Epistemological Detente?

Catherine Z. Elgin

What is the problem?

Epistemological reliabilists maintain that tenable beliefs owe their tenability to their being produced or sustained by reliable processes. Evidentialists maintain that such beliefs owe their tenability to their being backed by suitable evidence or reasons. Both groups hold that the two approaches conflict. In particular, they diverge over the correct way to answer the question, What makes S's belief that *p tenable*?' I will argue that they are wrong. Reliabilists and evidentialists give different answers because they are answering different questions. Roughly, an evidentialist takes 'What secures the tenability of *S*'s belief that *p*?'?' to be asking what reasons *S* has that support the conclusion that *p*, where a reason is an intersubjectivly accessible consideration that she could appeal to defend her view. A reliabilist takes the answer to the question to rest on the processes or mechanisms that either cause or sustain the belief. *S* need not be aware of the process, mechanism or its reliability; hence she need not be able to appeal to it, or to anything else, to defend her belief. If I am right, insofar as there is a dispute between evidentialism and reliabilism, it is about what sorts of issues epistemology ought to be addressing, not about who has the correct answer to a shared, univocal question. But, I will urge, there need be no dispute because both are reasonable questions.

Kornblith's process reliabilism

Hilary Kornblith is a naturalist. He holds that the best way to understand a phenomenon is through the lens of science. In particular, when we want to understand why someone believes what he does, we should look to cognitive psychology. This approach, Kornblith maintains, favors process reliabilism, which holds that beliefs are epistemically tenable if and only if they are

generated or sustained by reliable processes. Psychology is the science that discloses the processes that generate and sustain beliefs and other mental states, and it demonstrates their reliability.

Process reliabilists do not all speak with a single voice. One can be a process reliabilist without being a naturalist. Sanford Goldberg (2007, 2010), for example, holds that some social processes are reliable. They are guided by social rules and conventions rather than, or in addition to, natural laws. When an epistemic agent deploys such rules and conventions in her belief forming or sustaining activities, the results are, ceteris paribus, tenable beliefs. Nor need a process reliabilist with a naturalist bent privilege psychology. Alvin Goldman (1999, 2006), draws on psychology, neuroscience, and the social sciences. An important question for naturalism is what science or sciences to privilege. An important question for reliabilism is what processes, natural or not, are sufficiently reliable. Here I focus on Kornblith's position, and in particular on the papers in *Second Thoughts and the Epistemological Enterprise* (2019). I say nothing about whether the strengths and limitations I identify apply to other forms of process reliabilism.

Kornblith emphasizes that many beliefs and inferences are subliminal. It is no accident that we walk through doorways rather than crashing into walls, that we eat the salad but not the fork, that we reach for a corkscrew rather than an eggplant when we want to open the wine. Without reflection, or any need for reflection, we act on subliminal beliefs, and do so successfully. Many subliminal beliefs are transient. Although we do not consciously access them, we continually update our mental files to accommodate changing circumstances. A competent driver rapidly forms sequences of ephemeral beliefs that attune him to current traffic conditions and the resources he has for dealing with them. The beliefs he forms are not stored in longterm memory, since they exhaust their utility in the context in which they emerge. But by drawing on such inputs, the driver makes quick, nimble, effective inferences about the opportunities and obstacles that confront him.

Some subliminal beliefs persist. We believe, for example, that the floor is solid, that water is wet, that cats don't grow on trees. Enduring subliminal beliefs are so obvious and commonplace that we need not consciously entertain them. Some are obvious and commonplace to everyone; others only to experts. Even if a novice has to think through the steps and remind herself of their relevance, to an experienced EMT, CPR is so obvious that he springs into action without deliberation. Persistent subliminal beliefs and the inferences we unconsciously make on the basis of them constitute the substratum of our mental lives that we simply, unthinkingly take for granted. Although neither the subliminal beliefs nor the subliminal inferences rise to the level of awareness, according to psychology they constitute the underlying psychological processes that form the basis for conscious beliefs and actions. All of our conscious beliefs and inferences are products of subliminal processes.

One might wonder whether we should call such subliminal commitments beliefs and inferences at all. Process reliabilism has an answer. Since the commitments are products of the same sorts of processes that occur in conscious belief formation and fixation, and since they perform the same sorts of function in our cognitive economy as their consciously available counterparts, they are plausibly the same sort of thing. Some subliminal commitments can be brought to awareness; and when we acknowledge them, we seem to be reporting what we already believe or infer, not coming up with a new belief or inference. Despite the fact that we never considered it, we already believed that cats don't grow on trees. Hence, the process reliabilist maintains, allowing such subliminal takes on things to count as beliefs and such subliminal processes to count as inferences makes sense.

