

The Ethics of Belief, ed. Jonathan Matheson. Forthcoming.

The Commonwealth of Epistemic Ends

Catherine Z. Elgin

Neither the deliverances of the Oracle at Delphi nor those of Grandpa's trick knee are remotely credible. But some sense perceptions, experiences, introspective insights, and/or a priori intuitions apparently carry significant epistemic weight. Many epistemologists hold that some or all of them constitute our epistemic core. Core deliverances, if there are such things, (1) are intrinsically credible in isolation if anything is, and/or (2) interweave into fabrics of commitments that are intrinsically credible if anything is. If the items in the core are basic, nothing epistemologically more fundamental either can or need underwrite them. Not everyone considers all the items on my list basic. Some epistemologists deny that any deliverances are basic, holding rather that justification requires even core deliverances to cohere. Some include deliverances of testimony in the core. However these details work out, items on the list are standardly taken as clear, central cases of epistemically privileged justifiers. No non-core items are deemed necessary to secure either the justification of such deliverances or their capacity to confer justification. Although some hold that core deliverances depend on one another for credibility, it is widely believed that they do not depend on anything outside the core.

This is all highly schematic. It glosses over oceans of epistemological disagreement. I present matters so sketchily to highlight and problematize a pair of widely shared assumptions. First, *epistemic individualism*: it is states of an individual epistemic agent that constitute her epistemic core. *Her* perceptions, *her* experiences, *her*

introspective insights, *her* a priori intuitions, and/or testimony that *she* is privy to are supposed to justify her beliefs. Second, *attunement*: the core deliverances that justify her beliefs do so because they properly attune her to their objects. Perceptions as of red justify because they properly attune her to red things. A priori intuitions about natural numbers justify because they properly attune her to natural numbers. Unless a belief is about other people's attitudes or actions, it can be wholly justified regardless of what others (apart, perhaps, from those who provide testimony) say or think or do. For it is the believer's relation to the belief's content that secures justification. One need not be a foundationalist to take such a position. Commitment to epistemic individualism and proper attunement are endemic in epistemology. Nevertheless, I will argue that these commitments are mistaken. Epistemic success requires that agents be properly attuned not just to the objects of knowledge but also to one another.

Orwellian Worries

George Orwell's *1984* calls epistemological individualism into question. Perhaps separate bits of core knowledge are justified in isolation. But *1984* provides reason to believe that they are sources of justification only if the epistemic climate is favorable. To an extent, we already knew that. Descartes' malevolent demon personifies the possibility that conditions are not congenial to knowledge. The malevolent demon is a supernatural complication. Its natural counterpart is the possibility that the salient regularities that obtain in this spatiotemporal region of the cosmos might not be the real laws of nature. Given that possibility, we rely on familiar regularities at our peril.

Susceptibility to skeptical challenges discloses vulnerabilities in our epistemic condition. Traditional challenges reveal that epistemic success is possible only if

epistemic circumstances are naturally and supernaturally congenial. Then the regularities we take to be laws of nature at least approximate the genuine laws of nature. Our minds are not being manipulated by malevolent demons; nor our brains by mad scientists. And so on. Perhaps, as some think, we can demonstrate that circumstances are congenial; perhaps not. But minimally, they must *be* congenial for agents to be epistemically successful.

Read as an epistemological thought experiment, *1984* discloses a further constraint: only if epistemic circumstances are socio-politically congenial is epistemic success possible. For Winston Smith to know, or reasonably believe that stones are hard,¹ he must stand not only in the right relation to stones, but also in the right relation to his fellows. If this is so, epistemic justification depends on social conditions. Deliverances – even core deliverances – are sources of justification only if neither they nor agreement among them is a product of coercion. This holds whether they themselves are separately justified (as foundationalists claim), or their justification emerges from coherence among them (as coherentists claim), or their justification derives from their reliability (as reliabilists claim). For deliverances to yield justification, epistemic agents must be suitably related, not just to the objects of their beliefs, but to one another.

The epistemological problems exemplified in the novel are stratified. Issues that seem relatively benign at one level turn out to be problematic at a deeper level. But the deeper we delve, the more extensive and implausible the Party's interference is. Eventually we may reach a level where the machinations appear too extreme to disclose anything of epistemological interest. At the limit Orwell's scenario may be an epistemological reductio. If we balk before reaching the limit it is not. But even if we dismiss the most extreme machinations, the story discloses that uncongenial socio-

political conditions can create an epistemologically toxic environment, resulting in something like Vogel's semi-skepticism.² Epistemic success rests on socio-political underpinnings. Minimally, certain socio-political threats must fail to be realized; maximally, certain socio-political supports must be in place.

1984 describes not just a political dystopia, but an epistemological dystopia. The Party's machinations corrupt all standard epistemic resources. The Party regularly revises the historical record to align it with the current political agenda and destroys all previous records. Since, the Party declares, Oceania is now at war with Eurasia and has long been at war with Eurasia, all records claiming that Eurasia was recently an ally are mistaken. That being so, they must be corrected and the erroneous accounts erased. The Party leaves no trace of its tampering. So although citizens know that the record can be changed, they have no way to tell whether, in any particular case, the record has been changed, how it has been changed, or how often it has been changed. It is pointless to consult what purport to be historical documents in order to find out what happened, for there is ample reason to doubt that they have any bearing on the facts.

