De Re A Priori Knowledge

Cian Dorr
University of Oxford
cian.dorr@philosophy.ox.ac.uk

Suppose a sentence of the following form is true in a certain context: 'Necessarily, whenever one believes that the F is uniquely F if anything is, and *x* is the F, one believes that *x* is uniquely F if anything is'. I argue that almost always, in such a case, the sentences that result when both occurrences of 'believes' are replaced with 'has justification to believe', 'knows', or 'knows a priori' will also be true in the same context. I also argue that many sentences of the relevant form are true in ordinary contexts, and conclude that a priori knowledge of contingent *de re* propositions is a common and unmysterious phenomenon. However, because of the pervasive context-sensitivity of propositional attitude ascriptions, the question what it is possible to know a priori concerning a given object will have very different answers in different contexts.

1. Exportability

Suppose that it is necessary that whenever one believes that the governor of California is a bodybuilder, and some person *x* is the unique governor of California, one believes that *x* is a bodybuilder. Then we will say that the occurrence of 'the governor of California' in 'the governor of California is a bodybuilder' is **exportable for belief**.[1] In general:

> When O is an occurrence of a definite description ⌜the F⌝ in a sentence or open sentence S, S^x is the result of replacing O in S with the variable '*x*', and ψ is a propositional attitude verb like 'believe' or 'know', O is **exportable for ψing** iff the sentence ⌜It is metaphysically necessary that whenever one ψs that S, and *x* is the F, one ψs that S^x⌝ is true

Note that S can contain free variables, including the pronoun 'one'. For example, the occurrence of 'the father of *y*' in 'The father of *y* is

[1] I will be assuming a Russellian account of descriptions, according to which believing that the governor of California is a bodybuilder requires both believing that some governor of California is a bodybuilder and believing that there is exactly one governor of California. I doubt anything substantive will turn on this. Even if the English word 'the' does not work in this way, there is nothing to stop us from stipulatively introducing a new word which does.

Mind, Vol. 120 . 480 . October 2011
doi:10.1093/mind/fzr091 Advance Access publication 14 February 2012

the mayor of one's hometown' is exportable for knowledge iff (1) is true:

(1) It is metaphysically necessary that whenever one knows that the father of *y* is the mayor of one's hometown, and *x* is the father of *y*, one knows that *x* is the mayor of one's hometown

Here 'one', '*x*', and '*y*' are all to be understood as bound by the universal quantifier 'whenever'.[2]

2. Three generalizations about exportability

I will be taking it for granted that sentences in which the complements of 'believe' and 'know' are 'that' clauses are equivalent to the corresponding sentences with singular terms of the form 'the proposition that … '.[3] One believes or knows that P iff one believes or knows the proposition that P; one has justification to believe that P iff one has justification to believe the proposition that P; one knows a priori that P iff one knows the proposition that P a priori.[4] I will be using '*de re*' in such a way that sentences of the form 'For all *x*, the proposition that $\phi(x)$ is a *de re* proposition about *x*' are trivially true. So whenever there is an *x* such that one believes or knows that $\phi(x)$, one thereby believes or knows a *de re* proposition.[5] The main question I want to investigate in this paper is which *de re* propositions can be known a priori.

Some philosophers of mind have maintained that one can be in the extension of the ordinary language predicate 'believes that P' without

[2] See Lewis 1975 on adverbs of quantification.

[3] For a recent defence of these 'pleonastic equivalences', see Schiffer 2003.

[4] Bach (1997) rejects these equivalences, holding that one can believe that P without believing the proposition that P, just as one can fear that P without fearing the proposition that P. Bach seems to regard this surprising claim as following from his thesis that belief reports involving 'that' clauses are semantically incomplete (see Sect. 4 below). But I do not see why belief reports using 'the proposition that … ' should not be semantically incomplete in just the same way.

[5] Following standard philosophical usage, I will treat 'For some *x*, $x = a$ and *b* ψs that *x* is F' as interchangeable with '*b* ψs of *a* that it is F' and with '*a* is such that *b* ψs that it is F', ignoring possible subtleties arising from the semantics of the anaphoric pronouns in the latter two forms. Those who doubt that these three forms are in fact equivalent are invited to substitute a quantificational construction whenever I use one of the others.

being in a *belief state* with the *content* that P.[6] While one could understand 'belief state' and 'content' in such a way as to make this claim trivially false, this is not how those philosophers of mind understand these expressions. Instead, they use terms like 'belief state' and 'content' as technical terms of art, purporting to stand for entities and relations which play some important role in explaining the truth of ordinary language attitude reports. I will *not* be talking about these entities and relations. Thus, when I conclude that a certain *de re* proposition can be known a priori, I will be leaving it open whether it could (in the putative deep sense) be the content of a belief state which constitutes a priori knowledge. Because of this, some will think that I am ignoring the really important and interesting questions in the vicinity of my title. But even those whose primary interest is in these deep questions may find it rewarding to investigate the questions that can be asked in (relatively) ordinary English. By answering them, we may improve our understanding of the explanatory work the technical notion of content needs to do, and we may return to our arguments about this notion better able to distinguish between, on the one hand, premisses which are defensible when taken as claims about familiar relations like *believing* and *knowing* and, on the other, premisses which must be understood as claims (how supported?) about less familiar relations like *being-in-a-belief-state-whose-content-is*.

In trying to determine which *de re* propositions can be known a priori, claims about exportability for a priori knowledge will be of special importance. Suppose, for example, that we could show that the occurrence of 'the spouse of Hillary Clinton' in 'the spouse of Hillary Clinton is married to Hillary Clinton if any one person is' is exportable for a priori knowledge — in other words, that (2) is true:

(2) It is metaphysically necessary that whenever one knows a priori that the spouse of Hillary Clinton is married to Hillary Clinton if any one person is, and *x* is the spouse of Hillary Clinton, one knows a priori that *x* is married to Hillary Clinton if any one person is

Setting scepticism about the very idea of a priori knowledge aside, it is clear that some people do know a priori that the spouse of Hillary Clinton is married to Hillary Clinton if any one person is. From this,

[6] The technical terminology varies. See e.g. Loar (1988) on 'psychological content', Stalnaker (1988) on 'compatibility with a subject's beliefs', Lewis (1979) on 'the objects of belief', and Chalmers (2011) on 'endorsing'.

(2), and the fact that Bill Clinton is the spouse of Hillary Clinton, it follows that Bill Clinton is an *x* such that some people know a priori that *x* is married to Hillary Clinton if any one person is. A fortiori, it is *possible* to have a priori knowledge of the proposition, concerning Bill Clinton, that he is married to Hillary Clinton if any one person is. The fact that this proposition is only contingently true is no obstacle.

In general: suppose that the occurrence of 'the F' in 'the F is uniquely F if anything is' is exportable for a priori knowledge. Then if it is possible for *x* to be the F while someone has the trivial a priori knowledge that the F is uniquely F if anything is, it is also possible for someone to know a priori that *x* is uniquely F if anything is.[7] If it is not a necessary truth that *x* is uniquely F if anything is, this possible a priori knowledge is knowledge of a contingent proposition.[8]

Thus there is a direct route from claims about exportability for a priori knowledge to conclusions about the a priori knowability of various *de re* propositions, including contingent ones. Even those who do not reject the idea of contingent a priori knowledge altogether may find these conclusions repugnant. This intuitive resistance will be especially strong for those who think that (3) is logically equivalent to (4):

(3) Bill Clinton is an object *x* such that some people know a priori that *x* is married to Hillary Clinton if any one person is

(4) Some people know a priori that Bill Clinton is married to Hillary Clinton if any one person is

For while sentences like (3), in which externally bound variables occur within the scope of 'knows a priori', are hard to process, it might seem obvious that (4) is false. Without empirical evidence, how could anyone know that it is not the case that someone other than Bill Clinton is married to Hillary Clinton? As a matter of fact I *do* think that (3) and (4) are logically equivalent, although I will not be

[7] If, in addition, the occurrence of 'the F' in '*x* is uniquely F, if the F exists' is exportable for a priori knowledge, it will also be possible to know a priori that *x* is uniquely F *if x exists*.

[8] If something other than *x* is actually uniquely F, it is possible knowledge of an actually false proposition. It sounds odd to say of a false proposition that it is 'knowable a priori'. But this is easily explained, either pragmatically (by appeal to some kind of presupposition or implicature carried by these sentences) or semantically (by claiming that the notion of possibility expressed by '-able' in these sentences is contextually restricted so as to hold fixed the truth value of the proposition under consideration).

assuming this in what follows.[9] So prima facie, I have a strong reason to deny that the relevant occurrence of 'the spouse of Hillary Clinton' is exportable for a priori knowledge.

In the rest of the paper I want to develop a strategy for pushing in the opposite direction, by arguing for controversial claims about exportability for a priori knowledge from premisses about exportability for belief. The connection goes by way of the following generalizations, which I will defend and clarify in sections 8, 9, and 10, respectively:

> *Generalization 1*: When an occurrence of a description is exportable for belief, it is almost always exportable for propositional justification (the attitude reported by sentences of the form '*a* has justification to believe that P', understood so as not to entail '*a* believes that P')
>
> *Generalization 2*: When an occurrence of a description is exportable for justification, it is almost always exportable for a priori justification ('*a* has a priori justification to believe that P')
>
> *Generalization 3*: When an occurrence of a description is exportable for a priori justification, it is almost always exportable for a priori knowledge[10]

Armed with these generalizations, and pending clarification of the 'almost always', we will be in a position to support surprising claims about a priori knowability by deriving them from premisses about exportability for belief. But in order to understand the significance of this argumentative strategy, we will have to face up to a crucial fact about all the attitude-reporting sentences we have been concerned with: their pervasive context-sensitivity.

[9] See Sect. 11 below for some further discussion of this point.

[10] Claims like Generalizations 1–3 have sometimes been taken for granted in the literature on the contingent a priori. For example: in the course of arguing that contingent a priori knowledge is a rare phenomenon, Nathan Salmon concedes that in the case of a few special descriptions, like 'the stick that such-and-such a visual perception is visually veridically presenting to oneself', 'merely grasping its information value ipso facto places the user in a position to form *de re* beliefs concerning the referent qua the thing so described'. He takes it to follow from this that 'One can know a priori concerning a particular stick *S* that if such-and-such a visual perception is visually veridically presenting a certain stick to oneself, then *S* is that stick' (Salmon 1986, p. 180). This comes close to the assumption that if occurrences of the relevant description are exportable for belief, they are exportable for a priori knowledge.

3. The context-sensitivity of attitude reports

The case for positing some form of context-sensitivity in propositional attitude reports, at least those in which externally bound variables occur in the scope of the attitude verb, has been well known for a long time, so my presentation of it will be brief.[11] Here is an example due to Ernest Sosa (1970), perhaps the first to make the argument for context-sensitivity fully explicit:

> [C]onsider … the case of a prominent citizen of Metropolis who suffers from pyromania. Impelled by his pyromania, he disguises himself from time to time in order to start some fires, and becomes known to the community as 'The Metropolis Pyromaniac'. Now there are other arsonists in town, but the police always know the work of our pyromaniac by certain peculiarities of it. Eventually our man's wife begins to wonder whether anyone suspects that he has set any of those fires. But he is able to insist that no one does. In this he is right. At the same time, however, soon after the latest fire the chief of police is asked by the press whether anyone is suspected by the detective assigned to the case, and he is able to reply affirmatively. For the pyromaniac has left all the usual signs of his work and the detective suspects that he started the fire. (Sosa 1970, p. 894)

Let me put the argument in a way that takes more account of niceties about use and mention. Talking to his wife, the Pyromaniac utters (5):

(5) No-one suspects that I set any of these fires

Meanwhile, talking to the press, the chief of police utters (6):

(6) There is someone such that the detective suspects that he set the latest fire

Neither utterance seems to convey anything false, and neither bears any of the usual hallmarks of non-literal speech. But if we held that both (5) and (6) are true in the contexts in which they are uttered, and that the only relevant sources of context-sensitivity are the present tense, the first person pronoun, and the domain of the quantifiers, then we would have to conclude that (7) is true, as spoken by the Pyromaniac at the time in question:

(7) There is someone, distinct from me, such that the detective suspects that he set the latest fire

[11] The following authors, among many others, argue for the existence of such context-sensitivity: Lewis (1979), Schiffer (1979), Crimmins and Perry (1989), Richard (1990, 1993), Crimmins (1992), Oppy (1992), Recanati (1993), and Bach (1997).

This is absurd, given the facts of the case.

One might try to resist this argument by claiming that 'I' as it occurs in (5) is occupying a position that resists existential generalization, so that although (5) is true (relative to the Pyromaniac's context), (5′) is not:

(5′) For some *x*, I am *x* and no-one suspects that *x* has set any of these fires

There are, indeed, uses of 'I' which seem to resist existential generalization. For example, when Geoffrey Nunberg's condemned prisoner (Nunberg 1993) utters

(8) I am traditionally allowed to order whatever I like for my last meal

it is arguable that he would be making a mistake if he were to go on to assert

(8′) For some *x*, I am *x* and *x* is traditionally allowed to order whatever *x* likes for *x*'s last meal[12]

But the assimilation of (5) to such cases strikes me as deeply unpromising. True, (5′) is a bit stilted, so it may be hard to have an intuitive reaction to it. But surely if the Pyromaniac is within his rights to utter (5), he can just as well say something like (9):

(9) There are, in the world, people whom some detectives suspect to be arsonists; but I am not one of them

Indeed, the argument for context-sensitivity can be made without bringing in indexical pronouns at all, just by noting that

(10) There is no-one whom the detective assigned to the case suspects to have set this fire

is a sentence that someone in the world we have been imagining could perfectly well utter without conveying anything false, and without any of the usual hallmarks of non-literality. Since (6) and (10) are manifestly logically inconsistent, the only way it could be possible for both to be used literally without anyone speaking falsely is for them to be context-sensitive (in some non-obvious way).