In addition to strictly subliminal beliefs, human beings harbor unreflective conscious beliefs and make unreflective conscious inferences using them. Some are mundane. We are apt to

unreflectively believe that the coffee maker has to be plugged in before it will work; that umbrellas protect against the rain. We thus unreflectively infer that the unplugged coffee maker will make no coffee and that it is advisable to unfurl the umbrella when it starts to rain. Other unreflective beliefs may involve expertise. A pediatrician unreflectively believes that a child displaying a distinctive rash has chicken pox. A truck driver familiar with local traffic patterns unreflectively infers that leaving the highway and taking the longer route will get him to his destination more quickly than sitting out the traffic jam. Such unreflective beliefs and inferences are accessible to the epistemic agent, so they could in different circumstances be products of reflection. Another driver might, for example, weigh the pros and cons of sitting out the traffic jam as opposed to taking the longer route. She might arrive via reflection at the same decision that the trucker unreflectively made.

Failing to reflect is not in itself problematic. 'The vast majority of our beliefs are formed unreflectively. . . . New information we acquire is typically inferentially integrated with the rest of our beliefs without any need for reflective self-examination.' (Kornblith 2019:140) Besides often being unnecessary, reflection is time consuming. Opportunities can be squandered while we muse over the desirability of the various alternatives. Moreover, an unreflective process may be more reliable than a reflective one. Reflection sometimes yields conclusions that are inconclusive, inefficient, erroneous, or biased. The uses we make of them in deliberations may be flawed. Reflection, reliabilists insist, is no panacea.

Still, sometimes we do reflect. We adduce reasons for or against a consideration and, we think, draw the best conclusion we can given our assessment of the weight of the evidence. Moreover, we sometimes think that reflection is the proper way to resolve a question. Members of a jury are expected to weigh the evidence admitted in court. They may have no antecedent opinions about whether the defendant is guilty. Even if they have antecedent views, they are expected to set

them aside. Similarly, in deciding whether to accept a hypothesis, a scientist is expected to weigh the evidence her experiment yields rather than depending entirely on her antecedent beliefs about the phenomena. The driver who deliberated about whether to seek an alternative route rather than immediately, unthinkingly exiting the highway may have been epistemically responsible. Given her relative ignorance of the situation, she needed, it seems, to think things through. Reflection can sometimes settle issues that unreflective beliefs leave open.

So far, I've given a taxonomy of types of belief. Some are subliminal. They are inaccessible from the first person perspective. Hence an epistemic agent cannot adduce introspectively available reasons to support these beliefs; nor can she appeal to her subliminal beliefs to support conscious beliefs. Others are accessible but are held unreflectively. Arguably an agent could come up with a reason for thinking that the coffee maker won't work unless it is plugged in, but she sees no need to. Nor, apparently is her belief epistemically precarious because she provides no backing for it. Yet others are reflectively endorsed. The agent identified and weighed her reasons and drew the conclusion that they favor. Even though many of our beliefs do not reach the threshold for reflective endorsement, it might still seem that the epistemically best ones do.

Kornblith disagrees. The first person perspective, he maintains, is inherently unreliable. Recent psychology has shown that human beings are overwhelmingly subject to confirmation bias (2019:143). If someone who unreflectively believes that *p* is asked to marshal evidence pertaining to *p* or to consider whether she is justified in believing that *p,* she is apt to be strongly inclined to recall more evidence that confirms *p* than evidence that disconfirms *p.* That being so, Kornblith maintains, we should doubt that a belief we endorse on reflection is more reliable than its unreflective predecessor. Indeed, it may be less reliable.

Moreover, human beings confabulate. When asked why she believes that *p*, a person is apt to fabricate reasons that neither figured in the production of the belief nor figure in the processes that actually sustain the belief. We attribute 'various beliefs to ourselves, beliefs which we take to have been instrumental in producing the very belief we are questioning, but which we did not hold prior to the exercise of self-examination' (Kornblith 2019:142-143). This does not automatically make the beliefs for which we confabulate reasons untenable. If they were produced or are sustained by a reliable process they are tenable. But the confabulated support is spurious. It contributes nothing. Taken together, our propensity to confabulate and our susceptibility to confirmation bias lead to a disheartening result. We do not – indeed, cannot – from a first personal perspective know why we believe what we do. Psychology holds that the basis for our beliefs consists of a variety of subliminal inputs and processes that are unavailable to introspection.