News media publish whatever the Party dictates, with no attempt to verify and indeed no means of verifying the information they convey. Was the quarterly output 145 million pair of boots (as previously predicted) or 60 million pair (as the current, 'corrected' version reads)? Or were no boots produced at all? In all probability, no one knows. Yet the corrected version is promulgated as fact.³ Within a day, all mention of the original prediction will be expunged from every record. Although so-called 'written records' exist, they have, and are widely recognized to have, no stable relation to the truth. They just state what the Party currently wants everyone to think; and that can change abruptly. With history regularly rewritten, and the news media untrustworthy,

denizens of Oceania have no documentary evidence against which to check their memories or their current opinions.

Spies are everywhere. Deviating from the Party line is punishable by torture, death, and vaporization (where all evidence that a person ever existed is expunged). So it would be suicidal to attempt to check one's memories or perceptions of current events against those of one's compatriots. Since most would probably parrot the Party line anyway, the responses of one's peers would be presumptively unreliable. Given the pressure to conform and the dearth of objective evidence to the contrary, denizens of Oceania might well accede to what the Party tells them without even considering whether it is true. Of course, they might all, like Winston Smith, be seething with epistemological resentment. Without courting death, Winston has no way to find out. But his compatriots' overt behavior indicates that their doxastic default position is to acquiesce to whatever the Party tells them. Arguably then they do not make assertions. Rather than expressing beliefs about what they take to be the case, their utterances articulate the Party line on a given topic. Let us say then that they typically avow, rather than assert. Intersubjective accord thus indicates not that a contention is true or warranted, but that it is politically expedient.

Sense experiences are equally problematic. When what Winston apparently sees before his very eyes conflicts with what the Party reports, he has no way to determine whether he is dreaming, hallucinating, or actually perceiving. Did he really see Jones, Aaronson and Rutland after the Party reported their execution, or did he merely imagine that he saw them? How could he tell? If this were just a single suspect sighting, it would be no more epistemically problematic than Ken's taking himself to have seen an ivory gull in Gloucester. That the bird he saw was an ivory gull is unlikely given the species'

normal range. Since no other birder was around, he had no way to confirm his sighting. So Ken may never know whether he actually saw the rare gull he took himself to see.

But Winston's situation is different from Ken's in at least two respects. First, even if he had been surrounded by other observers, he could not have confirmed his sighting; for no one with a healthy sense of self-preservation would dare to admit to seeing the traitors so long as the Party says that they are dead. Second, and more significantly, Winston's predicament is utterly general. He can glean no genuinely intersubjective confirmation or disconfirmation for anything he takes himself to see. Nor can anyone else. When it comes to reporting observations in Oceania, the relevant difference is not between seeing and not seeing, or seeming to see and not seeming to see, but between what one is supposed to see and what one is supposed not to see. If Winston says that he sees what he is supposed to see, his fellows will corroborate his reports; otherwise they will not – regardless of what passes across anyone's visual field.

In real life, very young children take all their seemingly perceptual deliverances at face value. Initially, they do not realize that dreams are non-veridical. They come to distinguish seeing from dreaming by learning which of their seemingly visual representations admit of intersubjective support.⁴ Denizens of Oceania lack the reinforcement required to learn to distinguish genuine cases of seeing from impressions that do not reliably correlate with their surroundings. They have no reason to trust their senses.

My assessment of Winston's prospects for perceptual knowledge might appear unduly pessimistic. Perhaps through self-monitoring he could discover that his perceptual deliverances are reliable while his dreams are not. Coherence and consistency would be his guide. Then, without regard to what anyone else says, he would be in a

position to know what he sees, hears or whatever. On the face of it, this seems plausible. But things are not so straightforward.

A person has no hope of establishing for himself that perception per se is reliable for the simple reason that it is not. We learn from experience that peripheral vision is not as reliable as focal vision; that the apparent shape of the distant tower is not its real shape; that if we do not know how far away an unfamiliar object is, we cannot tell its size just by looking at it; and if we are sufficiently nearsighted or astigmatic, that deliverances of uncorrected vision are not remotely reliable. Nor are the deliverances of the other senses any more secure. So the objection must be that, on his own, an individual not only can learn to distinguish deliverances of dreams and hallucinations from perceptual deliverances but can also to distinguish unreliable perceptual deliverances from reliable perceptual deliverances.

Let us concede for the moment that a perceiver can tell whether he is a consistent judge of color. The items that looked red to him yesterday look red to him again today; variations in the apparent colors of enduring objects can be correlated with variations in the light source; and so on. One visual deliverance supports another. Although this might seem sufficient to enable him to draw the relevant distinctions, it is not. Some items that consistently appear red to him might be green. People discover that they are colorblind by learning that other people discriminate shades that they cannot tell apart. If an agent cannot appeal to other people's color perceptions to check his own, he has no way to tell whether his socks match. For all he knows, he might be colorblind. The same holds for the other secondary qualities. On our own, we do not know what we are missing.

