# The Sense of Natural Meaning in Conscious Inference

*Anders Nes*

## Introduction

The conscious lives of people are partly a matter of what thoughts run through their heads. Their stream of consciousness is *inter alia* a stream of thought. A story of what it is like to think, then, belongs to the story of what it is like to be a human person. I here try to contribute something to the telling of that story by concentrating on one aspect of our stream of thought, viz. that we infer things. Inference may be unconscious in various ways, but some of our thinking seems to bear its inferential character on its conscious sleeves. Sometimes, it seems to be part of the conscious character of our thinking that an inference is drawn. I here propose that this inferential conscious character is partly constituted by a sense of meaning; specifically, a sense of what Grice (1957) called 'natural meaning'. In consciously inferring that Q from a certain presumed fact that P, I sense the presumed fact that P as meaning that Q, where 'meaning that' expresses natural meaning. Conscious inference and fluent understanding of speech thus both can be said to involve a sense of meaning, viz. of natural and of linguistic meaning, respectively.

In the next section, I elucidate what I mean by 'conscious inference' by contrasting conscious inferences with some more intuitive, but otherwise quite similar cognitive transitions. I affirm the idea that conscious inference involves some form of 'connecting thought' concerning the relationship of the conclusion one is drawing to premises from which one is drawing it. Such connecting thoughts have been held, on various views, to be thoughts that the premises are good reasons for the conclusion (Thomson, 1965, p. 297; Grice, 2001, p. 41), that they support it (Audi, 1986, p. 241; Boghossian, 2014, p. 3ff.), or that they imply it (Broome, 2013, p. 229). In section 2, I present my alternative proposal that connecting thoughts take the form of a sense of natural meaning. Section 3 compares my proposal with the indicated, alternative views of connecting thoughts.

The scope of my discussion here is limited in several ways. I set aside practical reasoning (reasoning issuing in an intention or action) and reasoning in which we adjust our credence but

1

do not conclude with an all-out belief, judgement or acceptance. I shall mainly focus on non-suppositional reasoning of an everyday sort. Except where otherwise noted, or clear from context, 'consciousness' throughout means phenomenal consciousness, consciousness of the sort that constitutes our overall experience.

## 1. Conscious Inference

Conscious inferences, in the sense relevant here, are inferences that bear their status as inferences on their conscious sleeves. This is not to say that they must be highly articulate, or be made with explicit self-consciousness about drawing an inference. Inferences of the following, everyday sort would typically be conscious in the sense intended. Wondering when Tom will return to the hotel, I learn that Tom went to the museum, and that all who went to the museum will be late; I infer that Tom will be late. Seeing my partner's car parked in the driveway as I approach our house, I infer that she is home. In these cases, it is not obvious that I need to use, in inner speech, any such inferential marker as 'so', 'therefore' or synonyms thereof. It is not clear that I need to have an explicit self-consciousness about drawing an inference, i.e. that I need to have an occurrent thought to the effect that I am inferring that such-and-such. At least, I shall not here insist on these as requirements of conscious inference.[1] What nevertheless makes typical specimen of such everyday inferences conscious, in the relevant sense, emerges when we compare them with cognitive transitions issuing in judgements naturally described as 'intuitions' or 'gut feelings'. In an investigation into rapid, intuitive decision-making, psychologist Klein and colleagues (2010, p. 194) report the following incident:

> [A] firefighter led his men into a burning house, round back to the apparent seat of the fire in the rear of the house, and directed a stream of water on it. The water did not have the expected effect, so he backed off and then hit it again. At the same time, he began to notice that it was getting intensely hot and very quiet. He stated that he had no idea what was going on, but he suddenly ordered his crew to evacuate the house. Within a minute after they evacuated, the floor collapsed. It turned out that the fire had been in the basement.

---

[1] The sort of everyday inferences I have in mind correspond broadly to what Boghossian (2014, p. 2) dubs 'System 1.5 reasoning' (though I do not necessarily endorse his claim that such inferences are a matter of 'voluntary mental action').

The judgement that triggered the order to evacuate here is attributed by the firefighter to a 'sixth sense', indicating that it occurred to him with an intuitive, quasi-perceptual character.

Comparing our everyday inferences with such intuitive judgements facilitates identifying a certain conscious dimension the former have and the latter lack. It helps with this as there are a number of features, relevant to their classification as inferential and to their conscious character, that may be common between the two groups, against the backdrop of which certain differences stand out. Research by Klein et al. (ibid.) among others suggest it is often possible to identify cues to which intuitive judgements are reliably sensitive, such as, in the firefighter's case, the unexpected lack of effect of the water, combined with unexpected increasing heat and quietude.[2] The agent may be conscious of these cues, in the sense that they are (represented in) the contents of conscious mental states (which is not to say the agent is conscious of them *as cues*). Thus the firefighter notices the unexpected behavior of water, heat and sound, although he was not disposed, prior to systematic interrogation and empirical inquiry, to mention these as what may have elicited his intuitive judgement. Both Klein's firefighter and myself, when I judge my partner to be at home on the basis of seeing her car parked outside, may thus be said to be conscious of something on which we rely in reaching our judgement. The cognitive transition into the intuitive judgement need not, moreover, be merely accidentally related to any evidential status enjoyed by the cues relied on. Such a judgemental disposition may develop through a process that involves feedback from the environment, so as to ensure, if the environment is sufficiently stable, that it comes to constitute genuine expertize, where the cues relied on are valid and, in effect, constitute good evidence for the accuracy of the cued judgement (Kahneman & Klein, 2009). Further, the judgement may be arrived at through a computational process that combines, in a coherent way, the cue-representing state with memorial representational states (Simon, 1992). The reason for denying that the firefighter's judgement is a conscious inference, then, is not that he is not conscious of certain apparent facts that he relies on in reaching his judgement or that the transition is no sense inferential but, say, brutely associationist.