Things may be even worse than this suggests, for recently psychologists have found that humans are subject to a wide range of biases that undermine reliability (see Kahneman et al. 1982; Gilovich et al. 2002; Kahneman 2011). We ignore base rates. Our reasoning is subject to framing effects – that is, we draw different conclusions from the same information, depending on how that information is presented. We are remarkably inept at probabilistic reasoning. We regularly jump to conclusions. Insofar as these gaffes are unavailable to introspection, we do not know that we are committing them. The subliminal psychological processes that underlie all of our beliefs are buggy. So in any given case, it seems, we are in no position to ascertain whether a belief is reliably produced or sustained.

So What?

I do not want to quarrel with psychological findings that Kornblith reports. My concern is with whether they have the consequences he supposes. We should, no doubt, concede that

Descartes was too optimistic about the powers of the individual, reflective, introspective mind. We are all fallible; we are vulnerable to the cognitive infelicities that psychologists identify and probably to others as well. But to favor third personal psychological insights over the first person perspective requires more than acknowledging that the first person perspective, with its reliance on deliberation and introspection, is far from fail-safe.

Cognitive psychology concerns itself with human reasoning in general. Its evidence is statistical. It can report with confidence that such and such a proportion of human subjects ignore base rates, engage in confirmation bias, confabulate, or whatever. This does not, however, disclose, or purport to disclose anything about what happens in an individual case. It cannot show that Ben's belief that the economy is improving is due to or infected by confirmation bias. Indeed, it cannot even show that Ben himself is given to confirmation bias. For statistical evidence does not support exceptionless generalizations. If Ben unreflectively believed that the economy is improving and went on to adduce evidence for that belief, he might have been subject to confirmation bias. Cognitive psychology can tell us that. It can, presumably, tell us how likely it is that an average believer is to be subject to confirmation bias. Maybe it can tell us whether Ben fits the profile for an average believer or whether his divergence from the average makes him more or less likely to be susceptible to confirmation bias. But that is as far as it goes. It cannot reveal that in attempting to adduce reasons for this particular belief, Ben engaged in confirmation bias.

Moreover, the reasoning that Kornblith characterizes as corrupted by confirmation bias qualifies as biased only on the assumption that the agent who engages in that reasoning is seeking to figure out whether a given hypothesis is true. If so, a propensity to overemphasize confirming information while downplaying or neglecting disconfirming information is unfounded. But in the cases Kornblith discusses, the agents already unreflectively believe the hypothesis. They take its

epistemic status as settled. So in looking for evidence, they might be doing something different – investigating how stable the belief is (see Lapidow and Walker 2020). They adduce considerations that assume the truth of the hypothesis and consider the extent to which it remains true in varying circumstances. If this is what they are doing, focusing on confirming evidence is not a bias, it is a viable epistemic strategy.

Often when an agent deliberates, it is because he has no settled belief about the matter. He entertains evidence for and the evidence against a hypothesis because he is not at the outset sure what to believe. It is not clear how a propensity to confirmation bias could undermine this cognitive function. If Ben had no fixed opinion on the matter before he marshaled evidence, there seems no way for confirmation bias to get a foothold on his opinion.

Kornblith argues that because Ben does not have first personal introspective access to the underlying bases for his belief that the economy is improving, he is in no position to answer the question 'Why do you believe *that*?' or at least in no better position than a psychologist adopting an objective, third-personal perspective on the issue. But psychology's third personal perspective cannot answer the question either. It may be able to answer the question about why folks in general, or perhaps folks like Ben, believe *things like that*, but it cannot supply the etiology for Ben's belief in particular. So if Ben's appeal to his reasons is untenable, and cognitive psychology cannot, and does not purport to be able to answer the question, there is just no way to tell why Ben believes that the economy is improving.

Ben has what he takes to be reasons. He follows the stock market and has noticed an upward trend. That, he thinks, is evidence that the economy as a whole is doing better. Is that his reason? According to Kornblith, we should be skeptical, since Ben might be confabulating. If so, he did not arrive at his belief because of his information about the stock market; nor is does he sustain it

because of his awareness of the trend he discerned. Rather, he unconsciously made up a reason when one was demanded. The charge of confabulating is plausible if the subject did not have the information he adduces prior to being asked for it – that is, if he really made it up. But it is harder to sustain if he had the information but it did not figure in the etiology of his belief. As we have seen, the actual basis of his belief is something Ben cannot supply, for it lies in the murky, inaccessible, subliminal processes that gave rise to that belief. Maybe the route from those processes to his belief that the economy is improving ran through his opinions about the stock market; maybe not. Maybe the route went through details of his personal history that do not traverse consciously available reasons at all. Perhaps because he was brought up in a particular way, subject to a variety of psychological pressures, he is unreasonably apt to look on the bright side. If so he is strongly, subliminally, even irrationally disposed to think that things are looking up. We can't rule that out.

