

Non-Relativist Contextualism about Free Will

Marcus Willaschek

Abstract: Contextualist accounts of free will recently proposed by Hawthorne and Rieber imply that the same action can be both free and unfree (depending on the attributor's context). This paradoxical consequence can be avoided by thinking of contexts not as constituted by arbitrary moves in a conversation, but rather by (relatively stable) social practices (such as the practices of attributing responsibility or of giving scientific explanations). The following two conditions are suggested as each necessary and jointly sufficient for free will: (i) the agent is able to form considered practical judgements and to act accordingly, and (ii) the agent (or some agent-involving event) is the original cause of her actions. A contextualist reformulation of the second condition is developed according to which only contexts in which responsibility is attributed are relevant for the kind of original causation required for free will, which allows for a non-relativist contextualism about free will.

1. Introduction

John Hawthorne and, more recently, Steven Rieber have suggested contextualist accounts of free will according to which, even in a deterministic universe, actions can be truly said to be free, given an appropriate conversational context.¹ However, Richard Feldman (2004) has argued convincingly that Hawthorne's account is unsatisfactory in several respects, and as I will argue below, his criticisms equally apply to Rieber's proposal. In what follows, I am going to suggest a different contextualist conception of free will that avoids the weaknesses of Hawthorne's and Rieber's proposals by employing a kind of contextualism that is very different from the standard Lewis-Cohen-DeRose type of conversational contextualism Hawthorne and Rieber take as their model.² As will be explained in more detail below, the contextualism I have in mind differs from the standard form of epistemic contextualism in at least two important respects: First, according to the contextualism employed here, what depends on context are not primarily the truth-conditions of certain locutions (locutions containing expressions such as 'knows that p ' or 'is free to do F '), but the correctness of the attribution of certain normative statuses (statuses such as being epistemically justified in believing that p or being responsible for doing F). Some such attributions take the form of locutions containing expressions such as

'knows' or 'is free', but most do not. Second, on the view developed here, a context is not fixed by the moves in a conversation, but rather by social practices such as the practice of attributing responsibility or the practice of giving scientific explanations.³ As will become clearer below, this kind of 'social practice contextualism' does not face the objections that beset proposals such as Hawthorne's and Rieber's. In particular, it can avoid the implausible relativist consequence that the same action can be free relative to one context but unfree relative to a different context.

I begin with a brief look at Hawthorne's account of free will, Feldman's objections to it, and Rieber's alternative proposal (section 2) before offering my own analysis of the concept of free will in terms of two necessary (and jointly sufficient) conditions. While the first condition concerns the agent's ability to form *considered practical judgements* and to act accordingly, the second condition requires the agent (or some agent-involving event) to be the *original cause* of her actions. In section 3, I briefly elucidate the first of these conditions; in section 4, I explain the kind of contextualism to be employed in section 5, where I develop a contextualist reformulation of the second condition. I conclude, in section 6, by pointing out some of the advantages of the resulting contextualist-compatibilist analysis of free will. Let me emphasise from the outset, however, that it is not my aim in this paper to argue for the adequacy of the contextualist account of free will vis-à-vis incompatibilist objections to it. My aim is not to win over incompatibilists, but merely to show that a compatibilist conception of free will can gain in plausibility by incorporating (the right kind of) contextualism.

2. Hawthorne, Feldman and Rieber on Free Will

Hawthorne's suggestion, to which he himself does not subscribe, is centred around the following analysis of 'S does x freely': 'S does x freely only if S's action is free from causal explainers beyond S's control –Psst!—apart from those causal explainers that we are properly ignoring'.⁴ This analysis is contextualist, since which causal explainers we are in fact ignoring, and which we are properly ignoring, varies with context. On this account, locutions of the form 'S freely does x' are subtly indexical in that their truth-value varies with the 'context of attention'. In ordinary conversations about human actions we typically do not attend to the fact that most, if not all, events have sufficient causal 'explainers'. We ignore the causal ancestry of our decisions in the distant past as much as we ignore their proximal causes in our brains and concentrate only on those causal conditions we cannot properly ignore in the given context. Which ones are these? Hawthorne discusses three ways of distinguishing between properly and improperly ignored explainers (a 'consequentialist', a 'descriptive', and a 'transcendental' strategy) without taking a stand on the issue (cf. Hawthorne 2001: 71–2). In any case, assuming that this distinction can be drawn in some way or other, Hawthorne's contextualism implies that ordinary ascriptions of free will typically turn out to be true (since typically, in an ordinary context, *all* causal

explainers beyond the agent's control may properly be ignored). Only in the rare context where the causal ancestry of our decisions is under discussion (say, in a cognitive science lab or in a philosophy class on free will), do we cease to ignore (and *a fortiori* cannot properly ignore) all causal explainers of our actions, so that in these contexts none of our actions will be free.

This approach promises to combine the strengths of compatibilism and incompatibilism while avoiding the weaknesses of both. While incompatibilists are correct in saying that we cannot maintain ascriptions of free will and responsibility in light of deterministic explanations of human behaviour, compatibilists are correct in saying that the availability of such explanations does not affect the truth of our ordinary ascriptions of free will.

Richard Feldman has criticised Hawthorne's contextualist analysis of free will as having three severe shortcomings. *First*, according to Feldman, it is questionable 'that "free" is context-sensitive in the way Hawthorne's contextualist theory implies' (cf. Feldman 2004: 270–1). In particular, it seems that the truth-value of ascriptions of free will is not relative to context; otherwise, it should be obvious to any competent speaker of English that philosophical incompatibilism is itself compatible with the truth of ordinary ascriptions of free will (just as it is obvious that 'It's raining' is compatible with 'The sun is shining' when uttered at different times or places). But in fact, most people would hold that 'S doesn't do x freely', when uttered in a philosophy class, is incompatible with 'S does x freely' when uttered in an ordinary conversation (as long the reference of 'S' and 'x', respectively, is held constant).⁵ In other words, it is highly implausible that the very same action can be both free (in one context) and unfree (in a different context). *Second*, Feldman complains that contextualism concedes too much to the incompatibilist (cf. Feldman 2004: 271–2). According to contextualism, in the context of a philosophical discussion about free will, incompatibilism would be true and compatibilism false, since in this context we do not ignore the causal ancestry of our actions: 'The philosophical debate, which obviously occurs in philosophical contexts, is won by the incompatibilists. Indeed, the compatibilists lose without much of an argument in Hawthorne's account. They capitulate from the outset' (Feldman 2004: 272). *Third*, according to Feldman, contextualism fails to address some central worries that fuel the debate about freedom (cf. Feldman 2004: 272–3). In particular, Hawthorne's account states only a necessary condition for freedom. In order to come up with sufficient conditions, Hawthorne would need a principled way of distinguishing between actions without causal explainers beyond the agent's control that are free and such actions that are merely random, but, as Feldman notes, Hawthorne's contextualism is silent on this issue.