[12] In fact I doubt this would be a mistake. Following Hawthorne and Manley (2012, Ch. 6), I would prefer to reconcile the acceptability of (8) with an standard account of 'I' by claiming that in the relevant context, we are generous in ascribing *de re* contents to traditions.

Here is another example that avoids distracting issues about pronouns that are not functioning semantically like bound variables. Hume's *Treatise of Human Nature* was published anonymously in 1739. Here are two things we might want to say in connection with this episode:

(11) Every philosopher who wrote a book called *A Treatise of Human Nature* was believed by many of his early readers to be an atheist

(12) A philosopher who wrote a book called *A Treatise of Human Nature* was so successful in preserving his anonymity that only a few of his intimate friends had any idea that he had written a book

Taken separately, these seem like perfectly fine, non-misleading things to say, and neither displays any of the ordinary signs of not being meant literally. So there is at least a *prima facie* case that neither sentence expresses a falsehood relative to the context in which it is uttered. But if we held that both (11) and (12) are true *simpliciter* (ignoring context-sensitivity due to tense), we would have to conclude that (13), which follows logically from (11) and (12) taken together, is also true *simpliciter*:

(13) A philosopher who wrote a book called *A Treatise of Human Nature* was believed by many of his early readers to be an atheist, while only a few of his intimate friends had any idea that he had written a book

But I doubt we can live with the claim that (13) is context-insensitively true. If Hume's readers did not believe him to be the author of the book they were reading, on what grounds could they possibly have come to believe him to be an atheist? So the only way to do justice to the possibility of uttering (11) and (12) literally without asserting anything false is to posit context-sensitivity in these sentences.

Cases like these do not constitute a knock-down argument for the context-sensitivity of *de re* attitude reports. Indeed, several recent works have raised objections to the idea that one could ever establish the presence of context-sensitivity by appeal to such data.[13] Although none of our imagined utterances bears any of the *obvious* hallmarks of

[13] The most sweeping such objections are made by Cappelen and Lepore (2005). Many others have criticized the use of such arguments in the case of 'know' in particular: see e.g. Schiffer (1996), Hawthorne (2004), and Stanley (2004).

non-literality, it might be claimed that some of them are still non-literal, in the sense that the speaker does not assert or intend to communicate the proposition semantically expressed.[14] Alternatively, and to my mind much less plausibly, it might be claimed that ordinary people are just systematically mistaken about one another's *de re* beliefs, so that in some of our examples, the speakers not only semantically express, but believe and intend to communicate, certain false propositions.[15] Space precludes an adequate treatment of such invariantist views. Let me just suggest that as far as the purposes of the present paper are concerned, the dispute between the contextualist and the first, 'pragmatic', form of invariantism may turn out to be relatively unimportant. Invariantists of this sort will, I think, have to say that the practice of using *de re* attitude reports and their negations nonliterally is very widespread. They will need some theory about the shifting but systematic standards of assertability that govern this practice. I suspect that this theory will end up isomorphic in relevant respects to a contextualist semantics. Where contextualists like me reconcile seemingly inconsistent claims by evaluating them both as true relative to different contexts, invariantists will say that although at least one of the claims is literally false, both are still legitimate, since whichever is literally false is not intended literally. While we might wonder what non-arbitrary criterion the invariantist will use to decide which utterances to count as nonliteral, the puzzle-dissolving power of the two modes of reconciliation will probably be much the same.

4. Excursus: two models of context-sensitivity

How should we accommodate the context-sensitivity of attitude reports in a semantic theory? I see two main possibilities.

According to the *indexical* model, the semantics of attitude reports works like Kaplan's semantics for words like 'I' and 'today' (Kaplan

[14] Sosa (1970), for example, ends up tentatively favouring an invariantist view on which it is very easy for *de re* belief reports to be literally true, since (roughly speaking) any occurrence of any description is exportable for belief. While he notes that many of the sentences that come out literally true on this view would be misleading to assert, he gives less emphasis to the harder-to-explain fact that many sentences that come out literally false on this view are perfectly fine to assert.

[15] Braun (1988) defends an error theory of this sort with regard to our propensity to say things like 'The Babylonians did not believe that Phosphorus is visible in the evening'. On Braun's view, when ordinary people utter sentences like this, they express certain false beliefs, which they fall into because of a pervasive blindness to the validity of certain Leibniz's Law inferences.

1989). Our semantic theory will associate these sentences not with single propositions, but with functions from 'contexts' to propositions — where contexts are, minimally, entities such that whenever someone is uttering a sentence, there is a unique context in which they are uttering it.[16] Of course, to get sentences like those in section 3 to express true propositions in the contexts in which we were imagining them being uttered, either the function from utterances to contexts or the function from sentences and contexts to propositions will have to depend, somehow, on some mental states of the speaker, such as communicative intentions.

According to the *incompleteness* model, by contrast, the job of the semantic theory is just to associate each of the target sentences with a *set* of propositions, the *admissible interpretations* of the sentence. We could take these propositions, or the set of them, to be the semantic value(s) of the sentence. Or we could take the semantic value to be something that determines the set: perhaps a 'propositional radical' (Bach 1994), thought of as a structure containing 'gaps' that can be filled in different ways so as to yield different propositions.[17] When a sentence is used literally, the speaker will assert, and intend to communicate, at least one of the propositions with which it is semantically associated. By contrast with the indexical approach, there is no further question whether an utterance of a sentence is true — or whether the sentence is true in the context that the utterance is in — over and above the questions whether the various propositions that the utterer was asserting, or intending to communicate, or successfully communicating, were true.

The two models reflect different conceptions of the scope and purpose of semantic theorizing; in the present paper I would like to remain neutral between them. But this will take a little effort, because the two models require different apparatuses for talking about logical relations between context-sensitive expressions.

Since the indexical model lets us evaluate each context-sensitive sentence as true or false relative to any given context, we can characterize relations of consequence or inconsistency between context-sensitive expressions by quantifying over contexts. For example, even if we treat words like 'no' and 'some' as sources of context-sensitivity because of the phenomenon of contextual quantifier domain restriction, we can still capture a certain intimate relation

[16] Though perhaps it need not be *determinate* what the unique context of an utterance is.

[17] See also Recanati 1993, Soames 2005.

between these words, by saying that whenever 'No F is a G' is true relative to a context, 'Some F is a G' is false relative to that context.

Those who favour an incompleteness model for some kinds of context-sensitivity should not dismiss claims like this as nonsense. They can make sense of them by drawing a distinction between *uniform* and *non-uniform* admissible interpretations of a context-sensitive sentence. Consider

(14) Mary is ready, and everyone who is ready will be coming on the trip, but Mary will not be coming on the trip

One can utter (14) and thereby assert something true (and nothing false). But to do so, one has to rely more heavily than usual on one's hearers' capacity to figure out what one has in mind. One must intend to draw their attention, somehow, to two different properties associated with the expression 'is ready', rather than just one. We often succeed in doing this sort of thing.[18] Nevertheless, the obstacles to making oneself understood in such cases are higher than usual: there is some presumption in favour of uniform interpretations.[19] This is reflected in the fact that when we read (14) in abstraction from the circumstances that might make it a sensible thing to utter, it strikes us as bizarre: a certain class of inconsistent propositions is salient to us.

If we think of propositions as structured, we can say that for an admissible interpretation of a sentence to count as uniform, it must have repeated constituents corresponding to words that occur several times in the sentence. But uniformity also places demands on the treatment of different words. For example, if we decide on an incompleteness model for contextual quantifier domain restriction, we should say that the uniform interpretations of (15) are all necessarily false:

(15) No philosopher is going to be fired and some philosopher is going to be fired

[18] Stanley (2005, pp. 63–5) gives examples involving quantifier domain restriction that illustrate this.

[19] Perhaps we should make an exception to this for demonstratives proper: in interpreting a sentence like 'That is the same size as that', we assume by default that the two occurrences of 'that' will make different contributions to the asserted proposition(s). Or perhaps we should accommodate this distinctive behaviour of demonstratives by denying that uniform interpretation (in the relevant technical sense) requires assigning the same interpretation to different occurrences of a demonstrative.

In this sense, (15) is a 'logical falsehood'. This is how we will capture the semantic relation between 'no' and 'some' that proponents of an indexical model would characterize in terms of truth in a context.

Just as we can distinguish uniform and non-uniform interpretations of a single sentence, we can also classify an interpretation of one sentence as uniform or non-uniform with respect to a given interpretation of some other sentence. For example, we can say that no true interpretation of 'No philosopher is going to be fired' is uniform with respect to a true interpretation of 'Some philosopher is going to be fired'.[20] We can extend this sort of comparison to larger sets of sentences. At the limit, we get the useful notion of a uniform interpretation of the entire language; we can characterize an argument involving context-sensitive sentences as 'valid' iff it preserves truth on every uniform interpretation of the language.[21]

The phenomenon of non-uniform interpretation is an important one; the indexical model too needs a way of characterizing it. One way to do so is to relax the idea that every utterance is made in a *single* context, by allowing that in some cases — including our imagined utterance of (14) — the prevailing context changes midway through an utterance. One could introduce a notion of a proposition being expressed by a sentence relative to a temporal sequence of contexts, and thereby allow that someone who utters a logical falsehood (a sentence that is false relative to each single context) may nevertheless literally speak the truth.[22] A quite different approach is to posit structural ambiguity, treated as homonymy, in sentences like (14). In the sense of 'sentence' that matters to semantics — perhaps the linguists' LF — there are several different sentences that look and sound like (14), that differ by assigning different numerical indices to the occurrences of the context-sensitive word 'ready'. The LFs in which the two occurrences receive the same index are false relative to every context; those in which they receive different indices are true relative to some contexts. Our imagined utterance of (14) is an utterance of an LF in

[20] This need not be taken as a new piece of ideology: plausibly, the interpretation of S_1 as expressing P_1 is uniform with respect to the interpretation of S_2 as expressing P_2 iff the conjunction of P_1 and P_2 is a uniform interpretation of $\ulcorner S_1 \text{ and } S_2 \urcorner$.

[21] Note that this machinery is structurally isomorphic to supervaluationist semantics for vague languages. This is no accident: many supervaluationists think of vagueness as something that can be partially 'resolved' in ways that vary across contexts (see e.g. Lewis 1979, p. 153).

[22] One could introduce a stronger notion of logical falsehood requiring falsehood relative to every *sequence* of contexts; but even paradigmatic logical falsehoods like (14) are not logically false in this sense.

which the two occurrences receive different indices; relative to the context of the utterance, this LF expresses a truth. The presumption in favour of uniform interpretations translates, on this view, into a presumption against the multiplication of indices. Strictly speaking, it makes no sense to attribute validity or logical falsehood to coarsely individuated 'sentences' like (14). But when we are being less strict, we will normally want to understand a coarsely individuated sentence to have one of these properties iff the disambiguations that use as few indices as possible have the property.[23]

For our purposes, the upshot of all the pictures is quite similar. To preserve neutrality, I will, for the remainder of the paper, use 'sentence' to refer to the coarsely individuated items, and use 'context' in a way that is neutral between contexts as conceived by the indexical model and uniform interpretations as conceived by the incompleteness model. I will be wanting to say things like 'S_2 is true in the context evoked by utterance U of S_1'. Indexicalists should understand this to mean either that U is in a single context and S_2 is true in that context; or (if they go in for indices) that U is an utterance of a disambiguation of S_1 using as few indices as possible, and a disambiguation of S_2 using the same indices is true in the context U is in. Those who prefer an incompleteness model should understand it to mean that some uniform interpretation P of S_1 was asserted in U, and every interpretation of S_2 that is uniform with respect to the interpretation of S_1 as expressing P is true.

It is natural to suppose that context-sensitive sentences always have at least one context-sensitive syntactic constituent. This follows from a principle of compositionality according to which each admissible interpretation of a (syntactically disambiguated) sentence is determined, according to invariant structural rules, by an assignment of admissible interpretations to its syntactically atomic constituents. Advocacy of an

[23] The claim that there are *infinitely* many different sentences whose surface form looks like (14) seems highly artificial. It would be less artificial if we could posit just *two* sound-alike sentences: one in which a syntactic relation of 'coordination' holds between the two occurrences of 'ready', and one in which it does not. It is hard to see how one could give an indexicalist semantics for sentences individuated in this way. By contrast, it is easy to see how things would go on an incompleteness model: we simply say that the admissible interpretations of the coordinated disambiguation of (14) are all necessarily false, whereas some of the admissible interpretations of the uncoordinated disambiguation are true. On this kind of approach, we would need to allow for structural syntactic ambiguity in multi-sentence discourses even when there is no ambiguity at the level of individual sentences. The coordinated disambiguations of the discourse 'Mary is ready. Therefore, someone is ready' are valid arguments (their admissible interpretations are all pairs of propositions where the first entails the second); the uncoordinated disambiguations are not valid.

incompleteness model of context-sensitivity has recently tended to go along with rejection of this principle, under the guise of an embrace of 'unarticulated constituents' (Recanati 1993, Bach 1994, Carston 2002). But it need not; in fact the reasons that have been given for rejecting the principle strike me as weak. If we endorse the principle, and take attitude ascriptions to be context-sensitive, we will have to answer a further question about the source of their context-sensitivity. Is it due to context-sensitivity in propositional attitude verbs, or to context-sensitivity somewhere in their clausal complements? Or do these sentences contain some unpronounced syntactic constituents on which the context-sensitivity can be pinned (Schiffer 1979, Ludlow 1995)? These are hard questions; I hope to say more about them in future work. But the claims of the present paper should go through equally well no matter how they are answered.

