

Knowledge From Vice: Deeply Social Epistemology.

NEIL LEVY

Macquarie University and University of Oxford

neil.levy@philosophy.ox.ac.uk

MARK ALFANO

Delft University of Technology and Australian Catholic University.

mark.alfano@gmail.com

In the past two decades, epistemologists have significantly expanded the focus of their field. To the traditional question that has dominated the debate — under what conditions does belief amount to knowledge? — they have added questions about testimony, epistemic virtues and vices, epistemic trust, and more. This broadening of the range of epistemic concern has coincided with an expansion in conceptions of epistemic agency beyond the individualism characteristic of most earlier epistemology. We believe that these developments have not gone far enough. While the weak anti-individualism we see in contemporary epistemology may be adequate for the kinds of cases it tends to focus on, a great deal of human knowledge production and transmission does not conform to these models. Furthermore, the dispositions and norms that are knowledge-conducive in the familiar cases may not be knowledge-conducive across the board. In fact, dispositions that, at an individual level, count as epistemic vices may be epistemic virtues in common social contexts. We argue that this overlooked feature of human social life means that epistemology must become more deeply and pervasively social.

1. Introduction

In the past two decades, epistemologists have significantly expanded the focus of their field. To the traditional question that has dominated the debate — under what conditions does an individual's belief amount to knowledge? — epistemologists have added questions about testimony and its aptness to justify belief (Coady 1995), the conditions under which it is epistemically permissible for agents to make assertions (Turri 2016) and to whom they should make them (Alfano & Robinson 2017), the dispositions hearers need to embody epistemic virtues such as intellectual humility (Tanesini 2018; Whitcomb et al. 2015; Alfano et al. 2017), when and to whom we should extend our epistemic trust (Goldberg 2010, 2013; Alfano 2016), epistemic injustice in both the generation and the reception of claims (Fricker 2007), the nature and effects of vices such as epistemic insouciance (Cassam 2016, forthcoming), the nature and effects of propaganda (Stanley 2015), and more.¹ In short, epistemology is now concerned with epistemic responsibility and irresponsibility in all their forms.

¹ For a review, see Turri, Alfano, & Greco (2017).

This broadening of concern has coincided with and in part been driven by an expansion in conceptions of *epistemic agency*. For most of its modern history, epistemology has been resolutely individualistic, as exemplified by the British Royal Society’s motto: *nullius in verba*. Testimony and assertion, however, are essentially social (Coady 1995, chapter 11; Alfano 2016); they force us to extend our theorizing beyond what is happening in the head of a single individual. Recognizing this fact, epistemology has become social and in many cases explicitly anti-individualistic (Kallestrup & Pritchard 2012; Carter et al. 2014).

We applaud these developments. But we believe that the account of epistemic agency that has emerged from the tentative turn to anti-individualism (Palermos 2016) in contemporary epistemology is inadequate. In this paper, we argue that some of our most significant epistemic achievements rest on dispositions that cannot be understood within an individualistic framework. Nor can these knowledge-conducive dispositions be understood as epistemic virtues. The dispositions involved in the production and maintenance of *cumulative culture*, the knowledge on which we focus here, are more akin to vices than virtues.

In focusing on epistemic virtues and vices, we do not assume the truth of virtue epistemology in any of its forms. We do not, that is, take sides in debates over whether knowledge *is* ‘true belief out of intellectual virtue’ (Sosa 1991, p. 277; see also Zagzebski 1996, p. 272). Nor do we commit to particular accounts of the nature of these virtues. Rather, we focus on what we take to be a common view, cutting across virtue epistemology and its rivals, that there are intellectual dispositions that are (quite) reliably knowledge-conducive, and therefore worth cultivating (Olin 2017). We will argue that dispositions which are more commonly seen as epistemic vices than virtues are reliably associated with the acquisition of cumulative knowledge. We will, moreover, resist one possible response to this claim: reclassifying these dispositions as epistemic virtues. We will argue that these dispositions do not count as responsibilist epistemic virtues, because they are not competences in virtue of which agents deserve credit for their knowledge.²

² It is worth contrasting our approach with two recent defences of closedmindedness. Battaly (2018) argues that in some “one off” cases and in certain hostile environments, closedmindedness might be what she calls an effects-virtue. In contrast, we think that some apparent epistemic vices are *routinely* effects-virtues, but this fact isn’t sufficient grounds for reclassifying them as genuine virtues. Fantl (2018) advocates refusing to engage with apparently

Here is the plan for this paper: in the next section, we sketch the lay of the anti-individualist land in contemporary epistemology. Next, we introduce the phenomenon of cumulative culture and argue that it is an exemplary kind of knowledge. In consequence, epistemological theories that do not accommodate cumulative culture stand in need of revision. We then explore *how* cumulative culture is developed, passed on, and maintained within a community. The practices, norms, and institutions that make cumulative culture possible are, we argue, at odds with most contemporary accounts of knowledge and its generation, which are too committed to individualism about cognitive agency. In the following section, we explore possible replies on behalf of individualism about cognitive agency. One reply (inspired by the work of the so-called Parisian school of cultural evolution) holds that cultural knowledge embodies sufficient individual intelligence to ground epistemic credit. A second assigns credit to those who transmit cultural knowledge (if not to recipients). A third argues that cultural knowledge is not the kind of knowledge with which virtue epistemology is concerned, while a fourth attempts to distribute (minimal) credit across multiple agents. We will argue that all these replies fail. In the final section, we argue that cumulative culture is fostered by, among other things, interlocking epistemic dispositions that have traditionally (and perhaps rightly) been regarded as vices.

2. Contemporary anti-individualism

For most of its history, epistemology — like almost all branches of philosophy — has been individualistic (Goldberg 2010; Pritchard 2018). Knowledge and other epistemic statuses were held to be states of individuals. A number of factors gradually came to shake that individualistic consensus. Broader currents in the philosophy of mind helped to shake individualistic certainties. One such current was the rise of content externalism (Putnam 1975; Burge 1979), which emphasised how mental contents may be dependent on the social and environmental setting of agents. A second, more revolutionary, stimulus was vehicle externalism (Clark & Chalmers 1998), which suggested that the vehicles of mental states may themselves be distributed

compelling arguments when we know the conclusions to be false. For him, closeness is a way of retaining justification in certain, limited, contexts, rather than a disposition of the kind we have in mind, which is a disposition that conduces to the generation of new knowledge.

across agents and the world. Concerned as it is with the normative properties of mental states such as belief, ignorance, and knowledge, epistemology had to respond to these movements.³

In addition, developments internal to epistemology converged to shake individualistic assumptions. Many of the responses to Gettier's (1966) criticism of the justified true belief account of knowledge implied or presupposed that knowledge can be dependent on factors outside the individual. More recently, the recognition that a high proportion of our beliefs depend on the testimony of other agents broadened the scope of epistemology, from the normative status of beliefs to the properties of belief transmission (Coady 1995). For many philosophers, these currents have served as a spur to the development of home-grown anti-individualism within epistemology. One route to anti-individualism proceeds through Gettier cases and the like. Two agents, *A* and *A'*, might differ from one another in whether their true belief that *p* constitutes knowledge, due entirely to features external to their heads. For instance, *A*'s belief, but not *A'*'s, might be safe due to the proportion of fake barns in the vicinity. Another route might proceed via testimony. *H* and *H'* might differ from one another in whether their true belief that *p* constitutes knowledge, due entirely to features external to their heads. For instance, *H*'s belief, but not *H'*'s, might be safe due to the proportion of epistemically faulty testifiers in the vicinity (Kallestrup & Pritchard 2016).

