

STRAWSONIAN EPISTEMOLOGY.
WHAT EPISTEMOLOGISTS CAN LEARN FROM
“FREEDOM AND RESENTMENT”

Marcus WILLASCHEK
Goethe-Universität Frankfurt am Main

Summary

This paper investigates the possibility of applying P. F. Strawson’s method in his seminal paper “Freedom and Resentment” (1962) to issues in epistemology. In that paper, Strawson sheds light on our concepts of responsibility and free will by looking at practices in which we employ these concepts not directly—by judging that someone is free or responsible—but rather indirectly by adopting “reactive attitudes” such as resentment and gratitude that presuppose that their targets are free and responsible in relevant respects. In a similar way, we may be able to shed light on the concept of knowledge by looking at aspects of our communicative practice such as assertions and their acceptance. In accepting as true what others assert, the paper argues, we typically attribute knowledge to the asserter. This connection between acceptance and knowledge-attribution can then be used to argue that the concept of knowledge implicitly employed in our communicative practice is externalist in spirit, but also includes internalist elements.

0. *Introduction*

With his seminal article “Freedom and Resentment” (Strawson 1974 [1962], henceforth *FR*), P. F. Strawson added a completely new dimension to the traditional debate about free will. One central move in that paper consists in addressing the free will issue *indirectly*, through a description and analysis of a practice in which freedom and responsibility are attributed implicitly, namely the practice of responding to other people’s actions with “reactive attitudes” such as resentment and gratitude. The aim of the present paper will be to investigate whether a parallel move can be helpful in epistemology. After a recapitulation of Strawson’s paper, in which I

will offer a novel reading of Strawson's overall argumentative strategy and highlight some often neglected aspects of his argument (section 1), I will look at an epistemic practice that might serve as a basis for a "Strawsonian" approach to epistemology, namely simple forms of everyday communication in which we *accept* as true what others assert (section 2). Next, I will argue that in typical cases of accepting as true what others say, we implicitly attribute knowledge to the asserter (section 3). Finally, I will outline what a Strawsonian approach might look like with respect to the debate between internalism and externalism in epistemology by offering a "Strawsonian" argument for (a generic version of) externalism (section 4). Obviously, my aim will not be to give a comprehensive defense of externalism, but only to indicate how an approach modeled on Strawson's strategy in "Freedom and Resentment" might be fruitfully applied in epistemology.

1. "Freedom and Resentment"

Strawson's starting point in "Freedom and Resentment" is the traditional debate between compatibilists (Strawson: "optimists") and incompatibilists ("pessimists") about freedom and determinism (cf. *FR*, 1). While compatibilists claim that our actions and decisions can be free and we can be responsible for them even if determinism is true, incompatibilists deny this. According to Strawson, this debate (as it had been conducted in the 20th century until the time of Strawson's essay) suffers from the fact that the traditional versions of compatibilism are committed to implausible consequentialist accounts of punishment and reward (cf. *FR*, 2ff.): If determinism were true, it may seem that the only rationale for punishment and reward could lie in influencing people's *future* behavior. But this would ignore the essentially backward-looking dimension of punishment and reward which is associated with the idea that punishing or rewarding someone is unfair unless they *deserve* this on the basis of what they have *done*. As incompatibilists have rightly pointed out, traditional compatibilism cannot account for this central aspect of our moral practice. But according to Strawson, incompatibilists in turn have overreacted by positing contra-causal freedom as the basis of moral desert and responsibility. Strawson's aim is to offer a modified version of compatibilism that takes non-consequentialist notions of moral desert and responsibility on board but that remains neutral with respect to the truth of determinism (cf. *FR*, 4). He tries to achieve this aim, however, not directly by analyzing, or

constructing, appropriate concepts of desert and responsibility, but rather indirectly, by giving a structural description of our practice of *attributing* responsibility. And the most fundamental, and most common, way in which we attribute responsibility is not verbally by *saying* of someone that they are responsible for some act or event, but rather by taking certain *attitudes* towards people, which Strawson calls “reactive attitudes” and of which resentment and gratitude are the paradigmatic cases (cf. *FR*, 5). In taking such an attitude towards some agent we *hold* them responsible for what they have done and thus *implicitly* attribute responsibility to them.

Now Strawson has often been understood as trying to provide an argument *for* compatibilism and thus *against* incompatibilism (e.g. Wallace 1994, 96ff.). I will return to this point below. But given Strawson’s official goal of *mediating* between compatibilists and incompatibilists,¹ I take it that his primary aim is not to refute incompatibilism, but rather to offer the compatibilist a non-consequentialist account of responsibility that is acceptable to incompatibilists. Thus, Strawson’s official strategy, as I understand it, must be to show two things: First, our practice of attributing responsibility is *not consequentialist in character*, since it is organized around values such as interpersonal goodwill and mutual esteem and reacts to the moral quality of *past* actions. And second, both the legitimacy of responsibility-attributions within this practice and the legitimacy of the practice as a whole are independent from the truth or falsity of determinism.

Accordingly, Strawson must first provide us with a plausible picture of the practice of attributing responsibility. Strawson achieves this by giving what might be called a “structural description” of one particular aspect of that practice, that is, a description that identifies some of its fundamental moves, viz. the reactive attitudes, and states the conditions under which we do and the conditions under which we don’t take these attitudes towards ourselves and others. The description will be successful if we, as competent participants of the practice, can recognize it as adequately depicting some of the fundamental moves and guiding principles of our communal practice.

Strawson intends the resulting description to represent what he thinks of as a part of human nature.² For this reason, his account of the reactive

1. Cf. Strawson’s remark that his argument “is intended as a move towards reconciliation” (*FR*, 1).

2. “[A]ll these types of attitude alike have common roots in our human nature and our membership of human communities” (*FR*, 16).

attitudes has often been taken to be purely descriptive. However, the conditions for adopting reactive attitudes reconstructed by Strawson are not just conditions under which we typically *do* adopt such attitudes, but also conditions under which it would be *appropriate* to do so, as Strawson himself emphasizes.³ Therefore, Strawson can fruitfully be read as offering a reconstruction of (some of) the *norms* that govern our practice of attributing responsibility by describing the conditions under which it is appropriate to adopt the reactive attitudes.

Strawson first traces our practice of reacting to harmful and beneficial actions with feelings of resentment, gratitude and other reactive attitudes (cf. *FR*, 5f.). We expect a certain minimal amount of goodwill and consideration from others. If their actions are harmful and thus seem to disclose a lack of goodwill, we feel resentment; if they show more goodwill than what may be expected, we feel gratitude. Restricting our discussion to harmful actions, these natural reactions will be, and ought to be, inhibited in two kinds of cases (cf. *FR*, 6–10): (i) the action, although harmful, was not done from a lack of goodwill, but accidentally, inadvertently, from ignorance etc. (so-called “excuses” in the post-*FR*-literature)⁴; (ii) the agent was “incapacitated” in such a way that he or she was unable to fully engage in normal interpersonal relations either temporarily (e.g. the agent is a child or she was out of her mind) or permanently, e.g. in cases of certain mental disorders (so-called “exemptions”). If someone is permanently incapacitated and therefore not an apt target of the reactive attitudes, we often take what Strawson calls an *objective attitude* to that person, which means that we see her as an object of treatment, manipulation and social policy (cf. *FR*, 9).

