



U.S. National Library of Medicine
National Center for Biotechnology Information

NLM Citation: Levy N. Socializing Responsibility. In: Hutchison K, Mackenzie C, Oshana M, editors. *Social Dimensions of Moral Responsibility [Select Chapters]*. Oxford (UK): Oxford University Press; 2018.

Bookshelf URL: <https://www.ncbi.nlm.nih.gov/books/>



Socializing Responsibility

Neil Levy¹

Abstract

There is a near universal consensus that the bearers of moral responsibility are the individuals we identify with proper names. I suggest that if we take the exercise of agency as a guide to the identification of agents, we may find that agents sometimes extend into the world: they may be constituted by several individuals and/or by institutions. These extended agents may be responsible for morally significant outcomes. I argue that institutions or extended agents may also be responsible for the failure of individuals to satisfy the epistemic conditions on moral responsibility. Individuals may believe virtuously but falsely, due to the way in which cues to reliability are socially distributed. I conclude by suggesting that a focus on individual responsibility may have distracted us from the urgent task of reforming the institutional actors responsible for widespread ignorance about morally significant facts.

Introduction

With few exceptions, work on moral responsibility in the Anglophone world is resolutely individualist.¹ The individual is not merely the primary unit of analysis and bearer of value; for the most part, individualism is taken for granted to such an extent that philosophers are no more aware of their individualism than fish are of the water in which they swim. In this paper, I will suggest that concerns about agency may pull apart from concerns about agents, when 'agents' are identified with individuals (i.e., human organisms, with boundaries defined by their skin). Put another way, I will suggest that if we identify agents with the loci of agency, we will sometimes find that agents are not identical to individuals: they may be constituted by individuals, or parts of individuals, together with extra-individual objects and entities, including other individuals (cf. Rovane 1998).

Author Affiliation: 1 Oxford Uehiro Centre for Practical Ethics; Email: neil.nl.levy@gmail.com.

This is a draft of a chapter that has been accepted for publication by Oxford University Press in the forthcoming book *Social Dimensions of Moral Responsibility* edited by Katrina Hutchison, Catriona Mackenzie, and Marina Oshana due for publication in 2018.

© Neil Levy 2018.

Monographs, or book chapters, which are outputs of Wellcome Trust funding have been made freely available as part of the [Wellcome Trust's open access policy](#)

¹ Important exceptions to the individualistic consensus include Ciurria (2015), Doris (2015), and Vargas (forthcoming). My own work also rejects the consensus: with a focus on (so-called) Frankfurt-style cases, I have argued that the mere presence or absence of other individuals can make a decisive difference to agents' moral responsibility (Levy 2008; 2015).

Since it is agents, understood as loci of agency, who are (very plausibly) the bearers of responsibility, both our normative interests in the distribution of obligations and of blame and our explanatory and predictive interests in the causes of changes in the world should lead us to abandon individualist presumptions in favour of a focus on agency, whenever and wherever we find it.²

Changing our focus from biological individuals to a more inclusive way of understanding the loci of agency is, I claim, justified on grounds that are conceptual and normative. It is justified on conceptual grounds because we cannot understand individual agency except by situating individuals within broader contexts (especially, though not only, the broader contexts of the other agents who constitute the social groups to which they belong). We cannot, that is, hope to understand responsible agency by developing accounts of individual agency, on the one hand, and collective agency, on the other.³ Individual agency is dependent on the social context in which individuals are embedded, both for its genesis and its continued existence. Individuals owe their powers and their limitations to stable collectives as well as much smaller and perhaps transitory groupings (even groupings as small as a dyad) and only by understanding this context can we understand how their agency is realized, the constraints upon it, and the ways in which agents may reasonably be expected to exercise it.

The change of focus from biological individuals to the broader contexts in which individuals are situated is justified on normative grounds for several reasons but a principal justification is the corrective role it can play in identifying how (alongside other resources) institutions distribute individual-level blameworthiness. Agents who perform wrongful acts while satisfying standard individual level criteria for blameworthiness *very* often have fewer resources (psychological and financial) and less (in the way of resources and social standing) to lose from punishment than those who do not, and these facts about them are explained by how resources are socially distributed. They have less self-control, for instance (self-control is correlated with socio-economic status, and the neural mechanisms for this fact are slowly being elucidated (Hackman, Farah & Meaney 2010)). They are subjected to a double dose of unfairness: first they are positioned socially and psychologically such that they face greater incentives to crime and have less capacity to resist these incentives, then they are blamed and punished when they fail to resist. Addressing the enormous and growing disparities within developed countries requires the redistribution of resources of many kinds, but a heavy emphasis on “personal responsibility” is a barrier to the policies that would lead to such redistribution. If we come to see that individual responsibility is itself frequently distributed, and refocus our attention on the collectives and institutions that are the agents of the reproduction of inequality, we may be better placed to allocate blame and to enact better policies (both the official policies formulated by political and social organizations and the informal policies which govern our behaviour toward one another).

In the first half of the paper, my focus will be on the locus of agency, and on questions concerning how it is exercised. Roughly, this half focuses on what is often called the control condition on moral responsibility, and its purpose is to clear the way for refocusing our concern with responsibility from individuals to broader contexts,

² I am a moral responsibility skeptic (see Levy 2011). However, normative concerns other than concerns over blame and praise—concerns, for example, with who is capable of responding to reasons and for whom respect for agency requires the presentation of reasons—closely track necessary (though for me not sufficient) conditions for the ascription of moral responsibility. Throughout this paper, I will write as if agents could justifiably be held morally responsible for their actions, rather than allow my own views to obtrude. In so doing, I will attempt to identify a set of conditions which are genuinely important for agency on my view and which ought to be held to be central for responsibility by those who are not skeptics.

³ There is a rich and growing literature on collective responsibility; see Sepinwall (2016) for a review, focused on the responsibility of corporations. This literature has significantly advanced our understanding of how collectives can exercise agency and the conditions under which the collective might be said to be culpable, but is not intended to address my primary concern: how *individual* agency is dependent on its social context, including the way in which agents are embedded in collectives, and the extent to which individual powers of agency are dependent on, or even partially constituted by, this context.

especially the contexts provided by collectives. I will argue on conceptual grounds that the focus on individuals abstracted from their contexts is unjustified, showing that the kinds of capacities that have plausibly been held to be sufficient for moral responsibility may be possessed by supra-agential aggregates (constituted of several agents, of agents embedded in institutional contexts, or of agents and artifacts).⁴ In the second half of the paper, I turn to the other condition on moral responsibility that is often identified: the epistemic condition.⁵ If considerations concerning the control condition force us to give up an exclusive focus on agents understood as individuals, there is no principled reason why we ought to retain it when assessing whether the epistemic condition is satisfied, and by whom it is satisfied. Broadening our perspective allows us to broaden the range of individuals and institutions who are obligated to satisfy the epistemic condition and who may be held responsible if it is not satisfied. In the conclusion, I draw the two threads of the paper together, suggesting that we might best identify the agents responsible for failures to satisfy the epistemic condition by reference to the locus of agency, rather than preconceived ideas about who agents are. We may thereby find that the obligations to bring it about that individuals are better informed are themselves as truly social—falling on social institutions and governments, for example—as are the agents who bring about significant changes in the world.