While these anti-individualist developments are welcome, they remain limited. Anti-individualists hold that the epistemic properties of beliefs and cognitive processes associated with epistemic agency may be partially dependent on properties of the broader environment. But the states and processes themselves continue to be those of individuals — often of individuals who parcel out the credit for knowledge among themselves — and knowledge-transmission is conceived of as essentially dyadic: transmission from speaker to hearer. When third parties are glimpsed, they serve only as the background against which transmission and formation of beliefs occurs. Their reliability may make the epistemic environment friendlier for transmission by monitoring it (Goldberg 2010), policing it (Goldberg 2011), or strategically rewiring it (Alfano & Robinson 2017).

Within mainstream epistemology, there is ongoing debate over individualism, but few participants contemplate anything stronger than a weak anti-individualism.⁴ In the next section,

³ Note that while all internalist views have (hitherto at any rate, if not necessarily) been individualist, externalist views have only sometimes been anti-individualist. For instance, agent reliabilism, in some of its forms, is individualist.

⁴ Pritchard (2018) is a rare exception, but he contemplates a stronger anti-individualism only to dismiss it. Alfano & Skorburg (2017, 2018) are the other main exceptions. It should be noted, however, that there is a burgeoning literature on the epistemic dimensions of democracy, which moves beyond the dyadic paradigm typical in epistemology (see List & Goodin (2001) for a representative and highly influential sample).

however, we describe important and pervasive epistemic practices that cannot be accommodated by the limited anti-individualism prevailing in epistemology.

3. Cumulative culture as exemplary knowledge

In this section, we describe a type of knowledge that arises from cumulative culture. We believe that every epistemological theory has to acknowledge the epistemic bona fides of such knowledge, and in subsequent sections we argue that many contemporary theories fail this test.

Human beings are deeply social animals. Group living — initially, in the bands characteristic of hunter-gatherers, and then in increasingly large sedentary populations — has always been central to our survival (Dunbar 1992). At least as early as half a million years ago, *Homo heidelbergensis* depended on planned hunting for survival, and this kind of hunting requires cooperation. Alloparenting, the cooperative raising of offspring, is probably even older (Hrdy 2009). Our reliance on the division of labor, especially cognitive labor, may have increased across time but even for hunter-gatherer bands, knowledge production was a collective enterprise.⁵

Knowledge production has always been a collective enterprise for *Homo sapiens* because we are dependent on cumulative culture (Richerson & Boyd 2005; Boyd, Richerson and Henrich 2011). While many other species of animal have behavioral traditions (like the famous potato washing of Japanese macaques), only humans appear to have *cumulative* culture: culture that is not merely transmitted but which then becomes a platform for the next round of innovation. Each generation typically builds on the cultural achievements — artifacts, rituals, but also and centrally cultural knowledge⁶ — of its predecessors. Cumulative culture allows us to flourish in environmental conditions that are relatively resource-poor, in ways that we could not otherwise do. We can discover environmental regularities and successful practices by accumulating the results of experimentation and exploration across generations

Cumulative culture depends on methods of social transmission ranging from intentional, structured pedagogy to unreflective imitation to spying and corporate espionage. The different methods support different aspects of cumulative culture and may sometimes be in tension with one another (Sterelny 2012). For example, explicit pedagogy may assist innovative improvements

⁵ There are possible exceptions. Marlowe's (2010) ethnography of the Hadza, a forager people of East Africa, emphasises how little specialization there was (we thank a reviewer for drawing our attention to this work). However, there is certainly some degree of cognitive specialization among the Hadza. Marlowe himself recounts how he and some adolescent males became ill from eating tubers they had gathered, because they lacked the specialist knowledge of the women who usually gathered and prepared them.

⁶ This way of putting it is slightly misleading, inasmuch as artifacts and rituals are often partially constitutive of cultural knowledge.

to inherited cultural artifacts, practices, and norms, while (relatively) blind imitation supports faithful transmission of cultural knowledge between generations (Fridland 2018). The former may depend on the sorts of dispositions often celebrated as intellectual virtues in contemporary epistemology, but the latter — which is equally essential for the development of cumulative culture and the knowledge it embodies — presents a very different picture.

To illustrate, consider the dizzying array of cultural practices required for flourishing in the Arctic (Boyd, Richerson and Henrich 2011; Henrich 2017). Traditional Arctic life required knowledge of how to make special clothing, to manufacture and use special tools for hunting, to construct snow houses, to build fires for cooking and melting water (without access to wood!), and much more. Each of these skills can be broken down into others that are independent of one another, and each of them is difficult to acquire due to its complexity. Take clothing. The Inuit stayed warm and comfortable by making clothing from caribou skin, which has better insulation properties than seal or polar bear fur. But not just any caribou skin would do: it had to be harvested at the right time of year, and then prepared by repeated stretching, scraping, and moistening. Subsequently the hides had to be shaped in ways that maximized heat retention while allowing moisture to escape. A ruff of wolverine fur, especially selected for length, was then added. Footwear was equally specialized, consisting of five separate layers: three different layers of stockings, each with a different design, then two different kinds of boots. The know-how needed to make just one of these items of clothing is difficult to acquire, let alone the whole package. Moreover, individual elements of the package (for example, knowledge of how to prepare one of the layers of stockings) are often of little use by themselves: it is only when they play a role in the entire package that they make a significant difference to survival.

If each element of the package is difficult to acquire on its own, and many elements have little value on their own, how did the Inuit succeed in acquiring the whole package? Almost certainly, the answer involves multigenerational accumulation of knowledge, with many different group members each playing a small part in the acquisition of the propositional and procedural knowledge of the group as a whole. Cumulative culture, an irreducibly and deeply collective enterprise, is essential to the acquisition of the kinds of knowledge that allows human groups to flourish in highly diverse environments, from the deserts to the tundra and the tropics.

Two facts together strongly suggest that cumulative culture is central to ecological success: (1) when important skills are lost, they are very difficult to reacquire, and (2) well-prepared explorers have often found the conditions in which local peoples flourished fatally difficult. Nineteenth-century history provides illustrations of both these facts. (1) The Polar Inuit seem to have been struck by an epidemic some time in the 1820s that killed many older members of the

community (Boyd, Richerson and Henrich 2011). As a result, they lost important skills, and the group entered a decades-long decline in numbers. This decline was halted only when they reacquired the skills from a migrating group around 1862. They were not able to reinvent the lost skills of kayak and snow house design in the intervening decades, despite the possession of a suite of related skills and strong motivation to do so. (2) There are many instances of well-prepared explorers dying in the territories of indigenous peoples (or surviving only with the aid of the local people) because they could not acquire the skills needed for survival. In the mid-nineteenth century, for example, the ships commanded by Sir John Franklin, an experienced Arctic explorer, became stuck in the sea ice at King William Island (Boyd, Richerson and Henrich 2011). This island is regarded by the local Netsilik people as especially rich in resources. Nevertheless, everyone aboard perished.

While the Arctic environment is a harsh one, requiring specialized skills of great complexity, it is by no means a unique case. Food preparation techniques across the world illustrate the counterintuitive complexity of the cultural practices required for flourishing. The Yantruwanta people of the Lakes area of South Australia prepared the seeds of the Nardoo plant by roasting the seed cases before grinding the seeds and mixing them with water prior to cooking. When members of the Burke and Wills expedition ran out of food, they attempted to prepare Nardoo as they had seen the Yantruwanta people do. But they missed the crucial step of roasting the seed cases; as a result, the Nardoo leached thiamine out of their systems, hastening their deaths (Burcham 2008). A number of indigenous peoples in South America were heavily reliant on cassava, a highly toxic plant. They detoxified it in a variety of complex and time-consuming ways. The Tukanoans, for example, scraped, grated, and washed the roots to separate the fibre from the starch and liquid. The liquid is then safe for consumption after boiling, but the fibre and starch must sit for two days before they can be baked and eaten. When cassava was exported to West Africa at the beginning of the seventeenth century, it rapidly became a staple crop. However, knowledge of how to detoxify it was not exported along with the crop, and chronic cyanide poisoning remains a problem in some areas today (Henrich 2017).