In a recent paper, Steven Rieber identifies a different weakness in Hawthorne's account, namely that it is completely *ad hoc*—'unmotivated apart from its capacity to solve the puzzle [about free will]' (2006: 230). He then suggests a contextualist account of free will which is supposed to avoid that weakness. Rieber offers the following analysis: 'to say that an agent did F freely is to say that [t]he agent caused F and in so doing was the original cause of F' (2006: 230–1).

What makes his account contextualist is that the phrase 'original cause' is taken to be context-sensitive: 'Whether something counts as the original cause can depend on what we are thinking about. Generalizing, we can endorse the following rule: [. . .] In a context in which it is salient that something prior to A caused B, the sentence "A is not the original cause of B" is true' (2006: 232). Since this rule applies not only to sentences about free will, but to *all* sentences containing the expression 'original cause', Rieber takes his proposal to avoid the *ad hoc*-character of Hawthorne's account (2006: 236). Rieber then goes on to derive as a 'corollary' the rule: 'The sentence "If B is the product of a causal chain going back to something which is not A, then the original cause of B is not A" is true in any context in which it is uttered or considered' (2006: 233). The two rules, together with the proposed analysis of 'acting freely' in terms of original causation, are then used to explain why the sentence 'Emma raised her hand freely' (uttered or considered in a context where the causes of Emma's raising her hand are *not* salient) can be compatible with the sentence 'If Emma's raising her hand is the product of a causal chain going back to something other than Emma, then her raising her hand was not free' (uttered or considered in a context where the causes of Emma's raising her hand *are* salient).⁶

Now I find Rieber's central idea—to analyze freedom of the will in terms of original causes and then to relativize talk of original causes to context—very promising. However, all three objections Feldman had raised against Hawthorne's proposal apply equally to Rieber's. *First*, even though Rieber can point to a number of real-life examples in which 'original cause' is clearly used without the implication that there are absolutely no prior causes (2006: 249 fn. 15), it just seems mistaken to say that 'Emma's raising of her hand was free' (when uttered in a context where prior causes are *not* salient) is compatible with 'Emma's raising of her hand was not free' (when uttered in a context where prior causes *are* salient), as long as we are talking about the same action. *Second*, the incompatibilist still wins the philosophical debate hands down, since if determinism is assumed for the sake of the argument, the fact that for every action there is a chain of prior causes is always salient in the philosophical context. And *third*, Rieber, too, fails to provide necessary and sufficient conditions for freedom of the will. The fact that Rieber offers an 'analysis' (2006: 230) of what it means to say that 'an agent did F freely' (where 'acting freely' is meant to involve 'the sort of freedom relevant for free will': 2006: 247 fn. 3) may suggest the contrary. But his proposed *analysans* ('[t]he agent caused F and in so doing was the original cause of F') states at most a necessary condition—not a sufficient condition—for freedom of the will, particularly if 'original cause' is taken to be context-sensitive in Rieber's sense. Just think of small children or the insane. No doubt they will often qualify as the original causes of their actions in the context-sensitive sense of that term, but still we would hesitate to consider their actions free (in the sense 'relevant for free will').

I think that Feldman's criticisms are well-taken. Both Hawthorne's and Rieber's contextualist accounts of free will, despite their attractions, are quite unconvincing. It seems to me, however, that Hawthorne and Rieber do not

present the best possible version of a contextualist account of free will, and that is because they model their suggestions on what has been called ‘semantic’ (cf. Pritchard 2002) or ‘conversational’ (cf. Barke 2002) forms of contextualism about knowledge.⁷ In what follows, I will outline a different contextualist treatment of free will, one that avoids Feldman’s objections and provides us with a satisfactory account of free will that, while located squarely in the compatibilist camp, incorporates ideas usually considered to lie exclusively in incompatibilist territory.

3. A Preliminary Analysis

Freedom of the will is not a property of a mysterious entity called ‘the will’, but rather a quality of what people do: To enjoy free will is to be able to act in a particular way, namely freely (in one sense of the word ‘free’). The difficult task in analysing the concept of free will is to capture the specific sense of ‘free’ relevant for freedom of the *will* (as opposed to, in particular, freedom of *action*, which consists, roughly, in the ability to do what one wants to do). Here is a preliminary suggestion:

- (FW) P’s doing A is free (in the sense of ‘free’ relevant for free will) iff
 (FW1) P is able (a) to form a considered practical judgement about whether or not do A and (b) to act in accordance with that judgement, and
 (FW2) P (or some P-involving event) is the original cause of P’s doing A.

I will not try to show here that FW1 and FW2, properly understood, are jointly sufficient for freedom of the will. Let me just note that FW, assuming that we can make sense of FW2, is not obviously inadequate since, unlike Rieber’s analysis, it takes into account that free will requires, in addition to original causation, some kind of rational control over one’s decisions and actions. As we will see, FW1 is important not only in order to account for the fact that small children and the insane do not enjoy freedom of the will, but also because FW2, on the contextualist account of free will developed below, implicitly refers back to FW1. I now turn to briefly explaining and motivating condition FW1; in the subsequent section, I will then suggest a contextualist reformulation of FW2.

Freedom of the will is, among other things, freedom from determination by one’s own motives (instincts, impulses, desires, inclinations, etc.). To be determined by one’s motives, in the relevant sense, means that the motive upon which one acts is effective (issues in action) independently of one’s considered practical judgement—a judgement, that is, of the form: ‘All things considered, I ought to do A’. For the purposes of this paper, I want to leave open the question of which ‘things’ (desires, beliefs, values, facts, etc.) have to be considered in order to arrive at a considered practical judgement, and whether or not practical considerations have to be governed by rational, moral, or other normative standards. What matters is that human beings can, and often do, form categorical judgements about what, all things considered, they ought to do in a given situation, and that these judgements do not necessarily coincide with what they

find themselves immediately inclined to do in that situation. To be able to form such judgements and to act accordingly is, I want to suggest, a necessary condition for freedom of the will. For example, if someone wakes up thirsty after surgery and knows he ought not to drink for a couple of hours but, because of drowsiness after anaesthesia, *cannot* resist and drinks some water from his bedside table, this person lacks freedom of the will (at least with respect to drinking).

Both empirically and conceptually, cases like this are often difficult to distinguish from cases of weakness of the will, where the person in question could have resisted drinking the water but didn't. But even if it is difficult to tell these cases apart, there seems to be an important difference between freely failing to do, and not being able to do, what one considers to be the best thing to do: If the agent *could not* have acted on her considered judgement, even if she had sincerely tried, she lacks freedom of the will; if she could have done so (with an 'effort of will'), she enjoys at least *some* freedom of the will.⁸ This explains why addicts are commonly considered to lack freedom of the will: The addiction is supposed to lead to an irresistible motive, so that the addict cannot act against it even if she judges that, all things considered, she ought to.⁹ For similar reasons, affective states such as extreme fear or rage tend to be regarded as undermining freedom of the will: They motivate actions (e.g. acts of cowardice or aggression) the agent could not abstain from doing even if these actions went against her own considered practical judgement.