How, then, do intuitive judgements such as that of Klein's firefighter differ in their conscious character from typical specimen of everyday inferences of the sort exemplified above? That there is a difference is indicated by the fact that I would not describe, say, my judgement that my partner is at home, inferred from the visible presence of her car outside, as an intuition or

---

[2] For reviews, see Kahneman & Klein (2009), and Kahneman (2011).

gut feeling. Connectedly, I have in this case certain dispositions to justify and to explain my judgement. I will advert to her car being parked outside to account for why one should think, and why I do think, she is at home. Such dispositions typically do not apply to the cues underlying intuitive judgements. Likewise, if someone told me, as I judge that my partner is at home, that her car being parked outside is indicative, or good evidence, or a reason to think that she is at home, or something to that effect, this would not come as news to me. It would seem to be spelling out something I am already, often more implicitly and less articulately, taking to be so in making my judgement that she is at home. This again typically contrasts with cues underlying intuitions. To take a familiar example from the visual domain: to be told, e.g. in an introductory psychology class, that darker shading at the top of a figure (a visual cue for concavity) is good evidence of concavity, while shading at the bottom (a visual cue for convexity) is good evidence of convexity, in ordinary outdoors environments where light fall from above, typically comes as news, as something of which one was not already aware or at least had not yet worked out. The research by Klein, Kahneman, and others into so-called System 1 like processes, suggests that this point often extends to cues underlying intuitive judgements outside the narrowly perceptual domain.

These dispositional differences do not, however, exhaust the characteristic differences in consciousness between our everyday inferences and the cue-based intuition. Suppose Klein's firefighter under prior hypnosis had been told about the identity of the cues, their causal and evidential status, etc., for judgements of the sort that triggered his order to evacuate. Suppose moreover that the hypnotist had ensured this information would be retrieved only upon certain post-hypnotic cues, where these cues take the form precisely of questions raised or testimony put to him about the whys and wherefores of the latter sort of judgement. In this way, the above-mentioned dispositions may be bestowed upon him, but his judgement, as it is made, would intuitively still be a gut feeling.[3] A difference, it seems, between the everyday inferences and this hypnotized firefighter's judgement is that relevant dispositions would not be grounded in a certain characteristic feature of the consciousness involved in his reaching the resulting judgement. This feature we may call an inferential conscious character, the presence of which I will regard as definitive of conscious inferences.

---

[3] This argument parallels one by Strawson (2011, p. 319) against simple a dispositional account of the experience of understanding.

Now, since we are interpreting 'consciousness' as phenomenal consciousness, an objection might be pressed here. What is present in our everyday inferences but absent even in the hypnotized firefighter's judgement, it might be objected, is that a certain cognitive state with content, C, somehow concerning the role or status of the grounds/cues on which one relies, is *access* conscious, in Block's (1995) sense, in that this cognitive state renders C globally available for report, promiscuous inference, etc. Yet this difference in access consciousness (the objection continues) does not guarantee any relevant corresponding difference in phenomenal consciousness; in particular, C need not be the content of any phenomenally conscious state. In reply, it might be argued that either access and phenomenal consciousness are not dissociable in the suggested way, or, if they are, then a more elaborate hypnotism scenario could be constructed in which the dispositions towards report and promiscuous inference definitive of access consciousness are comprehensively reduplicated through hypnotism, but where the so-hypnotized firefighter's judgement remains a mere hunch for him, and as such relevantly differs in its conscious character from our everyday inferences. I will not, however, try to settle this debate here. Settling it would require a paper of its own, and doing so is not needed for most of my purposes here. When I speak of 'consciousness', 'conscious character', etc. below, my favored interpretation of these claims is in terms of phenomenal consciousness. Apart from one or two remarks explicitly about phenomenal consciousness, the rest could be re-interpreted as claims about access consciousness, if that were all one wanted to concede in this domain. Both sides can and should agree that (to put is loosely) our everyday conscious inferences involves some form of awareness that somehow concerns the relationship of the proposition one is drawing as conclusion to at least some of the things from which one is drawing this conclusion.[4] Let us adopt the label 'connecting thought' for this form of awareness, where 'thought' is understood in a broad, catch-all sense, embracing all broadly cognitive states or episodes. A number of writers on conscious inference, or the closely related topic of the basing relation, have held that it involves connecting thoughts, so understood, referred to under such terms as 'connecting beliefs' (Audi, 1986), 'meta beliefs' (Korcz, 2010), and 'linking beliefs' (Broome, 2013). In the next section, I suggest that connecting thoughts take the form of a sense of natural meaning.

---

[4] I say 'at least some' instead of 'all' as I want to allow that non-demonstrative (and perhaps even enthymematic deductive) inferences may rely on background assumptions of various kinds, not all of which need to be occurrently entertained as the inference is drawn.

## 2. The sense of natural meaning in conscious inference

Conscious inference has, I will propose, something in common with fluent understanding of speech. Both can be said to involve a 'sense of meaning'. I label it a '*sense* of meaning' only to suggest that it is an consciousness that is immediate, in that it is not itself reached by deliberation or conscious inference. In fluent comprehension, we have a sense of meaning in that we have a sense of what is said by utterances. This notion of 'what is said' is a notion of what Grice (1957) called 'non-natural meaning'. In conscious inference, we have a sense of meaning in that we have a sense of what certain presumed facts mean, where 'mean' expresses what Grice called 'natural meaning'. I will now briefly explain Grice's distinction. I then make some observations about the phenomenology of fluent comprehension, before outlining the broadly analogous proposal about conscious inference.