Cassava and Nardoo are not unusual among food crops in requiring extensive processing prior to consumption if they are to be eaten without ill-effects. Edible plants often evolve toxins to deter animals from eating them: it took centuries of selective breeding to produce the easily edible varieties we know today. Indigenous groups selectively bred these crops themselves, but they did not succeed in eliminating all toxins. They coped – and flourished – by developing complex methods of detoxifying the plants (Henrich 2017).

4. Cultural knowledge generation and epistemic agency

What kind of epistemic practices are required for groups of human beings to acquire cumulative cultural knowledge? At first glance, it might appear that the virtues and norms epistemologists typically value are just the virtues and norms needed for such acquisition: intellectual curiosity might motivate experimentation; open-mindedness might underlie the disposition to learn from others; and a knowledge (or justification) norm of assertion might ensure that knowledge acquired by certain members of one generation is trusted and thereby transmitted to the next (Craig 1999). We think the prospects for assimilating the epistemic processes underlying cumulative culture to the paradigms of knowledge generation and transmission that take centre stage in contemporary epistemology are gloomy, however, and the kinds of norms and dispositions that appear appropriate to the latter are inappropriate or even counterproductive in the former.

4.1. Transmission through diffidence and overimitation

Cumulative culture is adaptive because the kinds of problems presented by cassava, the Nardoo plant, and the harsh Arctic winter are resistant to the cognitive powers of smart animals like ourselves when we attempt to solve them on our own or over timespans shorter than generations. The connection between the plants that must be detoxified and ill-health, for example, is causally opaque; only some people develop the symptoms of cyanide poisoning, for instance, and then only after a prolonged period of reliance on cassava. Individual cognition is ill-equipped to detect such distal, stochastic causal regularities.

It is striking, in fact, how difficult even recent science has found the very problems that cumulative culture managed to solve. Corn has been a staple crop for many indigenous groups in the Americas for centuries. Corn contains niacin, but it is not released through ordinary cooking. Amerindian groups introduced an alkali into corn-based foods (for instance, ash from the right kind of wood), thereby releasing the niacin. By the 18th century, corn had become a staple food crop in some parts of Europe, but cooking techniques were not exported along with the crop, and pellagra (ranging in seriousness from hair loss through to dementia and even death) became a serious problem. In the late 19th century, the problem spread to the southern states of the United States and became endemic. Early 20th-century science advanced a number of hypotheses concerning the causes of pellagra. Though corn-based diets were widely recognized as responsible, the cause was typically held to be an infectious agent found in spoiled corn. It wasn't until 1914 that an American physician, Joseph Goldberger, disproved the infectious agent theory. However, Goldberger's theory was widely derided by the medical community of the time, and outbreaks of pellagra remained a serious problem in the United States into the 1940s

(Lanska 2010). No doubt modern science succeeded much more rapidly (and much more explicitly) than the mechanisms of cumulative culture in solving the problem, but it is striking how difficult the problem proved for twentieth century science.

Were Amerindian groups better scientists or more careful experimenters than the US scientists who failed to identify the causal mechanism underlying pellagra? Perhaps. But there is a great deal of convergent evidence suggesting that the kind of impressive cultural knowledge that enables the flourishing of indigenous groups across the world is owed to epistemic dispositions and norms that, by the yardstick of contemporary epistemology, look more like vices than virtues. Conversely, in many of these cases, agents who manifested the epistemic virtues (as commonly conceived) would not have been able to acquire or transmit the knowledge embodied in cumulative culture, and therefore would not have been able to flourish in their environments.

Cumulative culture is dependent on mechanisms of intergenerational knowledge transmission. We have multiple adaptations for such transmission, ranging from dispositions to accept testimony (Mascaro & Sperber 2009; Sperber et al. 2010; Harris 2012) to a disposition to copy the behavior of prestigious individuals (Henrich & Gil-White 2001; Chudek et al. 2012). At least some of these dispositions to acquire behaviors and beliefs from other individuals manifest extreme epistemic deference. We exhibit a *conformist bias* (Henrich & Boyd 1998): a disposition to copy the behavior and acquire the beliefs of the majority relatively unselectively and unreflectively. We are disposed to copy the practices of our culture, despite either failing to understand the need for such practices or even thinking that elements of such practices are superfluous. These are not epistemic dispositions seen only in indigenous groups. Rather, they are adaptations for culture that arise in normal human development the world over.

We use the verb ‘to ape’ to describe slavish imitation. In fact, no other animal – not even the other apes – apes as much as we do. There is experimental evidence that we are *overimitators*: whereas other animals copy behaviors when they see that these behaviors are instrumentally rational, human beings are disposed to copy even those elements that do not appear to be required for pursuing a goal. For example, Nagell, Olguin & Tomasello (1993) compared the performance of human children and chimps on the acquisition of a skill demonstrated to them. Experimenters used a rake, *tine side down*, to draw sweets that were otherwise out of reach toward themselves. Using a rake that way is highly inefficient: many sweets slip through the gaps in the tines. When they were given the opportunity to perform the task themselves, chimps flipped the rake so that the flat side acted as a more efficient tool, with fewer sweets escaping. But human children tended to imitate the actions demonstrated. Later experimental work demonstrated that this disposition to overimitate is selective: children

overimitate when the actions appear to be intended by the model, regardless of whether they see the point of the actions. Thus, while children who observed a model turning on a light by butting it with her head tended to do the same, rather than use their hands (Meltzoff 1988), imitation dropped significantly if the model's hands were occupied, suggesting that the decision to use her head was not a component of how things are supposed to be done (Gergely, Bekkering & Király 2002).

Why do we overimitate, compared to other animals? Because we are dependent on cumulative culture, and they are not. Because we live in such a wide diversity of environments, no small set of behaviors could come to be encoded in our genome, such that we would require only environmental triggers to develop the right responses. Instead, the right responses are those our group has gradually developed over multiple generations. Because the relations between problems and evolved solutions are complex and causally opaque, we cannot develop these solutions for ourselves from scratch. We often can't even grasp *what* the problem is to which they are solutions, let alone *why* they are solutions. Pellagra, for instance, may take decades to arise in those who rely on corn; even the relationship between vitamin C deficiency and the relatively rapidly developing scurvy was sufficiently opaque for there to be considerable controversy over its cause into the 18th century (Magner 1992). We therefore often do better to copy the behavior of successful members of our group than to embody the traits celebrated as virtues in contemporary epistemology. Relying on our own epistemic agency in these common and high-stakes contexts is liable to get us into trouble. Hence the conformist bias, the prestige bias, and our disposition to overimitate. It is noteworthy that these are not dispositions seen only in children: if anything, the disposition to overimitate (especially prestigious models) increases with age (Henrich 2017).

4.2. Overimitation and bad credit

In many contexts, human beings display an epistemic deference approaching servility which contrasts starkly with the behavior of chimps. It is them, and not us, who seem to be epistemically virtuous agents by the standards of contemporary epistemology. Human children are not intellectually openminded in their imitation behavior; in fact, they display resistance to updating their newly acquired behavior, even when it is shown that some steps are superfluous (Lyons, Young & Keil 2007). They seem intellectually incurious, inasmuch as they do not experiment to see whether steps in the process could be eliminated. Adopting the terminology popular in virtue epistemology, we might put the point by saying that children appear to deserve little credit for acquiring knowledge from the model. They do not acquire knowledge primarily through the exercise of their own epistemic agency. In contrast, chimps are active recipients of information

conveyed by demonstration, modifying and improving on the model shown. At minimum, the child seems to deserve less credit than the chimp, who is an active recipient of the knowledge transmitted.

