nature of that mental process—not on what the thinker *believes* (perhaps mistakenly) about that process.

In general, it remains open to the proponent of this internalist conception of rationality to take the following two views. First, competent inference is a *different* type of internal mental process from any process of incompetent inference that could result in the acceptance of invalid arguments. Secondly, competent inference in the presence of misleading higher-order evidence that one is reasoning fallaciously is *not* a fundamentally different process from competent inference in the absence of such misleading higher-order evidence; such higher-order evidence may just play a purely

¹⁷ For some highly illuminating discussion of cases of this sort, see Lasonen-Aarnio (2008) and Joshua Schechter (2011).

epiphenomenal role, without making any genuine difference to the nature of the underlying process of competent inference. If proponents of this internalist conception of rationality take these two views, then they can also claim that the fact that the inference is competent is always enough to make it rational—regardless of the presence or absence of higher-order evidence of this sort.

Clearly, there are *some* cases where evidence about the unreliability of the process that results in a mental event will defeat the rationality of that event. For example, if you had compelling evidence that your sensory experiences were not reliable perceptions of your environment, then it might not be rational for you to form beliefs directly on the basis of your sensory experiences. Similarly, if you had compelling evidence that your *memory* was an unreliable record of the past, it might not be rational for you to form beliefs directly on the basis of your apparent memories. Moreover, it seems, such compelling evidence against the reliability of your memory or your sensory experiences could defeat the rationality of forming such beliefs even if the evidence is in fact misleading, and your experiences and memories are in fact completely reliable.

It might seem that cases where one has compelling but misleading evidence to think that one is not reasoning properly must surely be similar. What principled basis is there for treating these cases differently from the cases where one has compelling but misleading evidence against the reliability of one's memory or one's sensory experience?

However, in the cases where one has evidence that one is not perceiving properly, the way in which the process would be unreliable if this evidence were not misleading is compatible with the *real* nature of the *internal* mental process. According to the version of internalism about perceptual belief that I favour, this internal mental process involves simply having an experience, and forming a belief (with a certain degree of confidence) directly on the basis of that experience. That internal mental process could operate in just the same way regardless of whether or not one is perceiving properly. So, strictly speaking, the defeating evidence concerns the *external* properties of this process—whether the experience derives from the external world in the right way or not—and not the process's internal character.

By contrast, misleading evidence that you are reasoning fallaciously can only lead you to think that your belief-forming process is unreliable by leading you to place at least some level of confidence in what is in fact a *false* hypothesis about the type of *internal* mental process that it is. That is, it would have to lead you to think that this internal mental process was not in fact a process of competent inference, even though in fact it *is* a case of competent inference. But, as I said, according to the version of internalism that I am advocating here, it is the real nature of this internal process—the fact that it is a process of competent inference—and not the higher-order beliefs that the thinker has, or even the beliefs that the thinker is justified in having, about the nature of that process, that is crucial to the rationality of the mental event that results from that process.

So, I propose, the *mental event* of making a competent inference is *necessarily rational*. Moreover, on the coherentist picture, if you are in the *enduring state* of accepting a valid argument precisely *because of* its coherence with all one's other beliefs, this too is necessary rational. It would not matter if you have compelling but

misleading evidence that your beliefs are not coherent; all that matters is whether they really are, as a matter of fact, non-accidentally coherent.

Admittedly, in many of these cases, you might be in the weird situation of rationally drawing an inference even though it is *not* rational for you to *believe* the inference to be rational. This would be like a failure of the “JJ thesis”: the thesis that whenever you justifiedly believe a proposition p , you have justification for believing that you justifiedly believe p . This JJ thesis is an analogue of the notorious “KK thesis”: the thesis that whenever you *know* a proposition p , you are in a position to know that you know p . This KK thesis has been trenchantly criticized—persuasively, as it seems to me—by such philosophers as Williamson (2000). If the KK thesis fails, as Williamson and others have argued, it does not seem strange that the JJ thesis can also fail.¹⁸

Strictly speaking, however, there is a disanalogy between the KK principle and the version of the JJ principle that is under consideration here. This version of the JJ principle concerns the justification or rationality of certain mental *events*, whereas the question of whether or not one *knows* a proposition p seems more closely akin to the question of whether an *enduring* belief-state is rational or justified.

This point helps us to understand how it can be that a mental event is justified even if the thinker is not in a position to have a rational or justified belief in the proposition that that event is justified. Typically, one could only pick out a mental event that occurs in one’s mental life demonstratively, on the basis of introspection or memory. As Williamson and others have convincingly argued, introspection and memory are both fallible methods of coming to know the truth. Since they are both fallible in this way, it seems that they cannot always justify one in having the very highest degree of confidence about exactly what mental events are unfolding in one’s mind.

So suppose that you are perfectly rational, and you competently draw a logically valid inference from p to q (that is, in effect, you form a maximally strong conditional belief in q conditionally on the assumption of p), but at the same time, you have some reason to suspect that you are not reasoning competently.