1. Agency: Individual and Extra-Individual

The past two decades have seen a revival of interest in debates over the control condition on free will and moral responsibility. Incompatibilists—those who maintain that free will and moral responsibility are incompatible with causal determinism—have developed new lines of attack against their compatibilist counterparts, while compatibilists have developed sophisticated accounts of agency to ward off these attacks (see Kane 1996; Fischer and Ravizza 1998; Clarke 2003; Steward 2012; Vihvelin 2013; Pereboom 2014 for some of the many highlights of the debate in its agency-focused incarnation). These debates have been fruitful: they have illuminated essential aspects of human agency and led to a better understanding of how it can be constrained and undermined. But the accounts of agency developed by compatibilists and libertarians alike share a commitment to individualism that dramatically and, I think unjustifiably, limits their scope. They are silent on how agency is supported and even partially realized by a scaffold that extends beyond our skins, and therefore on the extra-agential, deeply social, conditions which may make the difference between responsibility and excuse.

Let me illustrate this individualist commitment with a discussion of a recent debate between compatibilists and incompatibilists. Recent controversy has often centred on *manipulation cases* (some philosophers believe that manipulation arguments constitute the central difficulty facing compatibilism; see Mele 2006). In brief, arguments based on manipulation cases proceed as follows: a scenario is presented in which one agent covertly manipulates another in a manner that leaves the manipulated agent satisfying whatever set of conditions a compatibilist might adduce as sufficient for freedom, but which is supposed to be intuitively responsibility-subverting nevertheless. Having generated the intuition that an agent may be unfree despite satisfying plausible compatibilist conditions on free will, the incompatibilist then proceeds to argue that there is no relevant difference between being covertly manipulated to act in this way and being determined to act in the same way. Perhaps the manipulation case that has attracted the most attention is Derk Pereboom's Four-Case Argument

⁴ The idea that agency may supervene on individuals plus things in the world builds on the extended mind hypothesis; the hypothesis that mind may extend beyond the skull and into the world; see Clark and Chalmers (1998). Clark and Chalmers argue that external artifacts sometimes play functional roles in cognition analogous (or even identical) to the functional role of brain mechanisms and therefore should be seen as partially constituting minds. Agency, in turn, is dependent on cognition: action is guided by representations of the world. The view discussed here differs from the extended mind hypothesis not only in focusing on agency and not mind, but also in focusing on how other individuals and social institutions—as well as or instead of artifacts—may constitute agency.

⁵ In Levy (2011), I argue that the control and epistemic conditions are not independent: rather, the epistemic condition is at least in important part a condition that must be satisfied in order for the agent to exercise control. It will do no harm to treat them as if they were independent in this context, however.

(Pereboom 2001; Pereboom 2014), which has the virtue (among others) of making the parallel between manipulation and determinism very clear.

In Case One of the Four-Case Argument, Plum is manipulated into murdering Ms. White for rationally egoistic reasons, by neuroscientists who use very advanced technology to produce his moment to moment states. But Plum nevertheless satisfies plausible compatibilist conditions on free will. He does not, for instance, experience an irresistible desire, and his reasoning is consistent with his character (he is frequently manipulated into acting for rationally egoistic reasons, though he also often acts for moral reasons). He endorses his first-order desire to kill Ms. White and he satisfies conditions on reasons-responsiveness and reactivity (his reasoning processes would recognize reasons, including some moral reasons, to do otherwise, were they presented, and there are possible scenarios in which he would actually do otherwise in response to such reasons).⁶ Nonetheless, Pereboom claims that it is intuitive that Plum is not responsible in Case One: the manipulation undermines his freedom.

Pereboom then proceeds to construct three further cases. In the first two, Plum is manipulated into killing Ms. White as before and for the same reasons, but each case is less outlandish and closer to an ordinary case involving wrongdoing in a deterministic world than its predecessor. In Case Two, Plum is not locally manipulated, as in Case One. Rather, he has been created by neuroscientists who have “programmed” him to weigh reasons egotistically (though not exclusively so), such that in the circumstances in which he finds himself he is causally determined to form a desire to kill Ms. White and to act on it. As before, the desire is not irresistible, Plum endorses it and the mechanisms upon which he acts are moderately reasons-responsive. In Case Three, Plum has been indoctrinated from infancy in such a way that he often (though not exclusively) acts for egotistical reasons, such that he is determined in the situation in which he finds himself to form a desire to kill Ms. White and to act on it. In Case Four, Plum is not manipulated at all: rather, he is an ordinary agent in a deterministic world, who often (though not exclusively) acts for egotistical reasons, such that he is causally determined in the situation in which he finds himself to form a desire to kill Ms. White and to act on it.

Pereboom argues that if one accepts that Plum in Case One is not responsible—and one should, he maintains—then one should also accept that Plum is not responsible in Case Two either. After all, he points out, the difference between the cases consists merely in *when* the manipulation occurs. But if one accepts that Plum is not responsible in Case Two, Pereboom argues, one should also think he is not responsible in Case Three, since it differs from Case Two only in *how* Plum’s mental states were created. Indoctrination need not be less constraining than neural manipulation, after all, so it is hard to see why it ought to undermine moral responsibility any the less. But, finally, Pereboom argues that if one accepts that Plum in Case Three is not morally responsible for killing Ms. White, one should accept that Plum in Case Four—which features an ordinary agent acting in a deterministic world—also fails to be morally responsible. According to Pereboom, if Plum’s mental states and values arose by chance or without anyone intending them, he would be no more responsible than he is in a world in which he has been manipulated into having them. Accepting that Plum is not responsible in Case Four, however, is accepting that compatibilism is false: ordinary agents cannot be morally responsible in deterministic worlds.