Notice that Strawson, in effect, describes a default-and-challenge pattern. *Normally*, we react to other people’s harmful or beneficial actions with reactive attitudes, *unless* it is a case of an excuse or an exemption. No inquiry into people’s motivation and their capacities is required unless there is reason to assume that there is an excuse or exemption.⁵

3. “It is one thing to ask about the general causes of these reactive attitudes I have alluded to; it is another to ask about the variations to which they are subject, *the particular conditions in which they do or do not seem natural or reasonable or appropriate*; and it is a third thing to ask what it would be like, what it is like, not to suffer them. I am not much concerned with the first question; but I am with the second; and perhaps even more with the third” (*FR*, 6; emphasis added).

4. Cf. Watson 1987, Wallace 1994; also cf. Claudia Blöser’s contribution to the present volume.

5. This is an aspect of Strawson’s account that it shares with related work by H. L. A. Hart (Hart 1948) and J. L. Austin (Austin 1956).

Corresponding to the reactive attitudes (by which we react to goodwill, or lack thereof, against ourselves and those we love) there are impersonal, generalized, or vicarious analogues, such as moral indignation or moral admiration, by which we react to the moral quality of actions done to others (including those done to people completely unknown to us). Again, these vicarious attitudes are natural reactions and will be inhibited only in exceptional circumstances that are basically the same as the excuses and exemptions that also inhibit the personal reactive attitudes.

In sum, the first part of Strawson's overall argument consists in providing us with a structural description of the conditions under which we typically adopt, or refrain from adopting, the reactive attitudes, thereby articulating the norms that govern the *appropriate* adoption of these attitudes (and thus the appropriate attribution of responsibility). Given this description, we can see that our practice is not consequentialist in spirit, since it consists in reactions to *past* actions and to the agent's *motivation* (goodwill or lack thereof) quite independently of considerations of present or future utility. Moreover, we distinguish between situations in which it is fair, or appropriate, and those in which it is unfair, or inappropriate, to adopt reactive attitudes. So the second part of Strawson's argument must consist in showing that this practice, and hence the concept of responsibility it implicitly employs, is compatible with the truth of determinism. If it were not, this would mean that the truth of determinism would require us to give up the reactive attitudes and to adopt an objective attitude towards everyone (let's call this a "universal objective attitude").⁶

Strawson's argument for the claim that our practice of adopting reactive attitudes is compatible with the truth of determinism comes in three steps: First, Strawson observes that for *logical* reasons the truth of determinism cannot require us to suspend all reactive attitudes. Determinism is a completely general thesis, whereas, given Strawson's description of our practice, refraining from adopting reactive attitudes requires either an excuse or an exemption, each of which hold only under very specific conditions. It does not follow from determinism that each individual act was done accidentally, or inadvertently etc., nor does it follow that each human being is incapacitated (cf. *FR*, 11ff.). Therefore, it cannot be a consequence of determinism that we have to adopt an objective

6. Cf. *FR*, 11. Strawson himself does not distinguish terminologically between the objective attitude toward individuals and what I have called the universal objective attitude. As we will see in what follows, however, this distinction will help to clarify the structure of Strawson's overall argumentative strategy.

attitude in each individual case. The objective attitude is required only in cases of “psychological abnormality” and “it cannot be a consequence of any thesis which is not itself self-contradictory that abnormality is the universal condition” (*FR*, 11). This has been called Strawson’s “argument from excuses”.⁷

Second, Strawson argues that even if we wanted to, it would not be possible for us to give up the practice of adopting reactive attitudes altogether by adopting what I have called here a universal objective attitude, since the reactive attitudes are part of human nature. It is simply not up to us whether or not to engage in them; engaging in them is part of what we are as human beings. Conversely, human beings, as the natural beings they are, are not able to adopt a universal objective attitude for more than a limited time. This has been called the “naturalist argument” (cf. Wallace 1994, 96).

Third, even if we could adopt a universal objective attitude, belief in the truth of determinism *should not* (rationally) lead us to do so, since whether we should adopt a universal objective attitude depends not on the truth of an abstract thesis such as determinism, but on “the gains and losses to human life”, that is, on the practical, social, and political issue of how we want to live our lives and organize our societies (cf. *FR*, 13). Abandoning the reactive attitudes would mean a severe impoverishment of human life, since it would be incompatible with serious interpersonal relationships such as love or friendship as we know them. Even if we were willing to accept these costs, the truth of determinism in itself would not be a sufficient reason to do so. This has been called the “pragmatic argument” (cf. Wallace 1994, 99).

One may wonder whether the naturalist and pragmatic arguments are really necessary for Strawson to achieve his official goal, which he may seem to have reached already with the argument from excuses. After all, this argument, if successful, shows that, given our actual practice of attributing responsibility, the truth of determinism would not count as a reason to give up the reactive attitudes; hence, the concept of responsibility we implicitly employ in that practice is compatible with the truth of determinism. Moreover, this practice is not consequentialist in that it makes the appropriateness of the reactive attitudes depend on the goodwill, or lack thereof, displayed in an action and on whether the agent is a proper

7. Cf. McKenna 2005, 167; Wallace refers to this as Strawson’s “internal arguments” (Wallace 1994, 96).

target of reactive attitudes (and hence on fairness-conditions). With this, Strawson has shown what he had announced, namely that compatibilism is not committed to an impoverished consequentialist conception of punishment and moral praise and blame (as mere means of social policy) since it can incorporate the reactive attitudes with their thick conception of moral desert and responsibility. So there seems to be no need for the further steps in Strawson's argument.⁸

Many readers have therefore taken Strawson to argue for a stronger claim that might require the naturalist and pragmatic arguments, namely the falsity of incompatibilism. This argument would consist of two steps: (i) incompatibilism is at odds with our actual practice of attributing responsibility; (ii) it is neither possible (naturalist argument) nor rationally desirable (pragmatic argument) for us to give up this practice. So if either our practice or incompatibilism has to go, it must be the latter.—But even if there may be a basis in Strawson's text for such a reading, I think that it does not represent the philosophically most convincing reading of Strawson's overall argument.

The naturalist and pragmatic arguments should not be seen as self-standing arguments against incompatibilism, but rather as part of a complex dialectic into which they are entered only in order to fend off an objection against Strawson's general strategy. According to this objection, the argument from excuses only shows that the truth of determinism, because it would be a general truth pertaining to all actions and all agents in the same way, could not be a reason not to adopt reactive attitudes in *each individual case*, since this would require case-specific excuses or exemptions for each case. This, however, seems to leave open the possibility that, if determinism were believed to be true, we would, and should, adopt what I have called above a *universal* objective attitude—an objective attitude towards anyone, *independently* of whether any of the recognized excuses or exemptions applies to them. It is only in order to rule out this possibility, and thus to block an incompatibilist objection based on this possibility, that Strawson invokes the naturalist and the pragmatic arguments. Hence, they form an integral part of Strawson's overall strategy, even if they are not intended as independent arguments against incompatibilism.

It may be questioned, however, whether they succeed even in this limited task. The naturalist argument rests on the assumption that the reactive attitudes are an immutable part of human nature—a claim which can

8. Thanks to Andreas Maier for suggesting this point.

be, and has been, rejected: First, Strawson's naturalism about the reactive attitudes is an empirical assumption for which he does not argue. Second, there is empirical evidence against it based on the differences in the reactive attitudes between people from different times and cultures. And third, we may well be able to develop our moral sensibilities in such a way as to overcome at least "negative" reactive attitudes such as resentment and guilt (cf. Pereboom 2008). Even without going into the details of these objections here, I think it must be granted that they make it exceedingly difficult to defend Strawson's naturalist argument.⁹

But there still remains the pragmatic argument. Against it, it has been claimed that the truth of determinism is relevant after all for the question of how to live, since, if determinism were true, it would be *unfair* to engage in the reactive attitudes and the related practices of praise and blame, punishment and reward. This objection, however, obviously presupposes the truth of incompatibilism. And if the pragmatic argument is not meant as an independent argument against incompatibilism, but is put forward only in order to reject an objection to Strawson's central claim that our practice of adopting reactive attitudes is compatible with the truth of determinism, then it is the incompatibilist who must not presuppose the truth of his own position in countering the pragmatic argument. Conversely, since Strawson's aim is only to supply the resources for an improved compatibilism, his argument may presuppose the truth of that view. Moreover, that presupposition is not unargued for, but motivated by the previously given description of our practice of adopting the reactive attitudes. It therefore seems to me that the pragmatic argument is successful at least in the limited sense of countering the objection that the truth of determinism would make it rationally imperative to adopt a universal objective attitude.

