

My research is about the contrast between first-person and third-person perspectives on belief. For example, if you believe that it will rain, then from your first-person perspective it appears to be a fact about the world that it will rain. But if someone else believes that it will rain, then from your third-person perspective this appears merely to be a fact about that person's state of mind. Indeed, whether another person's belief even appears to you to support that it will rain will depend on your other beliefs concerning that person's reliability, track record, and so forth. Many philosophers neglect this contrast, and assume that a rational agent has a third-person perspective on her own beliefs. But I claim that an agent's true epistemic situation cannot be accounted for in this third-personal way.

Reliabilists famously deny the significance of the first-person perspective altogether, claiming that the justification of one's beliefs depends instead on the objective reliability of their sources. Although my work reinforces familiar objections to reliabilism, as well as other externalist views, it also introduces new objections to traditional internalist views. In my view, externalists and internalists alike have neglected or misunderstood the first-person perspective.

One of two key mistakes is to treat one's own 'internal' mental states on the model of an 'external' source of information. When one consults with an external source, one learns that, e.g., the fuel gauge reads 'full', or that some other person believes it will rain. From this third-person evidence one must then infer further conclusions about the world. Many philosophers, including traditional internalists, assume that one should have a similar third-person perspective on one's own mental states. In particular, they suppose that one must infer conclusions about the world from evidence about what mental states one is in. But I argue that this leaves us vulnerable to a perennial skeptical challenge. In its most general form, the challenge says that you cannot be justified in believing that your own cognitive faculties of perception, memory, and reasoning are reliable, since any attempt to verify their reliability inevitably will employ those very faculties. To verify the reliability of perception, for example, you must appeal to evidence known by perception. And to verify the reliability of your reasoning faculties, you must employ reasoning. When we consider more external sources, it seems epistemically circular to evaluate a source's reliability by consulting with the source's own testimony. And so it might seem that any attempt to evaluate the general reliability of your own internal faculties must be circular in the same way. But I think this worry rests on a confusion about the contribution of internal and external factors to your first-person perspective. External factors, such as another person's saying that *p*, can provide evidence on whose basis you might believe that *p*. But internal factors, such as your reasoning faculty's saying that *p*, typically constitute your believing that *p* on the basis of other evidence. With this distinction in place, I contend that many philosophers, going back to Descartes, have tended to narrow the scope of the internal to an extent that is both independently implausible and needlessly friendly to skepticism. By resisting the mistaken tendency to treat internal states like the external states of a measuring device, I think we can defend against the skeptic's charge that evaluations of one's own reliability are necessarily epistemically circular. I develop these points in both an ongoing historical research project on the Cartesian Circle and a series of papers in contemporary epistemology, including 'What's the Matter with Epistemic Circularity?', 'Perceptual Justification and the Cartesian Theater', and 'Is Memory Merely Testimony From One's Former Self?'.

A second key mistake of traditional internalism is at work in discussions of self-knowledge and higher-order evidence. A prevailing picture in these discussions holds that reflectively

evaluating one's beliefs, and revising them when one deems them irrational, is an essential component of full-blown rational agency. One familiar worry about this picture is whether rational agency grants us the kind of reflective access to and control over our beliefs that the picture demands, and I try to partially substantiate this worry in my 'Self-Knowledge and Moore's Paradox' and 'Inferential Justification and the Transparency of Belief'. But my core objection is that we shouldn't describe an agent as deliberately manipulating her beliefs to conform to rational norms, as one might arrange items in one's closet. Instead, a rational agent's attention can remain entirely directed outward at the world. For example, suppose that I believe that it will rain, and because of this refrain from believing that the picnic will take place as scheduled. Even though this is a paradigmatic example of following a rational norm, I need never consider the state of my beliefs or the rational requirements they impose. For from my first-person perspective, it appears to be true that it will rain. And this apparent fact about the weather can be reason enough to refrain from believing the picnic will take place. Building on this last point, I am developing a new view about higher-order evidence, which breaks with the prevailing picture in a more radical way. According to this view, even if I do reflectively judge that I should withhold belief about the picnic taking place, this need not give me any reason to withhold belief. For this judgment will itself be just another component of my first-person perspective. That is, from my perspective, it will be just another apparent fact about the world that I should withhold belief. And this will only give me a reason to withhold belief, I claim, if I take this apparent fact to itself provide evidence about whether the picnic will take place—a circumstance that I argue is less common than might be assumed. The resulting view is novel in that it rejects even highly restricted 'higher-order requirements' on justification, which is ordinarily considered an extreme externalist position, yet does so on the grounds that these requirements neglect the agent's first-person perspective.

My published work has so far taken the form of a series of interrelated papers. But I have plans to write a longer manuscript that unifies and further develops them, with an emphasis on their implications for the ideal of *intellectual autonomy*. We take it to be a platitude that one should think for oneself, and not merely go along with the opinions of others. For example, we think that I ought to see to it that my belief is consistent with my other beliefs, rather than with other people's beliefs. But when we consider ourselves from a third-person perspective, this ideal of autonomy can appear puzzling. Why take myself, of all people, to be the one whose beliefs I should aim to be consistent with? And why not aim in the same way to bring other people's beliefs into consistency with mine? Traditional defenses of intellectual autonomy, which conceive of autonomy merely in terms of 'tending one's own garden', cannot give satisfying answers to these questions. Instead, we must acknowledge that they cannot be answered from a third-person perspective. Even so, I claim that one can remain committed to consistency and other manifestations of autonomy from within one's first-person perspective, even while one cannot justify these commitments from outside it.