5. Context-sensitivity and exportability

For the occurrence of 'the author of the *Treatise*' in 'the author of the *Treatise* is an atheist' to be exportable for belief is for (16) to be true:

(16) It is metaphysically necessary that whenever one believes that the author of the *Treatise* is an atheist, and *x* is the author of the *Treatise*, one believes that *x* is an atheist

But if sentences in which variables occur free in the scope of attitude verbs are context-sensitive, then (17) too is context-sensitive, so it cannot be true or false *simpliciter*. At best, we can ask whether it is true relative to a given context.

Are there in fact contexts relative to which (16) is true? In this section I will argue that there are, by appealing to the already recognized context-sensitivity in

(11) Every philosopher who wrote a book called *A Treatise of Human Nature* was believed by many of his early readers to be an atheist

I will claim that (16) is true relative to the contexts that would naturally be evoked by an utterance of (11).

It will be easier to make the case if we allow ourselves to semantically descend. So, let us anchor ourselves in a favourable context by stipulating there was someone — namely, Hume — who many of the early readers of the *Treatise* believed to be an atheist. Was their

believing that the author of the *Treatise* was an atheist while Hume was the author of the *Treatise* sufficient for their believing him to be an atheist, or was there some further condition they had to satisfy?

Well, what could this further condition be, given that so many of the readers managed to satisfy it? Could it be that, in order to believe of Hume that he is an atheist, one must not believe that anyone *other* than Hume was the author of the *Treatise*? Or that one must not believe of Hume that he is *not* an atheist? These suggestions do have some plausibility. When the *Treatise* first appeared, some of its readers attributed it to one George Turnbull (Mossner 1954/1980, p. 125): it would be odd to apply the predicate 'believes of Hume that he is an atheist' to those who concluded from their reading that Turnbull was an atheist. Likewise, it would be odd to apply this predicate to an ignorant friend of Hume's who believes that the author of the *Treatise* is an atheist while believing that his friend Hume is no atheist. But these facts are best explained by the claim that explicitly considering a person one knew to belong to one of these categories would tend to evoke a different context. Suppose that there are five people in a room all of whom believe that the author of the *Treatise* is an atheist. One believes in addition that Turnbull wrote the *Treatise*; another is Hume's ignorant friend. If Hume asks us, 'How many people in that room believe that I am an atheist?', we should answer 'Five' or 'None'; the answers 'Three' and 'Four' seem just wrong. Again: suppose a person who has believed for a while that the author of the *Treatise* is an atheist comes to believe that Turnbull wrote the *Treatise*, or gets to know Hume personally, while continuing to believe that the author of the *Treatise* is an atheist. It would be very odd to tell Hume 'So-and-so used to believe that you are an atheist, but no longer does so'. Beliefs do not disappear as easily as that.

Some authors have suggested that to believe any *de re* proposition, you need to have a certain distinctive kind of cognitive mechanism: something like a 'mental file' in which you deposit information about the object in question, or a name for the object in one's language of thought.[24] Could this be right, by our present standards? If it is right, then the fact that so many of those who believed that the author of the *Treatise* was an atheist believed of Hume that he was an atheist depends on the contingent further fact that these people opened 'author of the *Treatise*' files, or introduced appropriate names into

[24] The 'file' metaphor goes back to Grice (1969); it has been taken up by many subsequent authors.

their languages of thought. I have little grasp of what this further fact could be. (What would one have to do *not* to have a mental file or name for the author of some book one has heard of?) There may be interesting questions in cognitive science in the vicinity, but surely our knowledge that most of Hume's readers believed him to be an atheist does not commit us to any particular answers to these questions.

It seems, then, that (16) is true relative to our present context: believing that the author of the *Treatise* was an atheist is sufficient for believing concerning the author of the *Treatise*, if there is one, that he or she is an atheist. Is it also necessary? I see no reason to think so. We could easily assert a truth by uttering (17):

(17) Hume was someone many people believed to be an atheist, including many readers of the *Treatise*, as well as others who knew him from sight but had never even heard of his books

(17) sounds fine; it has none of the feeling of flirting with contradiction that is characteristic of sentences whose only true interpretations are non-uniform. Thus, contrary to a suggestion made by Burge (1977), the pronouns in the sentences we are concerned with behave quite differently from so-called 'pronouns of laziness', as exemplified in cases like this: I say 'The heaviest object on the desk is a book, but it might have been a computer' and thereby assert something I could equally well have asserted by saying 'The heaviest object on the desk is a book, but the heaviest object on the desk might have been a computer'.

Of course, there are many contexts relative to which (16) is false. Might some of these be contexts in which nothing at all is exportable for belief? No: some descriptions have occurrences that must be exportable for belief just as a matter of logic. Consider for example the occurrence of 'the shortest person one believes to be a spy' in 'the shortest person one believes to be a spy is a spy'. This is exportable for belief in every context, since (18) is a logical truth:

(18) It is metaphysically necessary that whenever one believes that the shortest person one believes to be a spy is a spy, and x is the shortest person one believes to be a spy, one believes that x is a spy

But (18) is clearly a degenerate case. Are there contexts in which exportability for belief occurs only in cases like this, where it is forced to

occur by narrowly logical considerations? I doubt it. Consider another example. During our walk, we have an animated conversation which involves your gesticulating dramatically. Noticing a woman in a passing bus looking at you strangely, I make the following comment:

(19) Someone in that bus thinks you are waving at her

In doing this, I assert at least one true proposition that I justifiably believe. This proposition is, plausibly, a consequence of the proposition that someone in the bus is looking at you and thinks that the person she is looking at is waving at her. Or if not, at least it is a consequence of the stronger proposition that someone in the bus is looking at you while you are gesturing dramatically and thinks that the person she is looking at who is gesturing dramatically is waving at her. Or at least some such implication is true. If there were *no* implication of this sort — if, say, the proposition I asserted required someone in the bus to have a mental file referring to you, where the existence of such a file is not a necessary consequence of believing any descriptive proposition — I cannot see how my assertion could be as unproblematic as it is. Even if I somehow have excellent reason to think that human beings generally introduce mental files in circumstances like these, I am not committing myself to any such claim in asserting (19).[25]

[25] One kind of potential counterexample to the exportability of the relevant occurrence of 'the person one is looking at' involves people who believe concerning themselves that they are looking at people who were waving at them, but do not do so 'under a first person mode of presentation' — in Lewis's terms (1979), they fail to self-ascribe the property of looking at someone who is waving at one. For example, amnesiac Rudolf Lingens might think to himself, 'I bet that Rudolf Lingens is looking at someone who is waving at him', failing to realize that he himself is Rudolf Lingens. It is natural to say that Lingens believes in this case that the person he is looking at is waving at him; but the pressure to say that the person he is looking at is such that he believes that *she* is waving at him is much less than it would be if Lingens knew his own identity.

To address this worry, what we need is a way of understanding the claim that an occurrence of 'the person one is looking at' is exportable for belief on which it does not require those who fail to have the relevant *de se* belief to have the *de re* belief. But it is difficult to find a definition of exportability which has this effect, because it is difficult to find a form of words available in ordinary English that serves to express claims about *de se* belief in a context-insensitive manner. We can say 'Lingens does not believe that he himself is looking at someone who is waving at him'. But I doubt the use of 'himself' here makes a *semantic* difference to the range of available interpretations, as opposed to making a merely pragmatic difference by suggesting something about the particular propositions that the speaker intends to assert or communicate. ('He does not believe that he himself is F, but he believes that he is F' sounds pretty bad, which is evidence that it lacks consistent uniform interpretations.)

English does have expressions which seem to force a *de se* interpretation — for example, 'Lingens expects to ϕ' seems to require Lingens to believe that he will ϕ under a first-personal

It is interesting that occurrences of descriptions like 'the person one is looking at', which pick things out in terms of their relations to the subject's perceptual states, seem like especially strong candidates for exportability. This pattern is part of what leads some invariantists to adopt theories according to which the literal truth of a *de re* attitude report requires a relation of 'acquaintance' between subject and object, paradigmatically present in cases of direct perception. These invariantists will dismiss many of my examples as involving non-literal use; but even they should be open to the possibility that some perception-related descriptions have occurrences that are literally exportable for belief. From a contextualist point of view, of course, there is nothing especially deep about the tendency for occurrences of perception-related descriptions to be exportable: it merely reflects the kinds of communicative purposes which creatures like us tend to have when we use attitude reports.

Are there *any* contexts relative to which the truth of *de re* belief reports requires some perception-like relation between subject and object? I suspect not. My impression is that we can get occurrences of a description to be exportable in pretty much any context just by loading the description up with details about the properties of the object and its relations to the believer, even when these relations involve nothing like perception. Consider, for example, occurrences of the following description:

> the person who, in the eighteenth century, wrote the *Treatise*, the *Enquiries*, the *History of England*, and the *Dialogues on Natural Religion*; who wrote nothing else one has ever heard about; who is, in addition, the only person one has ever heard anyone else referring to using the name 'David Hume'; who has never been known under any other name than this; and with whom one has never had any kind of perceptual contact

Abbreviate this as 'the F'. Perhaps someone with more ingenuity than I will be able to describe possible circumstances in which it would be

mode of presentation. But, annoyingly, 'believes to ϕ' is ungrammatical, and 'expect' can only be used to report belief about times thought of as future. Thus, to find a definition of exportability which clearly prevents people like Lingens from serving as counterexamples to the exportability of descriptions involving 'one', we might have to go beyond ordinary English by introducing some device like Lewis's 'self-ascribes'. I have no doubt that this could be done, and I doubt it would have any effect on my arguments. But to keep the use of jargon to a minimum, I will officially stick to the definition from Sect. 1, and simply stipulate that for the purposes of understanding that definition, occurrences of 'one' in the sentences S and S^x are always to be interpreted in the distinctive *de se* way.

intuitively incorrect to describe one as believing of Hume that he is an atheist, even though one believes that the F is an atheist and Hume is the F. But any such circumstances would, I think, have to be quite unusual.

6. Exportability in sentences that uncontroversially express a priori knowledge

Even given Generalizations 1–3, there is no direct route from the claim that the occurrence of 'the author of the *Treatise*' in 'the author of the *Treatise* is an atheist' is exportable for belief in a given context to any conclusions about the truth, in that context, of sentences attributing a priori knowledge. To get to such conclusions, we will need a way of arguing for the exportability for belief of occurrences of descriptions in sentences for which it is relatively uncontroversial that the propositions they express can be known a priori — for example, sentences of the form 'the F is uniquely F if anything is'.

As a warm-up exercise, it will be useful to think about occurrences of descriptions in sentences of the form 'The F is F', for example the occurrence of 'the author of the *Treatise*' in 'the author of the *Treatise* wrote the *Treatise*'. There is good reason to think that this occurrence is exportable for belief whenever the occurrence in 'the author of the *Treatise* is an atheist' is. If Hume's early readers believed him to be an atheist, they surely also believed him to be an atheist who wrote the *Treatise*. So it is plausible that (16a) is true whenever (16) is:

(16) It is metaphysically necessary that whenever one believes that the author of the *Treatise* is an atheist, and *x* is the author of the *Treatise*, one believes that *x* is an atheist

(16a) It is metaphysically necessary that whenever one believes that the author of the *Treatise* is an atheist, and *x* is the author of the *Treatise*, one believes that *x* is an atheist who wrote the *Treatise*

And if (16a) is true, then since believing a conjunction requires believing each conjunct, so is (16b):

(16b) It is metaphysically necessary that whenever one believes that the author of the *Treatise* is an atheist, and *x* is the author of the *Treatise*, one believes that *x* wrote the *Treatise*

But the fact that one happens to believe that the author of the *Treatise* is an atheist, rather than merely believing that the author of the

Treatise wrote the *Treatise*, does nothing to improve one's claim to count as believing of Hume that he wrote the *Treatise*. Thus if (16b) is true, (16c) is true as well:

(16c) It is metaphysically necessary that whenever one believes that the author of the *Treatise* wrote the *Treatise*, and *x* is the author of the *Treatise*, one believes that *x* wrote the *Treatise*

And this is what needs to be true for the occurrence of 'the author of the *Treatise*' in 'The author of the *Treatise* wrote the *Treatise*' to be exportable for belief.

Everyone who read the *Treatise* presumably believed that the author of the *Treatise* wrote the *Treatise*. Thus (20) is true in any context where (16c) is:

(20) Everyone who read the *Treatise* believed of Hume that he wrote it

Given the historical facts, it would be bizarre to utter (20). Why should this be, if it admits of true interpretations? In short, because our hearers are more likely to take us to have asserted one of the false interpretations than one of the true ones. And why should that be? In this case there is a simple pragmatic explanation: since the true interpretations of (20) are consequences of uncontroversial background facts, such as the fact that everyone who read the *Treatise* believed that the author of the *Treatise* wrote the *Treatise* and the fact that Hume wrote the *Treatise*, to interpret people as asserting any of these propositions would be to interpret them as pointlessly stating the obvious. But this is only the beginning of the story. In section 7 below, we will encounter further factors which will lead us to overlook the possibility of interpreting sentences like (20) as expressing one of these relatively weak propositions, even when they are not so weak as to be obvious or uncontroversial.

The foregoing argument suggests a generalization: almost always, if any widest-scope occurrence of a description is exportable for belief, all of them are. Here is a schematic argument for this conclusion. We assume that the occurrence of 'the F' in 'The F is G' is exportable for belief; we want to show that the occurrence in 'The F is H' is too, for arbitrary H.