In conclusion, Strawson has achieved what he had set out to do, namely to provide the compatibilist with a non-consequentialist account of praise and blame. Given that compatibilism's dependence on consequentialism about praise and blame had been one of the central targets of the incompatibilist critique, this amounts to an important indirect argument for compatibilism.

If we now turn to the question of what epistemologists might be able to learn from "Freedom and Resentment", I take the central point to be

9. For a thorough analysis and critique of Strawson's naturalist argument cf. Russell 1992. For a defence of Strawson's naturalism cf. Magill 1997.

this: Instead of trying to analyze a contested concept directly, it may be more promising to look at the practices in which the concept is being used, including those in which it is used only implicitly. This helps to avoid a mistake Strawson attributes to traditional compatibilists and to incompatibilists alike, namely the mistake of “over-intellectualizing” our practices (cf. *FR*, 23). I think that this is an idea that can be fruitfully employed in epistemology as well. Instead of trying to analyze concepts such as knowledge and justification directly, let us look at the practices in which these concepts are being used, including those in which they are used only implicitly. In the next section, I will turn to the task of giving a structural description of a practice that may be able to do for epistemological issues what the practice of attributing reactive attitudes does for the debate about freedom and responsibility. It will be important that this practice is “primitive” in the sense of being independent from, and more basic than, the explicitly epistemological vocabulary in which philosophical theories are couched. As we will see, there are many structural similarities between the practice described by Strawson and the one we are going to consider now.

2. *Communication, acceptance, and knowledge*

Many human practices employ epistemic notions either explicitly or implicitly. On the more explicit end of the spectrum, there are practices such as science (or specific scientific disciplines), the law, and, of course, epistemology. On the other end of the spectrum, there are certain simple forms of everyday communication: Speaker A tells hearer B that it is raining outside, and B comes to believe this (that it is raining outside) *because* A said so. Let us say that in such a case, B *accepts* that p *from* A. If A does not accept that p, we can say that B *rejects* that p.¹⁰

Note that the notion of acceptance we are interested in here differs from the one sometimes contrasted with belief (where belief is typically thought of as involuntary and acceptance as voluntary)¹¹ in that it is a three-place-relation: A accepts that p from B. Consider the following definition:

10. Note that this means that “rejecting that p” does not imply believing that not-p, but is compatible with suspension of belief about p.

11. Cf. e.g. Lehrer 1990; Engel 1998, Tuomela 2000. I will ignore the belief/acceptance distinction here and define “acceptance of” in terms of belief.

(A1) B accepts that p from A =_{Def} (i) A asserts that p and (ii) B comes to believe that p because A has asserted that p.

This doesn't give us what we need, though, because it would allow for cases in which B uses A merely as a "reliable indicator", for instance if A asserts that it's raining by using the French words "Il pleut" and B, not understanding French, comes to believe that it's raining because he has noticed that every time A says these words it happens to be raining. Although cases like that are clearly possible, they seem to be relevantly different from typical cases of communication. In order to exclude cases like this, we might try the following definition:

(A2) B accepts that p from A =_{Def} (i) A asserts that p and (ii) B comes to believe that p because B takes A to have asserted that p.

But still there are counter-examples: B may take A to be lying about p but also to be systematically mistaken about whether or not p, so that A asserts that p just in case A believes that non-p and A believes that non-p just in case that p. B can then learn that p from A's assertion that p without "accepting", in any intuitive sense, that p from A. I therefore suggest working with the following definition of acceptance:

(A3) B accepts that p from A =_{Def} (i) A asserts that p and (ii) B comes to believe that p by taking A's word for it.

Here the colloquial expression "taking someone's word for it" serves as a technical term meant to exclude the kinds of indirect means of gathering information compatible with A1 and A2. If B comes to believe something by taking A's word for it, this means that without second-guessing and theorizing about A's intentions, B simply comes to believe that p because this is what A *asserts*. In order to avoid misunderstandings, let me stress that acceptance, by definition, only concerns propositions B doesn't already believe independently of A's assertion.

Although the practice I suggest for consideration is one of communication, and although asserting that p is a way of testifying to others that p, I am not particularly interested here in the epistemology of testimony. The point I will focus on is not whether and how B can acquire *knowledge* by relying on A's assertions; rather, my focus will be on the attitude B must have towards A in order to acquire *beliefs* on the basis of what A

asserts, independently of whether B's beliefs amount to knowledge or not. In particular, I will be interested in the epistemic status B must ascribe to A with respect to p in order for it to be correct for B to accept that p from A. While the epistemology of testimony is interested primarily in the *transmission* of information and knowledge in communication,¹² I will be interested in the concepts and standards of *epistemic appraisal* that lie in the background of any successful act of communication. Following Strawson's lead, I will offer a structural description of (some aspects of) our communicative practice by describing some of the fundamental moves that are possible in it (so far, I have mentioned assertion, acceptance, and rejection) and the conditions under which they are appropriate, thereby outlining the normative structure of that practice.

The first thing to note about the role of acceptance in our communicative practices is that, at least in fundamental cases, *acceptance is the default*. By "fundamental cases" I mean cases in which no prior experience has lead B to take A to be, for instance, dishonest or incompetent. Lacking a history of misinformation or deception linking speaker and hearer, people typically accept what they are told unless they take there to be some special reason against acceptance. Moreover, this is just what the norms that govern human communication require. If someone in my office tells me that the copy-machine is not working and I don't accept this, I need a reason for non-acceptance in order for my response to be appropriate. This is not just a demand of politeness, since, absent some reason for rejection, it would seem to be inappropriate not to accept what the person says even if I keep my reserve to myself. By contrast, I do not need a reason for *acceptance*—over and above the fact that this is what the person said to me. So just as adopting reactive attitudes is the default in our practice of attributing responsibility, acceptance is the default in our communicative practices. Note that this does not exclude the possibility of acceptance that is *not* appropriate by default. However, in this case B needs some special reason for acceptance. For instance, if B takes A to be a notorious liar, it will not be appropriate by default for B to accept that p from A. But if the given situation is such that A would not profit from a lie and B knows this, for instance, it might still be appropriate for B to accept that p from A.¹³

12. For a critique of the widespread assumption that the transmission of belief/knowledge lies at the heart of testimonial knowledge, cf. Lackey 2008, Ch. 2.

13. If B accepts that p from A even though B takes A to be a notorious liar, then the fact that in this case A would not profit from a lie does not serve as the basis for some inference

Notice that this does not commit us to any view about testimony as a source of knowledge; my claim here is not that one is by default epistemically justified in believing what one is told, but only that acceptance is the default according to the standards that govern everyday acts of communication. According to these standards, rejecting what someone tells me without a particular reason for this rejection would be inappropriate. This alone does not imply that acceptance without further reason is *epistemically* permissible or even required, since it might in principle be possible that our communicative standards are less demanding than our epistemic ones. (Although I will argue in what follows that in typical cases of acceptance we attribute knowledge to the asserter, and that the conditions under which it is appropriate to accept that p from A include A's knowing that p, it is important here that the standards for appropriate acceptance can be reconstructed independently of any prior epistemological commitments.)