### 2.1. Non-natural and natural meaning

Both natural and non-natural meaning can be construed as relations between something, X, and something, Y, which X means.[5] So construed, non-natural meaning is a relation between an utterance, gesture, symbol, or the like, and a certain content it expresses. The content in question need not be a complete proposition; it might be a gappy proposition or maybe just an object meant. Even when what something non-naturally means is stated by a propositional complement, of the form 'that P', it is not essential to the obtaining of the relationship of non-natural meaning that P be true. Thus, as Grice (1957, p. 377) noted

(1) Those three rings of the bell (of the bus) mean that the bus is full.

can be true, in the non-natural sense of meaning, even if the bus is not full. Non-natural meaning is a broad category within which subtypes can be distinguished. For example, both the notion of

---

[5] When we use 'meaning' as a count noun, we typically intend the things meant rather than the relation of meaning (as, we when we use 'uncle' as count noun, we intend the people who are uncles rather than the relation of being one). There is no competition, of course, between the count-noun use of 'meaning' and a use of it for the relation of X's meaning Y.

what is said by an utterance and that of what is implicated by an utterance are notions of non-natural meaning.

Claims of natural meaning, on the other hand, can characteristically be put on the following form:

(N) (The fact) that P means that Q.[6]

In (N), the truth of P is presupposed. A key contrast, moreover, with a statement of non-natural meaning such as (1) is that it is essential, for (N) to be a truth of natural meaning, that Q also be true. It is essential, in other words, to natural meaning that nothing naturally means something that is not the case. A well-worn example of claims of natural meaning is the following case of natural signs:

(2) That smoke is emerging from the town means that something's on fire there.

If smoke is not in fact is emerging from town, (2) fails to be true through an unsatisfied presupposition. If there is smoke but no fire, (2) is false because nothing means, in the natural sense, something that is not the case. If there is both smoke and fire, but no suitable connection between them, say because the smoke was caused by something else, (2) is still false: the smoke emerging from the town did not, in this case, mean that there was a fire there.

As in this familiar sort of example, the required 'suitable connection' will often trace to a causal link, or to P and Q being effects of a common cause, as when a barometer's falling reading means that unsteady weather is approaching.[7] Yet the suitable connection need not be traceable to causal connections. For example, one can say, in a recognizably natural sense of 'meaning':

(3) That the number two is even and a prime number means that there are even primes.

Relatedly, although talk of 'natural meaning' is apt to suggest that it is a relation that obtains in nature, independently of our conventions, claims of natural meaning can be made true by systems

---

[6] The parenthetical 'the fact' is optional.
[7] To avoid prolixity, I here and below occasionally elide tokens of 'the fact that' or 'that' in front of 'P' and 'Q'.

of conventions, rules or norms. It could, for example, truly be said about a member of a certain club

> (4) That N.N. has failed to attend at more than three meetings this year means that he is ineligible to stand for being chair.

This remark could be true simply because the body of regulations of the club are such that the indicated non-attendance implies ineligibility to stand for being chair, where this is so even if, in this case, no causal inferences about N.N.'s suitability may be drawn from his non-attendance (perhaps he has a perfectly good excuse of a sort that happens not to be provided for in the regulations). Although the familiar paradigm examples of truths of natural meaning concern natural signs, then, it must be borne in mind that such truths extend far more widely. Otherwise the suggestion that conscious inference is characterized by a sense of natural meaning would obviously fail to apply to many cases of inference.

The notion of natural meaning, thus briefly elucidated, is widely employed and well entrenched in common sense. I am here taking it at face value, leaving questions of its further character or possible analysis for another occasion. I shall, though, in section 4 below, endorse the assumption that natural meaning is a form of evidence.[8]

## 2.2. The Sense of Meaning in Fluent Comprehension

The phenomenology of language understanding, or of 'comprehension' for short, has often been invoked as an existence proof of cognitive phenomenology.[9] It can also serve as a heuristically useful analogy for the phenomenology of conscious inference. When we fluently understand what is said by everyday utterances in our first language, we have, I suggest, a sense of what is said therein (where this notion of 'what is said' is a one of non-natural meaning). This sense of what is said, in fluent comprehension, seems to have the following five features.

---

[8] See, e.g., Horwich, (1998, pp. 18f.) for this assumption.
[9] See, e.g., Strawson (2011, p. 294).

– *Phenomenality* The sense of what is said partly constitutes the phenomenology of comprehension, being part of what makes it phenomenologically different from incomprehension.

– *Immediacy* Although the sense of what is said may depend on unconscious inference, it seems 'immediate' at least in so far as it does not issue from conscious inference. In this respect, the sense of what is said has a quasi-perceptual character, whether or not it is perception-like in some further or deeper way.[10]

– *Content-Relationality* The sense of what is said seems to have a relational content. When I fluently understand your utterance, I do not merely, on the one hand, have an awareness of your utterance, and, on the other hand, grasp the content that happens to be what is said by your utterance; I am also aware of the latter content *as what is said by* your utterance. To see this, we may compare fluent comprehension with the following type of case. As the thought *Tomorrow's a holiday* occurs to me, a bystander co-incidentally utters words meaning just this in French. Although I know enough French to perceive her utterance as French, and make out its words, I do not understand it. Alternatively, we may consider the following, otherwise similar type of case. I entertain the thought *Tomorrow's a holiday* through fluently reading 'Tomorrow's a holiday' in an English newspaper, as the bystander again speaks as before in French. In either case, the phenomenology seems different from that of fluently comprehending the French utterance as saying *Tomorrow's a holiday*. In the second case, moreover, it is not opaque in my experience whether *Tomorrow's a holiday* is what the English inscription or what the French utterance is saying. This is not opaque in the way it is opaque to, say, Klein's firefighter what creates immediate danger and makes evacuation imperative. To account for how each of these two cases differs from one of fluently understanding the French utterance as saying *Tomorrow's a holiday*, it seems we need to attribute, to the latter understanding, an awareness whose content is relational in that *Tomorrow's a holiday* is given as what the French utterance is saying.