It is tempting to reply to the manipulation argument by trying to identify a difference between the manipulated agent and ordinary agents in deterministic worlds. McKenna (2008) calls this kind of move a *soft-line* reply to these cases. He argues that soft-line replies will ultimately fail: even if we succeed in demonstrating a

⁶ By building in these stipulations, Pereboom aims to ensure that Plum satisfies the conjunction of conditions on moral responsibility set down by a variety of leading theorists. The stipulation that Plum endorses the first-order desire on which he acts, for instance, ensures that he satisfies the conditions that Frankfurt (1971) influentially advances; the stipulation that Plum would respond and react to some reasons to do otherwise, including moral reasons, in counterfactual situations ensures that he satisfies Fischer and Ravizza’s (1998) conditions for moral responsibility, and so forth.

responsibility-relevant difference between an agent who features in a manipulation case and the kinds of agents whom compatibilists think are morally responsible for their actions in a deterministic world, the incompatibilist can simply build that condition into their revised manipulation argument. The only condition the compatibilist could cite that the incompatibilist cannot bring on board is the stipulation that the agent not be manipulated, but that stipulation looks ad hoc and question-begging without an account of *why* manipulation undermines moral responsibility but determinism does not.

McKenna therefore advocates a *hard-line* reply to manipulation cases: rather than deny that the agents who feature in such cases are relevantly similar to determined agents, he denies that manipulation necessarily undercuts responsible agency. This bullet-biting maneuver requires us to accept that one agent may manipulate another into performing an action for which that second agent is nevertheless fully responsible (which does not entail, of course, that the manipulator is not *also* responsible for the action; responsibility may be able to be shared without any particular agent's portion diminishing relative to a situation in which responsibility is hers alone). No matter what set of conditions we identify as sufficient for moral responsibility, that is, we must accept that manipulators can bring it about that agents satisfy these conditions (unless we stipulate a question-begging 'no manipulation' condition). But that fact should not shake our faith in those conditions. If they are independently plausible as sufficient conditions on moral responsibility, we should accept that agents who satisfy them *are* morally responsible, even if they are manipulated into satisfying them, McKenna maintains.

John Martin Fischer's enormously (and deservedly) influential account of morally responsible agency shares with McKenna's an emphasis on the rich set of internal capacities displayed by most actual adult agents. For Fischer (see Fischer and Ravizza 1998; Fischer 2006), agents are morally responsible when their actions are caused by moderately reasons responsive mechanisms. Roughly, a mechanism is moderately reasons responsive if it is capable of recognizing a sufficient range of reasons to do otherwise were they presented, and there is at least one possible world in which the mechanism would cause the agent to do otherwise in response to such a reason.⁷ This account is genuinely illuminating. We really do discover important things about agents by asking about counterfactuals in precisely the way Fischer suggests. We discover what reasons they are capable of recognizing and responding to, and whether there are distinctive gaps in the patterns of such responsiveness, and so on. But in the hands of McKenna, Fischer and other philosophers, the answers to these questions are gerrymandered to reflect individualist commitments.

For Fischer and for McKenna, an agent exemplifies a rich enough set of agential capacities and is therefore morally responsible when she herself – understood as the person bounded by her skin – remains capable of appropriate response. When manipulation, for instance, leaves her unable to respond in that way, she is not responsible. But there is another way of dividing up cases. We can ask not about *agents*, in a way that is guided by our intuitions about who or what is an agent, but about *agency*. Rather than asking, with Fischer (for instance), "is the agent appropriately responsive and reactive to reasons?", we might ask "is appropriate responsiveness and reactivity to reasons exemplified?". Asking that question will yield a different way of picking out agents, and one that might be better justified.

To pave the way for this kind of move, it might be worth comparing Pereboom's Four Case Argument with a different thought experiment, Searle's *Chinese Room* experiment (Searle 1980; 1990). That thought experiment features a native English speaker, with no knowledge of Chinese, who is locked in the room with a box of Chinese symbols and an instruction book. Native Chinese speakers write questions in Chinese and post them through a slot in the room. The agent inside the room uses her instruction book to produce (to her, meaningless) symbols on a piece of paper, which she then posts back to the people outside. The instruction book and database of symbols is so cleverly constructed that she is able to intelligently answer any question put to her in Chinese. But she understands neither the questions nor the answers she returns. Searle thinks it is obvious that the agent

⁷ Moderate reasons responsiveness also requires the satisfaction of an ownership condition for Fischer, because moral responsibility is intrinsically historical. I set that condition aside here.

inside the room does not understand Chinese; he takes the thought experiment to show that rule-governed manipulation of symbols can never be sufficient for genuine thinking; syntax alone can never be sufficient for semantics.

One popular reply (or set of replies; here I abstract from the difference between the Systems Reply and the Virtual Minds Reply) argues that though the agent inside the room does not understand Chinese, there is an agent who does understand it. That agent is identical to or supervenes on the system as a whole: the person inside the room, the input and output devices, and the internal processes together constitute an agent who understands Chinese (see Cole 2014 for discussion). This reply changes our focus from agents to agency: from a (possibly arbitrary) focus on loci of agency held to be bounded by skins to the identification of appropriate response functionally. This same move is available to us in response to manipulation cases, of course. Instead of, or as well as, asking whether Plum is responsible, we can ask whether there is agency that meets the conditions for responsibility. Is the system composed of Plum *and the scientists manipulating him* appropriately responsive and reactive to reasons? Asking that question will help free us from the individualist prejudice, according to which agency is always a property of those entities we intuitively identify as agents, and that, in turn, will allow us to identify the actors who play an indispensable role in shaping the physical and natural world.

What, besides a taken for granted presumption in favor of individualism, explains why philosophers have not asked this question? Perhaps they take it to be inappropriate, for reasons centered on pragmatic considerations. A morally significant event has occurred—Ms. White has been murdered, say—and we want to know if someone is to blame, should be punished or have obligations to make amends. There is no one other than Plum and other individuals (say the neuroscientists who manipulate him) of whom to ask this question. This worry has some force: sometimes there is no agent who can be held to account, for pragmatic reasons or because there is no mechanism for holding them to account. But there are cases in which this is not true. More importantly, understanding the true loci of agency and thereby freeing ourselves from the grip of the individualist prejudice clears the way for better policy: if agency is in fact exercised by supra-agential aggregates, then understanding the world, and having a chance of shaping it as we would like, requires us to focus on such aggregates as well as (and sometimes instead of) individuals.