While we think this point is significant, the lesson we wish to draw has implications beyond virtue epistemology: implications for everyone insofar as they countenance the idea that there are knowledge-conducive dispositions that, because they are knowledge-conducive, we have good reason to cultivate. We will pursue this question by continuing to talk about epistemic credit and agency, but we thereby aim to pursue the broader question too: insofar as these dispositions do not underwrite credit to the agent, the case for seeing them as epistemic virtues or as dispositions worthy of cultivation is weakened. They may be knowledge-conducive in important contexts (and therefore what Battaly calls effects virtues), but they also have features that make them unfit for cultivation. They are not, for example, appropriately seen as motivated by or expressions of a love of knowledge.

Sometimes, when we acquire ways of going on from others without understanding the point of some steps, we may reasonably hope that once we have mastered the practice we will understand their point (MacIntyre 1981).⁷ We will then be in a position to responsibly pass on the knowledge. But when we are acquiring a way of going on that is adaptive in a causally opaque environment, we may *never* be in a position to understand why it is adaptive. Such knowledge is acquired as a result of cultural evolution, not careful and explicit experimentation: some variants survive and others are eliminated because some variants are more successful over the long run (where the long run is often longer than any individual's life). No one may be in a position to understand what makes successful variations successful. In these kinds of circumstances, those who transmit the knowledge seem to deserve no more credit for transmission than the recipients do for acquiring the knowledge. Credit accrues neither to hearer nor to speaker, as credit accounts of testimony require (Lackey 2007; Greco 2007; Sosa 2007). It is difficult, we suggest, to understand the acquisition and transmission of the knowledge that constitutes cumulative culture as epistemically virtuous by the lights of virtue epistemology or as manifesting dispositions of epistemic agency that we might seek to cultivate through education.

How might virtue epistemologists and fellow travelers reply to this claim? They might argue that the acquisition and transmission of cumulative culture manifests epistemic virtue, despite all we have said. We can imagine strong and weak versions of this reply (corresponding to individualistic and weak anti-individualistic positions in virtue epistemology). Alternatively,

⁷ This seems to be how Aristotle conceives of the transmission of ethical knowledge in the *Nicomachean Ethics* (Kraut 2006).

they might concede that the acquisition and transmission of cumulative culture do not manifest epistemic virtue, but that the knowledge embodied in cumulative culture is not the kind of knowledge that virtue epistemologists aim to illuminate. It is, they may say, knowledge-*how*, and they aim to illuminate propositional knowledge: knowledge-*that*. In this section and the next, we will evaluate these responses.

We have been arguing that cumulative culture is a kind of knowledge dependent for its production and maintenance on dispositions for which agents do not deserve credit. But in so doing, we have set aside significant controversies in the field of cultural evolution. We have drawn heavily so far on what Sterelny (2017) calls the Californian program. But the Paris program associated with Sperber (1996), Morin (2016a), and others argues that imitation plays a smaller role in the transmission of culture than the Californians suggest. Rather, culture is reproduced through mechanisms that are reconstructive, with the learner contributing a great deal to the final product. The Parisians maintain that the reconstructive mechanisms are intelligent and strategic, rather than blind and deferential. In fact, it might even be claimed that the Californians themselves attribute more intelligence to cultural transmission than we have suggested: as Buskell (2016) and Boudry (2018) have pointed out, imitation California-style need not be reflexive, automatic, or one-shot.

Accordingly, those who hold that the acquisition and transmission of cumulative culture manifest epistemic virtue (sufficient for an individualist virtue epistemology) can cite the evidence of the Parisians in defense of the claim. For example, they may cite evidence that children and adults are far from indiscriminating in who they imitate (Mercier 2017). People prefer, for instance, to imitate prestigious members of their cultural group. Similarly, while people are very receptive to testimony, they prefer the testimony of those who appear to share their values and who have a track record of reliability over those who do not (Harris 2012). These dispositions are adaptive. It is adaptive to copy prestigious individuals, because prestige is correlated with flourishing and therefore imitation of the prestigious is a good heuristic for the acquisition of adaptive behaviour. Similarly, it is adaptive to prefer the testimony of those who are well-disposed toward us and who are reliable, in order to avoid being deceived by those who might exploit us or inadvertently mislead us. It might therefore be argued that our dispositions to accept testimony and to imitate are competences associated with epistemic agency, in virtue of which agents deserve credit for their knowledge acquisition.⁸ We doubt this response will succeed with regard

⁸ Compare the response of virtue epistemologists to Lackey's (2007) cases in which credit seems due to the source of testimony and not the recipient. Opponents concede that the hearer may not deserve much credit, but they maintain that they are sufficiently creditable to count as knowing in virtue of that fact (Greco 2007).

to deferential dispositions, because they are too indiscriminating to pass muster as epistemic competences. This is clearest in the case of prestige bias, which filters in all kinds of behaviors merely because they are the behaviors of prestigious individuals. Because the link between the behavior of the prestigious individual and their success is often causally opaque, we are disposed to copy all their behaviors, thereby ensuring acquisition of the right ones (along with others that are, at worst, not too maladaptive). Hence, for example, our disposition to buy products endorsed by prestigious individuals, despite the unlikelihood that there is any connection between the facts that make (say) a basketball star prestigious and his choice of watch, deodorant, or cryptocurrency.

It is less obvious that our disposition to filter testimony cannot constitute an epistemic competence. Proponents of the Paris program have shown that agents, even children, are more epistemically vigilant than might be thought (Sperber et al. 2010; Mercier 2017). We filter testimony by reference to cues that tend to correlate with accuracy. For instance, we are sensitive to cues that the testimony is consensual, and we tend to reject unfamiliar claims. These cues tend to correlate with accuracy on the plausible assumption (for which the Paris school provides evidence) that consensual claims are more likely to be true, in the first case, and that familiarity arises from repeated hearing (and therefore indicates widespread agreement), in the second. These dispositions might be thought to be the kinds of thing in virtue of which credit is due to the recipient of testimony. However, the way these dispositions are actually instantiated often leads to their deployment in ways that have us respond unthinkingly to cues for apparent consensus which are in fact no such thing, and markers of unfamiliarity which are epistemically irrelevant. For instance, repetition of testimony increases our disposition to accept it, to the same degree if it is repeated by independent individuals or the same individual multiple times (Weaver et al. 2007). Likewise, disfluency seems to be treated as a marker of unfamiliarity, and therefore to decrease acceptance, whether the disfluency is genuinely generated by novelty of content or something as irrelevant as a hard to read font (Song & Schwarz 2008; Oppenheimer 2008).

While this is not the place to rehash the debate between Parisians and Californians, we think it is important to note that we believe that the intelligence the former document is less apparent in the kinds of domains with which we are concerned. High-fidelity imitation is not necessary for the long-term maintenance of folk tales, painting styles, and myths. But unobvious technological innovations and sophisticated environmental knowledge — the kind of knowledge that depends on the detection of a signal in a noisy environment — do depend importantly on such imitation. Agents, adults as well as children, default to imitation when success depends or seems to depend on following a precise sequence of steps, when mechanisms are causally

opaque, and when the demonstrator is presented as an expert (Acerbi & Masoudi 2015). These are precisely the sorts of contexts in which reconstruction risks the loss of cultural knowledge. If agents intelligently reconstruct the causal process, they may easily leave out crucial steps, since their causal contribution is opaque to us. While the telling of folk tales may be adaptive, their precise content typically doesn't matter much, and reconstruction may be given free reign. But when the precise content is crucial, deference is required.