Here's where things stand. Cognitive psychology maintains that all human beliefs are generated by subliminal processes that are inaccessible to introspection. Although psychology can say in general what some of these processes are, it cannot identify the particular subliminal beliefs and inferences that give rise to an individual's specific conscious belief. Psychology has shown that many human beliefs are held unreflectively. People who harbor such beliefs appeal to no reasons or evidence to support them. According to process reliabilism, this is not a problem. So long as the subliminal and/or unreflective processes that cause or sustain beliefs are reliable, the beliefs are tenable. Psychology also has shown that human beings in general are prone to a variety of biases. When these biases are in play, the processes are not reliable. Although Kornblith highlights confirmation bias and confabulation, there are a host of others as well. Ordinarily, people do not realize that their reasoning is biased. If the processes in question are subliminal, then although they might know that human reasoning in general is subject to bias, they cannot know that

their reasoning in any particular case is biased. It might seem that an agent could simply check. If Jess suspects that her route to a particular conclusion embodied the gambler's fallacy, she can articulate and analyze her reasoning and see whether it does. But according to Kornblith, because she doesn't know (and can't know) how she actually arrived at her conclusion, the check is unreliable. She is caught in Catch-22. She has no grounds for thinking that the reasoning she is examining is actually the reasoning she used to arrive at her conclusion.

An evidentialist alternative

An evidentialist would take a different tack. Nor more than reliabilists, do evidentialists speak with a single voice. Here I present my own view. It is an amalgam of evidentialism, holism, and social epistemology.[1] Plenty of evidentialists disagree with me. Nor do I claim that the position that emerges suffices for knowledge. Although evidentialists and reliabilists often maintain that satisfying their standards yields knowledge, here I am concerned, as Kornblith is in the papers I discuss, with the standards for tenable belief.

Granted, human beings have a vast array of subliminal and unreflective beliefs, and engage in a vast array of subliminal and unreflective reasoning processes. But despite being subliminal and/or unreflective, they are susceptible of epistemic assessment. Investigations into heuristics and biases afford evidence that, even if the beliefs and processes themselves are not introspectively available, they are manifest sources of inconsistency, incoherence, irrationality (see Tversky and Kahneman 1982). That being so, we have evidence that some automatic reasoning strategies are flawed. We can and should devise ways to identify, and if need be, correct or compensate for the biases – to learn to attend to base rates, to represent probability as relative frequency, to reconfigure

[1]My view is also committed to non-veritism (see Elgin 2017). But non-veritism plays no role in the argument presented here. So evidentialists of a veritistic bent need not reject the position sketched in this paper.

the evidence to test for framing effects, and so on (see Gigerenzer, 2000). We are not passive victims to our default inadequacies.

Kornblith's process reliabilism is hamstrung because it keys epistemic standing to etiology. If an agent arrives at a conclusion via a route that is unreliable, his belief is untenable. Moreover, because we do not have epistemic access to the underlying psychological processes or their inputs, we are in no position to determine whether we took an unreliable route. If we attempt to adduce reasons for a belief, we might easily be confabulating or engaging in confirmation bias. So our self-assessments are untrustworthy. This is a serious problem if the tenability of a belief is hostage to its etiology. But an evidentialist need say no such thing. Often there is no need to inquire into underlying mechanisms. Epistemic agents have sufficient, readily accessible, surface level reasons to support their beliefs. Ben, we saw, thinks he came by his belief that the economy is improving because he has been following stock market trends. Maybe he is confabulating. Maybe his attempt to account for his optimistic take on the economy is due to confirmation bias. But now that he thinks about the issue, he realizes that he has plenty of additional evidence of an economic upturn. As David Lewis (1983:278-279) says, 'Sometimes evidence is not lacking. We who have lived in the world for a while have plenty of evidence, but we may not have learned as much from it as we could have done. This evidence bears on a certain proposition. If only that proposition is formulated, straightway it will be apparent that we have very good evidence for it.'[2] Having formulated his hypothesis, Ben recognizes that he has considerable, independent support for it. He is aware that fewer people are applying for food stamps, that new businesses are opening, that unemployment is down. More dubiously perhaps, he thinks that no one would pay $4.50 for a latte if times were tough. That, arguably, is weak evidence, but it is not nothing. Even if the belief that

[2]In the paper cited here, Lewis maintains that a fiction can be the seed crystal for synthesizing the evidence. But clearly a newly formulated factual hypothesis could do the same.

prompted his entertaining the issue was a product of confabulation or confirmation bias, he has ample independent evidence to back it up. The original path to his conclusion may have been unreliable, but the additional evidence compensates for the conclusion's dubious pedigree. He has fallback positions.