Next, let us turn to reasons for rejection, that is, a reason that defeats the default. These defeaters can be grouped in three broad categories, depending on whether they concern, (i) the proposition asserted, (ii) the person who makes the assertion, or (iii) the epistemic relation between that person and the proposition asserted.

First, there are reasons that have to do with how plausible p is given the rest of B's beliefs. For instance, if B holds beliefs that are incompatible with p, or beliefs that make it seem highly unlikely that p, then this will typically count as a reason against accepting that p.

The second category comprises those reasons for non-acceptance that have to do with the person making the assertion. If there is reason to believe that A might be lying, for instance because A has been dishonest before or because A would profit from my falsely believing that p, I thereby have a reason not to accept what A says. In the same category belong cases in which A, either temporarily or permanently, is not the kind of person whose utterances can have the force of full-blown assertions or, if they do have that force, do not have the kind of authority that warrants acceptance. Young children, for instance, or people who are drunk, belong in this class.

The third category has to do with the relation between A and what A asserts; in particular, whether or not *A is in an epistemic position to assert*

to the conclusion that p, but rather makes it appropriate for B, in this special case, to take A's word for it after all.

that p. Let me first try to explain the notion of “being in a position to assert” by way of examples: If I am looking outside the window at a tree in plain daylight and nothing obstructs my view, I am in a position to assert that there is a tree in front of the window. By contrast, if the shutters are closed or if it is too dark to see, I am not in a position to assert this, unless I have some other access to the fact that there is a tree in front of my window, for instance by remembering it or by inferring it from other things I know. Typically, A will be in a position to make assertions about what happened to him yesterday. By contrast, if A makes an assertion about something that happened to *me* yesterday in A’s absence, it is at least not obvious how A can be in a position to assert this.

We can try to specify somewhat more generally what it means for someone to be in an epistemic position to assert that *p* by appealing to the so-called “Knowledge Rule for Assertion” (KR). According to KR, one is in an epistemic position that makes it appropriate to assert that *p* only if one knows that *p*. Its most prominent proponent is Timothy Williamson, who has claimed that the following principle is constitutive of assertion: “One must: assert *p* only if one knows *p*.”¹⁴ That knowledge that *p* is a condition for appropriately asserting that *p*, and that therefore assertion is an implicit knowledge-claim, has been maintained by many other philosophers, too (cf. e.g. Brandom 1994, Williams 2001, DeRose 2002, Hawthorne 2004, Stanley 2005).

But KR has also been contested. Some of its critics have denied that KR is valid as a rule for assertion by citing counter-examples to KR (e.g. Lackey 2007, Weiner 2005). Some critics have held that other rules can explain all relevant data in a more convincing or theoretically elegant way (e.g. Douven 2006, Lackey 2007, Weiner 2005). I cannot enter into a discussion of KR and its alternatives here. What must suffice for our purposes is the observation that even most of these critics admit in one way or another that KR does hold for *most* cases of assertion and that *typically*, by asserting that *p* one conveys that one knows that *p* (cf. e.g. Brown 2008, 91, Douven 2006, 469).

14. Williamson 2000, 243. Williamson adduces three kinds of considerations in favor of KR. First, unlike some other candidates for rules governing assertion, KR can account for the fact that it is somehow improper to assert “Your lottery ticket will lose”, even if the chances of its winning are minimal and if later, after the drawing, it turns out that the ticket *has* lost (ibid., 249ff.). Second, “How do you know that?” can be a proper response to an assertion, which it couldn’t be if assertion wasn’t an implicit claim to knowledge (ibid., 251-3). Third, there is something paradoxical and inappropriate about Moorian assertions of the form “*p*, but I don’t know that *p*”, a fact that can be easily explained by appeal to KR. (ibid., 253ff.).

And indeed, it seems hard to deny that if I assert without any sign of reservation, say, that Naples, Florida, is south of Naples, Italy, and it later turns out that I didn't know this and merely guessed, my assertion was not in accordance with the norms that govern our communicative practice. This does not mean that there cannot be counter-examples to KR. Lackey for instance constructs three such examples based on the idea that it can sometimes be appropriate to assert that *p* if one takes all the available evidence to tell in favor of *p* but, for personal reasons, one cannot bring oneself to believe that *p* (and hence doesn't know it) (cf. Lackey 2007). But cases like that are clearly exceptional. They do not tell against the claim that typically, in the vast majority of cases, it is inappropriate to assert something (without indicating any reservation) unless one knows it to be true. So I think that we can learn from the discussion of the Knowledge Rule that at least the following negative default rule seems to hold:

(DKR) If *A* does not know that *p*, *A* will *typically* not be in an epistemic position to assert that *p* *unless* there is some exceptional reason why it is appropriate for *A* to assert that *p* even though *A* doesn't know that *p*.

Since *A*'s not being in an epistemic position to assert that *p* is a defeater for the default of acceptance, this means that *A*'s not knowing that *p* will *typically* be a defeater to the default of acceptance, too. If *A* doesn't know that *p*, it will be appropriate for *B* to accept that *p* from *A* only if there are special reasons that either make it appropriate for *A* to *assert* that *p* in the absence of knowledge that *p* or make it appropriate for *B* to *accept* that *p* from *A* even though *A* doesn't know that *p*.¹⁵

We can sum up our brief overview of possible defeaters by saying that *typically* it will be appropriate for *B* to accept that *p* from *A* unless (i) *p* is implausible, (ii) *A* is dishonest or incapacitated, or (iii) *A* doesn't know that *p* (where "typically" indicates that it can be appropriate to accept that *p* even in the presence of a defeater given special reasons for acceptance).

This brief presentation of possible defeaters is doubly ambiguous between (a) a normative and a factual reading and (b) an 'objective' and a 'subjective' account of defeaters. Both ambiguities are closely connected. According to a factual reading, while competent participants in our com-

15. For examples of the first type of reasons, see Lackey 2007; for an example of the second type, see section 3 below.

municative practice typically *accept* what others assert, they regularly do not do so in the presence of a defeater. According to the normative reading, by contrast, it is *appropriate* to accept what others assert unless there is a defeater. On a *subjective* account of defeaters, defeaters are reasons against acceptance only if, and insofar as, the hearer is, or at least ought to be, *aware* of them, while on an *objective* account, the factual presence of a defeater is sufficient to make acceptance inappropriate. Another way of bringing out the contrast between subjective and objective accounts of the relation between default and defeat is by saying that, on an objective account of default-appropriateness, B's acceptance is appropriate *unless* there *is* a defeater, while on a subjective account, B's acceptance is appropriate *until* B becomes *aware* of a defeater.

Given these distinctions, I want to suggest that all four possible combinations (factual/subjective, factual/objective, normative/subjective, normative/objective) are relevant for the purposes of a description of our communicative practice. (1) People generally accept (factual reading) what others assert unless they take there to be (some indication of) a defeater (subjective account). But since most people are correct most of the time in their beliefs about the presence or absence of defeaters, this also means (2) that people generally accept (factual reading) what others assert unless there really is a defeater (objective account). Further, (3) it is (objectively) appropriate for a participant in our communicative practice (normative reading) to accept what others assert unless there really is a defeater (objective account). But (4) it is (subjectively) appropriate for a participant in our communicative practice (normative reading) to accept what others assert, even in the factual presence of a defeater, until he or she becomes, or ought to become, aware of the defeater. In what follows, I will primarily be interested in the normative side of our practice. Where it will be necessary to distinguish terminologically between the subjective and the objective accounts, I will speak of *correctness* instead of objective appropriateness and *entitlement* instead of subjective appropriateness. The term "appropriate" will be used as neutral between these two readings. Note that even though (objective) correctness and (subjective) entitlement mostly go hand in hand, they can diverge: If B accepts what A asserts in the presence of a defeater of which B was not, and need not have been aware, B is (subjectively) entitled to accept that p from A even though the acceptance is (objectively) incorrect.