Thus one necessary condition for freedom of the will is the ability to act in accordance with one's considered practical judgement. This does not mean that only decisions in accordance with one's considered judgement are free; after all, we do many things out of our own free will without forming any practical judgement at all, considered or not. Moreover, if we do form a considered judgement and act against it, our decisions may still be free as long as we *could* have acted in accordance with it. (Otherwise, weakness of the will would collapse into lack of free will.) In sum, what is required for free will (in this respect) is that we are *able* to form considered practical judgements and that we are *able* to decide (and typically also to act) in accordance with these judgements.

The condition of being able to act in accordance with one's considered judgement is neutral with respect to the distinction between compatibilist and incompatibilist conceptions of free will. Let's say that an ability to do A is a *standard ability* if, for someone who has that ability at t and given some (appropriately specified) standard conditions for doing A, it is possible to do A and possible not to do A. (Standard conditions for acting in accordance with one's considered judgement, for instance, will include sobriety, the absence of threats, etc.) By contrast, an ability to do A is a *strict ability* if, for someone who has that ability at t and given the total past before t and the laws of nature, it is possible to do A and possible not to do A. Having the standard ability to act in accordance with one's considered judgement is compatible with determinism, while having the corresponding strict ability is not: Having the strict ability implies that even in cases where one does not do A, one could have done A

under identical conditions, which is something determinism, under any plausible reading of that term, does not allow. Having the standard ability to do A, however, only implies that even if one did not do A at t, one could have done A *under standard conditions*, which is perfectly consistent with one's not doing A at t being causally determined. Since FW1 is neutral with respect to the distinction between standard and strict ability, it is neutral, too, with respect to compatibilism and incompatibilism. Much more could be said about the ability to act in accordance with one's considered practical judgement, but for my present purposes the brief remarks of this section must suffice.¹⁰ Before I can turn to FW2 and to contextualizing the concept of original causation, I will first have to explain the kind of contextualism I am going to employ for that purpose and to distinguish it from the contextualism of Hawthorne and Rieber.

4. What's in a Context?

Rieber gives a number of convincing examples which show that we often use the expression 'original cause' without thereby implying that the original cause itself was uncaused, for instance: 'The original cause of the burning of the house was lightning' (2006: 232). This can be a true statement even though the lightning itself was caused by the interaction of air masses of different temperatures. According to Rieber, it is only when we attend to the causes of the lightning, and thus change the context, that the statement that the lightning was the original cause of the fire turns out to be false. As long as we are interested only in the cause of the fire, it is true to say that its original cause was the lightning. 'The phrase "the original cause" thus appears to be context-sensitive' (2006: 232).

Expressions such as 'original cause', 'ultimate cause', 'first cause', and the like, may be used to capture the incompatibilist idea that free will requires what Robert Kane has called 'ultimacy'—that the '*arche* (or sufficient ground or cause or explanation)' of a free action must ultimately lie within the agent so that the causal chain leading to the action cannot be traced back to something beyond the agent herself (Kane 1996: 35). Now I think that Rieber's basic idea is correct: Expressions such as 'original cause' are often used in a context-sensitive way, which fact can explain how something can truly be said to be an original cause even in a deterministic universe. However, as has been pointed out by Feldman, this kind of context-sensitivity (on which both Hawthorne and Rieber base their accounts) does not seem to be of much help for the compatibilist. The contextualism of Hawthorne and Rieber may be characterized by the following features: (1) a context consists in the moves in a *conversation* (with, perhaps, an inner monologue as the limiting case); (2) what is relative to context are the truth-conditions of a particular *locution* containing expressions such as 'is free' or 'acts freely'; (3) what determines which causes of a decision or action are relevant ('may not properly be ignored'), in a given context, for assessing whether that decision or action is free depends primarily on what is being *attended to*, what is being *considered*, or what is *salient* in that context, which in turn depends on what

has been *mentioned* in the conversation in question, or on what we are *thinking* about. I do not want to take a stand on the issue of whether expressions such as 'is free' or 'acts freely' are context-sensitive in this sense. In any case, it seems that this kind of context-sensitivity cannot be the basis for a plausible compatibilist account of free will. The main reason is that it would lead to an extreme relativism about free will, according to which it would depend on arbitrary and accidental moves in a conversation whether some action is free or not in a given context. This is highly implausible since these conversational moves need not have, and typically do not have, any bearing on the internal qualities of the agent and her action.¹¹

By contrast, the kind of contextualism I want to employ does not have this consequence. Before I develop this contextualist approach in some detail with respect to the concept of free will, it may help to distinguish it in more general terms from the contextualism of Hawthorne and Rieber:

(1) On the account I want to present, a *context* is fixed by two factors, (a) a social practice of attributing normative statuses and (b) the facts of the matter. The social practice includes, among other things, rules of relevance, that is, rules that determine what kinds of fact are relevant for deciding whether someone or something is the apt target of attributing a given normative status. Since these rules must be quite general, something else is needed in order to determine what in a particular case, is relevant and what is not, namely the 'facts of the matter'. An epistemological example may help to clarify this point: There is a well-established social practice of attributing normative statuses of the form 'being someone who knows that *p*'.¹² According to this practice, one of the requirements for knowledge is that, in order to know that *p*, one must be able to rule out *relevant* error-possibilities. What is a relevant error-possibility, however, does not (only) depend on what has been said, but rather on rules such as these: 'If an error is likely to occur in the given situation, it is relevant' or 'If the proposition in question is likely to be false, the possibility of an error is relevant' etc. But obviously these rules alone do not suffice to determine what is relevant for assessing a given knowledge-attribution; what is missing is what I call 'the facts of the matter', that is, those facts, which in conjunction with the rules of relevance, suffice to determine which error-possibilities are relevant. For instance, if I claim to have seen Peter riding on his bicycle yesterday, but Peter rarely rides a bicycle whereas his twin brother Paul often does, then the possibility that I mistook Paul for Peter is relevant and must be excluded in order for me to know that Peter rode a bicycle yesterday. In this sense, rules of relevance (here: likely mistakes are relevant) together with the facts of the matter (here: Peter's rarely riding a bike, Peter's having a twin brother Paul, Paul's often riding a bike) determine what is relevant for correctly attributing the normative status of knowledge. More generally: A social practice with its rules of relevance together with the facts of the matter constitute the particular context relative to which the attribution of normative statuses must be evaluated. (To anticipate: *one* kind of social practice pertinent to the contextualist account of free will developed here is the practice of attributing responsibility; a different kind of

social practice pertinent to that purpose is the practice of giving scientific explanations.)