Suppose that the system composed of Plum and the neuroscientists exemplifies moderate reasons responsiveness but that neither Plum nor any individual neuroscientist meets these conditions. This might occur if, for some reason, no individual was in a position to grasp the nature of the actions they helped to bring about, because they each responded to a narrow slice of information in such a manner that the ensemble of responses manifested moderate reasons responsiveness. We might think that in such a case, we would have moral responsibility without an agent that could be blamed, punished or under an obligation to make amends. But that's far from obviously the case. If agency is constituted by a set of agents, then there *is* an entity that might be blamed or otherwise held to account. Of course there are practical problems with holding it to account *without* also (unjustifiably) holding the individuals who compose it to account, but these problems are familiar and sometimes surmountable. Nations may be held to account even when some or all of the individuals who help to constitute them cannot justifiably be; so can corporations. There are sometimes ways of holding the supervening agent to account that does not harm the composite agents: we may dissolve the supervening agent, for instance. I don't deny that there are very hard problems that may arise, but these are problems we face in other cases. They are real problems, I claim: they arise only *because* agency can be composed of other agents, rather than always being dependent on the boundaries defined by skin.⁸

⁸ There is a further difficulty that may arise in this context that does not arise with regard to nations and corporations: individuals like Plum and the neuroscientists might fail to constitute an aggregate agent across time. They may constitute a fleeting agent: by the time that anyone thinks to hold them to account, they may have dissolved. I grant that this may happen, but it has its individual level analogue too (we may sometimes blame agents who are dead, for instance). A further worry is that the boundaries of the actual extra-individual agents who actually play an essential role in distributing powers to individuals are always shifting, making it hard to draw boundaries. This fact, too, may make it hard to identify

It is worth emphasizing that it may happen that *both* the supervening agent *and* the individuals who compose it exemplify moderate reasons responsiveness. That might be the case in some of the iterations of the Four Case Argument, as Pereboom (2001; 2014) presents it. In some of those cases, the following three propositions may all be true: (1) the supervening agent may be appropriately reasons responsive; (2) some or all the manipulating neuroscientists are themselves reasons responsive, grasping the nature of the action they cause Plum to engage in, and (3) Plum himself is not moderately reasons responsive. This kind of case presents us with practical problems, once again, but not with any particular theoretical difficulties. The supervening agent may in such a case be held to account and so can some, but not others, of the individuals who compose it. Again, there may be closely analogous cases involving nations or corporations, since some individuals who compose them possess more power and more information with regard to the goals of such composite entities than others. Such corporate agents may be required to offer recompense or apology. They may be capable of the expression of guilt or shame, or analogues thereof; they may be fined or we may refuse to cooperate with them (Bjornsson and Hess 2016). In addition, we may expect the individuals who partially constitute them, or who benefit from their wrongdoing, to feel shame or guilt.

In the preceding pages, I have offered a (very sketchy) hint as to how we might respond to well-known thought experiments in the philosophy of agency, with a view to attempting to shake the hold that the exclusively individualist conception of agents has on us. While I have suggested, with my asides about nations and corporations, that these hints may have real-world implications, it may be that there are few composite entities in the world that actually exemplify the rich set of agential powers required by accounts like that of Fischer. Very often, individuals act together while failing to constitute a supervening agent. While that fact entails that, in these cases, there is no collective entity that is the bearer of moral responsibility, the social institutions that distribute powers and resources to individuals—amorphous and inchoate as they may be—often constitute the external scaffolding on which *individual* powers often depend. We cannot understand individual agency unless we understand these institutions and how they confer both powers and limitations on the individuals who are embedded within them.

Indeed, the cases in which individuals owe their powers to institutions which are not sufficiently integrated or appropriately structured to themselves be responsible collective entities are especially interesting, for the following reasons: In those cases in which it is true that, *had* these institutions possessed sufficient structure and feedback mechanisms to appropriately be held to account, we would blame them and not the individuals who are embedded in, or dependent on, them, those individuals are themselves blameless. The lack of responsible agency at the collective level does not bring it about that individuals possess it. If an organization is so structured that the individuals who partially constitute it lack access to sufficient information about their options or the effects of their choices to be morally responsible as individuals (and they are not responsible for their lack of knowledge), then they are blameless regardless of whether there is some other entity that can be blamed appropriately.

Consider, for example, the persistence of racism in Western nations. There is no doubt that the racist attitudes of individuals (conscious and unconscious) play an important role in the sustaining of systematic inequality and injustices. But structural and institutional racism may be as, or more, important than racist attitudes in producing harms, and these kinds of racism do not depend on the persistence of racist attitudes (Haslanger 2004; Glasgow 2009). Nevertheless, institutions are maintained in existence by the actions of individuals. Non-racist individuals may sustain racist institutions, when (for instance) they do not possess sufficient information to grasp the systematic effects of their actions. University admission offices, for example, might apply criteria that seem ‘color blind’, without realizing that they filter admissions in ways that favor those people who possess

an appropriate target to call to account. However, these barriers will not be encountered in every case. More importantly, they should not distract us from the non-punitive responses that are the most important: changes in policies, redistribution of resources and alterations in social practices.

superior (that is, preferred, for arbitrary reasons) cultural capital, rather than superior ability or knowledge. Even the blameworthy racist individual may not be blameworthy for the systematic effects of the institutions they help to sustain: being blameworthy for a token racist action need not entail being blameworthy for the pattern of actions which it helps to constitute.

More generally, I hope that by shaking the grip of the individualist view of agency, I thereby make readers more sensitive to other ways in which agency may be deeply dependent on external scaffolding. It is with this mind that I turn from a consideration of the control condition to the epistemic condition. Knowledge and the capacities to acquire it are themselves socially distributed, and the institutions (causally or morally) responsible for this distribution are appropriate loci of responses that ensure that it is better distributed and that significant truths come to be known. Further, the institutions responsible for individual-level failures of the epistemic condition may themselves constitute genuinely supra-individual agents.

2. The Epistemic Condition Beyond Individualism

Moral responsibility requires the satisfaction of an epistemic condition: agents who exercise control over their behavior may be excused in virtue of the fact that they are (non-culpably) ignorant of important facts concerning that behavior. For instance, until recently physicians failed to treat peptic ulcers with antibiotics, and thereby may have failed to avoid significant harms to many of their patients. But very plausibly they were not blameworthy for doing so, because they acted in the light of what they justifiably took to be medical knowledge at the time. It wasn't until Warren and Marshall established in the mid 1980s that peptic ulcers were caused by bacteria that it became reasonable to expect physicians to use antibiotics as a first-line treatment and (therefore) to blame them if they failed to do so.⁹ Prior to the mid 1980s, physicians failed to be blameworthy for their actions when they caused avoidable harms to peptic ulcer patients, because they failed to satisfy the epistemic condition on moral responsibility.

Individual physicians are obligated to apply the knowledge that they have learned. They are very rarely in a position to test medical claims for themselves. They can collect data on treatment and patient outcomes, but for the most part this data is uninformative, for a number of reasons. First, general practitioners see too few patients with any particular condition to generate statistically significant results (given that differences between treatment types are anything short of spectacular). Second, there are too many confounds in clinical practice to easily generate useful data through this kind of process: patients who suffer from a condition may differ in too many ways for comparisons to be meaningful. Third, physicians cannot use placebo controls, and their close involvement with individual cases entails vulnerability to cognitive biases, like the confirmation bias. Researchers use double blind procedures to avoid these biases, but physicians cannot use them.

