

Freedom & Responsibility in Context,
by Ann Whittle. Oxford: OUP, 2021. Pp. x + 235.

BARBARA VETTER
Freie Universität Berlin
barbara.vetter@fu-berlin.de

This is the penultimate version of a review published in *Mind* 2022:
<https://doi.org/10.1093/mind/fzac046>

1. Introduction and overview

When there is seemingly irresolvable and faultless disagreement in philosophy, it is often natural to resort to contextualism: perhaps we are all speaking the truth, but talking past each other in some way. We might get, as Ann Whittle puts it, ‘the sense that both sides were right, at least some of the time, but neither position managed to capture it all.’ (p.ix). In epistemology, this is by now a well-rehearsed and popular move: to accommodate both the philosophically persuasive arguments for scepticism and the immense plausibility of ordinary knowledge ascriptions, it is argued that we speak in different contexts, and the sceptical conclusion is true in some philosophical, but false in all or most ordinary contexts. In the literature on free will, this kind of move seems to suggest itself with equal plausibility: to accommodate both the philosophically persuasive arguments for incompatibilism and the immense plausibility, irrespective of the question of determinism, of our ordinary ascriptions of freedom and responsibility, we might argue that we speak in different contexts, and the incompatibilist conclusion is true in some philosophical, but false in all or most ordinary contexts. However, in the literature on free will and responsibility, this line of argument has been conspicuously near-absent. In her rich and tightly argued book, Ann Whittle fills this gap and proposes a contextualist account of freedom, control, and responsibility.

Previous attempts at a contextualist resolution of these debates have focussed on the causal aspect of freedom. Hawthorne (2001) focussed on ‘causal explainers’ as the source of context-sensitivity, Rieber (2006) on the notion of ‘the original cause’. But there is an alternative and in many ways more natural way of locating the source of context-sensitivity when we talk about freedom and responsibility:

the notion of ability or ‘can’ involved in classic ‘could have done otherwise’ accounts but also in many post-Frankfurt accounts that appeal, for instance, to the abilities that constitute our sensitivity to reasons. The context-sensitivity of ‘can’ is well-established, and Whittle’s book builds on it to formulate accounts of freedom and responsibility that are context-sensitive as well.

At the foundation of Whittle’s argument is CAM: Contextualism about agentive modals, such as ‘can’ (ch.1). Rejecting a direct application of CAM to arguments about free will (ch.2), she proceeds to isolate the relevant notion of ‘can’, the ‘all-in can’, which is argued to be context-sensitive as well (ch. 3). With contextualism about the all-in can in hand, Whittle turns to her target notions of freedom and responsibility. The application is the most straightforward with respect to leeway conceptions of freedom, or ‘regulative freedom’, which is directly understood in terms of the ability to do otherwise. Whittle argues for and defends her contextualist view of regulative freedom in ch. 4-5. She then turns to moral responsibility and its precondition, ‘robust control’. Chs. 6-7 focus on control, arguing that the main competitors for understanding it all involve crucial appeal to abilities, and offering her own account of robust control. That account, or rather its basic idea that control requires certain abilities, is then used to argue for contextualism about moral responsibility (ch. 8). Ch. 9 ends the argument of the book by defending contextualism about moral responsibility against the objection that it makes responsibility a matter of luck.

Whittle’s argument proceeds in what we can call a ‘bottom-up’ directions: she builds her case for contextualism about freedom and responsibility from the case for contextualism about the notions involved in them, in particular, the ‘all-in can’. Context-sensitivity is transmitted, as it were, from *analysans* to *analysandum*, or from the necessary condition to that for which it is a necessary condition. Whittle is well aware that this kind of transmission does not always work, and provides arguments in each case for why it does. In addition, at each stage in the argument, a more direct semantic argument for contextualism (about the all-in can, about freedom ascriptions, about ascriptions of moral responsibility) is offered. As Whittle notes, however, the examples are often difficult to evaluate, since we are dealing with philosophical terms of art. Thus the argumentative weight rests firmly on the bottom-up argument and the transmission of context-sensitivity from agentive modals all the way to ascriptions of freedom and moral responsibility.