(2) What *depends* on context, according to the contextualism to be developed here, is the correctness of the attribution of normative statuses such as being someone who knows that *p* or being responsible for doing A. Such attributions may (but need not) take a linguistic form. For instance, responsibility often is attributed implicitly by blaming or praising someone (which, in the case of blaming, often hurts the most if done 'silently'). If such attributions do take a linguistic form, they need not employ expressions such as 'knows' or 'acts freely'. Many evaluations of human actions, for instance, carry the implication that the agent is responsible for them: If I praise someone for having been courageous, this praise would be regarded as undeserved if the agent had not been responsible for his behaviour (perhaps because he acted as he did only under a severe threat). If we assume that the kind of responsibility implicitly attributed through praise and blame presupposes freedom of the will (a claim I will not argue for, but simply take for granted), then implicit attributions of responsibility equally are implicit attributions of freedom of the will. However, both kinds of attributions, as implicit, will make no use of expressions such as 'being responsible', 'freedom of the will', 'acting freely' etc. The kind of contextualism I want to defend is not about linguistic expressions, but about normative statuses. Hence, it is unaffected by the various linguistic arguments that have been directed against 'conversational' forms of contextualism such as Lewis's (and hence Hawthorne's and Rieber's).¹³

(3) What makes certain causes relevant for determining whether a given decision is free is not what one is in fact attending to, or what has been mentioned in a conversation, or what one happens to be thinking about, but rather the rules of relevance that guide the social practice in question. Again, an epistemological example may help: In a court of law, a witness's claim to know that the defendant was at the scene of the crime at a particular time can be questioned in many different ways, corresponding to the rules of relevance for error-possibilities. Assuming that the witness is not intentionally lying, the defence might question the source from which the witness is supposed to draw her knowledge (testimony, perception). Let us assume that the witness claims to have herself seen the defendant at the scene of crime. Then the defence might question the witness's ability to tell the defendant apart from other, similar-looking people, since this would make it more likely that the witness mistook someone else for the defendant, and that error-possibility would then have to be ruled out in order for the witness to count as knowing that the defendant was at the scene of crime. But now imagine that the defence mentions the possibility that the witness may only have dreamt that she saw the defendant. Is this an error-possibility that has to be ruled out in order for the witness to count as knowing that the defendant was at the scene of crime? It depends. In the extremely rare case where there are reasons to believe that the witness might indeed only have dreamt what she reports having seen—for instance, because the witness is a small child, or is known for having reported the content of her dreams as real

experiences before—then it may be necessary to rule out this possibility. However, if no such reasons are present, the possibility that the witness only dreamt that she saw the defendant simply is not relevant, in a court of law, for the question of what she can be said to genuinely know. Therefore, in actual court procedure, a witness account cannot be undermined simply by mentioning the possibility that the witness might have been dreaming. This practice implicitly relies on a rule of relevance according to which error-possibilities which are highly unlikely and for which no reason is given for thinking that they may nevertheless obtain, are not relevant and thus do not have to be ruled out in order to attribute knowledge. If this is correct, then something must be wrong with standard ‘conversational’ contextualism about knowledge. Merely mentioning error-possibilities does not suffice to create a context in which these error-possibilities have to be ruled out in order to for someone to count as knowing.¹⁴ Similarly, I want to claim that merely mentioning a decision’s causes does not undermine the claim that the decision is free. Whether or not some causal factor is relevant for determining whether a decision is free is not a matter of the factual moves made in a conversation, but rather of the rules of relevance implicit in the practice of attributing responsibility.¹⁵ Let us now turn to FW2.

5. A Contextualist Account of Original Causation

FW2 requires that the agent (or some agent-involving event such as the agent’s decision to do A) be the *original cause* of her doing A. On the most natural reading of ‘original cause’, FW2 seems to say that an action is free only if it has a cause that is itself uncaused. Therefore, talk of original causes is typically regarded as the exclusive domain of incompatibilists. Before I will show how this idea can be given a compatibilist interpretation if we take into account the context-relativity of ‘original causation’, let me briefly say something about the possible candidates for being original causes mentioned in FW2, namely either *agents* or agent-involving *events*. This disjunction is meant to make FW2 acceptable to both agent-causal and event-causal libertarians. In what follows, however, I rely exclusively on the event-causal disjunct, since only then can FW2 be given a contextualist reading. Moreover, I will assume agent’s *decisions* as the agent-involving events that are the potential original causes of free acts, where ‘decision’ is to be understood in a rather undemanding sense. A decision need neither be the result of conscious deliberation nor a conscious mental event. Rather, by ‘decision to do A’, I will mean the forming of an attitude towards doing A that, internal and external impediments apart, moves one (or allows one to be moved) all the way to doing A.¹⁶ Minimally, this can consist in nothing more than letting oneself be drawn one way or another. I will assume that some such attitude is involved in every free action.¹⁷ However, talk of decisions in this sense is not meant to carry much theoretical weight. For present purposes, any account of human agency will do that acknowledges that actions are caused, among other things, by agent-involving events of some kind. (For instance, Davidsonian ‘primary reasons’

would do as well.) My central claim will be that these events, in whichever way they may be specified, can be original causes of the agent's actions even in a deterministic universe if the concept of original causation is contextualized appropriately.

For the purposes of this paper, I now want to distinguish between two (non-exclusive) *kinds* of contexts in which the question of whether a given decision is the original cause of an action may arise: that is, *evaluative* and *explanatory* contexts. Recall that a context is fixed, on the one hand, by a social practice with its rules of relevance and, on the other, by the facts of the matter. Now a context is *evaluative* if the social practice in question is centrally concerned with evaluating the prudence, morality, legality, and so on of human decisions and actions, evaluations that imply attributions of responsibility. For instance, we may find a decision to be prudent or imprudent, moral or immoral, legal or illegal, where the evaluative application of these distinctions presupposes that the agent is responsible for what she does.¹⁸ By contrast, a context is *explanatory* if the social practice involved is geared towards giving scientific *explanations* of human decisions and actions (e.g. physiological, psychological, social, or economic explanations). Thus, a criminal lawsuit would be a typical example of an evaluative context, and the execution of a scientific experiment in a laboratory a typical example of an explanatory context. Both kinds of context may overlap (i.e. a given situation may belong both to an evaluative and an explanatory context).

Both kinds of context bring with them different standards for determining which, if any, causes of a decision are relevant (and thus rule out original causation). In *explanatory* contexts, the rules of relevance will vary with the scientific discipline, research project, and empirical methods. While the physical causes of human actions may be relevant to the explanation of human decisions in neurobiology and psychology, they will typically be irrelevant in economics. Even the effects of gravitation on human decision-making processes may become relevant in a research project about space travel, while remaining irrelevant in most other contexts. Thus something may be an original cause in one explanatory context but only a subsidiary cause in another. In many explanatory contexts, such as in neurobiology, psychology, psychiatry, or sociology, our decisions will not count as original causes of what we do, since here the natural (and social) causes of human decisions (assuming there are such causes) will be particularly relevant. Hence, if determinism is true, statements of the form 'D is the original cause of P's doing A' will be false in many explanatory contexts.