Multiple sorts of accessible evidence in addition to empirical data support beliefs. Because an agent recognizes that a conclusion is obviously entailed by other beliefs she already endorses, she has pro tanto reason to accept it. Greta sees Harry across the room and readily infers that since he is in Cambridge, he's not in Hong Kong. So she concludes that Harry is not in Hong Kong. Manifest logical entailment provides all she needs. Her support might involve the conviction that any alternative to her conclusion is too far-fetched to be a viable alternative. Sure, if material objects were multiply realizable, Harry could be both in Cambridge and in Hong Kong at the same time, but multiple realization of material objects is too implausible to take seriously. So she dismisses that alternative. She may even defend her belief by arguing that if things had been otherwise, she would have heard about it (see Goldberg 2010). In other cases, the fact that a belief content meshes well with her background information while the alternatives do fail to mesh may suffice. Although all such reasons presumably emerge from subliminal beliefs and processes, they do not depend on the details of those processes. What goes on at the epistemically accessible, surface level affords all the support the agent needs.

All of these strategies are fallible. The evidentialist does not maintain that they are fail-safe or even that they are reliable. She does, however, maintain that they are viable approaches to the human epistemic predicament. They give reasons a role in reasoning. Ben can defend his beliefs. He can say, 'Here are my reasons. Here is why I think they are sufficient.' By offering reasons, he opens his belief to public assessment. He, in particular, is subject to criticism if his specific reasons

are not good enough. But neither he nor his belief is impugned by generalizations about the human propensity to commit various fallacies.

The discussion so far has focused on the isolated individual agent. This is a familiar approach, having its origins in Descartes, whose isolated ego sought to figure out what it could know on its own. But it has become increasingly evident that we are epistemically interdependent. So a preferable way to look at the issues is to recognize that there is a social dimension to epistemology. We learn from testimony, instruction, and conversation. Let's look at another case, to see how this helps. Kornblith sketches Mary's predicament (2019:148). I supply additional details.

Mary is a physician and a mother of two. She has a lot on her plate. One Tuesday morning, she gets the kids ready for school, reminding her daughter that she has violin lesson and her son that his book report is due. After dropping the kids off at school, Mary drives to the clinic. She is a conscientious driver, attentive and responsive to what is going on around her. As she feeds the kids breakfast and drives to work, Mary continuously registers her surroundings and updates her beliefs accordingly. It is no accident that she avoids stepping on the cat's tail as she walks across the kitchen, even though she did not consciously notice him underfoot. Nor is it an accident that she brakes to avoid hitting a seemingly suicidal squirrel. She has and acts on the full panoply of subliminal and unreflective beliefs that Kornblith discusses. All the while, Mary is brooding over a tricky diagnosis. Although she suspects that her patient has disease *A*, his signs and symptoms are ambiguous. Some are more strongly indicative of disease *B*. Both diseases are serious, and call for different treatments. She might, of course, be a victim of confirmation bias, subtly adducing more evidence for *A* than for *B*. We cannot rule that out. She mentally reviews the patient's symptoms and his test results, attempting to figure out what they indicate. That is, she deliberates. Still, she's not

satisfied. In reviewing the case, she asks herself, 'What am I missing?' That is, she introspects. Her introspection is selective. She focuses on those of her beliefs, hunches, suspicions, and doubts that seem relevant to the diagnosis. As a result of her deliberations, she decides she needs more information. She orders additional tests. She checks online for recent research that bears on the diagnosis. She looks only at reputable medical journals, ignoring blogs of medical quacks and skeptics. In so doing, she appeals to expertise. She presents her conundrum at the department's Difficult Diagnosis seminar, hoping that her colleagues will be able to point to a factor she missed.

Mary is highly educated in her field. When she asks herself 'What am I missing?' she knows what sorts of factors she should focus on. She appreciates that she has to attend to base rates, and also has to acknowledge that tricky medical cases are often anomalous. She is then sensitive to both the typical and the exceptional. Being aware of how her information and her perspective on it are or might be limited, she does what she can to compensate for the limitations. She augments her empirical information by running additional tests. She increases her understanding of the diseases by consulting recent research. Appreciating that she might be subject to confirmation bias or confabulation in favoring diagnosis *A*, she consults with colleagues whose judgment she trusts. Although they too are subject to biases, there is no reason to think that they would err in the direction she does. If a colleague unreflectively thinks that the patient suffers from disease *B*, then confirmation bias, if he is subject to it, would prompt him to identify evidence that unduly supports diagnosis *B*. By seeking their advice and appealing to their judgment, she broadens her epistemic reach and compensates for her limitations. When she consults with them she does not just ask what they think, she asks why. So she seeks their reasons, not just their verdicts. When they seek to offer insights, they may imaginatively adopt her perspective. 'If I were in your situation, I would . . .'. The capacity to do this indicates that they can engage in

relevantly constrained simulative introspection. They adopt her perspective, and see how things would look from there. But they draw on resources that their own perspective provides. The fruits of introspection are not thus limited to the contents or idiosyncrasies of an individual mind. Nor is Mary a passive recipient of these inputs. She argues with her colleagues, pits her reasons against theirs. Her final diagnosis is a thus product of joint inquiry. It lies at the intersection of multiple, divergent lines of reasoning.