There is a wealth of arguments and detailed engagements with other views in the literature that I cannot even try here to do justice to in this review. Suffice it to say that anyone working on freedom and responsibility will find arguments, insights, and thought-experiments worthy of their attention in this book; and that these arguments are developed carefully and with much consideration given to limitations and possible objections.

In what follows, I take a closer look at, and raise some concerns about, two crucial steps in the argument. First, I will consider the ‘all-in can’, the sense of ‘can’ that is relevant for freedom, and argue that it sits uneasily with Whittle’s foundation in the semantics of agentive modals. Second, I will take a closer look at the notion of ‘robust control’ and its role in the argument for contextualism about moral responsibility, and argue that Whittle’s own conception of control is in tension with her argument for contextualism about moral responsibility. Both aspects will also throw some doubt on the feasibility of the ‘bottom-up’ architecture of Whittle’s argument.

2. All-in ability, freedom, and the semantic foundations of contextualism

At the basis of Whittle’s argument is CAM, or contextualism about agentive modals – modals which ‘concern what it is possible for an agent to do, given the way that they are ...; and what it is possible for an agent to do, given the circumstances they find themselves in’ (p.7). CAM is as firm a foundation as one could have. It is part and parcel of the standard semantics of modals, a powerful linguistic tool that has proven its flexibility and explanatory capabilities over several decades now. Very roughly, and following the seminal statement in Kratzer (1977), modal semantics takes all modal statements to include an implicit ‘in view of’ clause whose complement is filled in by context. To use an example of an agentive modal, a skilled trombone player stranded on a desert island can still truly claim ‘I am able to play the trombone’, expressing the fact that she can play the trombone in view of her intrinsic features and previous training; but she might equally truly, and perhaps sadly, say ‘I cannot play the trombone’, denying that she is able to play the trombone in view of those facts *and* the facts about her being stranded on a desert island with no musical instrument in sight. (Whittle provides similar examples and an extended defense of CAM in ch. 1.) In what follows, I will not question CAM as the foundation of Whittle’s argument.

It might seem that CAM alone provides us with a simple and straightforward way of arguing for contextualism about free will, beginning with the incompatibilist’s Consequence Argument. The consequence argument invokes a ‘Transference Principle’: if an agent can ϕ , and the agent’s ϕ ing requires Q, then the agent can make it the case that Q. Some contextual resolutions of ‘can’ invalidate this principle: when my hands are tied, I still possess a general ability to raise my hand, where my raising my hand requires my hand being untied, even if I do not (at this moment) have an ability to untie my hand. But, Whittle argues in ch. 2, this contextualist route out of the consequence argument is too easy, for the sense of

‘can’ or ‘ability’ at issue is simply not what is at stake in the debate. (This mirrors Whittle’s own response to the new dispositionalists’ arguments in Whittle 2010.)

What sense of ‘can’, then, is at stake in the consequence argument, and the debate about free will in general? Whittle takes up a traditional way of characterizing this sense as the ‘all-in can’, or as ‘all-in ability’ (like Whittle, I will take ‘ability’ in a metaphysically deflationary way, as simply the noun that goes with ‘can’). The all-in can, hailing back to discussions in Austin and others, is characterized as ability plus opportunity, or as the case where the agent has the ability and ‘nothing prevents’ them from exercising it (p. 43). Whittle concedes to her opponent that it is consistent with CAM that the all-in can might have an invariantist semantics, and proceeds to argue specifically for contextualism about the all-in can in ch. 3.

However, in conceding to her opponent that there is a determinate (albeit, Whittle argues, context-sensitive) sense of ‘all-in can’, Whittle is conceding too much; or so I will now argue.