But in *evaluative* contexts, different standards of relevance are operative. Here, the causes of a decision are relevant (relevant to the question of whether the agent is a possible object of responsibility-attributing evaluations) only if they impair the agent's ability to form considered practical judgements and to act accordingly—or so I want to claim. I cannot offer a detailed defence of this claim in the limited space of this essay, so let me just briefly refer to R. Jay Wallace's extended argument for a very similar claim. Building on the work of John Austin and Peter Strawson,¹⁹ Wallace distinguishes between 'excuses' and 'exemptions', where only the latter exclude an agent (temporarily or perma-

nently) from the class of agents who can be responsible for what they do (Wallace 1994: 118). Wallace then gives a list of exempting conditions, i.e. conditions under which we exclude someone (briefly, for longer periods, or permanently) from the class of responsible agents: 'childhood, insanity or mental illness, addiction, posthypnotic suggestion, behaviour control, psychopathy, and the effects of extreme stress, deprivation, and torture' (1994: 166). Wallace does not claim that the list is complete, but it seems to be fairly comprehensive. He then argues that what these conditions have in common is that they 'impair the powers of reflective self-control' (1994: 166–180), which consist of '(1) the power to grasp and apply moral reasons, and (2) the power to control or regulate [one's] behaviour by the light of such reasons' (1994: 157).

Wallace's 'powers of reflective self-control' are closely related to the ability to form considered practical judgements that I suggested represents a necessary condition for free will. There is at least one important difference, however. Wallace gives pride of place to specifically moral reasons, whereas I have left open the question of which, if any, normative standards are essential for forming considered judgements about what to do. For the sake of the point under discussion, however, we may safely ignore this difference. Wallace's claim seems to be sound even if we replace his talk of 'powers of reflective self-control' with talk of 'the ability to form considered practical judgements'. The conditions that exempt someone from responsibility altogether (such as childhood, insanity, or addiction) are such that they impair either our capacity for practical reasoning (whether understood more narrowly, as the power to grasp moral reasons, or more widely, as the ability to form considered practical judgements), or the ability to act in accordance with one's considered (moral or practical) judgement, or both. (For the sake of brevity, I will sometimes refer to this set of abilities—the ones relevant for FW1—as the 'capacity of *practical reason*'.) And if I am right that these abilities are necessary for freedom of the will, we may conclude that the very conditions that deprive us of the status of being responsible for what we do deprive us equally of freedom of the will.

We can combine this result with another idea found in Austin and Strawson, of which Wallace does not make much. There is an important asymmetry between attributing responsibility on the one hand and withholding such attributions on the other. Whereas we need reasons *not* to hold someone responsible for what she does, we usually don't need reasons to hold someone responsible. Attributions of responsibility follow a default-and-challenge pattern: people are fully responsible for what they do (default)—unless there are specific circumstances that exclude or diminish their responsibility (challenge). This means, first, that exemptions must be based on factors that are generally considered to be *exceptional* (such as mental illness or addiction), since otherwise being responsible couldn't be the default position. It also means, second, that the causes of our decisions will typically not be relevant for establishing that someone is responsible for what she did—simply because this is something we don't have to establish but may presuppose in the absence of challenges. However, the causes of someone's decision will be highly relevant for establishing that someone is *not* responsible.

For instance, we may refrain from holding someone responsible because he suffers from mental illness if we believe that his decisions are causal effects of his illness. Of course, the causal ancestry of a person's decision is not the only reason to exempt that person from responsibility; but it at least features prominently among the many reasons for exemptions.

The claim that only causes that impair the agent's capacity of practical reason are relevant for attributions of responsibility may seem contentious. Incompatibilists may object that *all* causes of a decision are relevant. But even if this were granted, it remains a fact about actual practice of attributing responsibility that generally we don't have to inquire into the causes of someone's behaviour in order to find out whether that person is responsible for what she does; such causes become relevant only as 'exemptions' (in the sense defined by Wallace). Now, this practice allows for both a compatibilist and an incompatibilist interpretation: An incompatibilist will claim that *all* sufficient causes of our decisions are relevant for attributions of responsibility, since a decision cannot be free if it has a sufficient (non-agentic) cause; therefore, what our practice shows, according to the incompatibilist, is that we take it for granted that most of our decisions do not have sufficient (non-agentic) causes. A compatibilist, by contrast, may interpret our practice of attributing responsibility as showing that the truth or falsity of determinism is simply irrelevant for questions of free will and responsibility (cf. Strawson 1963). Both sides, however, should agree that, in our actual practice of attributing responsibility, (1) being responsible (more precisely, being someone who is the right kind of target for attributions of responsibility) is the default position (at least for adults), (2) the causal ancestry of our decisions becomes relevant for attributions of responsibility only if this default is challenged and (3) these challenges typically concern impairments of the capacity of practical reason. What compatibilists and incompatibilists disagree about in this area would then be whether *only* impairments of that capacity count as challenges to someone's (being an appropriate target of attributions of) responsibility or whether other challenges, such as the thesis of determinism, although hardly ever put forward in ordinary contexts, should count as valid challenges, too.

This is not the place to decide this issue. Rather, I want to insist only that a compatibilist interpretation of our practice cannot be ruled out offhand as implausible and that, if we put the three claims together, what we get is the following *rule of relevance for evaluative contexts*: The causes of human decisions are relevant in evaluative contexts (that is, relevant for the question of whether someone is responsible for what they do) only if these causes impair the agent's ability to form considered practical judgements and to act accordingly. Therefore, in an evaluative context it will be true that P's decision to do A is the original cause of P's doing A as long as there are no causes of that decision which impair P's ability to form considered practical judgements and to act accordingly. In other words, in evaluative contexts an agent's decisions will typically count as original causes of her actions, since, whatever causes of those decisions there may be, they will typically be irrelevant in these contexts. In explanatory contexts, by

contrast, the causes of human decisions will often be relevant; in these contexts, human decisions will not typically count as original causes of human actions.²⁰

It thus turns out that FW can indeed be accepted by the compatibilist if we contextualize the concept of original causation with respect to decisions as follows:

- (OCD_C) P's decision to do A is the original cause of P's doing A (in the sense required for freedom of the will) if and only if, *in an evaluative context*, P's decision is the original cause of P's doing A.

Since determinism does not imply that every human decision has causes that are relevant in evaluative contexts (i.e. relevant according to our actual practice of attributing responsibility, causes that impair our capacity of practical reason), determinism is compatible with someone's decisions being the original causes of her actions (in an evaluative context). In the remainder of this paper, I will now point out some of the advantages and consequences of a contextualist-compatible version of FW.

