

9. Can Selection Effects on Experience Influence its Rational Role?*

Susanna Siegel

In a chilling experiment conducted by Keith Payne, when primed with pictures of Blacks, participants more often misclassify a tool (pliers, wrench, or a drill) as a gun, compared with subjects who have been primed with pictures of Whites.¹ From this result, we know that whatever psychological state the prime puts the subjects in, it influences their answers on the classification task. But how is their visual experience affected? How do the pliers look to the subjects, when they see them?

The experiment does not speak directly to this question. There are several (combinable) options. A first option is *cognitive penetration*: subjects see the pliers, but nonetheless the pliers look to them like a gun, whereas to subjects given the Whites prime the pliers have their normal appearance. A different option is that the prime generates a *selection effect*, which can influence either the content of experience, or the role of experience. In *selection of features for experience*, subjects see the pliers, but due to the prime, they attend only to those features that pliers share with guns, such as being metallic and shiny. Here they have neither a pliers-experience nor a gun-experience, but rather an experience with more impoverished contents, and they end up jumping to the conclusion that it is a gun.

A different selection effect is *anti-selection of experience for uptake*. Here, when participants see the pliers, they look just like pliers, but due to the prime, this experience is ignored when subjects form the belief that the stimulus was a gun. Their pliers-experience is *anti-selected* for uptake into the usual kind of belief-formation process.² When subjects form the belief that the stimulus is a gun, their perceptual experience does not play its usual role in the formation

* For discussion, thanks to John Bennett, Ned Block, David Braun, Alex Byrne, Bill Child, Brendan Dill, Anya Farennikova, Jane Friedman, Paul Marcucilli, Alyssa Ney, Bernhard Nickel, Keith Payne, James Pryor, Nicholas Silins, participants in a seminar at Harvard, and audiences at the first Sheffield workshop on implicit bias, MIT, NYU, Oxford, Rochester, and Union College, where some of this material was presented. Special thanks to Eric Mandelbaum, Matthew McGrath, Declan Smithies, Scott Sturgeon, and Sebastian Watzl for extensive discussion.

¹ Payne 2001. The subjects are non-Black American college students, though similar effects are found for Black American college students.

² A structurally similar form of anti-selection would stop the process one step later: subjects would form a pliers-belief on the basis of the pliers-experience, but then the pliers-belief would be inferentially quarantined, in something like the way compartmentalized beliefs are. Could a mental state count as a belief, even if it was quarantined from all inferences? I leave this difficult question aside.

of this belief. Usually, the contents of experience are closely related to the contents of perceptual belief, but not in this scenario, where a *gun*-belief is formed in response to a *pliers*-experience.

Could cognitive penetration, selection for experience, or anti-selection for uptake explain the role of the Black prime in causing classification errors in the pliers/gun experiment?³ If so, either the resulting perceptual experiences or the resulting perceptual beliefs end up apparently confirming what the subjects expect or fear. These forms of cognitive penetration and selection effects challenge folk theories that construe perception as primarily a means of taking in information about what's currently in the immediate environment. According to the folk theory, perception is largely the result of initial inputs from the world to the senses, and any intervening perceptual processing disallows it from being influenced by the rest of one's psychological states, such as what one believes, wants, or fears. Integration with those other psychological states comes after experience, not before.⁴ In this sense, perception is primarily receptive, and that is why it contrasts so sharply with reasoning. We reason from information that we take ourselves to have already, and sometimes let our reasoning fall prey to our hopes or fears. But perception is different. Perception isn't subject to reasoning at all—neither good reasoning nor defective reasoning. The idea that experiences are primarily receptive is encouraged by the phenomenology of perception, in which we seem simply to confront the scene before us and take things in as they are.⁵

³ Cognitive penetration and selection effects are not the only options for explaining Payne's result, and in fact these hypotheses are better at illustrating types of effects on perception, than they are interpreting this particular experiment. An unpublished follow-up experiment by Payne probing the confidence levels of participants questions the assumption that the participants' keyboard response reflects what they at that moment take the target shown to be. The follow-up study replicated the main result and then asked participants to rate how confident they were in their responses, and found that confidence levels correlated almost exactly with performance: subjects reported little confidence in answers that turned out to be wrong, and high confidence in answers that turned out to be right. Crucially, in trials with stereotype-consistent errors, self-reported confidence levels were very low, and subjects reported feeling that their fingers (on the keyboard) could not keep up with the responses they wanted to make. Here participants systematically found themselves indicating an answer they thought was wrong. So perhaps they experience the pliers as pliers and believe they are pliers, but their finger-pressing behavior is not guided by what they believe.