These facts entail that for most knowledge claims, physicians find themselves at (or near) the end of chains of transmission for knowledge claims. They know very many of the things they know because they are the beneficiaries of testimony. In fact, almost everyone in science knows almost everything they know in very significant part in virtue of testimony. Consider the researchers who produce significant new knowledge. They use equipment designed by other researchers, which implement algorithms that they may be unable to understand. They use statistical techniques they could not have developed. They rely for the interpretation of their data on acquired knowledge they take for granted. Knowledge production in science is deeply distributed,

⁹ It is important to distinguish between medical researchers and general practitioners in making this claim. It is arguable that some medical researchers ought to have put the bacterial hypothesis to more stringent test much earlier; in retrospect, at least, we can see that the evidence against the hypothesis was weaker than was widely thought (see Zollman 2010 for discussion). However, individual physicians cannot and should not be expected to test established medical claims, except in very rare cases; rather, they discharge their epistemic obligations when they keep up with the current state of medical knowledge.

with each scientist dependent on very many others for their capacity to formulate hypotheses, test them and interpret their data.

Even within a single laboratory and with regard to a single research finding, knowledge production may be deeply distributed. A single paper, for instance, is typically the work of multiple authors, each of whom has a different set of skills. Some may have no interest in the topic of the paper at all: they are involved nevertheless because of their expertise with the equipment used or with the paradigm employed or in statistical analysis. Science has depended on the distribution of cognitive labor since its inception, but today the trend is toward deeper and deeper distribution of labor. Medical research is often carried out today by multiple laboratories in many countries, to ensure that a sufficiently large number of sufficiently diverse patients are enrolled and a sufficiently great range of expertise is brought to bear on the problem. For cutting edge physics, the distribution of cognitive labor is even deeper: the recent paper reporting the detection of the Higgs Boson particle had more than 5000 authors (Castelvecchi 2015).

It is in part *because* science is a deeply distributed enterprise that it is the paradigm of successful knowledge production. Individuals alone are not especially impressive at generating knowledge, but together they may explore the space of potential hypotheses much more successfully and—under truth-conducive conditions—cancel out one another's biases (in particular, the disposition of each of us to seek evidence in favor of a hypothesis and overlook evidence that conflicts with it). Diversity in anything that causes people to entertain conflicting hypotheses will ensure that one agent's confirmation bias will be compensated for by another's; each will be motivated to defend their hypothesis, and therefore produce arguments that undermine the other hypothesis. Deeper diversity (of cultural background, gender, life experience, political orientation) may allow for the detection of or compensation for those biases correlated with socio-cultural positioning. Whereas an all-male research group might overlook explanations of data that turn on women's agency, say, adding women to the research group may lead to a more thorough exploration of the space of hypotheses.¹⁰

The distribution of knowledge production is not a necessary evil, but is the key to our cognitive success (Mercier and Sperber 2011). Indeed, there is evidence to suggest that the cultural distribution of cognitive labor is an evolutionary adaptation: that is how we are designed (by nature) to come to understand, and thereby intervene in, the world. It is not just now, in the complex society in which we live today, or just with regard to the difficult questions that contemporary science faces, that knowledge production must be deeply distributed to be successful. In fact, we have *always* engaged in distributed knowledge production. Culture is central to our success as animals, and culture is in very important part the embodiment of knowledge accumulated by many individuals over many generations.

I use "culture" in the manner defined by Richerson and Boyd (2005, 5): culture is information capable of affecting individuals' behavior that is acquired from other members of their species through mechanisms of social transmission like teaching and imitation. In human beings (perhaps alone), culture is cumulative: new generations add to the stock of transmitted information. Cumulative culture often embodies knowledge that is inaccessible to individuals. Consider the very many cultural techniques required for survival in a harsh environment, like the Arctic. The indigenous people flourished in this environment where well-prepared expeditions could not, despite the advantages of 19th century science and the resources of one of the wealthiest empires ever known (Boyd, Richerson and Heinrich 2011). These expeditions were unable to acquire the knowledge needed for survival in this environment; it takes generations for such knowledge to accumulate and its range once acquired exceeds the resources of any single individual, such that it must be distributed across many people (so not only is its acquisition dependent on the work of many people over many generations, but it can be maintained only if there are sufficient numbers of individuals to sustain it). Distributed knowledge

¹⁰ Diversity is not sufficient on its own, of course. Groups may fail to outperform individuals at deliberation when dissenting individuals are unable to speak or be heard (Sunstein 2005). Too often, low status individuals are unable to express their views, or their opinions may be given little credibility (Fricker 2007).

acquisition is not just a fancy innovation that explains the relatively recent scientific takeoff; rather, it is one of the special tricks that explain human adaptive success (Richerson and Boyd 2005).¹¹

Given that our epistemic success is and always has been dependent on the deep distribution of cognitive labor, it is often the case that agents—even agents understood as identical to skin-bound individuals—satisfy or fail to satisfy the epistemic condition in virtue of their place in such distributed networks. Whether they are in a position to know that an action is harmful, for instance, is often dependent on what testimony they have received and the reception and assessment of testimony is dependent on social factors. Unsurprisingly (given that we are deeply cultural animals), human beings are disposed to accept the testimony of others. But our disposition to accept testimony is selective: we are sensitive to cues of lack of benevolence toward us and of incompetence in filtering testimony (Mascaro and Sperber 2009; Sperber et al. 2010). This disposition is obviously adaptive: the use of competence as a filter for reliability enables us to avoid acquiring false beliefs, while the use of lack of benevolence as a filter for reliability enables us to avoid exploitation by others. But cues for competence and for reliability are themselves dependent on social facts. We identify those to be trusted and mistrusted by reference to prevailing social norms and we identify competence by reference to socially transmitted credentials. Our early relatively promiscuous acceptance of testimony allows us to bootstrap our way to being appropriately selective in the testimony we accept: first we acquire, primarily from our caregivers, sensitivity to cues for reliability, and then subsequently we utilize those cues to acquire further testimony.