(i) Believing that the F is G is sufficient, given that *x* is the F, for believing that *x* is G (Assumption)

(ii) If (i), then believing that the F is both G and H is sufficient, given that *x* is the F, for believing that *x* is both G and H

Defence: if one believes that the F is both G and H and acquires the further belief that at most one thing is G, then given that by (i) one believes that *x* is G, one will be in a position to infer that *x* is both G and H. But it seems wrong to suppose that one's believing this could require any such inference from further premisses.

(iii) If (ii), then believing that the F is both G and H is sufficient, given that *x* is the F, for believing that *x* is H

Defence: believing a conjunction suffices for believing the conjuncts.

(iv) If (iii), then believing that the F is H is sufficient, given that *x* is the F, for believing that *x* is H

Defence: believing that the F is G does not enhance one's claim to count as believing that *x* is H. The case for this is most clear-cut in possible circumstances in which the F is *not* G: how could merely adding a false belief constitute one's coming to believe that *x* is H? But even if the F *is* G, and even if G-ness is the sort of property that might plausibly make something harder to have *de re* beliefs about, it is still hard to think of a scenario where the F's becoming G would leave intact the *de re* beliefs of those who believe that the F is G while eliminating the *de re* beliefs of those who do not.

I do not claim that each of the steps in this argument preserves truth in every case. For example, step (ii) will fail when the occurrence of 'the F' in 'The F is G' is 'degenerately' exportable, as with 'The shortest person one believes to be a spy is a spy'. But the argument seems strong enough to establish a presumption that all widest-scope occurrences of a given description are exportable for belief in a context if any are.

It would be convenient if we could strengthen this generalization further to the claim that *every* occurrence of a given description is exportable for belief if any are. Unfortunately there are clear counter-examples to this. For example, the occurrence of 'the author of the *Treatise*' in 'It is necessary that if any one person wrote the *Treatise*, the author of the *Treatise* did' (where the description is understood as taking narrow scope with respect to the modal operator) is not

exportable for belief in any ordinary context, since (21) is not true in any ordinary context:

(21) It is metaphysically necessary that whenever one believes that it is necessary that if any one person wrote the *Treatise* the author of the *Treatise* did, and *x* is the author of the *Treatise*, one believes that it is necessary that if any one person wrote the *Treatise x* did

For any *x*, the proposition that it is necessary that if any one person wrote the *Treatise x* did is at best a highly controversial claim of essentialist metaphysics. Having the trivial belief that it is necessary that if any one person wrote the *Treatise* the author of the *Treatise* did is certainly not sufficient, even by the loosest standards, for believing any such controversial claim.

We could retreat to a slightly weaker generalization: if any occurrence of a description that takes wide scope with respect to all *intensional* contexts is exportable for belief, all are. I am inclined to reject this too. Even in contexts where the occurrence of 'the author of the *Treatise*' in 'The author of the *Treatise* is an atheist' is exportable for belief, there is no strong pressure to treat occurrences of the same description in *negative* contexts, such as the occurrence in 'it is not the case that the author of the *Treatise* exists', as exportable for belief. Even after we have adopted the permissive standards of such contexts, it does not seem compulsory to apply the predicate 'believes of Hume that it is not the case that he exists' to someone who believes that *Treatise* came into existence through a random collision of atoms. And it would be even odder to apply this predicate to someone who believed that the *Treatise* was written by a committee.

We will therefore need some more discriminating generalization if we want to establish anything about occurrences of descriptions in sentences like 'The F is uniquely F if anything is', which uncontroversially express propositions that can be known a priori. I suggest the following:

Conditionalization: If an occurrence of a description in a sentence S is non-degenerately exportable for belief in a context, then so is the corresponding occurrence in any indicative conditional whose consequent is S

Given that the occurrence of 'the author of the *Treatise*' in 'the author of the *Treatise* is an atheist' is exportable for belief, Conditionalization

entails that the occurrence in 'If exactly one atheist and no-one else wrote the *Treatise*, then the author of the *Treatise* is an atheist' is too. Likewise, if I am right that the occurrence of that concept in 'the author of the *Treatise* wrote the *Treatise*' is exportable for belief in the same contexts as the one in 'the author of the *Treatise* is an atheist', Conditionalization entails that the occurrence in 'the author of the *Treatise* wrote the *Treatise*, if any one person did' is exportable across the same range of contexts.

What reason is there to believe Conditionalization? A line of thought I find attractive involves the idea that if one starts out believing both a conditional and its antecedent, and then ceases to believe the antecedent, one does not just on that account cease to believe the conditional. For example, suppose that you were one of those who believed that the author of the *Treatise* was an atheist, and in consequence — let us stipulate — believed of Hume that he was an atheist. Of course you also believed that *if the external world was real* he was an atheist. Now, having just finished Book I of the *Treatise*, you find yourself gripped by paralyzing doubts about the reality of the external world, other people included. You no longer believe that the author of the *Treatise* is an atheist. You no longer even believe that there are any authors, or atheists. Thus there is no longer any particular person whom you believe to be an atheist. But surely, assuming that you still believe that if the external world is real the author of the *Treatise* is an atheist, there is still someone — namely Hume — such that you believe that if the external world is real he is an atheist. Your new doubts need not have disrupted your old conditional beliefs about how things stand if the external world is real. But if being in this situation is a way to believe of Hume that if the external world exists he is an atheist, then more generally, believing that if the external world exists the author of the *Treatise* is an atheist while Hume is the author of the *Treatise* must be a way to believe of Hume that if the external world exists he is an atheist. It would be unacceptable, for example, to ascribe the conditional *de re* belief to you, while refusing to ascribe it to someone who doubted the existence of an external world from the beginning while still agreeing with you that if it does exist, the author of the *Treatise* is an atheist. So we have an instance of Conditionalization: if the occurrence of 'the author of the *Treatise*' in 'The author of the *Treatise* is an atheist' is exportable for belief, then so is the occurrence in 'If the external world exists, the author of the

Treatise is an atheist'. And if this argument is good, it clearly applies much more generally.[26]

7. The elusiveness of *de re* a priori knowledge

If I am right, sentences of the form 'For some *x*, it can be known a priori that *x* is uniquely F if anything is' are often true in ordinary contexts. But sentences like this tend to strike us as false. For example, it sounds very strange to claim that there is someone of whom it can be known a priori that he is married to Hillary Clinton if any one person is. In the present section, I will try to account for this oddness. My thought is that, even though the sentences in question have admissible interpretations that make them true, hearers tend to overlook

[26] Since I am not terribly confident that this argument for Conditionalization is sound, in this and several subsequent footnotes I will present an alternative route to conclusions about the possibility of *de re* a priori knowledge. The alternative route makes use of the idea that justification or knowledge might count as a priori 'modulo' a given proposition, even when it is not a priori *simpliciter*. One's knowledge (or justification to believe) that P is a priori modulo the proposition that Q iff it is independent of experience except in so far as experience is required to constitute one's knowledge (or justification to believe) that Q. As far as I can see, the considerations in favour of Generalization 2 and Generalization 3 which I will present in Sects 9 and 10 below generalize straightforwardly to the corresponding theses about a priori justification and knowledge modulo Q, for arbitrary Q. Together with Generalization 1, the relativized versions of Generalizations 2 and 3 give us a route from the premiss that an occurrence of a description is exportable for belief to the conclusion that it is exportable for a priori knowledge *modulo* any given proposition. For example: if we assume that the occurrence of 'the author of the *Treatise*' in 'the author of the *Treatise* is an atheist' is exportable for belief, we can conclude (assuming we have no need to worry about the 'almost always' provisos) that it is exportable for a priori knowledge *modulo* the proposition that the author of the *Treatise* is an atheist. But the proposition that the author of the *Treatise* is an atheist, like every proposition that can be known at all, can be known a priori modulo itself. So we can conclude that Hume is such that the proposition that he is an atheist can be known a priori modulo the proposition that the author of the *Treatise* is an atheist.

To get from conclusions like this to claims about a priori knowability *simpliciter*, we could appeal to the premiss that a priori knowability is closed under a kind of 'conditional proof': if it is possible to know a priori that P modulo the proposition that Q, it is possible to know a priori that if Q then P. This is plausible, at least if 'If Q then P' is interpreted as a material conditional. For suppose one's knowledge that P depends on experience only in so far as experience is required for one's knowledge that Q. Then the same is true of one's knowledge of logical consequences of the proposition that P, such as the proposition that if Q then P (a material conditional). But the fact that one's experience suffices for knowledge of Q seems irrelevant to the epistemic status of one's belief that if Q then P: if one's evidence for Q were weaker or nonexistent, that would not be enough to undercut the knowledge-constituting status of one's belief that if Q then P. (It might conceivably prevent one from being able to *entertain* the proposition that if Q then P, or to entertain that proposition under a certain favourable guise. But as I will emphasize in Sect. 9, this is not the sort of dependence on experience that undermines the status of a piece of knowledge as a priori.)

these interpretations: the propositions that jump to mind when we hear them are generally false ones. Anticipating this, speakers will realize that they have little chance of uttering one of these sentences without being taken to have asserted something false, at least without providing plenty of additional cues to help the hearer focus on the true propositions they would like to communicate. And since intending normally involves believing that one will succeed, this means that speakers will be unable to form the intentions which would be required for them to use the sentences in question to *assert* the true propositions.

Why should it be so hard to access the interpretations of the problematic sentences under which they are true? To begin with, we can observe that there are pragmatic pressures which make it unlikely that anyone who uttered a sentence of the form

(22) *a* is such that *b* believes that it is uniquely F if anything is

would be taken to have asserted only some consequence of the proposition that *b* believes that the F is uniquely F if anything is while *a* is the F. In many cases this can be explained simply by the fact that these propositions are part of the background knowledge common to speaker and hearer.[27] But even when they are not, they have other features which are liable to make us overlook them as possible interpretations of (22). For these propositions typically have a disjunctive character: they can be true in one or both of two conspicuously different ways. On the one hand, *a* might be the F, in which case *b* need only satisfy the very weak condition of believing that the F is uniquely F if anything is. On the other hand, *a* might not be the F, in which case *b* must bear some much more demanding relation to *a*. Testimony aside, most processes by which one might come to believe a proposition that was disjunctive in this way would involve coming to believe at least one of its two subcases. But speakers whose only ground for believing such a proposition was their belief that *b* believes that the F is the F while *a* is the F would have a strong motivation to convey the relevant information in some other way, rather than running the risk of leading their hearers to the erroneous conclusion that *b* satisfies some stronger condition. Anticipating this, hearers will conclude that speakers who utter (22) have grounds for believing that some relation between *b* and *a* obtains which would be sufficient for this to be an appropriate thing to utter whether or not *a* is the F. Anticipating *this*,

[27] Cf. the discussion of (20) in Sect. 6.

speakers will be willing to utter (22) only when they intend their hearers to infer that they believe one of these stronger propositions.

This style of explanation extends to logically complex sentences with constituents of the form '*b* believes that *x* is uniquely F if anything is'. Relative to contexts in which the occurrence of 'the F' in 'The F is uniquely F if anything is' is exportable, such sentences will express propositions corresponding to 'oddly shaped', 'gerrymandered' regions of logical space. Since it is a priori unlikely that a speaker would want to communicate one of these propositions, hearers will tend to pass over them as possible interpretations. This weak initial bias will then reinforce itself, as speakers anticipate it, and hearers anticipate speakers' anticipation of it, and so on. The end result will be that speakers will be unable to rely on hearers to come up with the problematic interpretations even when this is required by interpretative charity. Suppose, for example, that Sosa's Metropolis Pyromaniac attempted to confess his crimes to the detective in a roundabout way, by uttering

(23) Everyone in Metropolis who believes that these fires were set by a single person believes that they were set by me

The only uniform interpretations of (23) that are consistent with the background knowledge common to the speaker and his audience are consequences of the proposition that the speaker is the person who set the fires. But this kind of interpretation is called for so rarely that hearers will be unlikely to come up with it without a lot of coaching. Unless the detective is a philosopher of language, the Pyromaniac's attempt to confess his crimes by uttering (23) is likely to generate only bafflement.

This account of the pragmatics of belief reports carries over, *mutatis mutandis*, to justification, knowledge, a priori justification, and a priori knowledge. When we are first confronted with a sentence like

(24) It is possible for someone to know a priori of Bill Clinton that he is married to Hillary Clinton if any one person is

the interpretations on which it is true — those on which the occurrence of 'the spouse of Hillary Clinton' in 'The spouse of Hillary Clinton is married to Hillary Clinton if any one person is' is exportable for a priori knowledge — are much less likely to occur to us than those on which it is false. We will think instead of the states we would naturally take to be attributed by an utterance of '*b* knows of Bill

Clinton that he is married to Hillary Clinton if any one person is'. Finding that these are not states one can have a priori justification for being in, we will be disposed to reject the utterance of (24). Thus, those who want to use (24) to communicate a truth will need to work hard to provide their hearers with the tools to distinguish this truth from the falsehoods they will initially find salient.[28]

8. From belief to justification

It is finally time to begin the task of defending the generalizations introduced in section 2. Given context-sensitivity, they must of course be given a relativized form. Thus:

> *Generalization 1* (relativized): When an occurrence of a description is non-degenerately exportable for belief in a context, it is almost always exportable for justification in that context

To get a feel for the plausibility of this, consider an example. Suppose that I have justification to believe that the author of the *Treatise* is clever, and that it is necessary that whenever I believe this while *x* is the author of the *Treatise*, I believe that *x* is clever. How, under these circumstances, could I lack justification to believe that *x* is clever? Given that *x* is the author of the *Treatise*, it is impossible to believe that the author of the *Treatise* is clever without believing that *x* is clever. So if it is OK for me to believe the former proposition, it must be OK for me to believe the latter.