– *Fore-v-Background Structure* Even though the sense of what is said has this fairly rich, relational content, the different elements in this content do not typically seem to

---

[10] See Recanati (2004) for arguments congenial to the claim that our grasp of what is said by an utterance (as opposed to of conversational implicatures beyond what is said) is not achieved through conscious inference.

enjoy the same degree of phenomenal 'focus' or 'emphasis' in our overall experience; a certain structure of relative 'foregrounding' seem to apply across them. Typically, we focus more on the contents expressed than on the utterance expressing them, and more on either of these two than the relation of being-what-is-said-by, or the fact of its obtaining. The latter seems highly recessive in our experience.[11]

– *Doxicity* Although the sense of what is said is perception-like in some respects, it is also belief-like. It has a mind-to-world direction of fit, and involves a commitment on the part of the subject, on pain of her *mis*understanding, that that the relevant content, C, really be what is said by the relevant utterance U. It amounts to her *taking* U to mean C.

In the next section, I shall suggest these five features generalize to the sense of meaning I am attributing to conscious inference.

## 2.3. The Sense of Meaning in Conscious Inference

Conscious inference, I am proposing, also involves a sense of meaning. To put it slightly more precisely:

*The Proposal:* In consciously concluding that Q from premises $P_1$, …, and $P_n$, one has a sense of the presumed facts that $P_1$, …, and that $P_n$ as (together) naturally meaning that Q.

Let us apply this to our two examples of conscious inference in section 2. When I infer, learning that Tom went to the museum, and that all who went to the museum will be late, that Tom will be late, I have a sense of the presumed facts that Tom went to the museum, and that all who went to the museum will be late, as meaning that Tom will be late. When I infer that my partner is at home from her car being parked outside, I have a sense of the presumed fact that her car is outside as meaning that she is at home.

---

[11] Gurwitsch's (1964) theory of the theme v thematic field v margin structure of experience exemplifies one way of elaborating upon these rough and ready claims about experiential focus, emphasis, or foregrounding. Their proper theoretical elaboration must here be left for another occasion.

Some clarifications: First, the presumed facts that correspond to my premises need not be facts; the Proposal allows drawing conclusions from false premises. Second, The Proposal allows suppositional reasoning, for a presumed fact, relative to a thinker, need not be one she believes to obtain; she may imagine, pretend or suppose it to obtain. If she merely supposes that P, her sense (if she has one) of the presumed fact that P as meaning that Q is however embedded under the scope of the supposition. This is akin to how our emotional reactions to fictional portrayals of joys or misfortunes are embedded under the fictional pretense.[12] The so-embedded sense of P as meaning that Q does not, or at least should not, suggest to the thinker that Q actually obtains, only that it obtains in the scenario in which P is supposed to obtain. For reasons of space, I shall however have to set suppositional reasoning aside in what follows.

Third, the so called 'premises' in The Proposal are not necessarily exhaustive of anything that might be regarded a reason, ground or assumption of the conclusion. A conscious inference may trade on background assumptions of various sorts.[13] My inference that my partner is at home may well depend on innumerable such assumptions, e.g. that no one other than my partner would be driving her car, that she rarely leaves home without the car around this time of the day, etc. In so far as I rely on such background assumptions in inferring that she is at home, those assumptions would ipso facto be informing my sense that her car being parked outside (my occurrently entertained premise) means that she is at home. Fourth, in specifying the sense of natural meaning, The Proposal trades on a schema that is a generalization of (N) above, viz.

(N′) (The facts) that $P_1$, …, and that $P_n$ (jointly) mean that Q.

An alternative here would have been to posit a single conjunctive fact and retain the (N) form, i.e. to reformulate The Proposal in terms of '… has a sense of the presumed fact that $P_1$ and … $P_n$ as meaning that Q'. As far as I can see, nothing here hangs on the choice between these alternatives.

The Proposal largely boils down to the conjunction of two claims: that conscious inferences involve connecting thoughts, and that the contents of the latter represent a relation of natural meaning. The relational character of the contents thus ascribed secures that Content-Relationality generalizes, *mutatis mutandis*, to the sense of meaning posited in conscious inference. This is crucial to enable us to account for why it is not opaque to me, as I am inferring

---

[12] See Carruthers (2006) and several of the other contributors to Nichols (2006).
[13] Cf. note 4 above.

that my partner is at home, that this has something to do with her car being parked outside, as it is opaque to Klein's firefighter that the urgent need to evacuate has something to do with the unexpected behavior of heat, water, and fire. However, The Proposal does not, on its own, imply that the posited sense of meaning has any of the other four features attributed above to our sense, in fluent comprehension, of what is said. Conversely, it could be that conscious inferences involve connecting thoughts, and that the five features attributed to the sense of what is said generalize, *mutatis mutandis*, to these connecting thoughts, even if the connecting thoughts, *pace* The Proposal, do not represent a relation of natural meaning. In spite of this logical independence of The Proposal, on the one hand, and the idea that the posited sense of meaning has not only Content-Relationality but also the four other features, on the other, they seem to me to form an attractive package. Given Phenomenality, it enables an account of how a conscious inference is phenomenologically different from a gut feeling, say. Given Immediacy, we inter alia prevent The Proposal from engendering a vicious regress, where any conscious inference, and thus sense of meaning, requires a prior one. We should want Fore-v-Background Structure in so far as our main focus in inference tends to be on the conclusion, with less emphasis being placed on the presumed facts figuring as the premises from which we draw it, and with an even more backgrounded awareness of the relation of these facts meaning that the conclusion is true.[14] Finally, given Doxicity, we capture that the thinker takes the presumed facts that correspond to her premises to mean that her conclusion is true, and is committed, on pain of her making a mistake in her inference, to their indeed meaning that her conclusion is true.