These facts entail that false claims, even obviously false claims, can come to be accepted by agents who might be said to believe virtuously. Cues to reliability are themselves socially distributed: we learn who counts as trustworthy and whose testimony is unreliable. We acquire this information in the first instance from those who *show* us that they are competent and trustworthy (caregivers, in the first instance, in most cases), but then we acquire them through mechanisms of social transmission like imitation, explicit teaching and through subtle cues from those we trust (how they orient themselves in conversation, say, to whom they listen attentively and who they interrupt), as well as from the narratives with which we are surrounded. All too frequently, these cues are distributed so that many true claims are filtered out and false claims filtered in, for example when the true claims are promulgated by those who are denigrated as the source of testimony, such as women and minorities. When this happens, the speakers suffer epistemic injustice (Fricker 2007) at the hands of hearers who may be epistemically virtuous agents. They count as epistemically virtuous because they filter claims in response to the cues in a way that has the function of increasing reliability, and because they cannot reasonably expect to do better by ignoring testimony (again, we are all deeply reliant on testimony for almost all of our beliefs).

Consider, for illustration, Allen Buchanan's description of his lucky escape from the racism of the American South during the 1950s and '60s. Explicitly and by example, he was taught to regard black people as subhuman. Those he had trusted and relied on—family members, his pastor, his teachers, local government officials, were 'sources of dangerous error, not truth'. It was through the good fortune of leaving this environment that he came to be in a position to appreciate the falsehoods for what they are (Buchanan 2004: 95). Our epistemic capacities are designed (by evolution) to absorb local cultural knowledge, because we require such knowledge to cope with

¹¹ Kukla (2012) argues that there are deep problems with regard to contemporary science arising from the extent to which it is distributed across individuals. No one is in a position to understand or to take responsibility for the results of such research, she maintains. As a consequence, there is no one who we can hold accountable for its truth claims. Kukla thinks that appeals to testimony cannot explain how such claims can be justified, because no one is in a position to test these claims for reliability. It may be that a shift in focus from individuals to collectives as loci not only of agency but also of knowledge may help to alleviate these concerns. It is worth emphasizing that the problem is not new (taboos, for example, may sometimes embody genuinely distributed knowledge that no individual is in a position to answer for or even to understand) (Henrich and Henrich 2010); and that it has implications that extend well beyond concerns for the epistemic status of scientific claims. Distributed knowledge—justified for individuals or not—is central to our responsibility, including our individual responsibility. If no one can justifiably be held responsible when they cannot individually vouch for knowledge claims, there will be precious few cases in which anyone satisfies the epistemic conditions on responsibility.

our natural and social environment. But when toxic falsehoods are promulgated alongside adaptive knowledge, we are apt to absorb those too. Filters on testimony are insensitive to the differences between truth and falsehood when both are culturally reinforced, and individual epistemic virtue is insufficient to avoid potential disaster. In cases like this, individuals may perpetrate racist acts but fail to be morally responsible for them: their epistemic bad luck entails that they do not satisfy the epistemic condition on moral responsibility.¹²

These considerations generalize broadly. Consider the politically charged issues that divide people within and across countries. Should taxation be higher than it currently is? Is climate change a serious challenge to humanity? Should coreligionists be accorded a higher moral value than nonbelievers and heretics? Each of these questions is one on which some people are recognized as experts and accorded socially mandated credentials (economists, climate scientists, political pundits, priests, and so on). The great mass of humanity has no special expertise on any of these topics and no one is genuinely expert on more than a very few. These are not questions that we can easily evaluate for ourselves.¹³ In fact, the impression that many of us *do* evaluate these questions for ourselves is largely an illusion: we evaluate these questions using tools and biases (virtuous and vicious) that dispose us to take some sources of evidence seriously and dismiss others, and these tools and biases are socially produced and distributed. We must and should defer to others on these matters, except where we are capable of genuine expertise on them (where genuine expertise requires thousands of hours of learning). Epistemically virtuous believers come to have false beliefs on these topics through no fault of their own. They therefore non-culpably fail to satisfy the epistemic conditions on moral responsibility.

If we are justifiably to hold agents morally responsible for actions which are caused or enabled by false beliefs, they must be culpable for having those beliefs. Perhaps there are cases in which agents satisfy this condition, but there are very many, significant, cases in which they do not, and they do not because of the way in which knowledge claims are socially transmitted. Social animals like us, with cognitive faculties designed (by evolution) to utilize social cues for truth and falsity, are vulnerable to responsibility-undermining epistemic luck: when we are in an environment in which morally significant falsehoods are promulgated by those who count as epistemic authorities – that is, those individuals and institutions that pass the relevant tests for epistemic authority – we are unlikely to be able to come to understand that these claims are false. We lack a reasonable opportunity to come to better views: either these views are lacking from our environment, or they are advocated by individuals who fail the relevant tests for epistemic authority. Buchanan was the beneficiary of good epistemic luck, cancelling out the bad epistemic luck of his having been socialized in an environment in which epistemic authorities promoted (vicious) falsehoods: he found himself in a new, and more truth-conducive, social environment.¹⁴ In cases like these, there may be some individual agents who are blameworthy (perhaps agents who have deliberately brought it about that false claims would come to be accepted; see Oreskes and Conway (2010) for examples of individuals who may fit this description). However, the great mass of believers probably are not blameworthy (at most, they are due only a tiny fraction of the blame which might accrue to all of us insofar as we have a role in sustaining the institutions which distribute testimony and the cues of reliability which filter it).

¹² Many philosophers argue that agents like Buchanan's counterpart, who did not have the good luck of escaping his toxic environment, would satisfy the epistemic condition because they would (typically) be culpable for their ignorance (see, for example, Moody-Adams (1997); FitzPatrick (2008)). Responding to these critics would take us too far afield. I have responded to their worries elsewhere (Levy 2003; 2009).

¹³ Despite the depth of our reliance on the distribution of cognitive labor, we tend to overvalue individual reasoning and undervalue group deliberation (see Mercier, et al. 2015).

¹⁴ Note that these remarks do not commit me to denying that culpable ignorance is possible. Rather, they delimit the circumstances in which ignorance can be culpable: it is a necessary condition of culpable ignorance that the agent be presented with what she herself recognizes as a genuine opportunity to put her beliefs to the test. Culpable ignorance, on my view (owed to Smith 1983) requires a "benighting act"; such an act occurs when the agent culpably passes up such an opportunity.

They believe virtuously, but their beliefs are false due to the way in which knowledge is socially sustained and distributed.

The mechanisms by which knowledge is socially produced, sustained and transmitted often entail that individuals fail to satisfy the epistemic conditions on moral responsibility for false belief. When we acquire such beliefs via the virtuous utilization of mechanisms designed to equip us for lives in our culture, we often acquire widely shared false beliefs alongside genuine knowledge. When we act on these false beliefs, we will (at least often) fail to be morally responsible because we do not satisfy the epistemic conditions on moral responsibility. The explanation for our ignorance will essentially cite social and cultural practices, institutions, and norms.