So these are the outlines of a structural description of some basic aspects of our communicative practices: There are three fundamental moves that

can be made within this practice (assertion, acceptance, rejection) and two types of secondary moves (challenges and responses to challenges). The default response to A's asserting that p is B's acceptance of p, and there are three types of reasons that defeat this default, namely (i) that p is implausible, (ii) that A is dishonest in asserting that p or is in some sense "incapacitated" (a child, drunk, confused), and (iii) that A, in the given situation, is not in a position to assert that p. Given any of these defeaters, it will be inappropriate for B to accept that p from A unless there is a special reason to do so.

3. Acceptance as implicit knowledge-attribution

Obviously, the above is only a very general description of some aspects of our communicative practice to which much fine structure could be added if needed. I hope, however, that it has become apparent already that this practice is of great potential interest for epistemologists. Since the conditions for appropriate acceptance include the epistemic standing of the asserter, this practice must include standards of epistemic appraisal that, in themselves, are a novel and largely unexplored field for epistemological inquiry. Here, however, I will try to show that our communicative practices are also relevant for more traditional epistemological concerns. In order to do so, it will be important to link the structural description of our practice given so far to the central concept of traditional epistemology, the concept of knowledge. I will do this by arguing for (some qualified version of) the following claim:

(AIKA) Acceptance is an implicit knowledge-attribution.

By this I mean that in accepting that p from A, we commit ourselves to attributing knowledge that p to A. Put negatively, the claim is that it would be inappropriate to accept that p from A and at the same time, if confronted with the question of whether or not A knows that p, not to assent to the claim that A knows that p. If something like this claim is correct, this would constitute a strong and important analogy between acceptance and the reactive attitudes, which can play the theoretical role they play for Strawson precisely because they are implicit attributions of responsibility and free will.

Now AIKA clearly has some initial plausibility. This can be illustrated by cases such as Steward Cohen's airport scenario: A and B are at the

airport waiting to board a plane to New York. They wonder whether their flight will have a stop-over in Chicago, where they have an urgent business appointment at the airport. They overhear C saying to D: “I know that the flight stops in Chicago, I’ve checked the itinerary”. A says to B “The itinerary might contain a misprint or they might have changed the schedule at last minute. So C doesn’t *know* that the flight stops in Chicago. Let’s check with the airline”.¹⁶ Cohen has used this scenario in order to argue that knowledge-attributions are context-sensitive and that one aspect of the context is determined by how much is at stake for the attributor. While A may be correct to claim knowledge relative to a low stakes-situation, B is correct in denying knowledge to A given that there is a lot at stake for B and C. Whether or not this kind of contextualist interpretation is plausible, note that there is a minimal interpretation of the scenario according to which B doesn’t accept from A that their flight stops in Chicago *because* B doesn’t take A to *know* that their flight stops in Chicago. No matter what else one might want to say about the scenario, this minimal interpretation seems hard to deny. And this strongly suggests that acceptance is an implicit knowledge-attribution. Note that otherwise, it would have been an option for B to deny that A knows that the flight stops in Chicago, but still to accept that the flight stops in Chicago from A. But this does not seem to be an admissible move in our communicative practice. We can explain this if we assume that acceptance that p requires that one takes the asserter to know that p.

Against this, it might be objected that knowledge may be an unnecessarily strong condition for acceptance; rather, *warranted* belief might seem to be sufficient.¹⁷ In the airport-case, for instance, B takes A’s belief that the plane stops in Chicago to be insufficiently warranted, which is enough to explain why B doesn’t accept from A that the plane stops in Chicago. But note, first, that in a negative case like this, lack of warrant implies lack of knowledge. So the point of the objection can only be that in the positive case, warranted belief can be sufficient for appropriate acceptance even in the absence of knowledge. But this is an illusion. Let’s assume that B attributes warranted belief to A and that B accepts that p from A. From the fact that B accepts that p from A it follows (i) that B takes A to believe

16. Adapted from Cohen 1999.

17. I use the term “warrant” here to cover both internalist and externalist conceptions of the kind of epistemic standing a belief must have in order to count as knowledge (at least in a non-Gettier situation). It is important for me here to stay neutral with respect to internalist and externalist accounts of knowledge for reasons that will become apparent in section 4.

that p (ii) that B takes p to be true. But then, whenever B accepts that p from A on the basis of attributing the warranted belief that p to A, B will also attribute to A the warranted *true* belief that p. Even if, depending on the account of warrant one favors, warranted true belief may not always be sufficient for knowledge because of Gettier cases, it follows that at least for non-Gettier situations requiring warranted belief for appropriate acceptance comes to the same thing as requiring knowledge.

Nevertheless, there are counter-examples to AIKA: Imagine you know a person you consider to be *wise* (in the sense that their knowledge is wider and their judgment is better than most people's, including your own). Now this person tells you that in the 21st century there will be wars about water. You may think that it is impossible to *know* such a thing (because it concerns the contingent future), but still accept what the person says because you trust their judgment.

However, while this kind of example shows that AIKA cannot be correct as it stands, it is compatible with the weaker claim that A's lack of knowledge is a defeater for *acceptance by default*: If you don't take A to know that p, you may still accept that p from A, but in this case you need some reason to accept that p even though A does not know that p – in the example, the reason consists in the fact that the asserter is particularly wise. So the example seems to confirm, rather than disconfirm, that in order to be entitled by *default* to accept what others assert (i.e. for accepting that p from A to be correct in the absence of special reasons over and above A's asserting that p), one implicitly commits oneself to attributing knowledge to them.

Let's call this weaker principle DAIKA (with "D" for default):

(DAIKA) If B is default-entitled to accept that p from A, and B accepts that p from A merely on the basis of A's asserting that p, B implicitly attributes knowledge that p to A.

That B is default-entitled to accept that p from A means that it is "subjectively" (i.e. from B's perspective) appropriate to accept that p from A merely on the basis of A's asserting that p, that is, independently of any positive assumptions about A's person (his being generally honest, or particularly wise) and without investigating A's reasons for believing, and asserting, that p.¹⁸ Now here's an argument for DAIKA:

18. For a possible worry that the definition of acceptance might seem to rule out the possibility of non-default acceptance, cf. fn. 13 above.

- (1) It is correct (objectively appropriate) for B to accept that p from A only if *either* A knows that p *or* there is some special reason to accept that p from A (such as A being wise, or particularly competent, etc.) (premise)

This premise follows from the above description of our communicative practice, in particular from our discussion of the Knowledge Rule: If A doesn't know that p, then this defeats the default of acceptance, so that B needs some special reason for acceptance in case A doesn't know that p.

- (2) B does not have a special reason to accept that p from A (such as A's being wise or particularly competent, etc.) (assumption)
- (3) B's accepting that p from A is correct only if A knows that p (from 1,2).
- (4) B accepts that p from A (assumption)
- (5) B takes her accepting that p from A to be correct (assumption)

Note that "taking her acceptance to be correct" here is understood in the weak sense that B, if asked, would assent to the claim that her acceptance was correct (or deny that it was incorrect). Therefore, (5) will typically hold, since it is rationally impossible to come to believe that p while denying that the way by which one came to believe that p was correct; any reason to deny this would be a reason not to come to believe that p in this way (or, if it's too late for that, to give up the belief that p).