A perceiver fares better with respect to primary qualities. He can check his visual

deliverances against his tactile deliverances. He can correlate the way things look with the way things feel. If they align, he has, on his own, evidence of reliability. But suppose he discovers a systematic misalignment. Objects that feel circular look oval. That is, the perceiver's tactile deliverances are as of every point on a disk's periphery being equidistant from a central point, whereas his visual deliverances are as of the disk being longer in one direction than in another. How could he figure out which, if either, is right? The obvious (and correct) answer is to measure the disk along different axes. If all the measurements yield the same answer, it is a circle; otherwise it is not.⁵ Suppose his measurements reveal that the disk is a circle. Suppose he performs the measurements on numerous occasions on different objects and discovers that the misalignment is consistent. Things that feel circular to him generally look elliptical. Then he should conclude that either his visual perception of shape is defective or circles generally present the visual appearance of being elongated in one dimension. Without consulting others, he has no way to tell which is the case. He can find out through measurements which objects are circular; but by himself, he cannot discover not how circles should look. For there is no a priori reason why things that are equidistant from a given point should look equidistant from a given point.⁶

Of course, most of our perceptual deliverances are not just deliverances as of primary or secondary qualities. They are deliverances as of people, animals, plants, and other objects located in and moving about in space. The crucial question is whether such deliverances are trustworthy. To the extent that the reason to accept them does not turn on the agent's reliability with respect to secondary qualities, perhaps they do. If so, a life-long Robinson Crusoe (who was shipwrecked at birth and reared by wolves) would be in a position to place considerable trust in his mutually supportive perceptual deliverances.

The problem comes with establishing criteria for being the same thing – the same object, quality, relation or whatever. If the criterion were:

x and y are deliverances as of p just in case x presentations and y presentations are identical

then, at least assuming we are not prey to an analog of the private language argument, Life-long Crusoe would have no problem. But deliverances as of the same thing – be it a primary quality, a material object, an organism, or a person – vary. Things present different appearances in different lights, from different perspectives, against different backgrounds, and so forth. Many things change their appearances over time. Arguably an isolated individual could establish workable, stable criteria, but doing so not straightforward. Nevertheless, nothing I have said so far demonstrates that an isolated individual could not do this. For now, let us leave it at that.

Winston's predicament, however, is worse than Life-long Crusoe's. It is not just that he gets no positive reinforcement for his take on things, rather, he faces widespread peer disagreement. In the circumstances where he takes himself to see Jones, Aaronson and Rutland, his similarly situated compatriots avow that they do not. Where he thinks that he remembers that Oceania was recently allied with with Eurasia, his compatriots avow the contrary.

Recent discussion characterizes peer disagreement as disagreement that arises when equally able, similarly situated epistemic agents disagree about a topic. Epistemic peers are equally intelligent, equally well educated vis à vis the topic, equally well positioned to judge, and so forth. Conciliationists⁷ maintain that Anne's realization that her epistemic peer, Bob, disagrees with her undermines her justification for the claim on which they disagree. She has, they maintain, no reason to think she is in a better position

to judge the issue than Bob, so she should suspend judgment or at least lower her credence. The steadfast⁸ maintain that Bob's disagreement does not provide Anne with such an epistemic obligation. Anne, they maintain, should judge that even though Bob is generally her epistemic peer, in the case where they disagree, Bob must be mistaken. Bob's opinion then does not give Anne any reason to change her views.

If Winston's situation involves peer disagreement, some of his compatriots are his epistemic peers. We have no reason to think otherwise. They have the same education, the same sorts of experiences and the same evidence. They are equally well situated to judge visual appearances, news reports, and so forth. There is no indication that Winston is vastly more (or less) intelligent than his compatriots. The only apparent difference is his willingness to question the Party line – a willingness which his compatriots apparently lack. Since differences in skeptical or iconoclastic propensities do not disqualify people from being epistemic peers, at least as far as he can tell (and as far as we can tell), some of his compatriots qualify as his peers.

Winston's situation is more complicated than the ones epistemologists typically discuss. For he cannot tell whether his peers actually disagree with him. Although they may be parroting the Party line, they sometimes avow what they take to be a truth supported by evidence. He has no way to tell whether this is one of those times. If the disagreement is spurious, he obviously should ignore his interlocutors and hold fast to his belief. If it is genuine, and the steadfast position is correct, he should again remain firm. But if the disagreement is genuine and the conciliators are right, he should suspend judgment or lower his credence.

What should he do when he does not know whether the disagreement is genuine? We standardly take intersubjective agreement to sustain or enhance justification. The

greater the intersubjective agreement that p , the more likely it is that p . In Oceania, the situation is nearly reversed. The more people agree that p , the more likely it is that the Party wants them to take it that p ; and the more likely it is that the Party wants them to take it that p , the less reason they have to believe that p is true.⁹ In what we take to be normal epistemic circumstances, the steadfast appear intellectually arrogant. If Anne and Bob are epistemic peers, their situation seems symmetrical. It appears as reasonable for Anne to think that she has made a mistake as to think that Bob has. But given the pervasiveness of the Party's interference, Winston's situation seems different. Perhaps he should assume that his interlocutors are either disingenuous or misguided. He not only knows what he sees, he also recognizes that a well honed sense of self-preservation leads his peers to make their avowals on the basis of political expedience rather than warrant or truth. So, unlike Anne, he has positive reason to disregard them. Whatever may be the case elsewhere, in Oceania, it pays to be steadfast.