Modal semantics, as invoked in CAM, does not give us anything like a specific sense of ‘can’ that would correspond to the ‘all-in can’. All that modal semantics offers us are different modal bases, some of them including more and some fewer facts that complement the ‘in view of’ phrase. (In fact, modal semantics is more complicated than that – in addition to modal bases, it is often believed that we need ordering sources, and Kratzer’s recent work has taken a slightly different turn. But my argument holds, *mutatis mutandis*, for any version of the standard modal semantics, and I will here restrict myself to the simplest version. See Kratzer 1981 and Kratzer 1991 for more details. The case of agential modals is a little more complicated, see Kenny 1976, Mandelkern et al. 2017 and Jaster 2020.) Within this framework, we might understand the ‘all-in can’ as a sense of ‘can’ that takes *all* facts *into* account. But quantifiers like ‘all’ are notoriously context-sensitive: ‘all facts’ might be all facts about the agent’s intrinsic constitution, or about that and her immediate surroundings, or about the history of the entire world, depending on our domain of quantification. (Or if there is no restriction to domains and we really must include *all* the facts, we let the incompatibilist win by default.) That is contextualism alright, but it is at best a trivial contextualism about the all-in can. It simply collapses the all-in can into the ordinary ‘can’.

Whittle’s preferred understanding of the all-in can is in terms of the ‘nothing prevents’ clause. From the semantic point of view, that clause sits uneasily with the attempt to account for all contextual variation within the ‘in view of’ clause (or other semantic mechanisms that are shared by all modals). Be that as it may, the clause ‘nothing prevents’, too, has a quantifier in it. And again, that quantifier is context-sensitive in a way that trivializes the all-in can: it might quantify over factors that concern only agent’s intrinsic constitution, or factors that concern that and her immediate surroundings, or over anything in the history of the entire world.

Again, the all-in can seems to collapse into the ordinary ‘can’.

It is perhaps to avoid this kind of trivialization that Whittle’s argument for contextualism about all-in can appeals neither to the modal bases of all-in can statements, nor to the context-sensitivity of the quantifiers involved in characterizing them. Instead, ch. 3 argues at length for the context-sensitivity of ‘prevents’. (She also argues for the context-sensitivity of ‘ability’, but instead of modal bases, which give us more or less *specific* abilities, she invokes different specifications of the preagent, the ability’s exercise, yielding more or less ‘local’ abilities. See Kittle 2015 for a clear statement of the distinction.) The argument depends on theories of causation that are not uncontroversial. Moreover, of the cases that Whittle adduces as direct semantic evidence, the best example is one which, as she herself acknowledges, shows only the *extrinsicity* and not the context-sensitivity of ‘prevents’ (pp. 57-61). On the whole, contextualism about ‘prevents’ seems to me to be not nearly the firm semantic basis that CAM provides. But as we have just seen, CAM and the standard semantics of modals makes the argument about ‘prevents’ superfluous, for it already establishes a much more pervasive context-sensitivity for the all-in can, as traditionally characterized, than Whittle aims to provide.

The all-in can stems from a debate in the 1960s, which predates the rise of modern modal semantics. Given the thoroughly contextualist approach of the latter, it is difficult to make sense of this older notion. Those who wish to hold on to the all-in can might reject the relevance of modal semantics for the debate about free will; but that would be clearly at odds with Whittle’s overall project. Alternatively, we might characterize the all-in can not in terms of the underlying semantics but in terms of its role for the debate: the all-in can is that sense of ‘can’, whichever it is, that is relevant for free will. This, however, would be at odds with Whittle’s bottom-up strategy. If the all-in can was simply whatever sense of ‘can’ was relevant to free will, then we would have to establish contextualism about free will in order to argue for contextualism about the all-in can – we would have to proceed top-down, not bottom-up. In short, the combination of (i) a bottom-up strategy, relying on (ii) modal semantics and focussing on (iii) the all-in can seems unstable.

How else might a contextualist proceed? One route might be to simply stick to CAM and argue that the burden of proof is on the invariantist, who would need to provide *very* good arguments for the claim that talking about free will somehow, uncharacteristically, fix the kind of modal base that can go with our ‘can’ statements. (This is a strategy we might attribute to Kratzer 1977 and Jaster 2020.) Or else, the contextualist might give up on the bottom-up strategy and begin with evidence for the context-sensitivity of ‘free’, offering CAM as simply the best explanation for those data. Either way, she has her work cut out for her.