- (6) B is a competent participant in our communicative practice (assumption)
- (7) If (a) it is correct for B to do F only if C is the case, and (b) B does F, and (c) B takes her doing F to be correct (according to the standards of some practice P of which doing F is a part), and (d) B is a competent participant in P, then B is committed to the claim that C is the case.

That (7) is highly plausible can be seen if we apply it to a simple example: If it is correct for B to cross the street only when the lights are green, and B crosses the street, and B is a competent participant in the practice of road traffic, and B takes her crossing the street to be correct (according to the traffic rules), then B is committed to the claim that the lights are green. Of course, if B knowingly crosses the street when the lights are red, B is

not committed to the claim that the lights are green. But in this case she either is not a competent participant of the practice of road traffic (since she doesn't know one of its basic rules) or B cannot consistently take her crossing the street to be correct. Applied to the case at hand, since (3) corresponds to (7a), (4) to (7b), (5) to (7c) and (6) to (7d), it follows that:

(8) B is committed to the claim that A knows that p.

So in cases where B (i) accepts that p from A (ii) in the absence of any special reason that would make it correct to do so even if A did not know that p and where B (iii) is competent and (iv) takes her acceptance of p from A to be correct, B is committed to attributing knowledge that p to A. Since typical cases of default-entitled acceptance satisfy conditions (i) - (iv), we can say (with a grain of salt) that default-acceptance is an implicit knowledge-attribution.

4. *A Strawsonian argument in epistemology*

Finally, let us look very briefly and in general outline at an epistemological debate that may benefit from being considered from a Strawsonian perspective—the *debate between epistemological internalists and externalists* that has divided epistemologists since some of the early reactions to Gettier brought the option of externalism to general attention. Just like the compatibilist account of punishment, the externalist account of knowledge can seem to leave out something important, namely the fact that human beings are capable of reflective attitudes, and that knowledge, in many cases, seems to require not just de facto reliability, but that the subject has reflective access to the factors that turn a true belief into knowledge. So perhaps we can improve the externalist case by showing that our communicative practice (a) is fundamentally externalist in spirit but (b) also contains elements the recognition of which allows us to honor the fact that we are reflective beings.

In what follows, it will not be possible to discuss the pros and cons on both sides of the debate in any detail and to distinguish between the many different versions of internalism and externalism, including various intermediary positions. Since my aim here is the limited one of showing the possible relevance and fruitfulness of a Strawsonian approach in epistemology, it will be sufficient to work with a generic understanding of the

two competing views that brings out most clearly some of the intuitions motivating each side. Both internalists and externalists, in the sense of the distinction I am going to employ here, agree that knowledge requires more than just true belief, but they differ over what that additional factor consists in and whether the subject of knowledge must be able to have reflective access to it. Let's call whatever is sufficient (in a non-Gettier situation) to turn a true belief that *p* into knowledge a *warrant* for *p*.¹⁹ While internalists have variously cashed out the warrant-requirement in terms of reasons, justification, evidence, etc., externalists have favored accounts of warrant in terms of, e.g., reliability, information, and safety. We can ignore these differences for our present purposes and define generic versions of what has come to be called "access internalism" and "access externalism" as follows:

Access Internalism (AI): A knows that *p* only if there is adequate warrant for *p* and that warrant is reflectively accessible to A.

Access Externalism (EI): A knows that *p* only if there is adequate warrant for *p* (which warrant need not be reflectively accessible to A).

"Reflectively accessible" here means, very roughly, that A can become aware of the warrant just by reflecting on her own state of mind. Further important distinctions within the internalist and externalist camps are possible, depending for instance on whether reflective access is posited (or denied) with regard to the warrant itself, but also to the warrant's being adequate for knowing that *p*.²⁰ These distinctions will not matter for what follows, however.²¹

If the above description of our communicative practice is roughly correct, then an evaluation of AI and AE should start from two aspects of that practice we have highlighted above: First, acceptance is the (defeasible) default. Second, acceptance, at least if we are entitled to it by default, is an implicit knowledge-attribution. Note that we had reached this result without presupposing any specific analysis of the concept of knowledge. In particular, the argument for DAIKA did not depend on any commit-

19. Thanks to Thomas Grundmann for pointing out an error in an earlier definition of "warrant".

20. Cf. Alston 1989 for a classic statement of this distinction.

21. But note that the following argument for AE, if successful, will equally apply to a version of externalism that also denies reflective access to the adequacy of the warrant as a necessary condition for knowledge.

ment to either internalism or externalism. As I will argue now, the way we implicitly attribute knowledge in our communicative practice can be used as the basis for an argument against AI and thus for AE. We start from the following claim about the relation between entitlement (subjective appropriateness) and correctness (objective appropriateness):

- (1) If (i) B is default-entitled to accept that p from A, and (ii) B accepts that p from A merely on the basis of A's asserting that p, and (iii) none of the defeaters that would defeat B's entitlement in fact obtains, then B's accepting that p from A is correct (objectively appropriate) (premise).

This claim is meant as an analytic truth: Default-entitlement is defeasible, which means, roughly, that one's acceptance is subjectively appropriate until one becomes aware of a possible defeater. But if in fact there is no defeater, and hence the entitlement will not be defeated, acceptance merely on the basis of someone's assertion is not just subjectively appropriate (appropriate from the subject's perspective), but objectively appropriate (correct) according to the standards that govern our communicative practice.²²

- (2) The fact that A does not have reflective access to a warrant for p does not defeat B's default-entitlement to accept that p from A (premise).

This, of course, is a substantial premise. For the moment, I want to base it on the description of our communicative practice given above which (plausibly enough, it seemed) did not include A's lack of reflective access as a defeater to the default of acceptance. Now the adequacy of this description might be doubted, and I will return to this doubt soon. For now, I will assume that (2) is true as a claim about the standards that govern our communicative practice.

- (3) If (i) B is default-entitled to accept that p from A, and (ii) B accepts that p from A merely on the basis of A's asserting that p, and (iii) none of the defeaters that would defeat B's entitlement in fact obtains, then it is (objectively) correct for B to accept that p from

22. I assume here that in order to be default-entitled to accept that p from A, one must not falsely believe there to be a defeater to one's default-entitlement.

A even if A does not have reflective access to a warrant for p (from 1,2).

- (4) There are possible cases in which (i) B is default-entitled to accept that p from A, (ii) B accepts that p from A merely on the basis of A's asserting that p, (iii) it is (objectively) correct for B to accept that p from A, and (iv) A does not have reflective access to a warrant for p (from 3).
- (5) If B is default-entitled to accept that p from A, and B accepts that p from A merely on the basis of A's asserting that p, B implicitly attributes knowledge that p to A (= DAIKA).
- (6) There are possible cases in which (a) it is correct for B to attribute knowledge to A and (b) A does not have reflective access to a warrant for p (from 4,5).
- (7) A's having reflective access to a warrant for p is not a necessary condition for the correctness of B's attributing knowledge that p to A (from 6).

In order to turn this into an argument for AE and against AI, we need to connect the conditions under which it is appropriate to attribute knowledge to A to the truth of the claim that A knows that p. The most straightforward way of doing this is by assuming that:

- (8) It is correct (objectively appropriate according to the standards of our communicative practice) to attribute knowledge that p to A only if A satisfies all necessary conditions for knowing that p.

From (7) and (8) it follows that:

- (9) A's having reflective access to a warrant for p is not a necessary condition for A's knowing that p,

which in turn implies that:

- (10) AI is false and AE is correct.