Let me make that argument more explicit, in order to see how far it can be generalized. One premiss is a conception of propositional justification as permissibility: one has justification to believe that P iff one is *epistemically permitted*, given one's circumstances, to believe

[28] There are instructive similarities and differences between this pragmatic story about a priori knowledge attributions and the idea of 'elusiveness' promoted by contextualists about knowledge like Cohen (1988), DeRose (1995), and Lewis (1996). According to them, sentences like 'We know that we are not brains in vats' are unassertable despite being true in most ordinary contexts: the act of asserting them creates an extraordinary context, in which they are false. This is essentially what I would say about a sentence like (24). But there is the following difference: the mechanisms these authors posit for explaining the changes of context seem to be ones that would continue to operate even if we were fully convinced of the truth of the contextualists' semantic theories. For example, according to Lewis, the problem with 'I know that I am not a brain in a vat' is that asserting this sentence forces us *not to ignore* the possibility that we are brains in vats: this would render the sentence false in the context in which we uttered it, even if the proposition we were (foolishly) trying to communicate was a true one. By contrast, if we all attained a self-conscious understanding of the range of admissible interpretations of (24), I see nothing to prevent us from speaking truly by uttering it.

that P — that is, iff one's circumstances are such that believing that P under those circumstances is epistemically permissible *tout court*.[29] The other key premiss is that permission is closed under entailment: if one is permitted to have a certain property, one is permitted to have any other property entailed by that property. If the occurrence of 'the F' in 'ϕ(the F)' is exportable for belief, then the property of *believing that ϕ(the F) while x is the F* entails the property of *believing that ϕ(x)*. If permission is closed under entailment, it follows that if one is permitted, given one's circumstances, to (believe that ϕ(the F) while *x* is the F), then one is permitted, given one's circumstances, to believe that ϕ(*x*). But for a wide variety of descriptions which we might substitute for 'the F', the fact that *x* is the F will count as part of one's circumstances, in the relevant sense. For example, the fact that Hume is the author of the *Treatise* is part of my circumstances — it is not the sort of fact whose obtaining could even in part constitute my failing to live up to my epistemic obligations. If we denied this, we would have to regard the following as an acceptable speech: 'While you do have *justification* to believe that the author of the *Treatise* is clever, I am afraid you cannot actually *believe* this without violating your epistemic obligations, since you can do nothing about the fact that Hume is the author of the *Treatise*, and given that he is, there is no way for you to believe that the author of the *Treatise* is clever without believing unjustifiedly that he is clever.' But this is not an acceptable speech. If you did have justification for believing that the author of the *Treatise* was clever, it would have to be epistemically permissible not only for you to believe this under some circumstances or other, but for you to do so in a world that matches the actual world as regards Hume's being the author of the *Treatise*. When 'the F' is this sort of description, it will be metaphysically necessary that if one is permitted, given one's circumstances, to believe that ϕ(the F), and *x* is the F, then one is permitted, given one's circumstances, to (believe that ϕ(the F) while *x* is the F); and hence, by the closure of permission under entailment, to believe that ϕ(*x*). Thus provided that 'the F' is the sort of description whose satisfaction by an object counts as part of one's circumstances, those of its occurrences that are exportable for belief are also exportable for justification.

[29] I am thinking of this as a weak claim about the logic of justification, much weaker than the 'deontological conception of justification' criticized by Alston (1988). The latter claims that in order to believe something unjustifiably, one must be *blameworthy* for believing it and hence *able to avoid* believing it, in a sense of 'able' much stronger than mere metaphysical possibility given one's circumstances.

How is the distinction between facts that do and do not count as part of one's circumstances to be understood? One's epistemic obligations at a time have to do with one's beliefs at that time, and perhaps also at future times. If one has justification at *t* in believing that P, then there must not only be some epistemically permissible world or other in which one believes that P: there must be an epistemically permissible world in which one believes that P while the facts about subject matters other than one's belief state at *t* and subsequent times are just as they actually are. If so, the foregoing argument will fail only when 'the F' is a description that has to do with the subject's present or future beliefs. And in fact, it turns out that some cases of this sort are exceptions to Generalization 1. Recall from section 5 that the occurrence of 'the shortest person one believes to be a spy' in 'the shortest person one believes to be a spy is a spy' must be exportable for belief as a matter of mere logic, since (18) is a logical truth.

(18) It is metaphysically necessary that whenever one believes that the shortest person one believes to be a spy is a spy, and *x* is the shortest person one believes to be a spy, one believes that *x* is a spy

For this occurrence of 'the shortest person one believes to be a spy' to be exportable for justification, (18a) would also have to be true:

(18a) It is metaphysically necessary that whenever one has justification to believe that the shortest person one believes to be a spy is a spy, and *x* is the shortest person one believes to be a spy, one has justification to believe that *x* is a spy

(18a) is not a logical truth, and indeed there is no intuitive pressure to accept it. The following story is easy to make sense of: because of my paranoid tendencies, there are dozens of famous authors whom I believe, unjustifiably, to be spies. A trustworthy authority tells me I am not *completely* wrong: the shortest of these people is in fact a spy. However, the authors on my list are publicity-shy, and I have no justification to believe anything specific about their relative heights. So even after hearing from the authority, I still lack justification to believe anyone in particular to be a spy.

But this degenerate case is quite atypical. Since (18) is a logical truth, the explanation of its truth in a context has nothing to do with the distinctive nature of the relation expressed by 'believes' in that context. Relatedly, (18) is not the sort of thing that could play any role in

explaining why you count as believing any given person to be a spy. If we want to explore whether we can expect Generalization 1 to hold for typical descriptions involving the subject's mental life, we will do better to focus on an occurrence of such a description that is exportable for belief in some contexts but not in others. So consider the following case:

> *Colloquium*: I am going to give a talk at another department. One of the organizers sends me an email: 'We have one graduate student here whose views about metaphysics are close to yours. Unfortunately she has to teach a class during your talk. But you should definitely talk to her at the reception afterwards.' Immediately after the talk, a graduate student who was in the audience comes up and introduces herself, saying 'I am happy finally to meet someone else who shares my conviction that tables and chairs do not exist.' I reply: 'I am glad you could make it — I thought you had to teach during my talk.'

What I assert by making this reply is true. And it is plausible that relative to the contexts evoked by my utterance, (25) is true:

(25) It is metaphysically necessary that whenever one believes that the graduate student at department *y* whose views about metaphysics are close to one's own has to teach during one's talk, and *x* is the graduate student at department *y* whose views about metaphysics are close to one's own, one believes that *x* has to teach during one's talk

For the truth of what I asserted does not require me to have had any *other* way of thinking of the student at the contextually relevant time. For example, what I asserted could have been true even if I had not believed that *the student I was told about* had to teach during my talk, since it could be true even if I had entirely forgotten the conversation in which I came by my belief.

Now, suppose that I am terrible at metaphysics: I lack justification to believe that tables and chairs do not exist. Then the fact that the student in question is the graduate student whose views about metaphysics are close to mine cannot count as part of my circumstances, since in conjunction with the facts about the content of the student's views about metaphysics, which clearly *do* count as part of my circumstances, it entails that I violate my epistemic obligations. So our earlier argument for Generalization 1 does not apply in this case. Nevertheless, given (25), and given that I have justification to believe that the graduate student whose views are close to mine has to teach

during my talk, surely I must have justification to believe of her that she has to teach during my talk. The fact that I lack justification for my metaphysical beliefs is irrelevant to my justification for this humdrum *de re* belief. Although I hold the belief that she has to teach during my talk *partly in virtue of* my metaphysical beliefs, this belief is not *based on* or *grounded in* my metaphysical beliefs in the epistemological sense on which it would inherit their lack of justification.

This is a vindication of Generalization 1, but it is bad news for the conception of justification as permissibility which featured in the earlier argument that Generalization 1 holds for descriptions that are *not* about the subject's mental life. Given my external circumstances, the only possibilities in which I believe of the graduate student that she has to teach during my talk are ones in which I believe, impermissibly, that tables and chairs do not exist. Since permission is closed under entailment, it follows that it is not permissible for me to believe under these circumstances that she has to teach during my talk. On a conception of justification as permissibility under the circumstances, that would be enough for me to lack justification for this belief. But as we have just seen, it is not. Permissibility given one's circumstances may be sufficient for justification, but it is not necessary.

It is hard to get a grip on what is going on here. It is tempting to think that my metaphysical beliefs should somehow count as part of my circumstances when we are assessing the epistemic status of my belief of the student that she has to teach during my talk, even though they do not count as part of my circumstances when we are assessing their own epistemic status. But if we changed the example to one in which my belief that the student with views close to mine has to teach during my talk was somehow *inferred* from the unjustified beliefs in virtue of which her views were close to mine, we would not then want to treat these beliefs as part of my circumstances. There is no obvious way for an account of justification in deontic terms to allow for the dual status of these facts.

This opens up the possibility of Generalization 1 failing even for descriptions that have nothing to do with the subject's mental life. In a case where one has justification to believe that ϕ(the F) despite not being epistemically permitted to do so in one's circumstances (where x is the F, this fact has nothing to do with one's mental life, and the occurrence of 'the F' in 'ϕ(the F)' is exportable for belief), could it happen that one lacks justification to believe that $\phi(x)$?

If *Colloquium* is representative of cases where there is justification without epistemic permissibility, there is little to worry about here.

Suppose that necessarily, whenever one believes that *x*'s dissertation is about metaphysics and *y* is *x*'s dissertation, one believes that *y* is about metaphysics. In the imagined case, I have justification to believe, of the student, that her dissertation is about metaphysics, despite the fact I am not epistemically permitted to believe this given my circumstances, since the only way to do so requires me to believe, impermissibly, that there are no tables and chairs. This fact will prevent my being epistemically permitted to believe, of the student's dissertation, that it is about metaphysics. But it clearly does no more to prevent my having *justification* for the latter belief than it did to prevent my having justification to believe, of the student, that her dissertation is about metaphysics. Whatever it is about my metaphysical beliefs that makes their epistemic status irrelevant in evaluating my justification for the *de re* belief about the student makes them equally irrelevant in evaluating my justification for the *de re* belief about the dissertation.

This piecemeal combination of arguments and intuitions suggests to me that the 'almost always' proviso in Generalization 1 needs to be invoked only in degenerate cases like 'the shortest person one believes to be a spy'. Of course, in the absence of a full-fledged account of the nature of belief and justification, from which we could derive an explanation of the difference between the degenerate cases and the others, this will at best be a tentative hypothesis. So be it. I am not that worried: even if Generalization 1 turned out to be false in a wide variety of cases, because of possible cases of justification without permissibility, it would do little to undermine the idea that when one knows a priori that ϕ(the F), and the relevant occurrence of 'the F' is exportable for belief and independent of one's mental life, one normally knows a priori that $\phi(x)$. For in normal cases where one knows a priori that ϕ(the F), one's belief that ϕ(the F) *is* permissible given one's circumstances, in which case one's belief that $\phi(x)$ is also permissible given one's circumstances; and in view of the considerations to be considered in the next two sections, once one has granted this much it will be hard to resist the conclusion that one's belief that $\phi(x)$ constitutes a priori knowledge.

9. From justification to a priori justification

Let us now turn to

> *Generalization* 2 (relativized): When an occurrence of a definite description is exportable for justification in a context, it is almost always exportable for a priori justification in that context

I take it that one has a priori justification to believe a proposition iff one has justification to believe it, and this fact is independent of facts about one's experience. Granted, there are some well-known difficulties in understanding this definition so as to fit even paradigm cases of a priori justification. First, we will need to understand 'experience' in some non-obvious way if we want to avoid classifying as a priori introspective beliefs about one's own mental states, and for that matter ordinary beliefs retained in memory without a memory of the experiences on which they were initially based. Second, the fact that certain experiences are necessary for one even to be able to *entertain* a proposition — say, that everything red is red — should not prevent one from having a priori justification to believe it, even though, arguably, one can only have a priori justification to believe propositions one can entertain. These are delicate issues. Let us see how far we can get in making Generalization 2 plausible without having to resolve them.

To begin with, let us restrict our attention to descriptions which do not concern the subject's experience. Suppose that 'the F' is such a description, and that the occurrence of 'the F' in 'ϕ(the F)' is exportable for justification: necessarily, whenever x is the F and one has justification to believe that ϕ(the F), one has justification to believe that $\phi(x)$. Suppose further that one has a priori justification to believe that ϕ(the F) and that x is the F. Then the fact that one has justification to believe that $\phi(x)$ is entailed by two facts — that one has justification to believe that ϕ(the F), and that x is the F — each of which obtains independently of one's experience. If we could assume that the relevant notion of independence is closed under entailment (i.e. that whatever follows from facts independent of one's experience is itself independent of one's experience), we could conclude that one has a priori justification to believe that $\phi(x)$.

Not all notions of independence are closed under entailment. For example, although the fact that it is either raining or not raining is a consequence of the fact that $1 + 1 = 2$, we might want to claim that in a certain sense, the latter but not the former is independent of the fact that it is raining. We could attempt to understand the notion of the a priori in terms of some such hyperintensional notion of independence, thereby leaving room for the possibility that one lacks a priori justification to believe a certain proposition despite the fact that one would have had justification to believe it no matter what one's experience had been like. Do we really want to recognize this possibility? Here is one reason for not doing so: if, no matter what experiences one

had, one could acquire (doxastically) justified belief that P by attending to one's experiences and coming to believe P on the basis of them, then one should be able to acquire justified belief that P equally well by a process that cuts out the pointless step of attending to one's experiences. Admittedly, there are various reasons to resist this line of thought. Most obviously, one might resist it so as to leave room for a restrictive doctrine about metaphysical possibility that placed surprising a posteriori limits on the range of possible evidential states. But these considerations are unlikely to bear interestingly on Generalization 2. If one's imperfect grasp of the limits on metaphysically possible evidential states did not prevent one's justification to believe that ϕ(the F) from counting as a priori, why would it do so in the case of one's justification to believe that $\phi(x)$?