We therefore appear to have developed cues for switching to imitation from reconstruction. In addition to cues like the presence of experts or causal opaqueness, we use ostensive communication ('look at what I'm doing here') and others cues that indicate conventionality (Acerbi & Masoudi 2015). Whether cultural transmission is reconstructive or imitative depends on the domain studied — a fact Morin (2016b) comes close to acknowledging. Since the domains of causally opaque technology, such as food preparation techniques, and the detection of signals in noisy environments are crucial to our adaptive success, and it is in these domains that imitation reigns supreme, we make no apologies for resting our case for the rejection of individualism on these domains.⁹

4.3. Relocating epistemic agency in speakers rather than hearers?

If the recipients of testimony do not deserve credit, however, perhaps the speaker does. If cumulative culture consists largely in knowledge-how (an issue to which we turn below), the speaker or model may have as adequate understanding of the knowledge she transmits as we can reasonably expect from virtuous speakers. The propositional content of the knowledge transmitted is, roughly, $\langle \textit{this is the way to } \phi \rangle$, and since she thereby picks out the (or at least *a*) way to ϕ , she transmits it while grasping it well enough to deserve credit for doing so. She need not understand the underlying mechanics — the reason why the behaviors she refers to

⁹ Virtue epistemologists might appeal to another claim of the Parisians in defence of their view. The Parisians identify a range of "cultural attractors", which are species-typical cognitive dispositions that play an important role in stabilizing cultural traditions in particular forms. It might be claimed that the pervasive reliance on cultural attractors ensures that transmission of cumulative culture is individualistic. In the domains with which we are concerned, reliance on such attractors is risky and likely to be minimal. Such attractors risk distorting the knowledge transmitted when it is unintuitive (and it often is). In addition, we note that the cultural attractor story itself plausibly requires culturally specific, rather than species-typical, machinery to explain the content that develops. Take minimal counterintuitiveness theory (e.g. Boyer 2000; Atran 2002), often hailed as one of the great success stories for the approach. According to this account, minimally counterintuitive entities have an advantage when it comes to cultural epidemiology: they are more likely to be recalled and more likely to be transmitted. As Sterelny (2018) points out, however, the evidence strongly suggests that the institutional and ritual scaffolding of religious belief is central to its survival. These factors are not species-typical, of course — not in their contents, at any rates — and they are neither invented by individuals nor understood by them as belief transmission mechanisms. If we can generalize the lesson, the attractors themselves may turn out to be far more deeply cultural than the Paris school would be comfortable with.

demonstratively are a way to ϕ — any more than you need to understand the underlying physics in order to teach someone to ride a bike.

There are several reasons to find this credit-preserving line unsatisfactory. One is the fact that oftentimes people not only fail to grasp the underlying mechanics of the practices they transmit, but also have false beliefs concerning them. Moreover, these false beliefs may not be incidental to the practices but functional. Indeed, these false beliefs may be required *in virtue of* —not in spite of —the extent to which we are disposed to filter testimony by evidence of reliability: when it is difficult to see the point of a behavior or practice which is in fact functional, we might abandon it were it not supported by false beliefs. When its function is opaque even to careful observers (like the need to introduce an alkali when preparing corn), preservation of the practice may require that false beliefs about its role be transmitted along with the practice itself. In effect, the transmission of false beliefs may be required *to compensate for* the extent to which we approach being virtuous epistemic agents. Only by encasing cumulative cultural knowledge in a chrysalis of falsehood are we able to reliably transmit it from one generation to the next. Since the true function of certain elements of the practice may be unknown (may *never* have been known by anyone, if they emerged through blind cultural evolutionary forces), the chances that the rationale given will coincide with its real function are slim.

Moreover, those who transmit cumulative culture are far from paragons of epistemic virtue or responsible epistemic agency in doing so. Corresponding to the epistemic servility expressed by learners through overimitation, we can expect to find dogmatism and epistemic arrogance expressed by teachers through insistence on overimitation and similar displays of epistemic deference. In *The Antichrist*, Nietzsche characterizes such forms of cultural transmission of laws, norms, and practices, arguing that that they never describe

the uses, the reasons, and casuistry in the prehistory of a law: this would make it lose its imperative tone, the ‘thou shalt’, the presupposition of being obeyed. [...] What now needs to be prevented at all cost is any further experimentation, the continuation of values in a fluid state, scrutiny, selection, and criticism *in infinitum*. A double wall is set up against this: first: *revelation*, the claim that the reasoning behind the law does *not* have a mortal provenance, has *not* been slowly and painstakingly looked for and discovered, but instead had a divine origin, arriving whole, complete, without history, a gift, a miracle, simply communicated... And second: *tradition*, which is to say that the claim that the law existed from time immemorial, that it is irreverent to cast doubt on it, a crime against the ancestors.

Nietzsche here identifies two strategies for ensuring faithful cultural transmission: revelation and tradition. The disposition to employ such strategies is surely an epistemic vice in many cases, just as the disposition to respond to such strategies with deference is a vice. Yet precisely these

dispositions are often needed for intergenerational transfers of know-how, as well as the propositional attitudes that support it. Sometimes the functions assigned to behaviors are naturalistic but false, and sometimes the justification given is (as Nietzsche suggested) supernatural. The tabooing of certain foods for pregnant women in Fiji appears to be an instance of the former kind, invoking a naturalistic but false justification: the justification cited is that they might cause the child to have the properties of the fish that is avoided, while the real function of the taboo seems to be avoiding miscarriage (Henrich & Henrich 2010). There are many examples of the latter, supernatural, kind. For instance, in a number of cultures, important decisions (where to plant food crops, where to hunt, and so on) are made by consulting auguries, in ways that enhance the quality of the decision. For example, Naskapi hunters decide where to hunt by using the shoulder blade of a caribou, heated over hot coals so that it cracks and burns, as a kind of map. There is evidence in this and other cases like it that the divination is adaptive because it effectively randomizes behavior (Henrich 2017). Human beings are susceptible to a number of cognitive biases that lead us to see illusory patterns; outsourcing our choices to a randomizing device frees us from these biases and increases our chances of success. These facts make the analogy with the bike-riding case mentioned above tenuous at best. It is one thing to transmit a practice (and associated propositional knowledge: at minimum, the proposition < that *this* is the way to ϕ >) without grasping its mechanics; another to pass it on as a package encompassing false claims. It appears that neither the hearer nor the speaker deserves much credit for exercising epistemic agency to transmit knowledge in these cases.¹⁰

Finally, what are the prospects for dismissing the challenge to individualistic epistemology by showing that cumulative culture is knowledge-how rather than propositional knowledge? Our first response consists in pointing to the multitude of propositional knowledge that is, uncontroversially, partially constitutive of cumulative culture: *that corn should be prepared with wood ash* (or burned sea shells, or lime, depending on the local culture); *that Nardoo seed cases must be roasted prior to grinding the seeds*; *that pregnant women who eat certain fish will have children who look like fish*; *that the gods will manifest in the patterning of a burned caribou shoulder blade*. All these, and many more, adaptive cultural practices are given propositional

¹⁰ Nor are those who transmit cultural knowledge always in a position to knowledgeably assert that *this* is a way to ϕ : they may have never encountered the ecological conditions for which the practice is an adaptive response. Cultural practices may be responses to environmental fluctuations at temporal scales much longer than living memory (Shea 2009); they therefore may provide responses to problems that those who transmit them have never themselves encountered. Henrich (2017) recounts the story of Paralji, an elder of the Nangatara people, who saved members of his group from a devastating drought in the mid 1940s. When even the most reliable of the waterholes known to him failed, he used ceremonial songs as a guide on a 600 kilometre trek across the desert, across territory unknown to him. The route took them from waterhole to waterhole, thereby allowing the group to survive until the drought broke.

justifications by the people who engage in them. Uninformative (insofar as they often consist in nothing more than *this activity is required*) or false though many of these justifications are, they are not epiphenomenal: they are required for the success of the practice. Our second response consists in pointing out that knowledge-how itself entails knowing an answer to the question how one ϕ s. It is controversial whether knowledge-how *consists in* standing in a certain relation to a proposition (Stanley & Williamson 2001; Stanley 2011), possessing certain propositional attitudes (Bengson & Moffett 2012), or knowing a demonstrative ('*this* is how I ϕ '). While we take no stand on which of these views is correct, we note that all of them entail that the prospects are dim for invoking the knowledge-how/knowledge-that distinction to save epistemic individualism.¹¹

5. Knowledge from vice

We believe that the previous sections have established a strong case for holding that the formation and maintenance of cumulative cultural knowledge does not rely on dispositions in virtue of which agents deserve credit. In this section, we consider one final weak anti-individualist response before concluding. This response consists in attempting to show that credit can be spread over many individuals, each of whom deserves some credit. Palermos (2016) articulates such a view in the context of the use of epistemic artifacts designed by agents elsewhere and elsewhere. Analogously, we might think that the history of cultural practices will reveal a myriad of epistemic activities, none of which may have contributed very significantly to the adaptive design of the practice we see today, but each of which was deserving of *some* credit. Because there are many such epistemic activities, each of them deserving of some credit, together they account for the credit deficit. We might underwrite knowledge attribution by parceling out credit across many different agents, or perhaps by attributing knowledge not to any particular agent but to the group or community as a whole.