Anti-selection for uptake, however, may yet end up explaining Payne's central result. The follow-up result allows that the pliers-experience is anti-selected for belief, if the subjects have contradictory beliefs about whether the stimulus is a pair of pliers. If we looked ahead to how participants remember what they saw, and found that their memory correlated with their behavior rather than with what they would disown, if asked, at the moment of the experiment, one might find that their behavior was only the first manifestation of exactly the sorts of dispositions we'd take as evidence for believing that a gun was presented.

⁴ Something like this part of the folk theory is found in Fodor's idea that perceptual processing leading up to belief is dedicated exclusively to extracting information from those initial inputs, though his theory is directed at perceptual processes that may not be the exclusive determinants of conscious experience.

⁵ A number of writers have tried to characterize this aspect of perception, including Sturgeon 2000 on scene-immediacy, Martin 2002 on conviction, and Chalmers 2006 on edenic content.

The folk theory of perception goes with a folk-epistemology of perception, which in turn assumes a distinction between perceptual experience and perceptual belief.⁶ Roughly, perceptual beliefs are beliefs about what you're experiencing, which you form chiefly because you are having the experience. According to folk epistemology, in central cases of perceptual belief, where perceivers are not aware of any reason to discount their experiences, such experiences are the main person-level psychological state that determines the content of perceptual beliefs; and even though those experiences cannot be the product of rational (or irrational) person-level processes, nonetheless they can provide rational support for certain beliefs about what we see, as when we believe that there is mustard in the fridge upon seeing the mustard when we open the fridge door. Putting these two epistemic features together, perception can provide rational support without having any itself, and this winning combination makes it well suited to stop regresses of justification.

Most theories of the rational role of experience assume that experience conforms to the folk theory of perception, and are directed at cases in which experiences play the standard role in perceptual belief. The possibilities of cognitive penetration and selection effects on perception create the need to explain the rational role of experiences that don't conform to the folk theory of perception or don't play the standard role in perceptual belief.

We can see how the challenge to the folk theories unfolds by considering cases in which perception becomes an instrument of confirmation bias. Confirmation bias may be most familiar as a feature of suboptimal searches for evidence in hypothesis-testing, or imbalanced consideration of reasons in connection with forming or maintaining beliefs or fears. In these forms of confirmation bias, one searches for confirming but not for disconfirming evidence for a hypothesis, or one considers reasons that support what one suspects or fears, without considering reasons against it. But a different form of confirmation bias could be found at the level of perception. Here, a subject ends up perceiving, either at the level of experience or belief, what she already suspects or fears to be the case, and is systematically prevented from perceiving counter-evidence, leaving her with apparent confirmation from her perception of her fear or suspicion. This phenomenon can occur in the form of cognitive penetration, and the epistemic consequences in that case are increasingly widely discussed.⁷ Unlike cognitive penetration, however, the structure of selection effects on perception has barely been examined. We can distinguish between selection effects *per se*, and selection effects that unfold in a way that leads to confirmation bias. In experiences of objects, some factors

⁶ For complex forms of dissent to the distinction between experience and belief, see Glüer 2009 and Byrne 2009.

⁷ Discussion of epistemological consequences of cognitive penetrability of experience can be found in Huemer (2013), Lyons 2011, Siegel (2011 and 2013a), McGrath (2013a and 2013b). For discussion of the epistemological implications of cognitive penetration on perceptual processing generally (not focused exclusively on experience), see Fodor 1983, Churchland 1988, and Raftopolous 2009.

always determine which objects are being experienced and which features of those objects you focus on—such as where you happen to be looking and why.⁸ What you believe, want, or fear can easily influence what you are looking at. Likewise, experiences could be anti-selected for uptake due to all sorts of different factors, ranging from load-induced inattentiveness, to anti-selection that is at some level explained by its role in maintaining a prior fear, hope, or expectation. Here, the selection effects of interest are those that are generated by prior beliefs (including expectations), hopes, or fears, and that operate as a means of psychologically confirming or strengthening those prior states. These are selection effects that can lead to confirmation bias at the level of perception.