It does not follow, of course, from the fact that we may fail as individuals to be morally responsible for our actions, due to a failure to satisfy the epistemic condition that is itself due to forces working at a supra-individual level, that there is some supra-individual entity that *does* satisfy the epistemic condition. In fact, it will at best rarely be the case that such an entity exists. My primary aim is to not to identify targets to hold to account, but to clear the way for better practices. Focusing on the individual abstracted from her structuring, scaffolding and supporting context, as our preoccupation with individual-level responsibility encourages us to do, is an obstacle to responses that focus on restructuring the social environment (think, here, of how corporate resistance to commonsense measures to reduce sugar intake advocates that instead we take “personal responsibility” for our diets).¹⁵ Recognizing that the individual is embedded in complex social networks that distribute knowledge paves the way for a focus on the redistribution of resources, and on moving beyond blame and searching for evidence-based solutions. Traditional theories of responsibility focus us on the individual and her internal capacities (to control her actions, to regulate beliefs, to endorse her volitions, or what have you). More recent theories extend this concern to the collective and its powers to act and regulate itself. This leaves unexplored agents’ embedding in social networks, in power relations, in culture: the ways in which the collective form and sustain the individual. We cannot fully comprehend responsibility by dividing the territory into the individual, on the one hand, and the collective, on the other. We need as well to understand how the collective structures the individual. Identifying the social and cultural mechanisms responsible for our beliefs reduces the temptation to focus on individuals, and thereby enables us to focus our attention where it should be.

Conclusion

In the first half of this paper, I focused on the control condition on moral responsibility. I suggested that if we utilize *agency* as a criterion for identifying *agents*, we will sometimes find that agents are not always identical to individuals, where ‘individuals’ are skin-bound bodies. Instead, agents sometimes extend beyond individuals and into the world. Extended agency is not always collective agency: the agent may owe her powers and capacities to other individuals and to institutions without coming to constitute a higher-level entity. Sometimes, the locus of agency dissociates from the person, without thereby becoming collective. In the second half of the paper, I turned from the control condition to the epistemic condition. I argued that agents may fail to satisfy the epistemic condition on moral responsibility because of the way they are socially situated. They may exercise their epistemic faculties virtuously but acquire false beliefs because the socially produced cues to which these faculties are sensitive promote false beliefs and denigrate true. In this concluding section, I will draw the two threads of this paper together.

If the cues to reliability to which our testimonial mechanisms are sensitive dissociate from genuine reliability, whose fault is that? It may be no one’s fault: very often there is no agent responsible for a state of affairs. But many widely accepted false claims (that global warming is not occurring, or is not human induced, or is not a serious problem, for instance) arise and are maintained in very important part because of the actions of

¹⁵ One example among literally thousands: “Americans need to be more active and take greater responsibility for their diets”, according to the CEO of Coca Cola (Kent 2009).

individuals and of institutions that can reasonably be thought to possess and to exercise agency. I have already mentioned the “merchants of doubt” who may be individually responsible for knowingly promulgating false beliefs. But they could not have succeeded without the cooperation of institutions that fail to challenge them or are complicit in their behavior.

For a signal instance, take the media. The media gives the merchant of doubt a voice. Some media institutions may be identical to merchants of doubt, but they would not be taken seriously if many others did not accord the skeptics a degree of respectability out of proportion to their epistemic standing, due to the widespread commitment to a norm of “balanced reporting”. As this norm is often understood, reporting is balanced only if both sides are given an equal hearing. This is a norm that ought to be rejected or reinterpreted when we have good reason to think that one side is neither justified in its claims nor motivated by the search for truth or by a concern for the wellbeing of ordinary individuals. Of course, the media can no more assess the climate science than can other non-experts. But it need not do so in order to develop a rubric for giving a hearing to the sides: it can refer to the scientific consensus and weight voices accordingly. There is a strong scientific consensus on the reality of anthropogenic global warming: many studies have found that approximately 97% of relevant experts hold that it is occurring, that it is human induced and that it is a very serious problem (Cook et al. 2013; Carlton et al. 2015). Yet mainstream media organizations give both sides equal, or nearly equal, coverage (Boykoff and Boykoff 2004; Theel, Greenberg, and Robbins 2013). Other institutions—schools, universities and parliaments, for instance—may be complicit in the fact that this situation goes without significant challenge.

There may be a case for holding that some or all of these institutions, perhaps together with some individuals, constitute genuine loci of agency, which can be held responsible for the fact that many individuals fail to satisfy the epistemic condition. Someone may be to blame, or perhaps *something*; there may be genuine agents who ought to be held responsible and who have an obligation to correct the state of affairs, where these agents are not identical to any individuals.

Even if these institutions do not constitute agents who may justifiably be held responsible for the ignorance of so many, there is a strong case for thinking that they nevertheless play a significant role in the structuring of the agency of others—individuals and collectives—such that correcting the widespread ignorance will require reforming the institutions. The almost obsessive focus on the individual as the only locus not only of agency but also of responsibility has been counterproductive. It has focused us on arguing with one another, on identifying villains and on denigrating those who come to hold false beliefs. If we set aside this obsession with individuals, the way may be clear for us to identify the agents who are the most important actors on this stage and then in playing our individual and distributed parts in bringing them to play it better.¹⁶

References

- Björnsson G., Hess K. Corporate Crocodile Tears? On the Reactive Attitudes of Corporate Agents. *Philosophy and Phenomenological Research*. 2016. doi: [10.1111/phpr.12260](https://doi.org/10.1111/phpr.12260).
- Boyd R, Richerson P.J., Henrich J. The Cultural Niche: Why Social Learning is Essential for Human Adaptation. *Proceedings of the National Academy of Sciences*. 2011;108(26):10918–10925. PubMed PMID: 21690340.
- Boykoff M.T., Boykoff J.M. Balance as Bias: Global Warming and the US Prestige Press. *Global Environmental Change*. 2004;14(2):125–136.
- Buchanan A. Political Liberalism and Social Epistemology. *Philosophy and Public Affairs*. 2004;32(2):95–130.

¹⁶ I am extremely grateful to Katrina Hutchison, Catriona Mackenzie, and Marina Oshana for extensive comments on this paper; every page and every argument is substantially better for their input. I am also grateful to Wendy Carlton for capable copyediting and advice. Work leading up to this paper was supported by the Oxford Martin School and the Wellcome Trust (WT104848/Z/14/Z).