6. A Contextualist Analysis of the Concept of Free Will

By giving both conditions in FW a compatibilist reading, we get the following analysis of the concept of free will:

- (FW_C) P's doing A is free (in the sense of 'free' relevant for free will) iff
 (FW1_C) P has the standard ability (a) to form a considered practical judgement about whether or not to do A and (b) to act in accordance with that judgement, and
 (FW2_C) in an evaluative context, P's decision to do A is the original cause of P's doing A.

Given what has been said before, FW2_C is equivalent to 'There are no causes of P's decision to do A that undermine attributions of responsibility', which, on the account sketched out above, is equivalent to 'There are no causes of P's decision that impair P's ability to act according to her considered judgement'. Hence it may seem that, according to the contextualized notion of original causation, satisfaction of the second condition is a trivial consequence of satisfaction of the first. As a necessary condition for freedom of the will, FW2_C would then be superfluous. But this is not in fact the case. FW1_C requires the ability to act, under standard conditions, in accordance with one's considered practical judgement. Someone may have this ability and still not actually act in accordance with her considered judgement simply because the conditions aren't standard. For instance, I may have the ability required by FW1_C, but extreme emotional stress might prevent me from acting in accordance with my considered practical judgement. In that case, FW1_C would be satisfied, but my action still might not be free. FW2_C accounts for this, since the extreme emotional stress may be a causal

influence that is relevant for attributing responsibility. Hence, FW2_C is an independent necessary condition for freedom of the will. While FW1_C concerns a standard ability and thus depends, among other things, on what the agent *would* do under standard conditions (even if the actual conditions aren't standard), FW2_C concerns the specific causes of a decision in the *actual* world.

Before I turn to some advantages of this contextualist conception of free will, let me address an important objection: On the account suggested here, which causes are relevant in an evaluative context, and thus which actions are free, depends on the rules that govern the social practice of attributing responsibility. But these rules, it may be objected, might turn out to be inadequate. This might happen in one of two ways. First, we might become convinced by hard determinists that (a) all causes of our decisions are relevant in an evaluative context and (b) all our decisions have natural causes. It would then be a mistake to consider as relevant only those causes of our decisions that undermine our capacity of practical reason. Second, we might discover that there are specific causes we thought to be irrelevant for attributions of responsibility that in fact undermine our capacity of practical reason.

Concerning the first possibility, notice that the mere fact (if it is a fact) that there *are* such causes (a fact that would be implied by the truth of determinism) in itself is no reason to consider these causes to be *relevant*. Here, the critic would have to argue against our common practice that a distinction that is central to our practice of attributing responsibility—the distinction between causes that do and causes that don't undermine our capacity of practical reason—is in fact irrelevant. I don't believe that such an argument can succeed, but I cannot argue this point here.

Concerning the second possibility, I think that it has to be granted that there may well be as yet unrecognized specific causes of our decisions that undermine the capacity of practical reason. For instance, a sufficient level of the neurotransmitter serotonin is widely believed to be necessary for inhibiting aggressive behaviour. Since there are genetic disorders that result in an abnormally low serotonin level, people suffering from this disorder may not be able to control their behaviour (not be able to act on their considered judgement) in the same way as other people do.²¹ If these findings turned out to be correct, this might be a reason to acknowledge the genetic disorder in question as relevant for the attribution of responsibility. Perhaps, at least given certain further conditions, people suffering from this disorder will then no longer be made (fully) responsible for their acts of aggression—and something similar may be true about other as yet unrecognized causal conditions of human decisions, both genetic and environmental ones.²² In this respect, the standards governing the practice of attributing responsibility may need improvement and may develop over time. However, it is important to see that examples of this kind do not put in doubt, but rather presuppose the distinction between the relevant and irrelevant causes of our decisions. The point of the examples is not that *all* causes are relevant, but rather that some causes we thought to be irrelevant (or didn't know about in the first place) turn out to be relevant after all. As long as we accept that most people most of the time do have the ability to form considered practical

judgements and to act accordingly, these revisions do not undermine our practice of attributing responsibility as a whole, but merely call for local improvements. After all, the supposed fact that people with *abnormally* low serotonin have a problem with controlling their aggressive impulses in no way implies that people with *normal* serotonin levels cannot control their behaviour.

In closing, let me now highlight some of the advantages and implications of FW_C :

(a) As I noted earlier, both conditions in FW_C are compatible with determinism. But on FW_C , free will is equally compatible with indeterminism. Indeterminism allows for the possibility that our decisions have no sufficient causes at all; hence it allows *a fortiori* for their having no causes that are relevant in evaluative contexts.

(b) Given an appropriate context, an agent's decisions are 'self-originated' or 'spontaneous'. In evaluative contexts decisions can serve as 'unexplained explainers'.

(c) The concept of free will is *context-sensitive* in a way that does not make ascriptions of free will vary with context. Only the truth of claims about original causation varies with context, since a given claim may be true in an evaluative context but false in an explanatory context. But since FW_C ties ascriptions of free will exclusively to the rules of relevance that hold in *evaluative* contexts, this context-relativity does not transfer to ascriptions of free will. Even in a scientific laboratory, the question 'Was P responsible for doing A?' generates an evaluative context in which the causes of P's decision are relevant for answering that question only if they undermine P's ability to form considered judgements and to act accordingly. So we may concede Feldman's first point that 'free' does not seem to be context-sensitive but still insist that the question whether P's action was free depends on the context (at least indirectly).²³

(d) Therefore, Feldman's second objection to Hawthorne's account also does not apply. Feldman complains that Hawthorne grants too much to the incompatibilist by allowing that, in philosophical and explanatory contexts, ascriptions of free will would indeed be false. FW_C does not have this consequence. Rather, FW_C is thoroughly compatibilist, since even in contexts where we do consider the causes of someone's decisions, we can insist that these decisions are free as long as their causes are not relevant in evaluative contexts.

(e) Feldman's third objection to Hawthorne (that only a necessary, and not a sufficient condition for freedom is given) doesn't apply either, since the two clauses of FW_C may plausibly be assumed to be jointly sufficient (even though I have not argued for this claim here).²⁴

(f) FW_C , although thoroughly compatibilist, incorporates a central intuition that drives incompatibilist accounts of free will: the idea of original causation. It must be admitted that contextualising this idea will probably deprive it of much of the charm it originally had in the eyes of incompatibilists. But it seems to me that at least the pre-theoretical commonsensical idea that we are spontaneous agents can be accounted for in a satisfactory way even by the contextualised version of original causation. Isn't it highly plausible that by saying 'The original cause of the

fire was lightning' we do not commit ourselves to the truth of indeterminism? But then, it should be equally plausible that by saying 'I am the original author of my deeds' we are not committing ourselves to the truth of indeterminism either. It is only in the rarefied context of doing philosophy that we tend to forget that there are contextual constraints on the truth-value of our assertions. In ordinary contexts, however, these constraints form the unquestioned background against which our assertions and inquiries make sense. So even if the incompatibilist were correct in claiming that in philosophy we are interested in original causation not in a contextualised but in an absolute sense, it seems highly doubtful that this would reflect on the ordinary concepts of original causation and free will.