Activities deserving of credit, however small, are (minimally) virtuous epistemic activities; they must be expressions of cognitive agency to be candidates for inclusion in the broader narrative here envisioned. They are activities guided by the search for truth. Perhaps they consist in nothing more than noticing that an accidental deviation from a practice improves the result and therefore retaining the deviation. Even such a serendipitous improvement is the expression of minimal epistemic virtue, one might argue, inasmuch as its retention requires recognition of

¹¹ Elsewhere, one of us has defended the latter view, in arguing against the adequacy of the thoroughgoing intellectualism of Stanley or Bengson and Moffett; see Levy 2017.

the improvement and its cause, and is therefore creditworthy. However, it is far from clear that cultural practices owe their existence and form to a sufficient number of even minimally virtuous activities.

There is a great deal of epistemic activity to be found in the genesis of cultural practices. There is argument, experimentation, and the recognition that some variants of a practice are better than others. There is also, as we have seen, the filtering of testimony by reference to its plausibility and to the perceived character of testifiers. But there is a great deal of blind chance too, and — worse — much of the epistemic activity is ill-suited for credit attribution. Consider the role of argument, surely our paradigm epistemic activity. Argument certainly plays a role in the generation of cultural knowledge, but there are good reasons to think that the role it plays is itself deeply social in a way that makes it difficult to give the individuals involved credit, or to see the knowledge produced as the upshot of individual epistemic virtue or ability. Argument is *typically* and *adaptively* lazy and sloppy. Hence the apparently puzzling finding that we are bad individual reasoners across a number of contexts (Mercier and Sperber 2017). But the very dispositions that show up as glaring flaws in the lab may be strengths outside it, in groups of deliberators. We do much better in groups at tasks at which we do poorly individually (Mercier, Trouche, Yama, Heintz & Girotto 2015). A principal reason is that our biases contribute to an efficient division of cognitive labor (Mercier and Sperber 2011).

The myside (or confirmation) bias is an excellent illustration of how individual biases yield group epistemic success. The myside bias consists in our disposition to be easily satisfied with arguments for positions which we are motivated to accept, reserving our powers of critical scrutiny for arguments we are disposed to reject. In a sufficiently diverse group of deliberating agents, the myside bias ensures that a variety of arguments are generated and that each argument receives genuinely searching criticism. Group deliberation harnesses the power of asymmetrical individual scrutiny.¹² Interestingly, no evidence of confirmation bias emerges from studies of animal behavior (Mercier & Sperber 2017), which reinforces the hypothesis that it is an adaptation for culture. By harnessing our biases, group deliberation may allow us to identify significant truths. More intriguingly, it may allow the emergence and adoption of significant truths in ways that *circumvent* their identification. Group deliberation, that is, may not succeed in convincing partisan individuals but may nevertheless ensure that the results of deliberation come to be reflected in cultural practices (this might occur if there are institutional mechanisms — say the development of rituals — for implementing the results of deliberation without requiring that

¹² See Betz (2013) for supportive simulations and Alfano (2017) for a social virtue theoretic interpretation of these results.

people are convinced). Rituals and other social and institutional arrangements that conform to the conditions required for harnessing the wisdom of crowds — independence, decentralization, diversity, and aggregation (Surowiecki 2004) — would (and do) accomplish precisely this.

Group deliberation harnesses the power of argument, but the individuals who engage in it seem to deserve little credit for the process and even less for its results. If there is epistemic credit for, and if there is epistemic agency deployed in, the generation of knowledge, it is irreducibly social. It is at the level of the group as a whole that these processes are truth-conducive. At the individual level, they may entrench us in comfortable falsehoods. Our reliance on processes like this for the generation of the cumulative culture constitutes significant pressure toward a strong anti-individualism. The weakness of weak anti-individualism is justified by the claim that credit is shared between the individual who comes to have a belief and others who play a role in generating it. But in the cases we have examined, and which we claim to be central to the discovery and transmission of cumulative culture, there is precious little credit accruing to any of the individuals. Moreover, at least some of the individuals play an appropriate role in generating the knowledge without themselves coming to possess knowledge, since at least some of those who scrutinize arguments asymmetrically never come to accept the best-justified view. Sometimes, no one in a group understands how, or even *that*, a ritual or a narrative encodes a response to an ecological challenge because they have never encountered the challenge; once again, they do not seem deserving of credit for transmitting it.

Cumulative culture is not merely interpersonal; it is also and necessarily intergenerational. Adaptive cultural practices are generated by multiple individuals across decades and perhaps even centuries. It relies on individual and collective epistemic practices — argument, the recognition of success, deliberate experimentation — but also on cultural evolution working blindly.¹³ Insofar as there is epistemic credit to be apportioned, it accrues to many different players: individuals, groups, and the mechanisms of cultural evolution. Insofar as there are epistemic competences to be identified, some of those competences are those of the group. It seems difficult to hang on to even weak anti-individualism in the face of these facts.

In virtue of its interpersonal and intergenerational generation, there is a strong case for seeing cumulative culture as an example of group knowledge. However, cumulative culture does

¹³ No doubt these different forces may work simultaneously. Boyd, Richerson and Henrich (2011) recount how in the early decades of the 19th century, a prestigious Inuit hunter began to experiment with rudders for the kayaks the Inuit used, in order to stabilize them for the use of newly introduced firearms. The rudder he tried did not work; however, his prestige inspired other hunters to try their own hand at rudder making. According to contemporary reports, these rudders were so crude that the imitators hid them from view under the water out of “bashfulness”. These submerged rudders actually helped stabilize the kayaks. Thus, the innovation emerged out of deliberate experimentation, copying of prestigious models and chance.

not fit snugly within existing approaches to group justifiedness (and therefore, presumably, to counting it *as* knowledge in the first place). For example, according to Hakli (2011), a ‘collectively accepted group view that p ’ is epistemically justified just in case ‘the group can successfully defend p against reasonable challenges by providing reasons or evidence that are collectively acceptable to the group and that support p according to the epistemic principles collectively accepted in the epistemic community of the group’. As we have seen above, such reasons and evidence are unlikely to be forthcoming in the sorts of cases paradigmatic of cultural knowledge, in which blind deference is more common than reason-giving. Hakli also requires that the group, to be justified, must engage in open, plenary deliberation. Again, this does not seem to be how cumulative culture is transmitted. Likewise, Goldman (2014) argues for two principles connecting the justifiedness of a group’s belief to the justifiedness of the beliefs of each member of the group. First, the justifiedness of the group’s belief scales with the proportion of members of the group whose individual beliefs are themselves justified.¹⁴ Second, the belief-aggregation process that transforms individual beliefs into a group belief has high conditional reliability. We have already argued that the rituals and other transmission practices that are responsible for cumulative culture do in fact have high conditional reliability. However, we have also argued that the individual beliefs that serve as the basis for constructing group beliefs rarely enjoy high justificatory status.