3. Robust control, moral responsibility, and the transmission of context-sensitivity

Chapter 6-9 provide Whittle's contextualist account of moral responsibility, beginning with an ability-based notion of robust control (chs. 6-7), which is then invoked in understanding, and arguing for contextualism about, moral responsibility (chs. 8-9). The argument for contextualism about moral responsibility is roughly as follows (p. 172): moral responsibility requires robust control; robust control centrally involves abilities to do otherwise. Since ascriptions of those abilities are context-sensitive, so are attributions of robust control. And since attributions of robust control are context-sensitive, so are attributions of moral responsibility. Whittle recognizes that the last two steps do not follow deductively, but argues that they are eminently plausible: context-sensitivity is transmitted from ability ascriptions to robust control ascriptions, and from there to ascriptions of moral responsibility.

'Robust control', of course, is a technical term (unlike 'ability' and 'responsibility'). Whittle gives a highly original account of it, labelled ARC, which I will turn to in a moment. Her argument for contextualism, however, uses only the weakened premise that 'ARC is true. (Or, less controversially, abilities to do otherwise are central to an analysis of robust control.)' (p. 172) However, I will now go on to argue that a closer look at her own account, ARC, provides reason to think that *even if* ARC is true or abilities are central to an analysis of robust control, the context-sensitivity of 'can' does *not*, or at least not at all straightforwardly, transmit on to 'robust control' (and hence not to 'morally responsible'). So let us turn to robust control.

The account of control that Whittle gives is ecumenical in one way, and quite revisionary in another. She begins by discussing the two main competing notions of control: regulative control, usually understood as being able to do otherwise, and guidance control, understood in terms of reasons-responsiveness. She argues that these two types of account both crucially involve abilities if they are to be plausible; and, crucially, that either of the two kinds of control can be sufficient for moral responsibility without the other. Frankfurt-cases are classic cases of guidance control without regulative control; 'enabling' cases (where the abilities required for reasons-sensitivity are absent but would have been restored if needed) are cases of regulative control without guidance control; and both kinds of cases are argued to involve moral responsibility.

From this, Whittle does not conclude that we need to reject both accounts, or that we need to provide an account that disjunctively covers both kinds of control. Instead, she treats the abilities that both regulative and guidance accounts focus on, and which usually come together, as merely *ceteris paribus* characteristics of robust control, indicative of an underlying feature. In other words, she proposes to

give a ‘functional definition’ of robust control, along the lines of Menzies 1996 on causation:

The relation of robust control *typically* holds between an agent and their action when the agent has the ability to do otherwise [=regulative control] and performs the action ‘under their own steam’ [=guidance control]. (p. 143)

What, then, is the functional account of robust control? The full statement of (ARC), given on p. 144, is rather long, so let me summarize. An agent exercises robust control over their action A iff in A’ing, the agent exercises their ability to A, ‘typically... maintains abilities to do otherwise’ and ‘typically ... has the all-in, local ability to recognize the reasons that are pertinent to their action in that situation and, in the case of suboptimal actions, they typically have the all-in, local ability to have reacted to those reasons appropriately’. (The first two aspects yield a weaker kind of control, ‘minimal control’; the third takes us to robust control. Its second conjunct is designed to deal with an asymmetry in assigning responsibility for good and bad actions; I will ignore it in what follows, since it does not pertain to my own argument.)

Both in the general description of a functional definition, and in ARC itself, the adverb ‘typically’ appears. How are we to read it?

In the quotations from ARC, ‘typically’ seems to qualify the ascription of the abilities themselves. This goes well with some of the further discussion in ch. 7.2: the statement qualified by ‘typically’ is said to hold (without qualification) ‘if nothing interferes’ or ‘normally’. Plugging this into ARC, we get: an agent is in control iff (inter alia) they normally possess local-all in abilities to recognize (and, in some cases, to respond) to pertinent reasons, or they possess these abilities if nothing interferes. Let us call this the ‘generic reading’.