7. Conclusion

I have argued, in effect, that compatibilists and incompatibilists can agree on the following conditions for freedom of the will: (i) being able to form considered practical judgements and to act accordingly, and (ii) one's decisions being the original causes of one's actions. The second condition is acceptable to a compatibilist if we contextualise the notion of original causation. The contextualist accounts of free will proposed by Hawthorne and Rieber, however, appeared to be unconvincing, mainly because they are based on a conversational model of contextualism. In this paper, I have suggested a different contextualist account of free will based on a 'social practice' version of contextualism. It is my view that the resulting compatibilist conception of free will is both adequate to our pre-theoretic intuitions and to our actual practice of attributing responsibility. However, this is something for which I have not explicitly argued. My claim is merely that a compatibilist conception of free will can gain in plausibility if it incorporates a contextualised version of the idea of original causation.²⁵

Marcus Willaschek
 Institut für Philosophie
 Goethe-Universität Frankfurt am Main
 Germany
 willaschek@em.uni-frankfurt.de

NOTES

¹ Hawthorne 2001; Rieber 2006.

² Cohen 1988; DeRose 1995; Lewis 1996.

³ This means that the contextualism employed here is much closer to the views of David Annis or Michael Williams (cf. Annis 1978; Williams 1991, 2001) than to the contextualism of Lewis and others.

⁴ Hawthorne 2001: 68. This formulation is modelled on Lewis' structurally identical contextualist analysis of knowledge; cf. Lewis 1996: 545.

⁵ For a similar objection against contextualism about knowledge, see Schiffer 1996.

⁶ Cf. Rieber 2006: 234–5. One may wonder whether Rieber does not implicitly rely on a third rule according to which the sentence ‘A is the original cause of B’ is true if and only if A is a/the cause of B and the sentence is uttered in a context in which no cause of A is salient. Otherwise it is hard to see how Rieber’s analysis is supposed to explain why ‘Emma raised her hand freely’ ‘seems true in an ordinary context’ (Rieber 2006, 235).

⁷ For a critique of ‘semantic’ or ‘conversational’ contextualism about knowledge cf. Willaschek 2007.

⁸ Since freedom of the will comes in degrees, weakness of the will may gradually fade into a lack of freedom. Moreover, the distinction between resistible and irresistible motives may well be vague. Graduality and vagueness explain why weakness of the will and lack of freedom of the will are so difficult to distinguish empirically.

⁹ There are reasons to doubt that this picture of addiction is correct; for a survey of recent literature on the topic, see Yaffe 2002. What matters here, however, is not whether the picture is correct, but that its *prima facie* plausibility can be explained when we assume that freedom of the will is undermined by irresistible motives, where these are understood as motives on which we could not avoid acting, even if that action went against our considered judgement.

¹⁰ Looking more closely at the notion of practical consideration, we may for instance distinguish two aspects that are relevant in forming practical judgements: ‘critical reflection’ and ‘rational deliberation’. While critical reflection has to do with whether my motives are as they ought to be, rational deliberation aims at selecting general strategies and particular courses of action that satisfy my motives in the best way possible. Both rational deliberation and critical reflection employ *normative* standards by which we evaluate whether our motives are as they *ought* to be and what strategies and actions are *best*. The standards may be, for instance, those of prudential rationality, of social propriety, or of moral correctness. I’ll briefly return to the distinction between critical reflection and rational deliberation at the end of this essay, endnote 24.

¹¹ It might be objected that Hawthorne’s account does not have this consequence since, according to Hawthorne, causal explainers do matter for free will unless they are *properly* ignored. Whether it is proper to ignore a causal explainer or not does not depend on arbitrary conversational moves, but rather on the moral theory one employs. But this doesn’t help, since the ‘properly’ becomes relevant only where we in fact are ignoring a causal explainer. If we don’t ignore it, *a fortiori* we do not properly ignore it. And once a causal explainer of an action A is mentioned (whether properly or not, whether because it seemed relevant or simply by accident), we create a context in which ‘S did A freely’ is false. In this sense, Hawthorne’s account has the implausible consequence that whether some action is free or not depends on arbitrary moves in a conversation.

¹² In fact, there are several such practices which differ in several respects. For the purposes of this paper, we may ignore this complication. For more details cf. Willaschek 2007.

¹³ Cf. Cappellen and LePore 2005; Stanley 2004.

¹⁴ For a more detailed version of this argument cf. Willaschek 2007.

¹⁵ It may be asked whether there really is ‘the’ practice of attributing responsibility. One might hold that responsibility is attributed quite differently in legal as opposed to non-legal contexts, in the Western world as opposed to Asian cultures etc. This is certainly correct. However, for our purposes, responsibility is relevant only insofar as it is a necessary condition for deserving praise and blame (and in turn requires freedom of the will). While there may be subtle differences between the ways in which this kind of responsibility is attributed, e.g. in legal as opposed to non-legal contexts, at least in the

Western world these practices exhibit a common pattern which will be the topic of the following section. As will emerge there, it is this pattern which is relevant for understanding freedom of the will. Whether there exist sufficiently similar practices in other cultures is a question I want to leave open. What matters here is that the rules which structure our practices, although it may be difficult to make them fully explicit, are 'objective' in the weak sense of obligating all those who participate in these practices. In this they differ from the arbitrary moves in a conversation. Hence, although the contextualism I want to defend may possibly have relativist consequences when it comes to different cultures, it does not lead into the extreme relativism of conversational contextualism about free will.

¹⁶ Note that this notion of decision does not include what Alfred Mele calls 'distal decisions' (cf. Mele 1992: 143–4; Mele 2006: 32): Whereas a 'proximal decision' to do A immediately precedes doing (or beginning to do or trying to do) A, a 'distal decision' at t_1 concerns the question what to do at some later point in time t_2 .

¹⁷ This means that, while some decisions may themselves be free actions (typically the conscious and deliberate ones), on pain of regress decisions as such must not be conceived as (free) actions. Nevertheless, FW2 allows for an obvious sense in which decisions can be free, namely in the sense of being uncaused.

¹⁸ There may be a loose sense in which we can say of a child, for instance, that her behaviour was imprudent, even though we do not consider her to be responsible for what she does. What I have in mind, though, is a more restricted usage, whereby in saying that some act was, for instance, prudent or immoral, we evaluate not only the act in question, but also the agent. It is this usage that implies ascriptions of responsibility. The class of evaluative contexts obviously allows for further subdivisions, for instance according to the norms employed in evaluating human decisions and actions, so that one could speak of moral, legal, pragmatic contexts etc. In what follows I will ignore these differences for the reasons indicated in endnote 15.