Even if we do understand 'a priori' in terms of a hyperintensional notion of independence, we should recognize that there is an interesting distinction to be drawn between propositions one has justification to believe because of *contingent* features of one's experience, and propositions which one would have justification to believe no matter what one's experiences had been like. It is *prima facie* puzzling how propositions which could easily have been false could belong to the latter category; so a strategy for explaining why some do will be interesting, whether or not the explanandum is properly expressed using 'a priori'.

An interpretation of Generalization 2 on which 'almost always' is understood as a restriction to descriptions that do not concern the subject's experience would already be strong enough to vindicate the central claims of this paper. Still, it is interesting to consider to what extent we can expect this conclusion to carry over to cases where the fact that x is the F is not independent of the subject's experience, especially since many descriptions of this sort — 'the object that looks to one to be a stick', for example — seem especially apt to have exportable occurrences in ordinary contexts.[30]

Given our experience with Generalization 1, it will be no surprise to find that certain occurrences that are exportable for justification just as a matter of logic are exceptions to Generalization 2. Consider the

[30] Also, I anticipate that many of those who resist my attempts to argue that exportability for belief is a common phenomenon will, like Salmon (1986, p. 180), make an exception for some descriptions that concern the subject's experiences. Thanks to Scott Soames and Benj Hellie for pressing me to think about these cases.

occurrence of 'the shortest spy whom one has justification to believe to be a spy' in

(26) Either there are no spies whom one has justification to believe to be spies, or two spies are the same height, or the shortest spy whom one has justification to believe to be a spy is a spy

This is exportable for justification in every context, since

(27) It is metaphysically necessary that if one has justification to believe that (either there are no spies whom one has justification to believe to be spies, or two spies are the same height, or the shortest spy whom one has justification to believe to be a spy is a spy), and *x* is the shortest spy whom one has justification to believe to be a spy, then one has justification to believe that (either there are no spies whom one has justification to believe to be spies, or two spies are the same height, or *x* is a spy)

has to be true: if one has justification to believe the proposition that *x* is a spy, one has justification to believe any disjunction with this proposition as a disjunct. There is no particular reason to think that (27) has to remain true when we insert 'a priori' in front of the two underlined occurrences of 'justification', which is what would have to be the case for the occurrence in question to be exportable for a priori justification.

But these degenerate cases are obviously quite special. If we want to get a sense of how widely we can expect Generalization 2 to fail for descriptions that are not independent of the subject's experience, we should focus on a more typical case: 'the object that looks to one to be a stick', say. Let us anchor ourselves in a favourable context by stipulating that all occurrences of 'the object that looks to one to be a stick' that have widest scope, or scope within the consequent of a conditional, are exportable for belief and for justification. Suppose that while you are walking in the woods, you notice a perfectly ordinary stick *x*. *x* looks to you to be a stick. You have justification to believe both (A1) that the object that looks to you to be a stick looks to you to be a stick, and (A2) that the object that looks to you to be a stick looks to you to be a stick if any unique thing does. It follows from our stipulation about exportability that you also have justification to believe (B1) that *x* looks to you to be a stick, and (B2) that *x* looks to you

to be a stick if any unique thing does. Your justification to believe (A2) is a priori. Is your justification to believe (B2) a priori too? Here is an argument that it is:

(i) Either you have a priori justification to believe (B2), or your justification to believe (B2) is based on your empirical justification to believe some other propositions, or you have immediate empirical justification to believe (B2)

(ii) Your justification to believe (B2) is not based on your empirical justification to believe any other propositions

(iii) You do not have immediate empirical justification to believe (B2)

(iv) So you have a priori justification to believe (B2)

Understand 'immediate' and 'based on' in such a way as to render (i) trivially true. In support of (ii): among the propositions you have empirical justification to believe, the best candidate to be the basis for your justification to believe (B2) is (B1), the proposition that *x* looks to you to be a stick. But your justification to believe (B2) cannot simply be based on your justification to believe (B1). For, first, you can justifiably have a higher degree of confidence in (B2) than in (B1). To see this, notice that whereas we are easily gripped by attempts to raise sceptical worries involving scenarios in which (B1) fails ('I believe that *x* looks to be a stick — but how can I rule out the possibility that my visual experiences are all hallucinations produced by an evil demon, so that neither *x* nor anything else looks to me to be a stick?'), we find analogous attempts to raise sceptical worries involving scenarios in which (B2) fails completely ungripping ('I believe that *x* looks to me to be a stick if any unique thing does — but how can I rule out the possibility that it is some other object, distinct from *x*, that uniquely looks to me to be a stick?').[31] Second, given our stipulation about exportability, you will still have justification to believe (B2) even if you have lots of misleading evidence — such as evidence that you are hallucinating — that undermines your justification to believe (B1) while leaving your a priori justification to believe (A2) untouched. Thus, your justification to believe (B2) is better than it would be if it were based merely on inference from (B1).

[31] Such worries of course do become gripping when you have some independent way of thinking about *x* that does not involve its relation to your current experiences.

Are there any other propositions you have empirical justification to believe on which your justification to believe (B2) could be based? It might be suggested that it is based on your justification to believe propositions *about your beliefs*, such as the proposition that you believe that *x* looks to you to be a stick.[32] There are lines of argument here that are not obviously hopeless: 'I believe that *x* looks to me to be a stick; but if something other than *x* looked to me to be a stick, I would not be in a position to believe this or anything else concerning *x*; so it must be the case that *x* looks to me to be a stick, if any unique thing does.' But here again, I think we will find that your justification to believe (B2) persists even in the face of evidence that undermines your introspective justification for these self-ascriptions. If your therapist convinces you that you are systematically mistaken about your beliefs, your justification to believe (A2) will not be undermined. Given our stipulation that the relevant occurrence of 'the object that looks to one to be a stick' is exportable for justification, your justification to believe (B2) will not be undermined either. Moreover, if you somehow manage (through inattention or manipulation) actually not to believe any of the relevant *de re* propositions about your beliefs, this will not prevent your having justification to believe (B2).

This is not an exhaustive survey; but these results suggest that attempts to identify other empirically justified propositions as the basis for your justification for (B2) will founder in the same way: (B2) is much better justified than anything empirical on which its justification could be based. This leaves those who would resist the conclusion that you have a priori justification to believe (B2) with the option of rejecting (iii). But the idea that you have *immediate* empirical justification to believe (B2) seems implausible. When we think of our experience as giving us direct awareness of (or putative awareness of) certain facts about the world, we normally think of the facts in question as reasonably non-disjunctive — that *x* is thus and so, or that *x* looks thus and so, or that our sense-data are thus and so, or that we sense thus-ly and so-ly … The idea that we could be directly aware in experience of a conditional like (B2) would require a radical rethinking of the notion of direct awareness, unmooring it in an unsettling way from its anchoring in phenomenology. Such a rethinking might be warranted if it allowed us to maintain some sort of general prohibition on a priori justification to believe (certain kinds of)

[32] Recall that we are working with an expansive conception of 'empirical' justification that includes introspection.

contingent *de re* propositions. But we have already seen good reason to expect that such a priori justification will be commonplace, occurring whenever descriptions that do not concern the subject's experience are exportable for justification. Given this, it is hard to see any principled grounds for resisting the claim that one has a priori justification to believe (B2).

These considerations apply equally well to many other descriptions having to do with the subject's experience, like 'the object one sees to be a stick', 'the stick that one sees', and 'the stick veridically presented to one by visual perception *v*' (cf. Salmon 1986, p. 180). In general, if *x* is the F and the proposition that ϕ(the F) is the sort of proposition one might have a priori justification to believe, the proposition that $\phi(x)$ will not be the sort of proposition for which one could plausibly be taken to have *immediate* empirical justification, and it will be hard to imagine a story about *indirect* empirical justification that could explain why the basis for such a justification will always be present whenever one's a priori justification to believe that ϕ(the F) is. Degenerate cases like (27) will of course be exceptions to this rule. But nothing about the way they manage this suggests any worry about Generalization 2 that extends to any other kind of case.

10. From a priori justification to a priori knowledge

Finally:

> *Generalization* 3 (relativized): When an occurrence of a description is exportable for a priori justification in a context, it is almost always exportable for a priori knowledge in that context

Generalization 3 would be easy to argue for if a priori knowledge were just true belief for which one has a priori justification. But it is not. For one thing, to have a priori knowledge, one must have an *a priori justified belief*, not just belief for which one has a priori justification — one's belief must, in some notoriously elusive sense, be 'based on' the a priori justification one has. For another thing, there may be a priori Gettier cases, where we have true, a priori justified belief without knowledge.

According to a currently popular diagnosis (Sainsbury 1997, Williamson 2000, Sosa 2002), Gettier cases show that knowledge requires *safety from error*. A natural first attempt to articulate this appeals to the notion of what could easily have been the case:

> (R) One knows that P only if one could not easily have falsely believed that P

This suggests a way in which Generalization 3 might fail. Suppose that the occurrence of 'the president of Iraq' in 'The president of Iraq is the head of state of Iraq, if Iraq has a president' is exportable for a priori justification, and that you know a priori that the president of Iraq is the head of state of Iraq if Iraq has a president. Since your knowing this a priori requires you to have a priori justification to believe it, you must also have a priori justification to believe the proposition, concerning Jalal Talabani (the president of Iraq), that he is the head of state of Iraq if Iraq has a president. Call this proposition J. In a world where J was false, because someone else was president of Iraq, merely believing that the president of Iraq is the head of state of Iraq if Iraq has a president would not suffice for believing J. But you might still have believed J under these circumstances, for example by believing, falsely, that the Secretary-General of the Patriotic Union of Kurdistan is the head of state of Iraq if Iraq has a president. If this counts as something that could 'easily' have happened, (R) will entail that you do not know J, and *a fortiori* that you do not know it a priori. If so, you are a counterexample to the claim that the occurrence of 'the president of Iraq' in 'The president of Iraq is the head of state of Iraq, if Iraq has a president' is exportable for a priori knowledge.

But it is well known that (R) is too demanding. One might luckily acquire conclusive evidence which enables one to know that P, even if one also has, or could have had, some bad evidence that P on the basis of which one could easily have ended up believing falsely that P. (R) must be weakened somehow to allow for this possibility. The canonical fix for this kind of problem is to restrict attention to beliefs formed using the same 'method' or 'grounds' as one's actual belief:

(R*) One knows that P only if one has some grounds for believing that P such that one could not easily have falsely believed that P on those grounds[33]

But any weakening along these lines will no longer pose any obvious threat to Generalization 3. Suppose that one has a priori justification to believe J, as a necessary consequence of the facts that one has a priori justification to believe that the president of Iraq is the head of state of Iraq if Iraq has a president and that Jalal Talabani is the president of Iraq. Then, plausibly, one has grounds for believing

[33] Nozick (1981, p. 179) appeals to 'methods' in a closely related context. Sosa (2002) and Williamson (2000, p. 149) both suggest versions of the safety condition roughly equivalent to (R*). Comesaña (2005) argues that even (R*) is too strong.

J that one would lack under circumstances where one managed to believe J while someone else was the president of Iraq. Of course, one might *also* have some bad grounds for believing J, on the basis of which one could easily have believed it even if it were false. But a belief that is overdetermined by good and bad grounds in this way can still constitute a priori knowledge.

The lesson can be generalized. Suppose that one knows a priori on the basis of grounds G that ϕ(the F), and that necessarily, if one has a priori justification to believe that ϕ(the F) while x is the F, one has a priori justification to believe that $\phi(x)$. If one is in danger of believing falsely that $\phi(x)$, it must be because of some quite separate relation one bears to x, in addition to the relation of believing that ϕ(the F) while it is the F. It is hard to imagine a reasonable way of individuating grounds on which one's possible false belief that $\phi(x)$ derived from this other relation could count as being based on *all* the same grounds as one's actual true belief that $\phi(x)$.[34] Plausibly, your actual belief that $\phi(x)$ is based on at least one ground G′ which one can have only in a world in which x is the F and one has ground G for believing that ϕ(the F). If so, since one could not easily have had G without its being the case that ϕ(the F), one could not easily have had G′ without its being the case that $\phi(x)$. So (R*) places no obstacle in the way of one's knowing that $\phi(x)$.[35]

[34] We should make an exception to this for degenerate cases of exportability for a priori justification, analogous to the degenerate cases of exportability for belief and justification discussed in Sects 8 and 9. The only examples of this phenomenon I have been able to come up with are quite intricate. Say that x is *conditionally suspect* iff one has a priori justification to believe that if x exists, x is a spy. Consider the occurrence of the description 'the shortest conditionally suspect spy' in the open sentence 'If some spy is conditionally suspect, and no two spies are the same height, and x exists, then the shortest conditionally suspect spy is a spy'. This occurrence must be exportable for a priori justification, since if x is the shortest conditionally suspect spy, one has a priori justification to believe that if x exists x is a spy, and *a fortiori* has a priori justification to believe that if some spy is conditionally suspect and no two spies are the same height and x exists, x is a spy. But there is no reason to regard it as exportable for a priori knowledge.