The generic reading is, I think, the most natural reading of the definition as it stands, but it cannot be intended by Whittle. For it says nothing about what actually happens when an agent is in control of her action. When abnormal circumstances deprive me of my all-in abilities, it would still have to be true, on the generic reading, that I am exercising robust control simply because in normal situations I would have the relevant all-in abilities. This is not only implausible; it is also contrary to Whittle’s project, which specifically focusses on all-in abilities to avoid this line of thought. In ch. 2, Whittle considered the appeal to (roughly) general abilities in responding to the consequence argument, and objected that those were not the relevant abilities. A generic claim about all-in abilities comes dangerously close to a claim about general abilities (in fact, Maier 2013 has offered an analysis of general abilities along these lines) – just the kind of thing that Whittle has been trying to avoid since ch. 2.

So ‘typically’ must be meant to do something other than just put a generic operator on the ability ascriptions. And in fact, it does something different in Whittle’s more general description of the functional definition from p. 143 (quoted above). There it qualifies the relation between control and abilities to do otherwise. Whittle’s reference to Menzies (1996), together with a look at her own Whittle 2016, confirms this reading and clarifies the nature of the project. Take Menzies’ functional definition of causation:

[T]he causal relation is the intrinsic relation that typically holds between two distinct events when one increases the chance of the other event. (Menzies 1996, 101)

The definition relies on a relation between causation and probability-raising: in typical cases of the former, the latter also occurs. But probability-raising is used not to say what causation really is; rather, it is instead uses this typical feature to get a fix on the underlying phenomenon – whatever it is. Thus the definition also covers cases where causation occurs without probability-raising, as long as those are cases of the same underlying phenomenon.

If this is the idea of Whittle’s own functional definition, then, we use abilities to do otherwise and respond to reasons not to specify what robust control really is, but to get a fix on the underlying phenomenon, whatever it is. Thus we can cover cases, like Frankfurt cases or enabling cases, where those typical abilities are not present but the underlying phenomenon, whatever it is, is still operative. Whittle (2016) adds to this the idea that this underlying feature is what gives rise to *ceteris paribus* laws: it is the underlying mechanism that underwrites such connections as that between causation and probability-raising or, in our case, between robust control and certain abilities, and in many cases underwrites a lawlike generalization about them (*ceteris paribus*, causes raise the chances of their effects; *ceteris paribus*, when we have robust control, we have abilities to act otherwise and recognize reasons).

I think this is a promising and highly interesting view of robust control and moral responsibility. The problem is, however, that it does not sit well with the contextualist view of moral responsibility to be defended in chs. 8-9.

Consider a functional definition such as Menzies’ definition of causation above. Such a definition may very well include context-sensitive language (perhaps ‘increases the chance’ is an instance). But that is not to say that what the definition defines varies from context to context. To have a clear toy example, imagine geneticists giving a functional definition as follows:

(DefX) Gene X is the part of the human genome that is typically responsible for being tall.

‘Tall’ is highly context-sensitive; but it does not follow that ‘gene X’ as defined in DefX, is context-sensitive as well. To be sure, what DefX says may in principle vary depending on the contextual resolution of ‘tall’. We might contrive a context in which ‘tall’ will be read as ‘tall for a Basketball player’, and another in which it is read as ‘tall for a toddler’. But that is not the context in which the geneticists are speaking. They are looking for a relatively stable underlying mechanism, one which may be assumed to play a certain causal and evolutionary role. The context in which to read ‘tall’ will be adjusted to fit those interests: perhaps we should read it as ‘taller than the average human’, or ‘in the top 5 percentile of human adults with respect to height’ – whatever reading has the best chance of helping us pick out a part of the human genome.