Her exchanges with her colleagues are infused with moral and epistemic norms. When she confesses that she is in a quandary about the proper diagnosis, she is not just giving voice to her personal feelings of inadequacy. She grounds her doubts in publicly available reasons that her colleagues ought not, by their own lights, reject (see Scanlon 1998). She proffers reasons that should weigh with them as physicians, and asks them for medically relevant reasons that they think should weigh with her. That requires that they have shared or at least largely overlapping standards of relevance and significance. Because she trusts her colleagues' integrity and medical judgment, she appeals to them. She counts on them to raise objections if they think that she is wrong. If she suspected that they were incompetent or were out to sabotage her or her patient, she would not and should not consult them. But given the trust, she has reason to seriously entertain their insights and draw on their expertise.

In the end, she concludes that the patient suffers from disease *A*. That was what she originally suspected. But she is not back where she started. She has a better understanding of the medical situation with all of its uncertainties. As far as possible, Mary controlled for bias, compensated for vulnerabilities, looked for and assessed alternatives to what she was already inclined to think. She examined what she took to be her own thinking and the thinking of her colleagues, to insure that it was responsible and responsive to the evidence. Regrettably, she cannot

be completely confident that her diagnosis is correct. But she can be confident that it is epistemically responsible. She did the best that could be done to arrive at a correct diagnosis. She is aware of what the medically sound reasons are for that diagnosis and aware of how strong those reasons are. She can recommend a course of treatment to her patient and explain why she thinks it is the best thing to do. She was in no position to do this when she had only her own unreflective belief that he suffered from disease *A*.

Mary appealed to a variety of strategies that are, Kornblith maintains, unreliable – deliberation, introspection, intuition (in the form of clinical judgment), information gleaned from others who rely on introspection, deliberation, clinical judgment, mental simulation, and so forth. He is right, of course, that they are individually less than wholly reliable. Arguably they are jointly unreliable as well. But taken together they seem to yield results worthy of Mary's reflective endorsement. Mary's diagnosis is backed by reasons that she justifiably considers worthy of acceptance. The strategies she used yield results that are the best that can be done in the epistemic circumstances. If we were suffering from an ailment whose diagnosis was not obvious, we would want our physicians to act as Mary did.

Splitting the Difference

Belief, as it is commonly understood, has two components. One is the feeling or sense that things are as the belief content takes them to be. The other is serving as a basis for inference and action when one's ends are cognitive. L. Jonathan Cohen (1992) argues that the two should be prized apart. He reserves the term 'belief' for the first, and introduces the term 'acceptance' for the second.[3] To avoid confusion, let us call Cohen's sense of belief 'belief$_c$'. 'Belief' without a

[3]Elsewhere, I modify Cohen's conception of acceptance slightly. Cohen restricts acceptance to propositional items. I allow for a broader range of commitments (see Elgin 2017). For the discussion here, my divergence from Cohen does not matter.

subscript retains its everyday use. If Ben believes$_c$ that the economy is improving, he feels that it is so. If he accepts that the economy is improving, he is willing to posit that the economy is improving when making cognitively serious inferences and engaging in cognitively serious actions. I suggest that reliabilists are concerned with belief$_c$, whereas evidentialists are concerned with acceptance.

Belief$_c$ is passive. One finds oneself feeling, taking it, having a hunch that things are so. No doubt beliefs$_c$ have causes, but the causal trajectories are relatively opaque. An agent is quite capable of believing$_c$ that it is raining without filtering the belief content through a causal theory of perception. She simply feels it to be, or takes it to be so. Moreover, beliefs$_c$ are involuntary. They are not matters of choice; nor are they under the believer's direct control. This is the force of the claim that there is no will to believe. Nor is it possible to disbelieve$_c$ at will. If Joe finds himself with an unwelcome belief$_c$, he may be able to distract himself from considering the issue at all, but he cannot simply decide to stop feeling that things are as the belief content represents them to be. To be sure, he can, as Pascal recommends (2005:212-214), attempt to put himself in a position where his beliefs$_c$ will change for the better. Spending more time with religious adherents and less time with atheists and agnostics might eventually lead someone to believe$_c$ – that is, to feel – that God exists. And if such a person thought that believing$_c$ that God exists would be valuable, he could try to put himself in a position to be receptive to that belief$_c$. But his approach would be indirect. He could not get there directly by choosing to believe$_c$. Similarly, if a person thought it would be a good thing to believe$_c$ that the economy is improving, it might be a good idea to spend time with people like Ben, hoping to be influenced by their opinions. Arguably, such strategies are worth trying if we want to modify our beliefs. Still, whether one believes$_c$ depends on a host of

factors that are not available to the agent's scrutiny. In believing$_c$ the agent's stance toward a belief content is spectatorial.