[35] Sainsbury (1997) and Williamson (2000) conceive of the safety condition as attaching in the first instance to particular states or episodes of believing. Their idea is that one knows that P only if one is the subject of an episode of believing that P which could not easily have been an episode of believing something false. Since episodes of believing are not supposed to have their contents essentially — an episode which is in fact an episode of believing that P could have been an episode of believing some other proposition — this proposal is in some ways stronger, and in other ways weaker, than (R). As far as I can see, the remarks in the text about the 'same grounds' version of safety will apply equally well to the Sainsbury/Williamson proposal.

I have not given a positive argument for Generalization 3, only attempted to rebut an argument against it; it is hard to see how one could get much further without relying on some full-fledged analysis of knowledge. But given the intimate connections between a priori justification and a priori knowledge, there is enough of a presumption in favour of Generalization 3 to make it reasonable to expect those who would reject it to offer some sort of explanation of its failure. Appealing to the connection between knowledge and notions like reliability and safety seems by far the most promising strategy for providing such an explanation. If that strategy fails, we can reasonably regard it as unlikely that there is anything else about the nature of knowledge that could explain the existence of widespread exceptions to Generalization 3.

11. Conclusion: a priori knowledge and proper names

I have been making a case that the phenomenon of exportability is a fairly unitary one. Relative to a given context, if any widest-scope occurrence of a given description is exportable for any of the four propositional attitudes I have been considering (belief, justification, a priori justification, and a priori knowledge), then — usually — all widest-scope occurrences of that description, and all occurrences that take wide scope within the consequents of conditionals, are exportable for all four attitudes. I have tentatively suggested that the only exceptions to this are very special degenerate cases. If one also accepts that exportability is a widespread phenomenon, as argued in section 5, it will follow that a priori knowledge of *de re* propositions — including contingent propositions, and even propositions which could very easily be false — is similarly widespread. My hope is that this will also make such knowledge seem unmysterious.

My defence of Generalizations 1–3 bears little resemblance to a watertight argument. Most obviously, I have offered no explanation of the difference between the degenerate cases of exportability which are exceptions to the generalizations and the ordinary cases which are not, relying instead on the bare intuition that whatever explains the failure in the degenerate cases is unlikely to generalize very far. To explain what the difference amounts to, one would need to provide a metaphysical analysis of the propositional attitudes I have been concerned with — a theory that would answer the question 'What is it for a subject to bear this attitude to a given proposition?', and thus put

one in a position to say *why* a given occurrence of a given description is exportable for a given attitude. The project of providing such analyses is beyond the scope of this paper. I will be happy if I can clear the way for the project, by showing how certain claims about the epistemology of *de re* propositions that might seem objectionable if they were first encountered as consequences of some controversial analysis of the attitudes can be made plausible by appeal to premisses which do not depend on any particular analysis.

In the remainder of this final section I will briefly consider how much of what I have said about attributions of a priori knowledge with embedded externally bound variables carries over to attributions of a priori knowledge reports with embedded proper names. As I mentioned in section 2, my own favoured view is that attitude reports involving referring proper names are *de re*, in the sense that instances of the following schema are valid (true on all uniform interpretations):

> *Transparency*: For any x and y, if x is N, then: y ψs that $\phi(x)$ iff y ψs that ϕ(N)

(Here ψ represents an attitude verb and N a proper name). I would want to diagnose our negative reactions to sentences like

> (28) Although Superman is Clark Kent, Lois believes that Superman flies and does not believe that Clark Kent flies

as arising from the fact that we find their (necessarily false) uniform interpretations less natural than some of their (possibly true) non-uniform interpretations. This view of Transparency is defended in Dorr forthcoming. If I am right about this, our conclusions about ascriptions of *de re* a priori knowledge will carry over without further ado to ascriptions of a priori knowledge using referring proper names.

But this is not the place to argue for Transparency. So suppose it is false. Then there is no straightforward route from ⌜N is something that is knowable a priori to be F⌝ to ⌜it is knowable a priori that N is F⌝. However, our arguments about *de re* attitude ascriptions can be adapted without much fuss to attitude ascriptions involving embedded proper names. We can begin by defining a relativized notion of exportability:

> When O is an occurrence of a definite description ⌜the F⌝ in a sentence or open sentence S, S^N is the result of replacing O in S with the proper name N, and ψ is a propositional attitude verb: O is **N-exportable for ψing** iff the sentence ⌜It is metaphysically

necessary that whenever one ψs that S, and N is the F, one ψs that $S^{\ulcorner N \urcorner}$ is true

As far as I can see, the case for Generalizations 1–3 is not adversely affected when we replace all claims about exportability with claims about N-exportability. N-exportability for belief almost always suffices for N-exportability for justification (Generalization 1), which almost always suffices for N-exportability for a priori justification (Generalization 2), which almost always suffices for N-exportability for knowledge (Generalization 3). So if we want to establish that it is possible to know a priori that N is uniquely F if anything is, it will usually be enough to show that the occurrence of 'the F' in 'the F is uniquely F if anything is' is N-exportable for belief.

Even setting Transparency aside, there is good reason to think that attitude ascriptions involving embedded proper names are context-sensitive, in a way that carries over to claims about N-exportability. Consider:

(29) When she first saw him, Becky Sharp thought that Sir Pitt Crawley was a servant

The contexts in which (29) is true (in the fiction of *Vanity Fair*) plausibly include at least some on which widest-scope occurrences of descriptions like 'the person one is currently seeing' or 'the person one has just asked to carry one's luggage' are 'Sir Pitt Crawley'-exportable for belief. By contrast, these occurrences plainly are not 'Sir Pitt Crawley'-exportable for belief in contexts where (30) is true:

(30) Becky Sharp had no idea that she had just ordered Sir Pitt Crawley to carry her luggage

This argument is complicated by the fact that opponents of Transparency are liable to think that occurrences of names in attitude reports give rise to the same kind of scope ambiguities as occurrences of quantifiers. They will thus be able to attribute some putative cases of context-sensitivity in attitude reports involving names to this structural ambiguity. For example, it might be said that to interpret (29) as expressing a truth, we must assign it a structure in which 'Sir Pitt Crawley' takes scope outside the attitude verb, making (29) equivalent to a *de re* report:

(31) When she first saw him, Sir Pitt Crawley was such that Becky Sharp believed him to be a servant

One could attempt to use this move to avoid positing context-sensitivity in belief-reports where names take narrow scope. But this strikes me as an unpromising strategy. I will not try to do so here; but by accumulating examples, I think one could show that context-sensitivity of attitude reports involving names is quite fine-grained: for any given proposition P and any given attitude report S, either S is not used to assert P on the vast majority of occasions when it is uttered literally, or the negation of S is not used to assert the negation of P on the vast majority of occasions when *it* is uttered literally. If so, we could avoid positing context-sensitivity in narrow-scope readings of attitude reports with embedded names only at the cost of claiming that these readings are almost never the ones we intend.

In any case, once we recognize context-sensitivity in *de re* attitude reports, it seems arbitrary to go to such lengths to resist positing it in attitude reports involving names. And the considerations that support the thought that there are no contexts in which there are no non-trivial examples of exportability for belief also support the corresponding thought about N-exportability. The more details we load into the description 'the F', the harder it becomes to evoke a context where it would be inappropriate to apply the predicate 'believes that N is G' to someone who believes that the F is G while N is the F.

Occurrences of certain metalinguistic descriptions involving the name N are especially easy to hear as N-exportable. For example, many occurrences of 'the person to whom the name "Sir Pitt Crawley" refers', or 'the person one has heard of under the name "Sir Pitt Crawley"' (cf. Kroon 1987) seem to be 'Sir Pitt Crawley'-exportable for belief even in the relatively demanding contexts in which (30) is true. In contexts where occurrences like these are exportable, sentences like 'We know a priori that if the name "Sir Pitt Crawley" refers to a single person, it refers to Sir Pitt Crawley' will be true. There is no need to posit a special capacity for a priori knowledge about our own language, or a special role for linguistic knowledge in the definition of 'a priori', to explain the possibility of such a priori knowledge. However, the tendency of these metalinguistic descriptions to be N-exportable is only a defeasible presumption. If I tell you that I only recently learnt that Jones was the person I had been hearing of under the name 'Jones' all my life, you probably will not have much trouble finding a reasonable interpretation of my words.

The descriptions which have N-exportable occurrences in a given context may be quite a varied lot. In Kripke's celebrated Paderewski

example (Kripke 1979), (32) is a fairly natural thing to assert:

(32) Peter simultaneously believes that Paderewski has musical talent and that Paderewski lacks musical talent

I see no good reason to deny that (32) admits true uniform interpretations. If it does, then it is plausible that in the relevant contexts, both the descriptions 'the statesman named "Paderewski"' and 'the famous pianist named "Paderewski"' have a wide range of occurrences that are 'Paderewski'-exportable for all the attitudes we have been concerned with. So in these contexts, (33) is true:

(33) It is possible to know a priori both that if there is a unique statesman named 'Paderewski' it is Paderewski, and that if there is a unique famous pianist named 'Paderewski' it is Paderewski

On the other hand, (34) is plainly not true even in these contexts:

(34) It is possible to know a priori that if there is a unique statesman named 'Paderewski' and a unique famous pianist named 'Paderewski', they are the same person

The set of propositions that can be known a priori thus fails to be closed under logical consequence. Relative to this context at least, the two-dimensionalist project of assigning 'primary intensions' to sentences in such a way that all and only the a priori sentences are assigned necessary primary intensions is doomed to fail, since the set of sentences with necessary primary intensions is bound to be closed under logical consequence.[36] At least, the project is doomed to fail if we understand 'The sentence "P" is a priori' as equivalent to 'It is possible for someone to know a priori that P'. If this equivalence is rejected — as it is in most of the options taken seriously by Chalmers (2011, Sect. 9.i) — then whatever it means to call a sentence 'a priori', the question which sentences have this status is not part of the topic of this paper.

Given the prominent role they play in the best-known argument for the possibility of contingent a priori knowledge (Kripke 1972), names that are introduced by means of reference-fixing descriptions are worthy of special attention. Consider 'Bellerophon', which has been introduced (unofficially) as a name for the planet responsible for certain periodic fluctuations in the spectrum of the star 51 Pegasi.

[36] For an introduction and further references, see Chalmers 2006.

It would be strange to deny that astronomers believe many things about Bellerophon — for example, that Bellerophon is about half as massive as Jupiter. And given that we accept such attributions, there is pressure to treat various occurrences of 'the planet responsible for the fluctuations in the spectrum of 51 Pegasi', such as the one in 'The planet responsible for the fluctuations in the spectrum of 51 Pegasi is about half as massive as Jupiter', as 'Bellerophon'-exportable for belief.[37] For intuitively, finding out that Bellerophon exists was a fairly significant event in the history of astronomy. It is more plausible to date this event to the time when astronomers first found out that there was a planet responsible for the fluctuations, rather than the perhaps later time when they introduced a name for the planet, or opened a mental file for it, or did any of the other things one might be tempted to regard as necessary for believing that Bellerophon exists.[38]

Kripke famously used cases like this to argue for the possibility of contingent a priori knowledge. If we accept the appearance that relevant occurrences of 'the planet responsible for the fluctuations in the spectrum of 51 Pegasi' are 'Bellerophon'-exportable for belief, we should agree with him, in the light of Generalizations 1–3 (or to be precise, in the light of their analogues for 'Bellerophon'-exportability). We can in fact have a priori knowledge of the contingent astronomical proposition — call it B — that Bellerophon is responsible for the fluctuations if any unique planet is.

Kripke's argument has been widely criticized. A recurrent worry, first raised by Kripke himself (1972, n. 26), is that if we say that there is a priori knowledge in these cases, we will have to attribute extraordinary knowledge-extending powers to apparently trivial linguistic ceremonies. Here is a representative expression of the worry, by Donnellan:

> If a truth is a contingent one then it is made true, so to speak, by some actual state of affairs in the world that, at least in the sort of examples we

[37] I see no special reason to suppose that belief that the planet responsible for the fluctuations is F is *necessary* as well as sufficient (given the astronomical facts) for belief that Bellerophon is F. 'We only recently learnt that Bellerophon orbits very close to its star, but if there are intelligent beings in the 51 Pegasi system, they have probably known this as long as they have known the rudiments of astronomy' sounds like a perfectly fine speech, requiring none of the stretching characteristic of non-uniform interpretation.

[38] Similar considerations push us towards treating occurrences of 'the planet responsible for perturbing the orbit of Uranus' (Kripke 1972, n. 33) as 'Neptune'-exportable for belief, despite the fact that the current use of 'Neptune' involves no special connection to that description. But in this case we are pulled in the opposite direction by the fact that it sounds so odd to say that we are *now* in a position to know a priori that Neptune is responsible for perturbing the orbit of Uranus if any unique planet is. More grist for the contextualist mill!

> are interested in, exists independently of our language and linguistic conventions. How can we become aware of such a truth, come to know the existence of such a state of affairs, merely by performing an act of linguistic stipulation? (Donnellan 1979, p. 13)[39]

Clearly our present grounds for claiming that B is known a priori do not require us to attribute any special epistemic role to acts of linguistic stipulation. If the relevant occurrence of 'the planet responsible for the fluctuations in the spectrum of 51 Pegasi' is 'Bellerophon'-exportable for a priori knowledge, then we knew B a priori before we introduced the name.