But steadfastness comes at a considerable price. It plunges Winston into arrant dogmatism. The Party does not always deceive; it does so intermittently. Sometimes ordinary citizens of Oceania avow what they actually see or remember. In some cases of disagreement, Winston's peers are apt to be correct. But, according to the steadfast, whenever Winston finds himself disagreeing with a peer, he should judge that he is right. Moreover, the Party does not restrict its deception to political issues. To reinforce its thought control, it sometimes deceives about mundane matters. So Winston cannot restrict his dogmatism to certain topics. Every time he disagrees with anyone about anything, he should assume he is right. He could then never discover that he is colorblind. Having no reason to take himself to be a worse judge of colors than Neville, he should not let Neville's apparent disagreement undermine his conviction that his socks

match. Nor, evidently, does it matter how many others apparently agree with Neville. Winston should think that even if people are normally good judges of color, in this case, they are wrong or disingenuous. Neither the fact of apparent disagreement nor the number of those who apparently disagree with him should carry any weight. Since the Party may have its thumb on the scale, this position is not unreasonable. Maybe Winston has little reason to consider his peers wrong, but he has plenty of reason to suspect that they are disingenuous. Winston need not, and perhaps cannot, recognize anyone as his epistemic superior. So in the face of apparent disagreement he should always take himself to be right. His compatriots might conceivably be sources of new information, although its reliability would be highly doubtful. But they can never give him reason to revise his opinions.

Dogmatism is one worry; epistemic insecurity is another. The Party's machinations strip away valuable checks on a person's beliefs. If Winston is reliable in some domain, his take on things in that domain is likely to be correct. But if he cannot tell that he is reliable, this avails him little. He may be steadfast, dogmatic, and complacent about his take on things. But steadfastness sacrifices resources he needs to recognize and correct errors – at least those that do not reveal themselves through obvious inconsistency. He would be epistemically better off if he could adopt a higher-order perspective and recognize when and to what extent he is reliable. By himself, he is in no position to discover systematic or compensating errors. If he considers himself infallible, he is wrong; but if he considers himself fallible, he should be epistemically insecure.

Winston has no check on his memory or current sensations, hence no basis for distinguishing between remembering and seeming to remember, or between seeing and

seeming to see. So introspection is undermined as well. If one cannot distinguish between remembering and seeming to remember, seeing and seeming to see, and so forth, one cannot reliably identify one's own mental states. Is a given representation a memory or merely a fantasy? Is it a belief or merely wishful (or fearful) thinking? Does it reflect what is politically expedient or what is likely the case? Such states do not come already labeled. Memory and seeing involve being causally connected to the world in a law-governed way. Fantasy, wishful thinking, and hallucination have a different etiology. If Luke remembers having seen the Eiffel Tower, then his current mental representation as of his having seen it is suitably connected to his actually having seen it. If he imagines that he saw it, his current mental representation as of having seen it lacks that causal connection. If it is in principle impossible to determine whether the requisite causal connection obtains, the distinction between the real and the apparent wanes. Strictly from the inside, there seems no way to learn to tell the difference. This is not to say that every apparent memory must be checkable. But enough of them must be checkable to stabilize the distinction between the two sorts of representations. 'An "inner process" stands in need of outward criteria.'¹⁰ Strictly from the inside, there seems to be no way to differentiate the real from the merely apparent; and in Oceania, there are no external sources of confirmation.

One might object that Oceanians have the same sorts of beliefs about mundane matters as anyone else; they are just more circumspect about voicing them. Perhaps. But this assumption is not obviously correct. Oceanians might internalize the Party's strictures to such an extent that their attitudes are more reflective of political expedience than of evidence or truth. Beliefs should be sensitive to changes in the world in the sense that, as Tamar Gendler says, 'when we gain new all-things-considered evidence – either

as a result of a change in our evidential relation to the world, or as a result of a change in the (wider) world itself – the norms of belief require that our beliefs change accordingly¹¹. An un suicidal citizen of Oceania *is* sensitive to changes in the wider world. Her take on things *is* responsive to the way the world is, and it changes as the world changes. But the changes she is (most) sensitive to are not changes in the evidence as to whether *p* is the case; they are changes in the evidence that the Party wants people to take *p* to be the case. Possibly her belief-forming mechanism has a political override. She is then sensitive to evidence that bears on the truth of the content of her attitude unless that sensitivity clashes with her sense of the political expedient. Possibly, because her sense of the politically expedient so dominates, she fails to form beliefs, instead forming a hybrid attitude that interweaves doxastic and prudential elements. This may be one place where Orwell's scenario reveals itself to be a reductio, if human beings could not survive without a bedrock of mundane beliefs that are impervious to political meddling. But if a basic function of belief is to promote survival by giving us a way of recognizing and avoiding danger, and if the greatest danger comes from bucking the Party, it would not be unreasonable to expect people's attitudes to be more attuned to the Party's mandates than to other facts.