A full-dress defense of this package is beyond the scope of this paper. In what follows, I will concentrate on a salient respect in which this package differs from other characterizations that have been offered of connecting thoughts, by Thomson, Grice, Audi, Boghossian and Broome, namely its reliance on the notion of natural meaning. Since these alternative characterizations agree in regarding connecting thoughts as beliefs or at least belief-like, as per Doxicity, and in assuming connecting thoughts in inference not themselves to issue from conscious inference, as per Immediacy, I will take these features for granted as common ground. Nothing will hang on Phenomenality or Fore-v-Background Structure. I will now note some virtues of invoking natural meaning here, before comparing The Proposal with the alternatives in the next section.

---

[14] On the 'eclipsing' power of the conclusion in inference, see James (1890, pp. 243f.). The typically backgrounded or recessive standing of the relation of natural meaning is, of course, compatible with the possibility of its being called into question and thus becoming foregrounded.

The Proposal enables us to understand why we so commonly express inferences in terms of 'means', where the latter expresses natural meaning. Thus I could put my inference about Tom's return as follows: 'Tom went to the museum. All who went to the museum will be late. This means that Tom will be late'. My inference about my partner I could express so: 'That her car is present there means that she is at home'. If The Proposal is true, this familiar way of making an inference explicit spells out something that already, typically less articulately, is sensed to be so in drawing the inference.

The Proposal also allows us to account for why we commonly assess actual or potential inferences in terms of 'means', where the latter again expresses natural meaning. Consider the following interchange:

(5)  A:  Clinton's visiting Iowa. So she'll run.
     B:  That Clinton's visiting Iowa doesn't mean that she'll run.

It is hard not to hear B's contribution here as an objection. Correlatively, it would be hard to make coherent and literal sense of this concessive reply from A:

(6)  A:  True, that Clinton's visiting Iowa doesn't mean she'll run. Still, she's visiting Iowa. So she'll run.

A's reply here seems absurd in a way akin to that in which such Moorean absurd assertions as 'I believe P, but not P', and 'P, but I do not know that P' seem absurd. We naturally suspect the speaker to be ironic, to abruptly change her mind, to intend to report or give voice to merely sub-conscious attitudes or acts, or perhaps to harbor deviant beliefs about evidence, logic, or psychological states. This is easily explainable on The Proposal, as the inference expressed in A's second and third sentence in (6) involves (given that it is conscious), a commitment that obviously conflicts with the concession in her first sentence.

The Proposal also elucidates the widely felt impropriety of asserting 'lottery propositions' on the basis of probabilistic evidence alone.[15] Suppose ticket T has a 0.99999 chance of being a loser. If this is all you have to go on, many have found an impropriety about the assertion

---

[15] See (Williamson, 2000) for an influential discussion. Note that several who deny Williamson's claim that knowledge is required for proper assertion, such as Douven (2006), agree that lottery propositions are unassertible.

(7) T is a loser.

Given that you know that T has a 0.99999 chance of being a loser, it would be hard to account for this impropriety if the following inference were felt to be okay:

(8) T has a 0.99999 chance of being a loser. It's a loser, then.

The Proposal can account for our reluctance about this inference, as it is matched by our reluctance about the following:

(9) That T has a 0.99999 chance of being a loser means that it's a loser.

At least, it remains odd to suspend judgement on (9) whilst explicitly going in for the inference in (8).


## 3. Natural Meaning vs. Reasons, Support, and Implication

A number of writers, as we have seen, have taken conscious inferences to involve connecting thoughts. On alternative proposals, these are thoughts that the premises are good reasons for the conclusion (Thomson, 1965, p. 297; Grice, 2001, p. 41), that they support it (Audi, 1986, p. 241; Boghossian, 2014, p. 3ff.), or that they imply it (Broome, 2013, p. 229).

 Let us first consider the idea that a 'reason for' relation is represented, i.e. that consciously inferring that Q from that P involves a connecting thought that P is a (good) reason to think that Q. This fits in with our observation, in section 2, that conscious inference characteristically involves a disposition to advert to one's premise by way of accounting for why one should think that Q. A related observation, from Thomson (1965, p. 297), is this:

> There is a perfectly good sense of 'accident' in which if you said 'The philosophy department has a new chairman, so his last name begins with a B or M or L', knowing full well how things stand—that is, knowing that your 'premise' is not in fact a reason for thinking your conclusion

true—you would have to be joking, for it would an accident, a stroke of luck, if you were right. […] No wonder it is right to say: you are not reasoning unless you believe your premise to be a reason for your conclusion.

Thomson here in effect points to another 'Moorean' absurdity in this area, viz. the putative weirdness of consciously inferring that Q from that P whilst conceding that P is no reason to think that Q. If inferences involve connecting beliefs that one's premise is a good reason for one's conclusion, it would explain this. These two arguments, from dispositions to justify and the power to explain a certain absurdity, could *mutatis mutandis* be offered on behalf of the suggestion that connecting thoughts represent a support relation between the presumed fact that P and the conclusion that Q.

The Proposal can explain the Moorean absurdity adduced on behalf of these two alternatives, given a plausible assumption. This is the assumption that there is something similarly weird about denying either that that P is a reason to think that Q or that that P supports that Q, whilst consciously committing oneself to the claim that the presumed fact that P means that Q (a commitment implicit in inferring that Q from that P, on The Proposal). To combine such a denial with the latter commitment does seem odd. At least, there is, on the face of it, something odd about judging either of the following true.