In this case, you might certainly have all sorts of doubts at the meta-level. For example:

1. While you are making the inference, you might have some doubts about the *psychological* question of whether you really are inferring competently. (Given my interpretation of “competence”, this psychological question is a question about what *dispositions* you are manifesting on that occasion, and introspection is not an infallible guide to the truth about such questions.)
2. While you are making the inference, you might also doubt the general proposition that there are two propositions x and y such that y is a logical consequence of x and you are currently inferring y from x . (Again, the fallibility of introspection seems to allow for the rationality of such doubts.)
3. After you have made the inference, you might doubt whether you should trust your *memory* that you have just inferred q from p (as opposed to having inferred q'

¹⁸ For a more detailed argument against the JJ thesis, see my post “A Refutation of the JJ Principle” on the epistemology weblog Certain Doubts (2 November 2009) <http://el-prod.baylor.edu/certain_doubts/?p=1516>.

from p'). Given the fallibility of memory, it seems that you cannot be absolutely certain that you are not wrong about the content of your past thoughts.

So the kind of case that we are considering may involve many doubts at the higher-order level. Nonetheless, Proposal 2 commits me to denying that any of these higher-order doubts will “trickle down” from the meta-level to the first-order level. If in spite of these higher-order doubts, you really do competently infer q from p , then I say that this inference is perfectly rational.

As we have seen above, the fact that *perfect* rationality would involve thinking in a certain way W does not entail that it is *good advice* to give to an imperfect thinker like you or me that we should think in that way W . If you and I are overwhelmingly likely to think in a less-than-perfectly rational way, it might be *good advice* to urge us to think in a way that guarantees that we will only be *slightly* irrational, rather than *egregiously* irrational. So it may be good advice to give to imperfect thinkers like us to urge us to pay attention to the evidence that suggests that we are not reasoning competently.

Nonetheless, it seems to me, a perfectly rational thinker would continue to draw the inference even if she had (misleading) evidence that she was reasoning incompetently, and even if she entertained serious doubts about whether or not she really was perfectly rational.

In this way, I believe that the conception of justified inference that I have sketched in this essay can be defended against objections. For this reason, it seems to me that there are some grounds for regarding it as a promising account of the epistemology of inference.¹⁹

References

- Audi, R. (2001). *The architecture of reason*. Oxford: Clarendon Press.
- Christensen, D. (2010). Higher-order evidence. *Philosophy and Phenomenological Research*, 81, 185–215.
- Edgington, D. (1995). On conditionals. *Mind*, 104, 235–329.
- Friedman, J. (2011). Suspended judgment. *Philosophical Studies*, Online First (25 June), 1–17. doi:10.1007/s11098-011-9753-y.
- Hall, N. (1994). Correcting the guide to objective chance. *Mind*, 103, 505–517.
- Harman, G. (1986). *Change in view: Principles of reasoning*. Cambridge, MA: MIT Press.
- Jackson, F. (1985). On the semantics and logic of obligation. *Mind*, 94, 177–195.
- Joyce, J. M. (2009). Accuracy and coherence: Prospects for an alethic epistemology of partial belief. In F. Huber & C. Schmidt-Petri (Eds.), *Degrees of belief*, Synthese Library Vol. 342 (pp. 263–297). Berlin: Springer.
- Kyburg, H. (1961). *Probability and the logic of rational belief*. Middletown, CT: Wesleyan University Press.
- Lasonen-Aarnio, M. (2008). Single premise deduction and risk. *Philosophical Studies*, 141, 157–173.
- Lewis, D. K. (1986). A subjectivist's guide to objective chance. Reprinted with postscripts in Lewis, *Philosophical papers* (Vol. II, pp. 83–113). Oxford: Clarendon Press.
- Moore, G. E. (1939). Proof of an external world. *Proceedings of the British Academy*, 25, 273–300.
- Plantinga, A. (1993). *Warrant: The current debate*. Oxford: Clarendon Press.

¹⁹ Earlier versions of this paper were presented at Cornell University and Yale University, and at a conference on the Epistemology of Inference at Brown University. I am grateful to those audiences for their helpful comments. The first draft was written while I held a Research Fellowship from the Leverhulme Trust, to whom I should also record my gratitude.

- Schechter, J. (2011). Rational self-doubt and the failure of closure. *Philosophical Studies*. doi:[10.1007/s11098-011-9823-1](https://doi.org/10.1007/s11098-011-9823-1).
- Turri, J. (2010). On the relationship between propositional and doxastic justification. *Philosophy and Phenomenological Research*, 80, 312–326.
- Wedgwood, R. (2002a). Internalism explained. *Philosophy and Phenomenological Research*, 65, 349–369.
- Wedgwood, R. (2002b). Practical reason and desire. *Australasian Journal of Philosophy*, 80, 345–358.
- Wedgwood, R. (2006). The internal and external components of cognition. In R. Stainton (Ed.), *Contemporary debates in cognitive science* (pp. 307–325). Oxford: Blackwell.
- Wedgwood, R. (2007). *The nature of normativity*. Oxford: Clarendon Press.
- Wedgwood, R. (2011). Primitively rational belief-forming processes. In A. Reisner & A. Steglich-Petersen (Eds.), *Reasons for belief* (pp. 180–200). Cambridge: Cambridge University Press.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: Clarendon Press.
- Wright, C. (1985). Facts and certainty. *Proceedings of the British Academy*, 71, 429–472.