To say that the agent is passive with respect to them is not to say that belief$_c$ contents are simple. A good deal of education, training, or experience may be required for someone to be capable of believing$_c$ that things are a certain way. Mary's medical education equips her with the capacity to believe$_c$ that the tests reveal that her patient's white blood count is low. Lissa's musical training equips her with the capacity to believe$_c$ that *Opus 17* is a fugue. Jeremy's teaching experience equips him with the capacity to believe$_c$ that his students are clueless about the law of cosines. Those of us without the appropriate backgrounds lack the capacity to believe$_c$ these things. We lack the requisite epistemic resources. We could, of course undertake the training that will enable us to believe$_c$ them, but it may be a long and arduous trek.

Beliefs$_c$ can be sophisticated or unsophisticated, simple or complex, accurate or distorted. Beliefs$_c$ can be assessed as true or false, accurate or inaccurate, correct or incorrect. When, for example, a student says she has a feeling that she aced the exam, she expresses a belief$_c$ that she aced it. Agents may be assessed as fortunate or unfortunate for having specific beliefs$_c$. But because beliefs$_c$ are involuntary, agents ought not be praised or blamed for having them. Moreover, being the way the world presents itself to an agent, a belief$_c$ is not the sort of thing that has or needs a reason. If the agent is properly attuned, her beliefs$_c$ are reliable. Often this is the case. To some extent reliability is a product of evolution. To some extent it is a product of education, experience, and even luck. Often enough the way the world presents itself to an agent is accurate. If their accuracy is not a fluke, the beliefs$_c$ that emerge are reliable.

To accept that *p* is to be willing to use *p* as a basis for cognitively serious inference and action. That is, to accept *p* is to *treat p* as being the case. Some of those inferences are assertoric.

Rather than taking the form *if p then q*, they have the form *since p, it follows that q.* If Mary accepts that the patient's white blood count is low, she is prepared to use 'His white blood count is low' in inferences about what might be ailing him and to use it in inferences to rule out other diagnoses; she is willing to use it as a reason to prescribe certain medications and refrain from prescribing others. Acceptance is voluntary. One decides that a contention is worthy to serve as a basis for inference or action. If she decides to accept it, she deems it choice worthy. The decision can be, and arguably should be, based on reasons. Mary might decide to accept the result of the blood test, or she might decide that more evidence is required. If she decides to accept it, she ought to be able to defend her acceptance. She should be able to answer 'What makes your conclusion acceptable?' where the answer is keyed to publicly available evidence and publicly shared standards. As a result, both her choice, and her decision to make that choice open her up to praise and blame. In accepting that *p*, one adopts an agential stance. One is, and may be held, responsible for what one accepts.

Elsewhere I have argued that acceptance rather than belief or belief$_c$ is the central epistemic attitude. I will not rehearse those arguments here. Rather I will emphasize only that acceptance and belief$_c$ have different epistemic profiles. Although we often accept what we believe$_c$ and believe$_c$ what we accept, the two can diverge. Despite having studied hereptology, Megan cannot help but believe$_c$ that garter snakes are dangerous. She accepts that they are not – she even lets her children handle them – but despite her best efforts they still strike her as dangerous. Arguably, many cases of irrational belief qualify as irrational precisely because the agent who harbors them believes$_c$ but does not or ought not accept them. Acceptance without belief is common in the sciences and, I have argued, in philosophy (see Elgin 2025). An epistemic agent may be willing to use a contention in her inferences and actions without ever quite believing$_c$ it. To believe$_c$ that *p* is to feel that the world is such that *p*. It is then to feel that *p* is true. If the agent is rational, this seems

to require feeling that *p* will never justifiably be rejected. Many philosophers and scientists, who accept their favored accounts without reservation, would still be reluctant to draw the inference that those accounts will never be rejected for good reason. In more ordinary cases, we often accept a contention thinking that it is good enough to perform the function we want performed in particular circumstances, but recognize that it is not, or that we do not consider it, strictly true.