(10) That P means that Q. But that P is no reason to think that Q.

(11) That P means that Q. But that P does not support that Q.

It is a good question why it should be odd, as it seems to be, to judge either of the above true. One plausible line of reply is this. First, natural meaning seems to be a form of evidence, as proposed e.g. by Horwich (1998, 18f.). Second, if P is evidence that Q, P is a reason to think, and supports, that Q. If these implications are not only true but obviously so, we can understand why there is an obvious problem about being committed to P's meaning that Q whilst denying either that P is a reason to think that Q or that P supports that Q. Likewise, if natural meaning obviously is a form of evidence, our proposal can explain why conscious inference of Q from P is

associated with a disposition to advert to that P as something that supports, or is a reason to think, that Q.[16]

However, there are certain absurdities, or anyhow oddities, which our proposal can account for but not the two current alternative specifications of connecting thoughts. They include the oddity of consciously inferring that Q from that P whilst consciously suspending judgement on whether that P means that Q. This oddity may not be quite as absurd as the counterpart where, instead of suspension of judgement, you have outright disbelief that P means that Q, but still seems quite odd. This oddity is explained by The Proposal but not by the alternative specifications of the connecting thoughts. Although P's meaning that Q seems to be *sufficient* for P to be evidence that Q, and as such be reason to think or support that Q, it is not *necessary*. For one thing, it does not seem to be an absolutely general requirement, for a certain fact that P to be evidence that Q, that Q be true. It seems there can be misleading evidence, whereas it is essential to natural meaning that P means that Q only if Q. Even when Q is true, the fact that P can be evidence that Q without those facts meaning that Q. A DNA match at the scene may be evidence that N.N. is the culprit, and N.N. may indeed be the culprit, but if the DNA material comes from N.N.'s unknown twin brother who co-incidentally passed the scene sometime before the crime, the DNA match does not mean that N.N. is the culprit. Also, our example (9) above, repeated here:

(9) That T has a 0.99999 chance of being a loser means that T's a loser.

---

[16] It has been objected that (10) and (11) have acceptable instances. Any known fact that P, the objection goes, trivially means that P, yet it need not be, and typically is not, a reason to think or something that supports that P. However, if this is an exception to (10) and (11), it is matched by a corresponding exception to the putative Moorean absurdity posited by Thomson. If the known fact that P is not (trivially) a reason to think or something that supports that P, but inferences by reiteration from known truths are okay, someone could, non-absurdly, contend: 'P. So P. But P is no reason to think/doesn't support that P'. Thus, for my purposes, this objection leaves it open for me to retreat to the claim that (10) and (11) are unacceptable just so far as there is a Moorean oddity to be explained here in the first place.

More generally on this point: Let us grant that inferences by reiteration from known truths are okay. I am committed, then, to holding any known fact that P trivially to mean that P. I take no stand on whether any known truth that P is also a reason to think or something that support that P. If they are, there is no problem. If they are not, this is a problem for the idea of Thomson, Grice, Audi and Boghossian that such a relation always is taken to obtain in conscious inference. True, it is also a problem for our conjunction (i) that natural meaning implies evidence and (ii) that evidence implies reason to think/support. My response would be to apply a revision to either (i) or (ii) bringing in some suitable non-triviality restriction in the antecedent. A full elaboration of these issues must await another occasion.

is doubtful, and remains so even if T happens to be a loser. However, the high chance of T being a loser arguably is evidence, and a reason to think, that it is a loser. There is, then, nothing incoherent about holding that P supports, or is a reason to believe, that Q, whilst suspending judgement on whether P means that Q.

Let us turn to the suggestion, recently defended by Broome, that the connecting thought, or 'linking belief' in his terms, affirms an implication between one's premises and the conclusion. Broome (2013, p. 229) writes:

> If you reason, you must think of your conclusion as arising somehow from your premises. So you must have a belief that links the premises and conclusion. In the case of belief reasoning, implication is the weakest relation between premises and conclusion that would allow you to think of the conclusion as arising from the premises.

Broome (ibid.) proposes that you believe the implication to hold by believing a 'conditional proposition': 'When the premises are P, Q, R and so on, and the conclusion is T, you believe that, if P, Q, R and so on, then T'. Broome does not specify just what sort of implication a linking belief should be a belief in, in particular whether it is material, strict, or counterfactual implication. Let's briefly consider these options in turn.

Material implication is a very weak relation. Connectedly, it would be acutely irrational to conclude that Q from that P whilst rejecting or even suspending on whether P materially implies that Q, since that rejection (/suspension) commits one to reject (/suspend) on Q, inconsistently with one's concluding judgement that Q. Material implication seems too weak, however. This can be seen by considering a case where you already know a would-be conclusion that Q to be true on independent grounds, e.g. where you are a detective knowing Q from legally inadmissible evidence, now considering independent grounds for Q. In this case, you may rationally believe any arbitrary proposition, P, to materially imply Q. Yet, if you take that P to provide no evidence or reason to believe that Q, you would not, except absurdly, infer that Q simply from that P. In other words, if you do, non-absurdly, infer that Q from a certain proposition that R, the fact that you are thereby thinking of 'your conclusion as arising somehow from your premise' (as Broome puts it) is not explained by the fact that you hold a belief in the material implication 'if R, Q', for you believe such an implication 'if P, Q' to hold for any proposition P whatsoever.