I take this argument to be valid. Premises (1) and (5) follow from the above description of our communicative practice, which I will take for granted now. This leaves only premises (2) and (8) in need of defense against possible objections.

(2) is the claim that lack of reflective access to a warrant is not a defeater to the default of acceptance. This claim can be based on the observation that B does not need to ascertain that A has access to a warrant in order for B to be entitled to accept that p from A. In fact, if B is entitled *by default* to accept that p from A, there is *nothing* A needs to ascertain over and above the fact that A asserted that p.

Against this, it might be objected that, in accepting that p from A in the absence of defeaters, we *presuppose*, and are entitled to presuppose, that A has access to a warrant for p. This objection might be backed by the claim that by asserting that p, A typically presents herself as knowing that p. Since knowledge (according to the objector) requires the reflective accessibility of a warrant, we assume, and may assume, that A has access to a warrant whenever A asserts that p. In effect, the objection claims that (a) it is a necessary condition for it to be appropriate to accept that p from A that A has access to a warrant and that (b) if B is default-entitled to accept that p from A, then B is also default-entitled to assume that A has access to a warrant for p.

Note that this objection presupposes the truth of AI. This is not a problem for the objector, however, since the objection is directed against an argument against AI. So in this stage of the dialectic, the objector is allowed to assume the truth of AI. What is a problem for the objector, however, is that the claim on which the objection is based (namely that we require access to a warrant in every case of appropriate acceptance) is completely *ad hoc* as a claim about our communicative practice. A neutral description of our practice that is not already informed by an internalist conception of knowledge in no way suggests that access to a warrant is a necessary condition for (default-appropriate) acceptance. To see this, consider how the supposed fact that A's access to a warrant is a necessary condition of appropriate acceptance could play out in actual cases of acceptance. First of all, the supposed default assumption (that if A asserts that p, A has access to a warrant for p) will make a difference in the practice only if there are possible defeaters to this default, i.e. if there are cases in which B needs to ascertain, and A needs to manifest, that A has access to a warrant for p. Without the possibility of situations in which the default assumption can be defeated, that assumption would be completely gratuitous. (Otherwise, we might equally posit that acceptance is appropriate only if A's assertion has divine approval, which, however, may be indefeasibly assumed to be present in each case of acceptance.)

Now the obvious, and in fact the only, candidates for situations in which the supposed default assumption of A's access to a warrant is being challenged, and possibly defeated, are situations in which A is required to offer a reason, or some evidence, for the truth of his assertion that p. By offering a reason or evidence in favor of p, so the objector may claim, A shows his access to a warrant for p. If A is unable to offer a reason (or some evidence) for p, this shows that A lacks such a warrant.

But even if this much is granted, it does not follow that access to a warrant is necessary for appropriate acceptance *in the default case* (where no reason-giving is required of A). It does not follow because the reason-giving that is actually required of A for B's acceptance to be appropriate always has the character of a *response* to a challenge. If there is no challenge, it is entirely unclear what an adequate response would have to look like (cf. Williams 2001, Ch. 14). For instance, if A asserts that there is coffee in the kitchen, and B objects that there is no more coffee in the coffee maker, A might respond that there is coffee in a thermos. If A asserts that the wall is blue, and B objects that it looks blue only because of the blue light shining on it, A might respond that the wall looks blue even when the blue light is switched off.

In both examples, it is the objection, or challenge, that indicates the direction an appropriate response will have to take. Without an objection, reason-giving is not only not required, it is even entirely unclear what would count as an adequate reason or as sufficient evidence. If A asserts that the wall is blue and B who can't see the wall right now, but has seen it before, objects that it is not blue, typically it will be sufficient for A to point out that he can see the wall right now and that it *looks* blue. But, as we have just seen, this response will be insufficient if the challenge is that the wall really is white, but there is a blue light shining on it. Again, a different response will be needed if the challenge is that the wall is yellow, but seen in green light, or that A's color vision is malfunctioning, and so on.

It is important to note that the mere *possibility* that the wall is yellow, or that A's color vision is malfunctioning, will not suffice as a challenge that undermines the appropriateness of acceptance. If there is no reason to assume that A's color vision may in fact be malfunctioning, the mere logical and physical possibility of a malfunction will not make it inappropriate for B to accept from A that the wall is blue. In order for there to be a real challenge to the default appropriateness of B's accepting that p from A, there must be some indication that something is wrong with A's assertion that p (cf. Willaschek 2007). In case such a challenge comes

up, what is required of A in order for it to be appropriate for B to accept that p from A is that A can respond appropriately to the very challenge that has been raised. If there is no challenge, however, there is *nothing* by way of reason giving or offering evidence that is required of A in order for it to be appropriate for B to accept that p from A. And therefore, even if it is granted that A's giving a reason for p shows that A has reflective access to a warrant for p, from the fact that accepting that p from A in the presence of a relevant challenge is appropriate only if A gives a reason (or offers evidence) for p, it does not follow that A needs to have access to a warrant in cases where it is appropriate *by default* for B to accept that p from A (i.e. in cases where there are no challenges). Rather, since no reason-giving is required in the case of default-appropriate acceptance, and since it is not clear what kind of reason would be required in the absence of a challenge, a general requirement of reflective access (according to which A needs reflective access to a warrant for p in order for it to be appropriate to accept that p from A) cannot be motivated by a description of our communicative practice.

This means that while the suggestion on which the objection against (2) is based (that we attribute access to a warrant by default) may be *compatible* with our communicative practice, it lacks independent support from the neutral description of that practice given earlier. But then, the objection fails, since (2) is entered as a claim about the norms that govern our communicative practice. As we have now seen, as such (2) is plausible and well-motivated, while its denial would be completely ad hoc.

But there still remains premise (8), according to which it is appropriate to attribute knowledge that p to A only if A satisfies all necessary conditions for knowing that p. Against this it might be objected that it can be *pragmatically* appropriate to attribute knowledge that p to A even if A does not satisfy all necessary conditions for knowing that p. According to this objection, a description of a practice in which a concept such as knowledge is attributed can give us no more than the conditions under which it is appropriate to *attribute* that concept. These conditions, however, will depend on the pragmatic point and the practical restrictions of the practice in question. In the case of acceptance, for instance, they may reflect the need for easy and generally successful communication, which would only be hindered by too strict conditions for acceptance. Therefore, it might be appropriate for us to accept what others assert, and thus to attribute knowledge to them, even in situations in which it is possible, although unlikely, that in fact they do not possess knowledge. Hence,

the conditions under which communication will generally be successful (and thus the conditions under which attributions of knowledge will be pragmatically appropriate) may not be the same as the conditions under which such attributions are strictly speaking true. Something similar may be said about our practice of attributing responsibility by adopting reactive attitudes: When it comes to attributing knowledge or responsibility, we may be a little lax in our everyday practices—a laxity that is harmless given the purposes of that practice, but fatal in philosophy.

Now I think that already the idea that our standards in philosophy may be somehow constitutionally stricter than those in “ordinary” practices is misguided. After all, in ordinary contexts, someone’s life or well-being may depend on whether someone really knows that what they are asserting is true, which rarely is the case in philosophy. So it is not clear why standards for acceptance and knowledge that are strict enough for practical purposes should not be strict enough for philosophy. Something similar can be said for other concepts as well. For instance, it is not clear why standards for attributing responsibility in everyday life should be less strict (e.g. in not requiring indeterminism) than those in philosophy, given that it is only the former that really have an impact on people’s lives. Moreover, this objection drives a wedge between the “correct” philosophical account of a concept and the “lax” way in which we use it outside philosophy that threatens to deprive the philosophical analysis of the concept of all practical relevance. After all, why should we care how philosophers analyze some concept such as knowledge or responsibility if this analysis is unrelated to the way in which we actually use this concept—and use it correctly, relative to the standards of our actual practices? And finally, if it is granted that in our actual practices we do *not* require some condition (e.g. indeterminism, reflective access to a warrant) in order to apply a given concept, what could be the basis for the claim that nevertheless the *correct* account of that concept does include this condition? To invoke one’s own intuitions about hypothetical cases hardly seems sufficient for this. In conclusion, it seems that the conditions under which we actually attribute a concept, and do so appropriately according to the standards governing the relevant practice, are the best access we have to the truth-conditions of statements in which we explicitly apply that concept, so that the alleged gap between truth-conditions and conditions of appropriate attribution on which the objection depends does not really exist.