The Party's thought control is so great that it even convinces Winston that $2+2=5$. At this point, adherents of the a priori might insist that Orwell has gone too far. Unfortunately, it is not obvious that he has. Winston is tortured into believing that $2+2=5$. He does not take to it easily. According to BonJour, 'In the most basic cases [a priori] reasons result from direct or immediate insight into the truth, indeed, the necessary truth, of the relevant claim.'¹² Presumably ' $2+2=4$ ' is a basic case. If so, thinkers typically have a direct or immediate insight into its necessary truth. Nevertheless, it is

not implausible that through a sufficiently intense and painful process of operant conditioning, a person could be brought to block the direct and immediate insight that $2+2=4$ and substitute the (less painful) impression that $2+2=5$. The mere fact that $2+2$ *could not equal 5* does not insure that someone could not be conditioned to think it does. Nor does it insure that someone could not be brought to have a direct and immediate impression that $2+2=5$. No doubt one could maintain that Winston and other torture victims lost their moorings to the extent that they no longer know what ' $2+2=$ ' *means*. So they do not mean what we mean by ' $2+2=x$ '. If 'insight' is a success term, then they lack insight. Nevertheless, Winston has what he takes to be a direct and immediate insight. In Oceania, seemingly a priori deliverances are unreliable too.¹³

Torture is memorable. So we may suppose that Winston remembers being tortured into taking $2+2$ to equal 5. If so, he has reason to reject his direct and immediate deliverance. But more subtle forms of mind control – perhaps hypnosis, propaganda, or delicate forms of operant conditioning – might have the same result. The Party uses a variety of methods that are evidently widely effective. There is no reason to believe that every interference with a priori deliverances leaves a phenomenologically salient mark.

In all the cases I have discussed, things have been so arranged that there can be no independent confirmation of one's judgments. It is impossible or ineffective, for Winston to ask anyone, 'Did you see that?' 'Do you remember that?' 'Does your calculation agree with mine?' 'Didn't yesterday's newspaper say something different?' 'Weren't we taught the contrary in school?' It is fruitless for him to consult contemporary news media or historical records, for they are bound to toe the current Party line.

Winston, and presumably all other even mildly reflective members of the society, are aware of what the Party is doing. Although the Party purports to be telling the truth,

this is a blatant pretense, which no one who is paying attention could believe. Nor, does it matter. Citizens behave as the Party wants them to, regardless. They realize that they cannot trust the media, or their senses, or their compatriots, or their calculations. But their higher-order awareness of the corruption of their epistemic resources does little to improve their epistemic lot. Possibly their higher-order awareness enables them to frequently suspend judgment rather than harbor false beliefs. But they know very little. For the recognition that the Party regularly promulgates falsehoods does not enable them to determine what, in particular, is true.

Some of the contentions a citizen of Oceania accepts are no doubt accurate. Sometimes, she actually sees what she takes herself to see. Sometimes, she remembers events that actually occurred as she remembers them. News media occasionally report something that in fact occurred. Some of her calculations are correct. Moreover, some coherent constellations of considerations she accepts are presumably accurate as well. If she believes the members of such a constellation she has locally justified, locally reliable true beliefs. But because such islands of accuracy are surrounded by a vast sea of unjustified contentions, they are not trustworthy. Given the overarching epistemic circumstances, luck plays too great a role in their being true and in their being justified.

Ordinarily, we assume that the coherence and mutual supportiveness of a suitable number and range of deliverances of sensation, introspection, a priori insight, inference and testimony suffice for justification. It would, we think, be a miracle if all these sources pointed to the same falsehood. Orwell's lesson is that this need not be so. Convergence on the same false or unwarranted conclusion would be a miracle only if the sources were suitably independent of one another. But if mutual support is generated perniciously, the independence assumption is undermined. Minimally, to afford

justification, deliverances must not products of collusion or coercion. More strongly, we may also need reason to believe that they are not.

A vast Orwellian conspiracy is not required to undermine the epistemic standing of familiar deliverances. More subtle arrangements that silence certain voices or deflate their credibility can produce the same effect on a more modest scale.¹⁴ By omitting or downplaying the significance of particular perspectives, prejudice and stubbornness can skew matters to the point where intersubjective agreement does not supply justification. If naysayers are silenced, then the fact that everyone who speaks agrees is not a good reason to think that a contention is true. If the credibility of naysayers' testimony is deflated, then the fact that the balance of what is taken as evidence weighs heavily in favor of a hypothesis is at best a weak reason to believe it. The negative epistemic consequences do not just undermine the opinions of those who have been silenced. They affect the entire epistemic community. Without access to the discredited opinions, community members do not know what they are missing. Moreover, the skewing of the evidence that results from deflating credibility undermines the trustworthiness of what evidence they have, for it artificially inflates the credibility of the voices that are heard. Such a situation is epistemically unjust.

1984 represents an extreme case. But epistemic injustice is not unusual. Here is an example: In the 1940s and early 1950s Barbara McClintock published a series of papers contending that she had discovered sequences of genetic material that change position on the chromosome of corn. She was not taken seriously. As a result, she said, 'I stopped publishing detailed reports . . . when I realized, and acutely, the extent of disinterest and lack of confidence in the conclusions I was drawing from my studies.'¹⁵ Historians of science disagree about whether sexism figured in the scientific community's

disregard of McClintock's findings.¹⁶ Whatever the reason, she was effectively silenced. Her discovery was not taken up until the 1960s when François Jacob and Jacques Monod discovered that the same transposition occurs in bacteria. Plainly the understanding of genetics was retarded by the failure to give McClintock's discoveries their due. Even if it is not Orwellian in motivation or scope, such epistemic injustice deprives the community of inquiry of data and skews the evidence it has.