In contrast, P strictly implies Q iff Q holds in any possible world, within a certain contextually restricted range, where P holds. A fairly tight restriction on a range of worlds must be imposed here, if the resulting view is to have a chance of applying to many everyday non-demonstrative inferences. Narrowing down merely to, say, the nomologically possible worlds will not do. Take the case where I infer that my partner is at home, on the basis of her car being in the driveway. If you will, you may add to my basis here various other, more or less backgrounded common or garden assumptions about my partner's daily routines and circumstances. In drawing this inference, I may quite reasonably think that, for all I know about the car (and about her routines etc.) there are non-miraculous ways in which her car can be there and she not at home. Thus, if I really am to be committed to some strict implication from my grounds to her being at home, a narrower restriction on the relevant range of worlds must be imposed. Perhaps the most promising such restriction would be to the worlds that are 'nearby' or 'close' to actuality, where 'nearby' and 'close' are understood as they are in 'safety' conditions on knowledge (cf., e.g., Sosa, 1999; Williamson, 2000). The idea would be that, in inferring that Q from that P, one is taking it that in any nearby Q world, P holds, or, as it is sometimes put, that it could not easily be that Q without P. Let's say that P 'safely implies' Q just in case any nearby P world is a Q world. I shall return to this proposal in a moment.

Another interpretation is that Broome's 'conditional proposition' is a counterfactual. Sticking as closely as possible to Broome's formulation above, the idea would be that in inferring T from premises P, Q, R, and so on, you believe that 'if [*it were that*] P, Q, R, and so on, then [*it would be that*] T'. On the standard Lewis-Stalnaker semantics for counterfactuals, this proposal however runs into much the same problem as that found for material implication, since, on that semantics, it suffices for the truth of 'If it were that P, it would be that Q' that P and Q both are true.[17] If you believe that Q on independent grounds, then for any other proposition P you believe to be true, you may rationally believe that if it were that P, it would be that Q.

We can circumvent this problem by revising Broome's conditional into the following: 'it would not have been the case that (P, Q, R, and so on), unless T'. Note that, even on the assumption (which could be questioned, but which I will make) that 'unless' here is semantically equivalent to 'if it weren't that', this conditional is not logically equivalent to that in the last paragraph, as contraposition is invalid for counterfactuals (Lewis, 1973, pp. 35f.). This

---

[17] The standard semantics might of course be resisted. For present purposes, though, the most obvious alternative, adopted e.g. in (Sosa, 1999), construes 'If it were P, it would be Q' as expressing safe implication, and so matches the option distinguished above.

counterfactual also avoids the problem identified for our first strict implication, where the relevant range of worlds was merely narrowed down to nomologically possible worlds. When I infer that my partner is at home, on the basis of her car being parked in the driveway, it is not crazy to suppose that I am committed to thinking that her car wouldn't have been there unless she were at home. At least, this commitment is compatible with allowing that the car could, compatibly with the laws of nature, be home without her. Another interesting feature of the present counterfactual is that it matches Dretske's (1971) condition for conclusive reasons, this being that P is a conclusive reason that Q iff P wouldn't have been the case unless Q. Let's assume Dretske's condition here is at least a sufficient condition for being a reason, and thinkers implicitly recognize its sufficiency. Then, if the present counterfactual is the content of the linking belief, it would also explain the absurdity, observed by Thomson, of inferring that Q from that P whilst denying that P is a reason to think that Q. Since the condition 'it wouldn't be that P, unless Q' is structurally similar to so-called sensitivity requirements on knowledge, I will say that P 'sensitively implies' Q iff it wouldn't be that P unless Q.[18]

We have found two contenders for being an implication represented in connecting thoughts: safe and sensitive implication. How do these relate to each other and to our proposal that connecting thoughts represent natural meaning? To begin with the last question: as Dretske (1988, p. 54-9) argues, claims of natural meaning can very often be given a gloss in terms of sensitive implication, and vice versa. For example, when a low barometer reading means unsteady weather ahead we can give a gloss on this by saying 'The barometer wouldn't have been this low unless unsteady weather was impending'. So perhaps natural meaning can be analyzed in terms of sensitive implication. On the other hand, Luper (2006), critically responding to Dretske, argues, in effect, that natural meaning may be accounted for in terms of safe as opposed to sensitive implication. One possible upshot, then, is that whether safe or sensitive implication turns out to be the best candidate for an implicational content in connecting beliefs, that implicational content would be analytically equivalent to the content assigned to connecting beliefs by The Proposal.

However, it seems that natural meaning and *safe* implication come apart, and that, when this happens, The Proposal fits better with intuitive verdicts on inferences. Let LOOK be the proposition that a certain animal before me has the stripy, equine, and otherwise zebraish look

---

[18] Dretske (1970; 1971) is a proponent of sensitivity requirements. For critical discussion, see Sosa (1999) and Williamson (2000).

typical of zebras, and FAKE be the proposition that the animal is a mule (or other non-zebra) disguised to have the zebraish look.[19] Here LOOK safely implies not-FAKE, since in all the close LOOK worlds, the zebraish-looking animal is indeed a zebra. However, it seems wrong to say that LOOK means that not-FAKE. How can the fact that the animal has the zebraish look already mean that the zebraish look does not result from a disguised mule? This indicator does not indicate its own non-misleadingness. Relatedly, an inference from LOOK directly to not-FAKE is intuitively hard to endorse. It would clearly raise heckles were someone to argue: 'This animal has a zebraish look. So: it is not a non-zebra made out to have a zebraish look'. The Proposal explains why this inference is hard to swallow.