Functional definitions, to be successful in picking out underlying mechanisms, thus appear to introduce their own contexts. They do not pick out slightly different mechanisms in different contexts of utterance. Even if we can vary the contexts a little, we are more likely to pick out the same mechanism across the contexts. (The part of the human genome that is typically responsible for being tall-for-a-Basketball-player might also be typically responsible for being taller-than-average.) There may be some variation in what the definition picks out across different contexts of utterance (perhaps different parts of the genome are responsible for being tall-for-a-toddler and being tall-for-an-adult), but that should only prompt us to make the functional definition more precise. If we do not, we make it shift between different mechanisms with no guarantee that those will underwrite a coherent set of *ceteris paribus* laws. In short: in functional definitions, the context-sensitivity of expressions in the *definiens* does not simply translate into corresponding context-sensitivity of the *definiendum*.

By the same token, when ARC is read as a functional definition, the presence of context-sensitive language in the *definiens* does not give us reason to think that the defined term, ‘robust control’ will be context-sensitive as well. It may be that the very giving of the functional definition creates a fixed context for the interpretation of the context-sensitive terms in it; or that some contextual resolution is required to give the definition any chance of success at all; or that different admissible resolutions simply yield the same, context-insensitive reference for the defined term. We cannot assume without further argument that context-sensitivity carries over from the *definiens* to the *definiendum* in a functional definition; and in fact, given the very purpose of functional definitions, we may even have reason to think that it should not.

Whittle’s argument for contextualism about moral responsibility, we have seen, invokes the transmission of context-sensitivity from ‘can’, through ‘robust control’, to ‘morally responsible’. If I am right, then her own account of robust control stops the transmission, and weakens the case even for the argument that does not invoke

the account itself. Abilities to do otherwise may well be ‘central to an analysis of robust control’, but without that analysis (or rather, its *analysandum*) thereby inheriting the context-sensitivity of the relevant ability ascriptions. Thus, once again, it is Whittle’s bottom-up strategy that is called into question, this time through her own interesting theoretical contribution to the notion of robust control.*

*For helpful discussions, I would like to thank Nathaniel Baron-Schmitt, Maike Albertzart, Max Goetsch, David Heering, and Geert Keil. I also gratefully acknowledge funding from the DFG (Deutsche Forschungsgemeinschaft) within the Centre of Advanced Studies ‘Human Abilities’, and the ‘Capacities and the Good’ project (grants number 409272951 and 439616221).

References

- Hawthorne, John. 2001. ‘Freedom in Context.’ *Philosophical Studies* 104:63–79.
- Jaster, Romy. 2020. *Agents’ Abilities*. Berlin/New York: de Gruyter.
- Kenny, Anthony. 1976. ‘Human Abilities and Dynamic Modalities.’ In Juha Manninen and Raimo Tuomela (eds.), *Essays on Explanation and Understanding*, 209–232. Dordrecht: Reidel.
- Kittle, Simon. 2015. ‘Abilities to do otherwise.’ *Philosophical Studies* 172:3017–3035.
- Kratzer, Angelika. 1977. ‘What ‘must’ and ‘can’ must and can mean.’ *Linguistics and Philosophy* 1:337–355.
- . 1981. ‘The Notional Category of Modality.’ In H. J. Eikmeyer and H. Rieser (eds.), *Words, Worlds, and Contexts. New Approaches in Word Semantics*. Berlin/New York: de Gruyter.
- . 1991. ‘Modality.’ In Arnim von Stechow and Dieter Wunderlich (eds.), *Semantik: Ein internationales Handbuch der zeitgenössischen Forschung*, 639–650. Berlin/New York: de Gruyter.
- Maier, John. 2013. ‘The Agentive Modalities.’ *Philosophy and Phenomenological Research* 87:113–134.
- Mandelkern, Matthew, Schultheis, Ginger, and Boylan, David. 2017. ‘Agentive Modals.’ *Philosophical Review* 126.

- Menzies, Peter. 1996. 'Probabilistic Causation and the Pre-emption Problem.' *Mind* 105:85–117.
- Rieber, Steven. 2006. 'Free Will and Contextualism.' *Philosophical Studies* 129:223–252.
- Whittle, Ann. 2010. 'Dispositional Abilities.' *Philosophers' Imprint* 10.
- . 2016. 'Ceteris Paribus, I Could Have Done Otherwise.' *Philosophy and Phenomenological Research* 92:73–85.