At a minimum, these reflections on *1984* reveal that epistemology needs a non-interference requirement. Only if beliefs and belief generating mechanisms are free of Orwellian, or quasi-Orwellian interference can their products be epistemically warranted. Arguably, they show something more. If my point about the differentiation of representations and the fixation of belief is correct, then to be capable of having beliefs at all requires the support of a non-coercive community, not merely the absence of interference by a coercive one.

The epistemological is political

Habermas maintains that political wellbeing requires uncoerced conversation.¹⁷ *1984* reveals that epistemic well being does too. Agreement among free and equal inquirers enhances the epistemic standing of a claim; coerced agreement does not. If inquirers are free, they can take up any perspective they like, and examine the issue from that perspective. If they are equal, all inquirers have an equal opportunity and an equal right to venture hypotheses, to raise objections and counter-hypotheses, and provide reasons for them. Under such circumstances, there are no political impediments to discovering what is the case. But if intersubjective agreement results from coercion, collusion, credibility inflation or deflation, none of this holds. If inquirers cannot

examine an issue as they see fit, if they cannot raise objections or gain a hearing for them, or if their views are given undue weight, agreement provides small reason to think that a conclusion is true or that it stands up to serious testing.

'Free and equal' here is a political requirement on a community of inquiry; it is not a claim that all are equally knowledgeable or equally intelligent. Some hypotheses are plainly untenable; some perspectives are obviously skewed; some methods are demonstrably unsound. These are quickly and rightly dismissed. Nevertheless, to block the Orwellian threat, the opportunity to venture a hypothesis, and the right to have it assessed on its merits must be real. As Mill insists, 'The beliefs we have the most warrant for have no safeguard to rest on but a standing invitation to the whole world to prove them unfounded.'¹⁸ The epistemic value of uncoerced conversation lies, I suggest, not so much in the particular agreements that it generates as in its propensity to uncover and correct errors.

It might seem that all that is needed is a non-interference requirement. For an agent to be in a position to know a fact or to understand a constellation of facts, she must stand in a suitable relation to those facts. The Orwellian scenario is a reminder that other people can undermine her capacity to stand in such a relation. That scenario is enough to discredit the assumptions about individualism and attunement that I mentioned at the outset. I think, however, that the epistemic community plays a more substantive role. It constitutes epistemic norms and specifies the conditions for their satisfaction.

With relatively few exceptions, facts are independent of anyone's beliefs about them. We do not need intersubjective agreement (nor would such agreement help) to make beliefs true. Moreover, the reliability of our modes of access to the facts may be beyond our control. For the purposes of argument, let us suppose that they are.¹⁹ Will's

visual perception yields a deliverance as of a black dog. He straightway believes that a black dog is in front of him. He is passive with respect to this belief in that it is a product of involuntary belief-forming mechanisms. He is, in Kantian terms, heteronomous. Just as Kant's heteronomous subject acts on whatever inclinations she happens to have, a heteronomous doxastic subject like Will forms beliefs on the basis of whatever deliverances he happens to have. This does not entail that he believes the content of every deliverance; but it is not up to him which ones he believes. Rather than saying that he forms beliefs, it would be more accurate to say that belief contents just strike him as correct. If asked why he believes there is a black dog in front of him, he could honestly reply, 'That's just the kind of guy I am'. The heteronomous doxastic subject is under the sway of whatever belief-forming mechanism is in effect. If the mechanism is reliable, his belief is apt to be true; if not, it is not. If the mechanism is reliable, his doxastic position is objectively secure. But from a subjective perspective, he seems unacceptably vulnerable. Since his belief contents just strike him as correct, he is in no position to criticize them or reflectively endorse them.

One might argue that such a subject has resources that I did not recognize. He has second order beliefs. He could reflectively endorse his first-order beliefs, and be right to do so, if his mechanism for forming second-order beliefs was reliable. Although true, this is unhelpful. If he is heteronomous with respect to his second-order beliefs – if, that is, his belief that a first order belief is credible just strikes him – the problem recurs. To attempt to solve it by appeal to third-order beliefs which would enable him to endorse his second order beliefs which would enable him to endorse his first order beliefs, sets off a disastrous regress. Is there any way to vindicate second-order considerations without embarking on a regress?

Given that I've posed the problem in Kantian terms, it is no surprise that I extract an answer from Kant. One formulation of the Categorical Imperative is that those maxims are acceptable that an agent can endorse as a legislating member of a kingdom of ends. These maxims are not just laws that the members of the kingdom of ends are subject to, they are laws that they *make themselves* subject to. According to Kant, in the moral realm, legislators enact the laws that bind them. I suggest that the same holds in the epistemic realm. What gives certain second-order claims their epistemic authority is that they express standards, rules, or principles that epistemic agents can on reflection endorse. Thinking of themselves as reasonable and rational, they are prepared to accept those second-order considerations as specifying constraints on what they ought to accept.