Notoriously, LOOK does not *sensitively* imply not-FAKE. For the counterfactual 'If it were that not not-FAKE, it wouldn't be that LOOK' is false. When we consider these cases, then, sensitive implication seems a better bet than safe implication for connecting beliefs. Yet, when we turn, conversely, to cases with sensitive but not safe implication, doubts emerge over the propriety of an inference over the sensitive implication. Consider this vignette—a variant upon one in Lewis (1973, p. 35f.):

> Boris wanted to go to a party, but preferred to avoid Olga, who was having a crush on him. Olga, of course, was especially keen on attending if Boris attended. Something that easily could have happened is this: Olga did not attend, while Boris (perhaps being informed of her nonattendance) attends. Something that even more easily could have happened is this: Boris, finding his reluctance towards co-attending with Olga to be outweighed by the party's attractions, decides to attend whether or not Olga does, and Olga attends too.

Now, that Olga didn't attend (supposing she didn't) does not safely imply that Boris didn't, for there is a close world in which she fails to attend but Boris attends. On the other hand, on the standard semantics, were it that Boris attended, Olga would have attended, for the closest world in which he attends, she does too. Therefore, that Olga didn't attend (supposing she didn't) sensitively implies that Boris didn't. How are we, then, to feel about this inference: 'Olga didn't attend. So Boris didn't either.'? It seems reasonable to resist endorsement here on grounds of unsafety: 'Look, even supposing Olga didn't attend, it easily could be that Olga didn't attend but Boris did!' Notice that, on the same grounds of unsafety, it is reasonable to resist the claim that

---

[19] This famous case is derived from Dretske (1970).

Olga's non-attending (supposing she didn't attend) means that Boris didn't attend. In other words, it seems reasonable to remark here: 'Even if Olga didn't attend, that does not necessarily mean that Boris didn't attend', thus expressing suspension of the natural meaning claim. The Proposal, then, which requires commitment to that claim, can explain our doubts about this inference, in a way an assumed commitment to sensitive implication cannot.

Of course, these brief reflections do not rule out that the content of the connecting thought could turn out to be equivalent to an implication, e.g. to safe-*and*-sensitive implication, or perhaps some other form of implication. It suggests, though, that if P implies that Q, on a given notion of implication, but P does not mean that Q, or vice versa, it is the notion of implication that needs to be shifted to get a good fit for the content of the connecting belief. It suggests, in other words, that a plausible, implicational content for connecting thoughts will not be competitor to The Proposal as much as a truth-conditionally equivalent specification of what is taken to be so in connecting thoughts. This provisional assessment needs of course to be pitted against a wider range of cases. I only purport here to have made a prima facie case for natural meaning as at least a contender for the content of connecting thoughts.

## Acknowledgements

## References

Audi, R. (1986). Belief, reason, and inference. *Philosophical Topics* 14, 27–65. (Page ref. to reprint in his *The Structure of Justification*. Cambridge: Cambridge University Press, 1993.)

Block, N. (1995). On a confusion about a function of consciousness. *Behavioral and Brain Sciences* 18, 227–47.

Boghossian, P. (2014). What Is Inference? *Philosophical Studies* 169, 1–18.

Broome, J. (2013). *Rationality through Reasoning.* Oxford: Wiley-Blackwell.

Carruthers, P. (2006). Why Pretend? In S. Nichols (Ed.), *The Architecture of the Imagination: New Essays on Pretence, Possibility, and Fiction* (pp. 89–110). Oxford University Press.

Douven, I. (2006). Assertion, knowledge, and rational credibility. *Philosophical Review* 115, 449–85.

Dretske, F. (1970). Epistemic operators. *Journal of Philosophy* 67, 1007–23.

Dretske, F. (1971). Conclusive Reasons. *Australasian Journal of Philosophy* 49, 1–22.

Dretske, F. (1988). *Explaining Behavior.* Cambridge, MA: MIT Press.

Grice, P. (1957). Meaning. *Philosophical Review* 66, 377–88.

Grice, P. (2001). *Aspects of Reason.* Oxford: Oxford University Press.

Gurwitsch, A. (1964). *Field of Consciousness.* Pittsburgh, PA: Duquesne University Press.

Horwich, P. (1998). *Meaning.* Oxford: Oxford University Press.

James, W. (1890). *Principles of Psychology.* London: Holt.

Kahneman, D. (2011). *Thinking Fast and Slow.* New York, NY: Farrar, Straus and Giroux.

Kahneman, D. and Klein, G. (2009). Conditions for intuitive expertise: a failure to disagree. *American Psychologist* 64, 515–26.

Klein, G., Calderwood, R. and Clinton-Cirocco, A. (2010). Rapid decision making on the fire ground: the original study plus a postscript. *Journal of Cognitive Engineering and Decision Making* 4, 186–209.

Korcz, K.A. (2010). The epistemic basing relation. In E.N. Zalta (Ed.), *Stanford Encyclopedia of Philosophy (Spring 2010 Edition)*.

Lewis, D. (1973). *Counterfactuals.* Oxford: Blackwell.

Luper, S. (2006). Dretske on knowledge closure. *Australasian Journal of Philosophy* 84, 379–94.

Nichols, S. (Ed.) (2006). *The Architecture of the Imagination: New Essays on Pretence, Possibility, and Fiction.* Oxford University Press.

Recanati, F. (2004). *Literal Meaning*. Cambridge: Cambridge University Press.

Simon, H.A. (1992). What is an "explanation" of behavior? *Psychological Science* 3, 150–61.

Sosa, E. (1999). How to defeat opposition to Moore. *Philosophical Perspectives* 13, 141–53.

Strawson, G. (2011). Cognitive phenomenology: real life. In T. Bayne and M. Montague (Eds.) *Cognitive Phenomenology* (pp. 285–325). Oxford: Oxford University Press.

Thomson, J.J. (1965). Reasons and reasoning. In M. Black (Ed.), *Philosophy in America* (pp. 282–305). Ithaca, NY: Cornell University Press.

Williamson, T. (2000), *Knowledge and Its Limits.* Oxford: Oxford University Press.