Is there *any* context in which we could truly say 'We extended our knowledge of astronomy by introducing the name "Bellerophon"'? I believe so. Suppose that, although Fred realizes that the fluctuations in the spectrum of 51 Pegasi are caused by a planet, he mistakenly thinks that 'Bellerophon' was introduced as another name for Pluto. It is quite natural to describe Fred as believing that Bellerophon orbits the Sun, and as *not* believing that Bellerophon is responsible for the fluctuations in the spectrum of 51 Pegasi. Once we start thinking about people like Fred, we evoke a different class of contexts, in which occurrences of 'the planet responsible for the fluctuations' are not 'Bellerophon'-exportable. By the standards of these contexts, our coming to know that Bellerophon was responsible for the fluctuations was not such a noteworthy achievement — presumably we did not know this until we had introduced the name, and would not have come to know it if we had picked a completely different name. So the linguistic ceremony *did* help us extend our astronomical knowledge. On the other hand, it is not clear that 'We know a priori that Bellerophon is responsible for the fluctuations if any unique planet is' is true relative to these contexts. For Fred to come to believe that Bellerophon is about half as massive as Jupiter, he would have to come to believe the relevant truths about the use of the name 'Bellerophon': it would not be enough if he merely came to suspend judgement as between these truths and his former false beliefs about the use of the name. If so, then what seems to be crucial about the name-introducing ceremony is that it gave *us* true beliefs about our use of the name.[40]

[39] Salmon (1987), Soames (2003, p. 411), and many others express similar concerns. See Jeshion 2001 for citations.

[40] Given what we have said about Fred, it is quite plausible that we could have believed propositions like B even before performing the ceremony, if we knew enough about our habits to predict that we would choose 'Bellerophon' rather than any other name.

If these sociological beliefs are justified only empirically, and are required given the rest of our circumstances for us to count as believing B, there is no reason to count B as justified a priori, rather than as based on the same empirical grounds as the sociological beliefs.

When Kripke's opponents worry about the generation of astronomical knowledge by linguistic ceremonies, they typically have in mind a different picture, on which what was required for us to believe and know propositions like B was merely that we introduce *some name or other* for the planet responsible for the fluctuations. This is stranger than the picture we ended up with by focusing on people like Fred. It is odd, for example, to suppose that we can deliberately *delay* acquiring knowledge just by postponing introducing a name, while knowing exactly how the name-introducing ceremony will proceed when we get around to it. This oddness persists even if we relax the view to allow that the same knowledge can also be attained by purely internal operations, like introducing a new name in the language of thought, or opening a new mental file, assuming that these are operations one could in principle deliberately refrain from performing.[41] I doubt that there are legitimate contexts that count such operations as *necessary* for believing B: there will always be some descriptive proposition, perhaps very rich and detailed, belief in which is sufficient (given the

[41] Hawthorne and Manley (2012, Ch. 2) try to make the idea that name-introductions could bring such epistemic advantages seem less wild by citing other cases where they think the introduction of new vocabulary broadens the range of propositions one is in a position to believe and know. They give the example of introducing 'and' into a language that previously had only 'or' and 'not', by laying down the standard introduction and elimination rules. I do not find their case convincing. Consider the epistemic effects of stipulatively introducing a new *predicate*. Let the meaning of 'bloog' be such as to render 'Necessarily, all and only blue spheres are bloogs' true. Did I just come to know that all bloogs are blue? My first impulse is to say no: I have known this as long as I have known that all blue spheres are blue. I can get myself into a mood where I will answer 'Yes' by imagining someone like Fred, who, say, falsely believes that 'bloog' means 'red sphere', and thinking about the naturalness of describing Fred as 'failing to believe that all bloogs are blue'. But once I am in this mood, it is not so tempting to say that I would still have known that all bloogs are blue if I had chosen some other word, and it *is* tempting to say that my epistemic advantage over Fred is based on my empirical evidence that 'bloog' was introduced as it actually was. I see no relevant difference between this and introducing a new connective by saying something like 'Let all sentences of the form "Necessarily, $P \uparrow Q$ iff not-P and not-Q" be true', or between this procedure and introducing a connective by laying down introduction and elimination rules. Perhaps things work differently when there are sentences involving the new vocabulary (not in the scope of attitude verbs or similar contexts) that cannot be shown, using the stipulated rules and axioms, to have the same truth value as any sentence in the old vocabulary: e.g. the introduction of second-order quantifiers and variables into a previously first-order language. But this is not the case in Hawthorne and Manley's example, or when we introduce names by descriptive reference-fixing.

astronomical facts) for belief in B. Nothing I have said so far, however, rules out the possibility that operations like the introduction of a name can make the difference *in some particular cases* between belief in B and its absence. For example, there may be a context in which the predicate 'believes that Bellerophon is F' does not apply to *everyone* who believes that the planet responsible for the fluctuations is F, but does apply to all of those who in addition possess names that were introduced using the description 'the planet responsible for the fluctuations', or deploy mental files containing that description.

If there are contexts like this, is it true in them that the belief in B that we acquire by introducing a new name or file constitutes a priori knowledge? We can argue that it does, by appealing to considerations similar to those we appealed to in arguing for Generalizations 1–3. First, the mere act of introducing a name or file associated with some description is not the sort of thing for which it would make sense to demand *epistemic* justification. It might have been frivolous, unwise, a waste of energy, for astronomers to bother introducing a name for the planet responsible for the fluctuations of 51 Pegasi; but if they thereby came to believe B, this belief was not on that account epistemically unjustified. Thus from the standpoint of epistemic justification, facts about one's supply of names and files seem to be just as much a part of one's 'circumstances' as astronomical facts; so by the same reasoning I appealed to in arguing for Generalization 1, the astronomers must have had justification to believe B. Moreover, since the fact that one has introduced a name seems to be independent of one's experiences in the relevant sense, this justification was a priori in character. And finally, it is no easier to see in this case than in the case of an exportable description how the a priori justified belief in B could fail to constitute a priori knowledge.

We could take this as showing that there is after all nothing mysterious or problematic about acquiring a priori knowledge of astronomical truths just by introducing new names. Or — if we are impressed by the idea that name-introductions are devoid of epistemic value — we could take it as a reason to think that name-introductions can never make a difference to the range of propositions one counts as believing. I will not attempt to adjudicate between these two possible reactions.[42]

[42] Thanks to John Hawthorne, Benj Hellie, David Manley, Michael Nelson, Kieran Setiya, and Jessica Wilson; to the editor and referees for *Mind*, who have subsequently been identified to me as Berit Brogaard and Carrie Jenkins; and to audiences at Cornell and USC. Special

References

Almog, Joseph, John Perry, and Howard Wettstein (eds) 1989: *Themes from Kaplan*. Oxford: Oxford University Press.

Alston, William P. 1988: 'The Deontic Conception of Epistemic Justification'. In Tomberlin 1988, pp. 257–99.

Bach, Kent 1994: 'Conversational Impliciture'. *Mind and Language*, 9, pp. 124–62.

——1997: 'Do Belief Reports Report Beliefs?' *Pacific Philosophical Quarterly*, 78, pp. 215–41.

Braun, David 1988: 'Understanding Belief Reports'. *Philosophical Review*, 107, pp. 555–95.

Burge, Tyler 1977: 'Belief *De Re*'. *Journal of Philosophy*, 74, pp. 338–62.

Cappelen, Herman and Ernest Lepore 2005: *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism*. Malden, MA: Blackwell.

Carston, Robyn 2002: *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Oxford: Blackwell.

Chalmers, David 2006: 'Two-dimensional Semantics'. In Lepore and Smith 2006, pp. 574–606.

——2011: 'Propositions and Attitude Ascriptions: A Fregean Account'. *Noûs*, 45, pp. 595–639.

Cohen, Stewart 1988: 'How to be a Fallibilist'. In Tomberlin 1988, pp. 91–123.

Comesaña, Juan 2005: 'Unsafe Knowledge'. *Synthese*, 146, pp. 393–402.

Crimmins, Mark 1992: *Talk About Beliefs*. Cambridge, MA: MIT Press.

Crimmins, Mark and John Perry 1989: 'The Prince and the Phone Booth: Reporting Puzzling Beliefs'. *Journal of Philosophy*, 86, pp. 685–711.

Davidson, Donald and Jaakko Hintikka 1969: *Words and Objections: Essays on the Philosophy of W. V. Quine*. Dordrecht: D. Reidel.

DeRose, Keith 1995: 'Solving the Skeptical Problem'. *Philosophical Review*, 104, pp. 1–52.

Donnellan, Keith S. 1979: 'The Contingent *A Priori* and Rigid Designations'. In French, Uehling, and Wettstein 1979, pp. 13–27.

thanks to Scott Soames, for advice and encouragement over the many years this paper has been in the works.

Dorr, Cian forthcoming: 'Transparency and the Context-Sensitivity of Attitude Reports'. In Garcia-Carpintero and Martí forthcoming.
French, Peter, Theodore E. Uehling, Jr and Howard K. Wettstein (eds) 1979: *Midwest Studies in Philosophy II: Contemporary Perspectives in the Philosophy of Language*. Minneapolis: University of Minnesota Press.
Garcia-Carpintero, Manuel and Genoveva Martí forthcoming: *Thinking and Talking about Nothing*. Oxford: Oxford University Press.
Grice, H. P. 1969: 'Vacuous Names'. In Davidson and Hintikka 1969, pp. 118–45.
Grimm, Robert and Daniel Merrill (eds) 1988: *Contents of Thought*. Tucson: University of Arizona Press.
Hawthorne, John 2004: *Knowledge and Lotteries*. Oxford: Oxford University Press.
Hawthorne, John and David Manley 2012: *The Reference Book*. Oxford: Oxford University Press.
Jeshion, Robin 2001: 'Donnellan on Neptune'. *Philosophy and Phenomenological Research*, 63, pp. 111–35.
Kaplan, David 1989: 'Demonstratives'. In Perry, Almog, and Wettstein 1989, pp. 481–566.
Keenen, E. L. (ed.) 1975: *Formal Semantics of Natural Language*. Cambridge: Cambridge University Press.
Kripke, Saul 1972: *Naming and Necessity*. Cambridge, MA: Harvard University Press. Revised edition 1980.
——1979: 'A Puzzle About Belief'. In Margalit 1979, pp. 239–83.
Kroon, Frederick 1987: 'Causal Descriptivism'. *Australasian Journal of Philosophy*, 65, pp. 1–17.
Lepore, Ernest and Barry C. Smith (eds) 2006: *The Oxford Handbook of Philosophy of Language*. Oxford: Oxford University Press.
Lewis, David 1975: 'Adverbs of Quantification'. In Keenen 1975, pp. 3–15.
——1979: 'Attitudes *De Dicto* and *De Se*'. *Philosophical Review*, 88, pp. 513–43.
——1996: 'Elusive Knowledge'. *Australasian Journal of Philosophy*, 74, pp. 549–67.
Loar, Brian 1988: 'Social Content and Psychological Content'. In Grimm and Merrill 1988, pp. 99–110.
Ludlow, Peter 1995: 'Logical Form and the Hidden Indexical Theory: A Reply to Schiffer'. *Journal of Philosophy*, 92, pp. 102–7.

Margalit, Avishai (ed.) 1979: *Meaning and Use*. Dordrecht: Reidel.
Moser, Paul (ed.) 2002: *The Oxford Handbook of Epistemology*. Oxford: Oxford University Press.
Mossner, E. C. 1954: *The Life of David Hume*, second edition (1980). Oxford: Oxford University Press.
Nozick, Robert 1981: *Philosophical Explanations*. Cambridge, MA: Harvard University Press.
Nunberg, Geoffrey 1993: 'Indexicality and Deixis'. *Linguistics and Philosophy*, 16, pp. 1–43.
Oppy, Graham 1992: 'A Semantics for Propositional Attitude Ascriptions'. *Philosophical Studies*, 67, pp. 1–18.
Recanati, François 1993: *Direct Reference: From Language to Thought*. Oxford: Blackwell.
Richard, Mark 1990: *Propositional Attitudes: An Essay on Thoughts and How We Ascribe Them*. Cambridge: Cambridge University Press.
——1993: 'Attitudes in Context'. *Linguistics and Philosophy*, 16, pp. 126–48.
Sainsbury, R. M. 1997: 'Easy Possibilities'. *Philosophy and Phenomenological Research*, 57, pp. 907–19.
Salmon, Nathan 1986: *Frege's Puzzle*. Cambridge, MA: MIT Press.
——1988: 'How To Measure the Standard Metre'. *Proceedings of the Aristotelian Society*, 88, pp. 193–217.
Schiffer, Stephen 1979: 'Naming and Knowing'. In French, Uehling, and Wettstein 1979, pp. 28–41.
——1996: 'Contextualist Solutions to Scepticism'. *Proceedings of the Aristotelian Society*, 96, pp. 317–33.
——2003: *The Things We Mean*. Oxford: Oxford University Press.
Soames, Scott 2003: *Philosophical Analysis in the Twentieth Century, vol. 2: The Age of Meaning*. Princeton: Princeton University Press.
——2005: 'Naming and Asserting'. In Szabó 2005, pp. 356–82.
Sosa, Ernest 1970: 'Propositional Attitudes *De Dicto* and *De Re*'. *Journal of Philosophy*, 67, pp. 883–96.
——2002: 'Tracking, Competence and Knowledge'. In Moser 2002, pp. 264–86.
Stalnaker, Robert C. 1988: 'Belief Attribution and Context'. In Grimm and Merrill 1988, pp. 140–56.
Stanley, Jason 2004: 'On the Linguistic Basis for Contextualism'. *Philosophical Studies*, 119, pp. 119–46.
——2005: *Knowledge and Practical Interests*. Oxford: Oxford University Press.

Strawson, P.F. 1974: *Subject and Predicate in Logic and Grammar*. London: Methuen.
Szabó, Zoltán Gendler (ed.) 2005: *Semantics Versus Pragmatics*. Oxford: Oxford University Press.
Tomberlin, James (ed.) 1988: *Philosophical Perspectives 2: Epistemology*. Atascadero: Ridgeview.
Williamson, Timothy 2000: *Knowledge and its Limits*. Oxford: Oxford University Press.