We believe that the prospects for showing that the dispositions that underlie cumulative culture are the kinds of things in virtue of which agents deserve credit are dim. Equally dim are the prospects for rethinking them as virtues. They are not the kinds of dispositions which we would or should seek to cultivate in one another, even if they are effects virtues — even if, that is, they do tend to lead to a preponderance of good epistemic effects in certain environments. It is a striking fact about these dispositions that parents, educators, and philosophers proverbially warn *against* them. They exhort children to think for themselves, not to blindly follow what they’re told. In a more elevated idiom, we might think of Kant: ‘*Sapere Aude!* “Have courage to use your own understanding!” — that is the motto of enlightenment’ (Kant 1784/1991, p. 54). Following this kind of sentiment, people seek to inculcate in one another dispositions that do not merely lead to truths, but that lead to understanding. They think of confirmation bias as, precisely, a bias, and close-mindedness as a vice.

We think that none of this is a coincidence. While we maintain that these supposed vices are knowledge-conducive much more often and more reliably than others have recognized, there is *something* to the identification of opposing dispositions as virtues. It takes hard work to

¹⁴ Lackey (2016) endorses a similar principle.

consider a range of hypotheses and to be open-minded. It takes hard work to think for oneself. We fail at these attempts at epistemic virtue *routinely*. If we're right, these failures may be happy: if we succeeded very much more often, we would do less well epistemically. But that doesn't entail that our successes are not themselves epistemically virtuous or expressive of cognitive agency: it is, perhaps, only by great effort that we achieve a cognitive sweet spot, where we follow the crowd and are closedminded just enough.

Of course, this suggestion is speculative. Nevertheless, there doesn't seem to be any good reason to inculcate closedmindedness and deference: we're quite good enough at these things as it is. At minimum, we must resist the suggestion that these dispositions are virtues in any strong sense, because they do not express the kinds of motivations that virtue epistemologists (both of the responsibilist and of the agent-reliabilist varieties) take be partially constitutive of virtue. The motivations of the closeminded and the overly deferential might be every bit as bad as virtue epistemologists have typically thought; for all that, these dispositions are reliably knowledge-producing under many cultural conditions.

Given that the paradigm cases that feature in traditional epistemology are indeed common — individuals really do acquire knowledge on their own, and knowledge really is transmitted from speaker to hearer — perhaps we should adopt a multilevel account, akin to the multilevel selection account favoured by many evolutionary biologists. Epistemic competences are shared between individuals and groups, irreducibly collective, and even non-agential, and these competences work together simultaneously, in mutually reinforcing (and no doubt sometimes conflicting) ways. Epistemology ignores the irreducibly collective and the non-agential at the peril of neglecting some of our most significant cognitive achievements.¹⁵

¹⁵ We are grateful to two extraordinarily helpful reviewers, and to the editors of *Mind*, for comments that enabled us greatly to improve this paper. We are also grateful to Étienne Brown for comments. Neil Levy was supported by the Australian Research Council. Mark Alfano was supported by a subaward agreement from the University of Connecticut with funds provided by Grant No. 58942 from John Templeton Foundation. The contents of this paper are solely the responsibility of the authors and do not necessarily represent the official views of UConn or John Templeton Foundation.

References

- Acerbi, A. & Mesoudi, A. (2015). If we are all cultural Darwinians what's the fuss about? Clarifying recent disagreements in the field of cultural evolution. *Biology and Philosophy* 30: 81–503.
- Alfano, M. (2016). The topology of communities of trust. *Russian Sociological Review*, 15(4): 30-56.
- Alfano, M. (2017). Twenty-first century perspectivism: The role of emotions in scientific inquiry. *Studi di Estetica*, 7(1): 65-79.
- Alfano, M. & Robinson, B. (2017). Gossip as a burdened virtue. *Ethical Theory and Moral Practice*, 20: 473-82.
- Alfano, M., Stey, P., Robinson, B., Christen, M., Yu, F., & Lapsley, D. (2017). Development and validation of a multi-dimensional measure of intellectual humility. *PLoS ONE*, 12(8): e0182950.
- Alfano, M. & Skorburg, J. A. (2017). The embedded and extended character hypotheses. In J. Kiverstein (ed.), *The Routledge Handbook of Philosophy of the Social Mind*. Routledge.
- Alfano, M. & Skorburg, J. A. (2018). Extended knowledge, the recognition heuristic, and epistemic injustice. In D. Pritchard, J. Kallestrup, O. Palermos, & J. A. Carter (eds.), *Extended Knowledge*. Oxford University Press.
- Atran, S. (2002). *In gods we trust: The evolutionary landscape of religion*. New York, NY: Oxford University Press.
- Battaly, H. (2018). Can Closed-mindedness be an Intellectual Virtue? *Royal Institute of Philosophy Supplement* 84: 23-45.
- Bengson, J., & Moffett, M. A. (2012). Nonpropositional intellectualism. In J. Bengson & M. A. Moffett (Eds.), *Knowing how: Essays on knowledge, mind, and action* (pp. 161–195). Oxford: Oxford University Press.
- Betz, G. 2013. *Debate dynamics: How controversy improves our beliefs*. Dordrecht: Springer.
- Boudry, M. (2018) Replicate after reading: on the extraction and evocation of cultural information. *Biology & Philosophy* 33: 27.
- Boyd, R, Richerson, P.J. & Henrich J. (2011). The Cultural Niche: Why Social Learning is Essential for Human Adaptation. *Proceedings of the National Academy of Sciences* 108: 10918-10925.
- Boyer, P. (2000). Functional origins of religious concepts: Ontological and strategic selection in evolved minds. *Journal of the Royal Anthropological Institute* 6: 195–214.
- Burcham, P.C. (2008). Toxicology Down Under: Past Achievements, Present Realities, and Future Prospects. *Chemical Research in Toxicology*, 21: 967–970.

- Burge, T. (1979). Individualism and the Mental. *Midwest Studies in Philosophy* 4: 73-121.
- Buskell, A. (2016) Cultural longevity: Morin on cultural lineages. *Biology & Philosophy* 31: 435-446.
- Carter, J. A., Kallestrup, J., Palermos, S. O., & Pritchard, D. (2014). Varieties of externalism. *Philosophical Issues*, 24(1): 63-109.
- Cassam, Q. (2016). Vice epistemology. *The Monist*, 99: 159-80.
- Cassam, Q. (forthcoming). *Vices of the Mind*. Oxford University Press.
- Chudek, M., Heller, S., Birch, S. & Henrich, J. (2012). Prestige-biased cultural learning: bystander's differential attention to potential models influences children's learning. *Evolution and Human Behavior* 33: 46-56.
- Clark, A. & Chalmers, D. (1998). The Extended Mind. *Analysis* 58(1): 7-19.
- Coady, C. (1995). *Testimony: A Philosophical Case Study*. Clarendon Press.
- Craig, E. (1999). *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford University Press.
- Dunbar, R. (1992). Neocortex size as a constraint on group size in primates. *Journal of Evolution*, 22(6): 469-93.
- Fantl, J. (2018). *The Limitations of the Open Mind*. Oxford University Press.
- Fricker, M. (2007). *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford University Press.
- Fridland, E. (2018). Do as I say *and* as I do: Imitation, pedagogy and cumulative culture. *Mind & Language* 4: 355-377.
- Gergely, G., Bekkering, H. Kiraly, I. (2002). Rational imitation in preverbal infants. *Nature* 415: 755.
- Gettier, E. (1966). Is justified true belief knowledge? *Analysis*, 23(6): 121-23.
- Goldberg, S. (2010). *Relying on Others: An Essay in Epistemology*. Oxford: Oxford University Press.
- Goldberg, S. (2011). The Division of Epistemic Labor. *Episteme* 8: 112-125.
- Goldberg, S. (2013). Self-trust and extended trust. *Res Philosophica*, 90(2): 277-92.
- Goldman, A. (2014). Social process reliabilism: Solving justification problems in collective epistemology. In J. Lackey (ed.), *Essays in Collective Epistemology*. Oxford University Press.