- Carlton J. S., Perry-Hill R, Huber M., Prokopy L. S. The Climate Consensus Extends Beyond Climate Scientists. *Environmental Research Letters*. 2015;10:094025.
- Castelvecchi D. Physics Paper Sets Record with More than 5,000 Authors. *Nature News*. 2015;15(May)
- Ciurria M. Moral Responsibility Ain't Just in the Head. *Journal of the American Philosophical Association*. 2015;1(4):601–616.
- Clark A., Chalmers D. *The Extended Mind*. *Analysis*. 1998;58(1):7–19.
- Clarke, R. 2003. *Libertarian Accounts of Free Will*. New York: Oxford University Press.
- Cole, D. 2014. "The Chinese Room Argument." *Stanford Encyclopedia of Philosophy*, edited by E. N. Zalta. < Available at: <http://plato.stanford.edu/entries/chinese-room/> >
- Cook J., Nuccitelli D., Green S. A., Richardson M., Winkler B., Painting R., Way R., Jacobs P., Skuce A. Quantifying the Consensus on Anthropogenic Global Warming in the Scientific Literature. *Environmental Research Letters*. 2013;8(2):1–7. A.
- Doris, J. M. 2015. *Talking to Our Selves: Reflection, Ignorance, and Agency*. Oxford: Oxford University Press.
- Fischer, J.M. and Ravizza, M. 1998. *Responsibility and Control: An Essay on Moral Responsibility*. Cambridge: Cambridge UP.
- Fischer, J.M. 2006. *My Way*. New York, NY: Oxford UP.
- FitzPatrick W.J. Moral Responsibility and Normative Ignorance: Answering a New Skeptical Challenge. *Ethics*. 2008;118(4):589–613.
- Frankfurt H. Freedom of the Will and the Concept of a Person. *Journal of Philosophy*. 1971;68(1):5–20.
- Fricker, M. 2007. *Epistemic Injustice*. Oxford: Oxford University Press.
- Glasgow J. Racism as Disrespect. *Ethics*. 2009;120(1):64–93.
- Hackman D.A., Farah M. J., Meaney M. J. Socioeconomic Status and the Brain: Mechanistic Insights from Human and Animal Research. *Nature Reviews Neuroscience*. 2010;11(9):651–9. M.J. PubMed PMID: 20725096.
- Haslanger, S. 2004. "Oppressions: Racial and Other." In *Racism in Mind* Levine, edited by M.P. and Pataki, T., 97-123. Ithaca, NY: Cornell University Press.
- Henrich J., Henrich N. 2010. "The Evolution of Cultural Adaptations: Fijian Food Taboos Protect Against Dangerous Marine Toxins. *Proceedings of the Royal Society B: Biological Sciences*. 1701;277:3715–3724. PubMed PMID: 20667878.
- Kane, R. 1996. *The Significance of Free Will*, New York: Oxford University Press.
- Kent, M. 2009. "Coke Didn't Make American Fat." *Wall Street Journal*, October 7.
- Kukla R. "Author TBD": Radical Collaboration in Contemporary Biomedical Research. *Philosophy of Science*. 2012;79(5):845–858.
- Levy N. Cultural Membership and Moral Responsibility. *The Monist*. 2003;86(2):145–63.
- Levy N. Counterfactual Intervention and Agents' Capacities. *Journal of Philosophy*. 2008;105(5):223–239.
- Levy N. Culpable Ignorance And Moral Responsibility: A Reply To FitzPatrick. *Ethics*. 2009;119(4):729–741.
- Levy, N. 2011. *Hard Luck: How Luck Undermines Free Will and Moral Responsibility*. New York: Oxford University Press.
- Levy, N. 2015. "Frankfurt in Fake Barn Country." In *The Philosophy of Luck*, edited by Duncan Pritchard and Lee Whittington, 79-92. Hoboken: Wiley, 2015.
- Mascaro O., Sperber D. The Moral, Epistemic, and Mindreading Components of Children's Vigilance Towards Deception. *Cognition*. 2009;112(3):367–80. PubMed PMID: 19540473.

- Mckenna M. A Hard-line Reply to Pereboom's Four-case Manipulation Argument. *Philosophy and Phenomenological Research*. 2008;77(1):142–159.
- Mele. A. 2006. *Free Will and Luck*. Oxford: Oxford University Press.
- Mercier H., Sperber D. Why do Humans Reason? Arguments for an Argumentative Theory. *Behavioral and Brain Sciences*. 2011;34(2):57–111. PubMed PMID: 21447233.
- Mercier H., Trouche E., Yama H., Heintz C., Girotto V. Experts and Laymen Grossly Underestimate the Benefits of Argumentation for Reasoning. *Thinking and Reasoning*. 2015;21(3):341–355.
- Moody-Adams, M. M. 1997. *Fieldwork in Familiar Places: Morality, Culture, and Philosophy*. Cambridge, MA: Harvard University Press.
- Oreskes, N. and E. M. Conway. 2010. *Merchants of Doubt*. London: Bloomsbury Press.
- Pereboom, D. 2001. *Living without Free Will*. Cambridge: Cambridge UP.
- Pereboom, D. 2014. *Free Will, Agency, and Meaning in Life*. New York: Oxford University Press.
- Richerson, P.J. and R. Boyd. 2005. *Not by Genes Alone*. Chicago: University of Chicago Press.
- Rovane, C. 1998, *The Bounds of Agency*. Princeton: Princeton University Press.
- Searle J. Minds, Brains and Programs. *Behavioral and Brain Sciences*. 1980;3(3):417–57.
- Searle J. Is the Brain's Mind a Computer Program? *Scientific American*. 1990;262(1):26–31. PubMed PMID: 2294583.
- Seppinwall A. J. Corporate Moral Responsibility. *Philosophy Compass*. 2016;11(1):3–13.
- Smith H. Culpable Ignorance. *Philosophical Review*. 1983;92(4):543–71.
- Sperber D., Clément F., Heintz C., Mascaro O., Mercier H., Origg G., Wilson D. Epistemic Vigilance. *Mind & Language*. 2010;25(4):359–393.
- Steward, H. 2012. *A Metaphysics for Freedom*. Oxford: Oxford University Press.
- Sunstein C.R. Group Judgments: Statistical Means, Deliberation, and Information Markets. *New York University Law Review*. 2005;80(3):962–1049.
- Theel, S. Greenberg, M. and Robbins, D. 2013. “Media Sowed Doubt In Coverage Of UN Climate Report” < Available at: <http://mediamatters.org/research/2013/10/10/study-media-sowed-doubt-in-coverage-of-un-clima/196387> >
- Vargas, M. (forthcoming). “Implicit Bias, Responsibility, and Moral Ecology.” In David Shoemaker (ed). *Oxford Studies in Agency and Responsibility*, Vol. 4. Oxford: Oxford University Press.
- Vihvelin, K. 2013. *Causes, Laws, and Free Will: Why Determinism Doesn't Matter*. Oxford: Oxford University Press.
- Zollman K.J.S. The Epistemic Benefit of Transient Diversity. *Erkenntnis*. 2010;72(1):17–35.