I therefore think that the objection against premise (8), too, can be rejected with good reason. But then, our Strawsonian argument for AE will

be successful: In our communicative practice, we often implicitly attribute knowledge by accepting what others assert and the conditions under which we do this tell strongly in favor of an externalist account of knowledge.

However, this externalist result must be supplemented by pointing out that there is also an *internalist* element in our practice: In order for B to accept that p from A, in case of a challenge A will often be required to answer the challenge by explicit reason-giving. Sometimes it will be sufficient if someone—A or somebody else—answers those challenges. But in many cases—for instance, if the challenge can only be answered by stating what A himself sees, hears, or remembers—A may be the only person who can give an adequate response to the challenge. And if A indeed can give such a response, A will in many cases have shown to have reflective access to a warrant for p. To be sure, this does not mean that access to a warrant is a necessary condition for knowledge in general. But it might indicate that access to a warrant is a necessary condition for knowledge in many situations in which an assertion, or knowledge-claim, has been challenged. So there is some truth in internalism, even though it is restricted to a special, and, when viewed from the perspective of our communicative practice, rather atypical kind of situation. (This incorporation of an internalist element into a generally externalist conception of knowledge is meant to parallel Strawson's attempt to improve on traditional compatibilism by incorporating a seemingly incompatibilist element, namely a non-consequentialist conception of praise and blame.)²³

5. *Conclusion*

In this paper, I have suggested that Strawson's indirect and practice-based approach to freedom and responsibility in his classic paper "Freedom and Resentment" can serve as a model for a similar approach in epistemology. I have given a structural description of one fragment of our communicative practice, namely the practice of accepting what others assert. I have argued that, according to the standards that govern that practice, acceptance is the default while rejection requires particular reasons. One

23. As will be obvious, I have not followed Strawson in favoring "reconciliation" over partisan argument. While Strawson, as I have argued, is primarily interested in offering an improved version of compatibilism and only secondarily, if at all, in arguing against incompatibilism, I have presented an argument for externalism and only added a conciliatory gesture in the end. This may be more a difference in emphasis, though, than in substance.

of these defeaters is lack of knowledge on the side of the asserter, which means that appropriate acceptance, at least in the default case, is an implicit knowledge-attribution. Therefore, the conditions under which we appropriately accept that *p* allow us to analyze the concept of knowledge that we implicitly employ in the knowledge-attributions that go with acceptance. Finally, I have presented an argument for access externalism on the basis of the connection between acceptance and knowledge-attribution. The main idea behind that argument was that, when we accept that *p* from *A* in default situations, that is, without requiring *A* to give any reasons, we implicitly attribute knowledge to *A* without requiring *A* to have access to a warrant. This can be taken to show that the concept of knowledge we implicitly employ in our communicative practice does not include access to a warrant as a necessary condition. Access externalism should be supplemented, however, by an account of the internalist elements in our communicative practice, which become apparent whenever challenges to an assertion are raised and have to be answered – often by the asserter herself. Obviously, the above argument for externalism is not meant to finally settle the long-standing debate between internalist and externalist accounts of knowledge. I hope, however, that it illustrates the potential of a Strawsonian account in epistemology.²⁴

REFERENCES

- Alston, William P. 1989: *Epistemic Justification: Essays in the Theory of Knowledge*. Ithaca, N.Y.: Cornell University Press.
- Austin, John L. 1956: "A Plea for Excuses". *Proceedings of the Aristotelian Society* 57, 1–30.
- Brandom, Robert 1994: *Making It Explicit*. Cambridge, MA: Harvard University Press.
- Brown, Jessica 2008: "The Knowledge Norm for Assertion". *Philosophical Issues* 18, 89–103.
- Cohen, Stewart 1999: "Contextualism, Skepticism, and the Structure of Reasons". *Philosophical Perspectives* 13, 57–89.
- DeRose, Keith 2002: "Assertion, Knowledge, and Context". *The Philosophical Review* 111, 167–203.

24. For valuable comments on earlier versions of this paper I would like to thank Michel de Araujo Kurth, Sebastian Boll, Mikael Janvid, Andreas Maier, Hannes Ole Matthiessen, Andreas Müller, Michael Williams as well as audiences in Cologne and Milan.

- Douven, Igor 2006: "Assertion, Knowledge, and Rational Credibility". *The Philosophical Review* 115, 449–485.
- Engel, Pascal 1998: "Believing, Holding True, and Accepting". *Philosophical Explorations* 1, 140–151.
- Hart, H. L. A. 1948: "The Ascription of Responsibility and Rights". *Proceedings of the Aristotelian Society* 49, 171–194.
- Hawthorne, John 2004: *Knowledge and the Lotteries*. Oxford: Clarendon.
- Lackey, Jennifer 2007: "Norms of Assertion". *Nous* 41, 594–626.
- 2008: *Learning from Words: Testimony as a Source of Knowledge*. Oxford: Oxford University Press.
- Lehrer, Keith 1990: *Theory of Knowledge*. London: Routledge.
- Magill, Kevin, 2000: "Blaming, Understanding, and Justification. A Defence of Strawson's Naturalism about Moral Responsibility". In: Ton van den Beld (ed.), *Moral Responsibility and Ontology*. Dordrecht: Kluwer, 183–198.
- McKenna, Michael 2005: "Where Frankfurt and Strawson Meet". *Midwest Studies in Philosophy* 29, 163–180.
- Pereboom, Derk 2008: "Revising the Reactive Attitudes". In: Michael McKenna & Paul Russell (eds.), *Free Will and Reactive Attitudes: Perspectives on P. F. Strawson's "Freedom and Resentment"*. Aldershot: Ashgate Press, 219–234.
- Russell, Paul 1992: "Strawson's Way of Naturalizing Responsibility". *Ethics* 102, 287–302.
- Stanley, Jason 2005: *Knowledge and Practical Interests*. Oxford: Oxford University Press.
- Strawson, Peter F. 1974 [1962]: "Freedom and Resentment". In: *Freedom and Resentment and other Essays*. London: Methuen, 1–25.
- Tuomela, Raimo 2000: "Belief versus Acceptance". *Philosophical Explorations* 3, 122–137.
- Wallace, R. Jay 1994: *Responsibility and the Moral Sentiments*. Cambridge, MA: Harvard University Press.
- Watson, Gary 1987: "Responsibility and the Limits of Evil". In: Ferdinand Schoeman (ed.), *Responsibility, Character, and the Emotions*. New York: Cambridge University Press, 256–286.
- Weiner, Matt 2005: "Must We Know What We Say". *The Philosophical Review* 114, 227–51.
- Willaschek, Marcus 2007: "Contextualism about Knowledge and Justification by Default". *Grazer Philosophische Studien* 74, 251–272.
- Williams, Michael 2001: *Problems of Knowledge*. Oxford/New York: Oxford University Press.
- Williamson, Timothy 2000: *Knowledge and Its Limits*. Oxford: Oxford University Press.