Kornblith notes that reliabilists typically refrain from offering epistemic advice, while evidentialists offer advice freely (2019:17). This insight supports my contention that reliabilists and evidentialists are engaged in different epistemic projects. The reason reliabilists do not offer advice about what to believe or how to arrive at tenable beliefs is that on their view, advice can get no purchase. If beliefs$_c$ just happen to epistemic agents – if, due to epistemically inaccessible causes, people just find themselves feeling that *p* is so – advice would be idle. 'Try to form reliable beliefs', like 'buy low, sell high', is not operationalizable. It tells us where we want to end up, but gives no clue how to get there. Evidentialists can offer plenty of pointers. They set standards of acceptability and criteria that have to be satisfied for those standards to be met. Accepting is something agents do, and we are open to inputs about what we should do to achieve our epistemic aims.

Kornblith is skeptical of evidentialist advice, as it seems to require adopting a first personal perspective. That perspective is not purely individualistic, however, since others can simulate it and thus bring to bear public standards that enable them to criticize, amplify, or extend the conclusions any individual draws. Taken literally, advice that starts with 'If I were you' is worthless. If I were you, I would be in exactly your predicament, and would be just as bewildered as you are. That's no help. But the phrase, as it is actually used, means something more like 'If I were in your predicament, but had my epistemic resources, here's what I would do'. That makes

sense and is sometimes helpful. It provides another (quasi-first-personal) way to look at the predicament. It is thus a strategy an evidentialist would favor, since it is a way of bringing to bear additional evidential resources to resolve a dilemma. It is by no means fail-safe. The resources may be inadequate, irrelevant, or biased. But the opportunity to draw on additional resources, taken from alien perspectives, is an epistemic asset.

Detente?

Where does this leave us? Given my full throated defense of evidentialism's account of Mary's and Ben's epistemic successes, it might seem that I have provided reasons to dismiss Kornblith's process reliabilism. Not so. Evidentialism, as I have characterized it, affords no good answers, and often no answers at all, to the questions Kornblith wants to ask. What psychological processes generate and sustain an epistemic agent's belief that *p*? Are these processes reliable? If, as is evidently the case, human beings are susceptible to a variety of heuristics and biases that undermine the reliability of the inferential processes that we deploy, how is it that many of our beliefs are accurate? That is, how do we compensate for or overcome or avoid subliminal biases? How should epistemology account for tenable beliefs that are not products of and often are not even susceptible of reflection? These are legitimate, important epistemological questions. Kornblith's naturalistic process reliabilism affords plausible answers.

Evidentialists might reply that they are concerned exclusively with the epistemic status of beliefs that are answerable to evidence. That gives them a good reason to ignore the phenomena that Kornblith focuses on. This may be true. But it is an argument for a division of labor, not an argument for dismissing process reliabilism. If we recognize that the question 'What secures the tenability of your belief? ' can be disambiguated so that under one interpretation it means 'What are

the underlying sources of your belief and how reliable are they?' and under another it means 'What intersubjectively accessible and acceptable reasons are there to support your belief?' we can see how both reliabilism and evidentialism contribute to epistemology. I suggest then that we recognize that both projects are worth doing, and that neither discredits nor undermines the other. Once we realize that they are attempting to answer different questions, the recognition that they provide different answers is neither surprising nor alarming.

References

Cohen, L. Jonathan (1992). *An Essay on Belief and Acceptance*. Oxford: Clarendon Press.

Elgin, Catherine Z. (2025), *Epistemic Ecology.* Cambridge MA: MIT Press.

Elgin, Catherine Z. (2017). *True Enough.* Cambridge MA: MIT Press.

Gigerenzer, Gerd (2000). *Adaptive Thinking.* Oxford: Oxford University Press.

Gilovich, Thomas, Dale Griffin, Daniel Kahneman (2002). *Heuristics and Biases*. Cambridge: Cambridge University Press.

Kahneman, Daniel (2011). *Thinking Fast and Slow.* New York: Farrar, Straus, and Giroux.

Kahneman, Daniel, Paul Slovic, and Amos Tversky (1982). *Judgment Under Uncertainty: Heuristics and Biases,* Cambridge: Cambridge University Press.

Kornblith, Hilary (2019). *Second Thoughts and the Epistemological Enterprise.* Cambridge: Cambridge University Press.

Lapidow, Elizabeth and Caren M. Walker (2020). 'The Search for Invariance: Repeated Positive Testing Serves the Goals of Causal Learning' *Language and Concept Acquisition From Infancy Through Childhood* ed. Jane B. Childers. Springer: 197-219.

Lewis, David. (1983). 'Postscripts to "Truth in Fiction"' *Philosophical Papers, vol. 1.* Oxford: Oxford University Press, pp. 276-280.

Pascal, Blaise (2005). *Pensées.* Indianapolis: Hackett Publishing Co.

Scanlon, T. M. (1998). *What We Owe To Each Other.* Cambridge MA: Harvard University Press.