The suggestion then is that what blocks the regress is that the second-order endorsement is a product of agency. Epistemic subjects are epistemic agents; they take their beliefs, practices and so forth to be answerable to certain norms because they think that cognitively acceptable beliefs, practices, and so forth ought to be answerable to those norms. They thus make the epistemic laws that bind them. And because they are agents, they can both start and end a justificatory path.²⁰

Why should epistemic agents be construed as joint legislators rather than as autocrats? Why shouldn't each agent decide by herself and only for herself what epistemic principles merit her purely personal reflective endorsement? The answer is this: if an agent reflectively endorses an epistemic principle, she considers it reasonable that her cognitively serious actions, such as inferring and asserting, accord with that principle. She repudiates the gambler's fallacy because she recognizes that in committing the fallacy she makes herself vulnerable to Dutch Books. She endorses modus ponens because she recognizes that it is truth preserving. Such an agent has no reason to think

that as an epistemic agent she should be subject to principles that similarly situated epistemic agents are not subject to. She does not, for example, think that the desirability of avoiding Dutch Books stems from a personal predilection. Because she takes the principles she reflectively endorses to be reasonable and rational in the epistemic circumstances, she thinks they should be binding on similarly situated epistemic agents. But she recognizes that epistemic agents should be subject only to principles that they consider worthy of reflectively endorsement. By her own lights then, only such principles as she can justify to the similarly situated epistemic agents – those who constitute her community of inquiry – merit her reflective endorsement. Because the principles that govern their epistemic practices must be ones that its members can justify to one another, the community of inquiry serves as a stay against idiosyncrasy or bias in reflective endorsement. It not only fails to interfere, it actively fosters epistemic success.

Does it follow that Life-long Crusoe is incapable of epistemic success? Possibly. He evidently lacks the resources to rule out certain sources of error. One is confirmation bias. Someone given to confirmation bias weighs evidence that supports his convictions more heavily than evidence that undermines them. If Life-long Crusoe does so, he may have no way to discover his error, particularly if the bias is slight. Then many, if not all, of his beliefs are fated to be unjustified. Suppose, however, that he is not prey to confirmation bias. Then he is, we may assume, as reliable as the rest of us. It might seem therefore that he is just as capable of knowing, understanding, and being justified as we are. But we have a resource that he lacks. We assess our beliefs in light of the standards that the community of inquiry has designed to (among other things) filter out confirmation bias. So we have a reason to believe that beliefs that satisfy the standards are not the result of confirmation bias. Life-long Crusoe has no such reason. His

grounds for his belief are thus epistemically impoverished as compared to ours. Satisfying the standards of a community of free and equal epistemic agents enhances the epistemic standing of a belief.

Whether we should conclude that Life-long Crusoe is incapable of knowledge (or some other epistemic success) is not clear. That depends on where thresholds are set. If first order reliability suffices, and Life-long Crusoe is not prey to confirmation bias (or kindred epistemic failings), he knows. If epistemic success requires having reason to believe one has avoided the pitfalls, he does not. But even if he knows, though, his situation is epistemically precarious. It is, from his perspective, just by luck that he has managed to avoid the pitfalls.²¹

An epistemic agent can be mistaken about which principles the members of her community of inquiry can justify to one another. In reflectively endorsing p , she commits herself to p 's satisfying standards that are justifiable to the relevant community of inquiry. If it does not, her endorsement of p is an error. By her own lights, she ought not accept it. This raises a more serious worry. What if the community standards are wrong?

A community of inquiry sets and reflectively endorses its epistemic standards, and its members justify their inferences and conclusions to themselves and one another by appeal to those standards. Because the community has ample evidence of human fallibility, the standards and criteria of application they reflectively endorse not considered fixed or final. They are subject to rejection or refinement if they do not serve the community's evolving epistemic ends.²² The claim of current standards to epistemic authority lies in their having thus far withstood the standing invitation to prove them unfounded. This is no guarantee, but it affords some reason to countenance them.

Nevertheless, one thinks, the principles might be misguided. Being unrefuted is not the same as being irrefutable.

Maybe we would do better then to insist that some other feature – perhaps reliability or bearing a suitable relation to foundational claims – is constitutive of epistemic acceptability. Is there any reason to think that the sort of reflective endorsement I have described is anything more than indicative of acceptability? In answering this, we should remember that we are epistemic *agents*. Epistemic resources are things we work with, not just things we admire or credit. We use them to apply, extend, and improve upon what we take ourselves to know and understand. The considerations that have withstood Mill's challenge are considerations that epistemic agents consider as sound a basis for reasoning and epistemically responsible action as any they have. Maybe they are not really (in some sense of 'really') a sound basis for reasoning and action. Perhaps an agent's beliefs are incorrect, or her methods are biased. Perhaps her community shares her errors and biases. This could be so. Although her system has gone through a fairly rigorous process of testing and correction, there is not guarantee that all flaws have been eliminated.

Still, what is an agent to do? Either she assesses her deliverances on the basis of standards she reflectively endorses or she does something else. If she forgoes assessment entirely, she is heteronomous. She thinks and acts on whatever strikes her, subjecting her deliverances to no critical filter. Should her beliefs turn out to be true, that is, from her perspective, a lucky accident; she has nothing she can call a reason for them. If she deliberates and acts on the basis of standards she does not reflectively endorse, her epistemic situation is equally bleak. Even if satisfying alien standards is objectively more truth conducive than satisfying her own standards, she has no reason to think so. By her

own lights, the fact that a belief that satisfies those standards is true is still a lucky accident. To assess her beliefs on the basis of alien standards thus would be epistemically irresponsible. She has no reason to trust them. Still, to be complacent about whatever standards she finds herself with would be irresponsible as well. The mere fact that the standards are hers is no reason to credit them. But if her standards have survived serious, sustained testing and correction by a community of free and equal epistemic agents – if, despite their best efforts, those agents have not been able to prove them unfounded – relying on those standards in the current epistemic circumstances is not irresponsible. Fallible and potentially flawed as they are, the standards are at least as good as any available alternative.²³

¹George Orwell, *1984*, New York: New American Library, 1961, p. 81.