¹⁹ Cf. Austin 1956/7; Strawson 1963.

²⁰ As mentioned before, these kinds of contexts may overlap. In a court of law, expert witnesses are heard in order to determine whether the defendant is legally sane; in a cognitive science lab, the brain processes involved in moral problem-solving may be explored. Here, the context that determines the 'direction of inquiry' also determines the rules of relevance. If we are interested primarily in questions of responsibility, as is the case in a court of law, the rules will be those of the evaluative context, so that only causes that impair the agent's ability to decide in accordance with her considered judgement will be relevant. If we are interested primarily in scientific explanation, as in a cognitive science lab, then all causes of the subject's decisions will be relevant.

²¹ Cf. e.g. Brown *et al.* 1979; Carey and Goldman 1997; Meyer-Lindenberg *et al.* 2006. Both the central claim about the correlations between genes, serotonin and aggression and the details of these correlations are highly controversial. I do not want to take a stand on this issue. I use the possible influence of serotonin on aggressive behaviour merely as an example of a certain *type* of causal condition that undermines reflective self-control. My point is that there is no reason to deny that some such conditions may have to be recognized as relevant in our practice of attributing responsibility.

²² I argue for this in Willaschek 2002.

²³ It may be objected that the context-relativity of the notion of an original cause is just as paradoxical as the context-relativity of the concept of free will. There are two important differences, however. First, context-relativity about original causation, on the view proposed here, does not imply context-relativity about responsibility, whereas context-

relativity of free will does (on the assumption that responsibility requires free will). Hence, the latter is much more paradoxical than the former. Second, the context-relativity of original causation may be assimilated to a difference in perspective: something can be an original cause in one respect, but not in another, which is not paradoxical at all. This move is not available for free will, since, arguably, what drives the interest in free will is the interest in responsibility. Therefore, if we were to say that a decision was free in one respect but not in another, the next question would be whether the decision was free in the respect required for responsibility, so that, with respect to the kind of free will required for responsibility, nothing would be gained by making that distinction in the first place.

²⁴ A major objection to the two conditions' being jointly sufficient might arise from cases of manipulation or brain-washing that undermine a person's freedom of the will but that may seem not to be excluded by FW_C . A possible response could consist in a more detailed account of the ability to form considered practical judgements. For instance, one might distinguish between two aspects of that ability, namely 'practical deliberation' (roughly, means-ends-rationality) and 'critical reflection' (about the ends one pursues); cf. endnote 10 above. Manipulation might then be said to undermine freedom of the will precisely to the extent in which it undermines the ability to critically reflect on one's ends; cf. Willaschek 2009.

²⁵ I would like to thank Steven Gross and Joel Anderson for extremely valuable philosophical comments and suggestions as well as linguistic help concerning earlier versions of this paper.

REFERENCES

- Annis, D. B. (1978), 'A Contextualist Theory of Epistemic Justification', *American Philosophical Quarterly*, 15: 213–19.
- Austin, J. L. (1956/7), 'A Plea for Excuses', *Proceedings of the Aristotelian Society*, 57: 1–30.
- Barke, A. (2002), *The Closure of Knowledge in Context*. Paderborn: mentis.
- Brown, G., Goodwin, F. K. and Ballenger, J. C. (1979), 'Aggression in Human Correlates with Cerebrospinal Fluid Amine Metabolites', *Psychiatry Research*, 1: 131–39.
- Cappellen, H. and LePore, E. (2005), *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism*. Malden, MA: Blackwell.
- Carey, G. and Goldman, D. (1997), 'The Genetics of Antisocial Behavior', in D. M. Stoff, J. Breiling and J. D. Maser (eds), *Handbook of Antisocial Behavior*. New York: Chichester; Weinheim: Wiley, 243–54.
- Cohen, S. (1988), 'How to be a Fallibilist', *Philosophical Perspectives*, 2: 91–123.
- DeRose, K. (1995), 'Solving the Sceptical Problem', *Philosophical Review*, 104: 1–52.
- Feldman, R. (2004), 'Freedom and Contextualism', in J. K. Campbell, M. O'Rourke and D. Shier (eds), *Freedom and Determinism*. Cambridge, MA: MIT Press, 255–76.
- Hawthorne, J. (2001), 'Freedom in Context', *Philosophical Studies*, 104: 63–79.
- Kane, R. (1996), *The Significance of Free Will*. Oxford: Oxford University Press.
- Lewis, D. (1996), 'Elusive Knowledge', *Australasian Journal of Philosophy*, 74: 549–67.
- Meyer-Lindenberg, A., Buckholtz, J., Kolachana, B., Hariri, A., Pezawas, L., Blasi, G., Wabnitz, A., Honea, R., Verchinski, B., Callicot, J., Egan, M., Mattay, V. and Weinberger, D. (2006), 'Neural Mechanisms of Genetic Risk for Impulsivity and Violence in Humans', *Proceedings of the National Academy of Sciences*, 103: 6269–74.
- Mele, A. (1992), *Springs of Action*. New York: Oxford University Press.
- (2006), *Free Will and Luck*. New York: Oxford University Press.

- Pritchard, D. (2002), 'Two Forms of Epistemological Contextualism', *Grazer Philosophische Studien*, 64: 19–55.
- Rieber, S. (2006), 'Free Will and Contextualism', *Philosophical Studies*, 129: 223–52.
- Schiffer, S. (1996), 'Contextualist Solutions to Scepticism', *Proceedings of the Aristotelian Society*, 96: 317–33.
- Stanley, J. (2004), 'On the Linguistic Basis for Contextualism', *Philosophical Studies*, 119: 119–46.
- Strawson, P. F. (1963), 'Freedom and Resentment', *Proceedings of the British Academy*, 48: 1–25.
- Wallace, R. J. (1994), *Responsibility and the Moral Sentiments*. Cambridge, MA: Harvard University Press.
- Willaschek, M. (2002), 'DNS—Doch nicht schuldig? Zum Zusammenhang zwischen genetischer Disposition und persönlicher Verantwortung', *Jahrbuch für Wissenschaft und Ethik*, 7: 243–57.
- (2007), 'Contextualism about Knowledge and Justification by Default', *Grazer Philosophische Studien*, 74: 251–72.
- (2009), 'Der eigene Wille. Zum Zusammenhang zwischen Freiheit, Selbstbestimmung und praktischer Identität', in J.-C. Heilinger, C. G. King and H. Wittwer (eds), *Individualität und Selbstbestimmung*. Berlin: Akademie-Verlag.
- Williams, M. (1991), *Unnatural Doubts: Epistemological Realism and the Basis of Scepticism*. Oxford: Blackwell.
- (2001), *Problems of Knowledge*. Oxford: Oxford University Press.
- Yaffe, G. (2002), 'Recent Work on Addiction and Responsible Agency', *Philosophy & Public Affairs*, 30: 178–221.