- Greco, J. (2007). The Nature of Ability and the Purpose of Knowledge. *Philosophical Issues*, 17: 57-69.
- Hakli, R. (2011). On dialectical justification of group beliefs. In H. B. Schmid, D. Sirtes, & M. Weber (eds.), *Collective Epistemology*. Ontos Verlag.
- Harris, P. (2012). *Trusting What You're Told: How Children Learn from Others*. Cambridge, Mass.: Harvard University Press.
- Henrich, J. & Boyd, R. (1998). The evolution of conformist transmission and between-group differences. *Evolution and Human Behavior* 19: 215-242
- Henrich, J., F. Gil-White. (2001). The Evolution of Prestige: Freely Conferred Deference as a Mechanism for Enhancing the Benefits of Cultural Transmission. *Evolution and Human Behavior* 22: 165-196.
- Henrich, J. & N. Henrich. (2010). The Evolution of Cultural Adaptations: Fijian Food Taboos Protect against Dangerous Marine Toxins. *Proceedings of the Royal Society B: Biological Sciences* 277 (1701): 3715-3724.
- Henrich, J. (2017). *The Secret of Our Success*. Princeton: Princeton University Press.
- Hrdy, S. B. (2009). *Mothers and others: The evolutionary origins of mutual understanding*. Cambridge: Harvard University Press.
- Kallestrup, J. & Pritchard, D. (2012). Robust virtue epistemology and epistemic anti-individualism. *Pacific Philosophical Quarterly*, 93(2): 84-103.
- Kallestrup, J. & Pritchard, D. (2016). From Epistemic Anti-Individualism to Intellectual Humility. *Res Philosophica* 93: 533-552.
- Kant, I. (1784/1991). An Answer to the Question: 'What is Enlightenment?', in Kant, I. 1991. *Political Writings*. Cambridge: Cambridge University Press.
- Kraut, R. (2006). How to justify ethical propositions: Aristotle's method. In R. Kraut (ed.), *The Blackwell Guide to Aristotle's Nicomachean Ethics*. Wiley-Blackwell.
- Lackey, J. (2007). Why We Don't Deserve Credit for Everything We Know. *Synthese* 158: 345-61.
- Lackey, J. (2016). What is justified group belief? *Philosophical Review* 125(3): 341-96.
- Lanska, D.J. (2010). Historical aspects of the major neurological vitamin deficiency disorders: the water-soluble B vitamins. In Finger, S., Boller, F & Tyler, K.L. 2010. *Handbook of Clinical Neurology: History of Neurology*. Elsevier, pp. 445-476.
- Levy, N. (2017). Embodied Savoir-Faire: Knowledge-How Requires Motor Representations. *Synthese* 194: 511-530.

List, C. & Goodin, R. (2001). Epistemic democracy: Generalizing the Condorcet jury theorem. *Journal of Political Philosophy* 9:277-306.

Lyons, D.E., Young, A.G. & Keil, F.C. (2007). The hidden structure of overimitation. *Proceedings of the National Academy of Sciences* 104: 19751-19756.

MacIntyre, A. (1981). *After Virtue*. University of Notre Dame Press.

Magner, L.N. (1992). *A History of Medicine*. New York: Marcel Dekker.

Mascaro, O. & Sperber, D. (2009). The moral, epistemic, and mindreading components of children's vigilance towards deception. *Cognition* 112: 367-80.

Meltzoff, A. N. (1988). Infant imitation after a 1-week delay: Long-term memory for novel acts and multiple stimuli. *Developmental Psychology*, 24: 470-476.

Mercier, H. (2017). How Gullible Are We? A Review of the Evidence From Psychology and Social Science. *Review of General Psychology* 21: 103-122.

Mercier, H. & Sperber, D. (2011). Who do humans reason? Arguments for an argumentative theory. *Behavioral and Brain Sciences* 34: 57-111.

Mercier, H. & Sperber, D. (2017). *The Enigma of Reason*. Harvard: Harvard University Press.

Mercier, H., E. Trouche, H. Yama, C. Heintz, and V. Girotto. (2015). Experts and Laymen Grossly Underestimate the Benefits of Argumentation for Reasoning. *Thinking and Reasoning* 21: 341-355.

Morin, O. (2016a). *How Traditions Live and Die*. Oxford: Oxford University Press.

Morin, O. (2016b). Reasons to be fussy about cultural evolution. *Biology and Philosophy* 31: 447-458.

Nagell, K., Olguin, R. S., & Tomasello, M. (1993). Processes of social learning in the tool use of chimpanzees (*Pan troglodytes*) and human children (*Homo sapiens*). *Journal of Comparative Psychology*, 107: 174 -186.

Nietzsche, F. (1888 / 2004). In A. Ridley & J. Norman (eds.), J. Norman (trans.), *The Anti-Christ, Ecce Homo, Twilight of the Idols, and Other Writings*. Cambridge University Press.

Olin, L. (2017). Is every epistemology a virtue epistemology? In A. Fairweather & M. Alfano (eds.), *Epistemic Situationism*. Oxford University Press.

Oppenheimer, D.M. (2008). The Secret Life of Fluency. *Trends in Cognitive Science* 2: 237-241.

Palermos, S.O. (2016). Spreading the Credit: Virtue Reliabilism and Weak Epistemic Anti-Individualism. *Erkenntnis* 81: 305-334.

Pritchard, D. (2018). Extended Virtue Epistemology. *Inquiry* 61: 632-647

- Putnam, H. (1975). The Meaning of “Meaning”. In *Mind, Language and Reality. Philosophical Papers, vol. 2*, 215-271. Cambridge, UK: Cambridge University Press.
- Richerson, P.J. & Boyd, R. (2005). *Not by genes alone*. Chicago: University of Chicago Press.
- Shea, N. (2009). Imitation as an Inheritance System. *Philosophical Transactions of the Royal Society B* 364: 2429-2443.
- Song, H., & Schwarz, N. (2008). Fluency and the detection of distortions: Low processing fluency attenuates the Moses illusion. *Social Cognition* 26: 791-799.
- Sosa, E. (2007). *Apt Belief and Reflective Knowledge, Volume 1: A Virtue Epistemology*. Oxford University Press.
- Sperber, D. (1996). *Explaining Culture*. Oxford: Blackwell Publishers.
- Sperber, D. Clément, F, et al. (2010). Epistemic Vigilance. *Mind & Language* 25: 359-393.
- Stanley, J., & Williamson, T. (2001). Knowing how. *Journal of Philosophy* 98: 411-444.
- Stanley, J. (2011). *Know how*. Oxford: Oxford University Press.
- Stanley, J. (2015). *How Propaganda Works*. Princeton University Press.
- Sterelny, K. (2018). Religion Re-Explained. *Religion, Brain & Behavior* 8: 406-425.
- Sterelny, K. (2012). *The Evolved Apprentice*. MIT Press.
- Sterelny, K. (2017). Cultural Evolution in California and Paris. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences* 62: 42-50.
- Surowiecki, J. (2004). *The Wisdom of Crowds*. Random House.
- Tanesini, A. (2018). Intellectual humility as an attitude. *Philosophy and Phenomenological Research* 96: 399-420.
- Turri, J. (2016). *Knowledge and the Norm of Assertion: An Essay in Philosophical Science*. Open Book Publishers.
- Turri, J., Alfano, M., & Greco, J. (2017). Vice epistemology. *The Stanford Encyclopedia of Philosophy* (Winter 2017 edition), URL = <https://plato.stanford.edu/archives/win2017/entries/epistemology-virtue/>.
- Weaver, K., Garcia, S. M., Schwarz, N., & Miller, D. T. (2007). Inferring the popularity of an opinion from its familiarity: a repetitive voice can sound like a chorus. *Journal of Personality and Social Psychology* 92: 821-833.

Whitcomb, D., Battaly, H., Baehr, J., & Howard-Snyder, D. (2015). Intellectual humility: Owning our limitations. *Philosophy and Phenomenological Research*, 94(3): 509-39.