²Jonathan Vogel. 'Are There Counterexamples to the Closure Principle?' in *Doubting: Contemporary Perspectives on Skepticism*. Kluwer, 1990, pp. 13-27.

³Orwell, p. 41.

⁴Bernard Williams, *Truth and Truthfulness*, Princeton: Princeton University Press, 2002, pp. 135-140.

⁵This assumes that he can trust his measurements. For now, let us simply grant that.

⁶This is a variant on the Molyneux problem.

⁷See David Christensen, 'Epistemology of Disagreement: the Good News' *Philosophical Review*, 116: 187-217.

⁸Earl Conee, 'Rational Disagreement Defended,' in *Disagreement* ed. Richard Feldman and Ted Warfield. Oxford: Oxford University Press, 2010, pp. 69-90.

⁹I am grateful to Jonathan Adler for this point.

¹⁰Ludwig Wittgenstein, *Philosophical Investigations*. Oxford: Blackwell, 1953, #580. See also, Wilfrid Sellars, 'Empiricism and the Philosophy of Mind,' *Science, Perception, and Reality*. London: Routledge and Kegan Paul, 1963, pp. 127-196.

¹¹Tamar Gendler, 'Alief in Action (and Reaction)'. *Mind and Language* 23 (2008), 565.

¹² Laurence Bonjour, 'In Defense of the *A Priori*', in *Contemporary Debates in Epistemology*, ed. Matthias Steup and Ernest Sosa. Malden: Blackwell, 2005, p. 99.

¹³ One might think that closely related beliefs, such as the belief that $1+1+1+1=4$ suffice to indicate to Winston that he is wrong. So systematicity considerations would tell against Orwell. But this need not be the case. All that the conflicts can reveal to Winston is that he's going wrong somewhere. He might, nevertheless, be incapable of figuring out which of his beliefs is incorrect. Remember what it is like to not understand how to do a math problem. You may fully realize that something has gone wrong somewhere without having a clue which of the claims is incorrect.

¹⁴See Miranda Fricker, *Epistemic Injustice*, Oxford: Oxford University Press, 2007; and Mary Kate McGowan, 'On Silencing and Sexual Refusal' *Journal of Political Philosophy* 17 (2009), 487-494,

¹⁵Barbara McClintock, Letter to J. R. S. Finchman, May 16, 1973,

<profiles.nlm.nih.gov/LL/Views/Exhibit/documents/biographical.html>, consulted June 10, 2009.

¹⁶See Evelyn Fox Keller, *A Feeling for the Organism*. New York: W. H. Freeman, 1983; Nathan C. Comfort. *The Tangled Field: Barbara McClintock's Search for the Patterns of Genetic Control*. Cambridge: Harvard University Press, 2001.

¹⁷Jurgen Habermas, 'Truth and Society: The Discursive Redemption of Factual Claims to Validity,' *On the Pragmatics of Social Interaction*. Cambridge: MIT Press, 2001, 85-104.

¹⁸John Stuart Mill, *On Liberty*, Indianapolis: Hackett, 1978, p. 20.

¹⁹This is a simplification, since we can learn to refine our sensibilities and thereby gain epistemic access to aspects of things that were originally inaccessible. Moreover, learning enables us to recognize things to which we were once oblivious.

²⁰I am indebted to Jonathan Adler for this point.

²¹Thanks to Alvin Goldman and Sanford Goldberg for pressing me to address this issue.

²²An example may bring this out. 'Inanimate objects are identical when their parts are identical' is a principle that many philosophers consider correct. Being a universal claim, it should hold for particles of a viscous fluid. However, if it holds, then 'F=ma' does not. Molecules in a viscous fluid move at different rates. In prototypical applications of 'F=ma', forces act on objects like billiard balls that have sharp boundaries. But in viscous fluids, the 'forces' on the 'object' are the effects on momentum of molecules moving in and out of that 'object'. So preserving 'F=ma' requires continually redefining what constitutes a particle, letting different molecules comprise it at different times. Although physicists are willing to concede that 'F=ma' does not hold at the quantum level or at relativistic speeds and distances, it is an extremely valuable law for characterizing the behavior of middle sized items in this neighborhood of the cosmos. These include the viscous fluids flowing around here. Fluid mechanics thus characterizes its particles so as to comport with the law. Rather than insisting that all component molecules of a particle be the same from one instant to the next, they let the individual molecules come and go, but keep the average enclosed mass constant. See Mark Wilson, *Wandering Significance*, Oxford: Clarendon Press, 2006, pp158-9). 'F=ma', evidently, is a so central a law of physics that scientists are willing to make drastic revisions in the criteria for the identity of a fluid particle over time in order to preserve it. In this case, the tension is acute. One way or another a major revision in antecedently plausible principles is needed. Either scientists must revise a very reasonable metaphysical commitment about the identity of an object over time, or revise a fundamental law of physics. The recognition of the clash leads to a revision in the principles epistemic agents are willing to accept on reflection.

²³I am grateful to Jonathan Adler and Jonathan Matheson for helpful comments on an earlier draft of this paper.