My aim is to examine the rational impact of two varieties of selection effects of this kind. First, I'll present a case in which the same kind of objects are repeatedly selected for experience, where the selection leads to confirmation bias. This type of selection effect is a selection of objects rather than features for experience, but I'll argue that analogous cases involving selection of features warrant the same treatment. These cases pull us in the direction of allowing that experiences you *don't* have can affect the rational status of the experiences you *do* have. If patterns of experience can be epistemically significant in this way, then ultimately what needs explaining is how the rational role of those experiences can be influenced by their pattern, and by their relationship to the prior psychological states that help select which objects or features are experienced and which are not.

I defend the position that patterns of experiences can be *epistemically downgraded* by certain selection effects. Experiences are epistemically downgraded, if the rational support they offer for believing their contents are reduced or eliminated, relative to a baseline that I'll assume is the normal case. Normally, seeing a car drive by gives you excellent reason to believe that a car just drove by, and I'll assume that the experience you have in seeing the car plays a major role providing this rational support.⁹ Since the perceivers in the cases I consider are unaware of the selection effects and their causal role in confirmation bias, one might think that only externalist theories of rational belief (such as reliabilism) could explain why certain selection effects lead to epistemic downgrade. But I'll sketch an internalist explanation of epistemic downgrade in these cases.

Besides the selection of objects or features for experience, the second variety of selection effect I consider is anti-selection of experience for uptake. Anti-selection for uptake creates fragmentation in the mind that is similar in some

⁸ Here I'll assume that if you attend to a feature F of an object, then F will characterize the way the object looks to you. This assumption can be taken as a restriction on the type of feature-attention at issue. It rules out unconscious attention to features, for example, as well as "feature illusion" in which you attend to an object's redness (for example) but experience it as blue instead of red.

⁹ For discussion of the different ways in which experiences could provide rational support for beliefs, see Siegel and Silins 2013.

ways to standard cases of compartmentalized beliefs.¹⁰ Here the challenge is to explain the rational powers of psychological states that are isolated from one another. In cases of compartmentalized belief, it is controversial how, if at all, beliefs in separate compartments bear rationally on one another. Similarly, in cases of anti-selection for uptake, an account is needed of which other mental states, if any, the anti-selected experiences rationally bear on. I suggest that anti-selected experiences can retain the same rational bearing on mental states that lie outside its “compartment” as it would have in a more unified, less compartmentalized mind, and that this conclusion seems especially plausible in the case where anti-selection for uptake facilitates confirmation bias.

The existence of these two varieties of confirmation-bias-inducing selection effects is much less controversial than cognitive penetration. Selection effects are compatible with the idea that all of the psychological processing leading up to perceptual experience is modular. Many experimental results that provide suggestive but not decisive evidence for cognitive penetration are explicable in principle by selection effects.¹¹ But here I set aside controversies about the psychological reality or prevalence of confirmation-bias-inducing selection effects and cognitive penetration. My methodology is to define hypothetical versions of each kind of confirmation bias, discuss its epistemological consequences, and leave the empirical inquiry into whether any of these effects actually happen (and if so, how much) for a separate discussion. Of course, these facts are not epistemologically irrelevant, since we have a special interest in our own epistemic situation, and for that we need to know the empirical facts. But the epistemological significance of those facts can be pursued in a relatively aprioristic way. Once we have answers to the empirical questions about which types of effects on perception we are subject to, we can see what different epistemological theories tell us about our actual epistemic situation. In addition, distinguishing different types of effects on perceptual experience and perceptual belief may help pin down competing interpretations of some experimental results.

The discussion proceeds as follows. In section 1, I introduce the main epistemological problem that arises from confirmation-bias-inducing selection of objects or features for experience. In sections 2 and 3, I sketch an internalist way to respond to the problem. I discuss anti-selection for uptake in section 4, and its significance for the epistemology of mutually unintegrated mental states in section 5. Throughout, I use “experience” as shorthand for “perceptual experience,” though I focus on visual experiences.¹²

¹⁰ On compartmentalized belief, see Lewis 1982.

¹¹ Suggestive but indecisive results include Barrick 2002, Balci and Dunning 2006, Hansen 2008, Levin and Banaji 2006, Proffitt 2006, among many others.

¹² In my terminology, perceptual and visual experiences are phenomenal states, and so can be had in cases of hallucination or illusion, as well as veridical perception.

1. SELECTION FOR EXPERIENCE AS A MEANS OF CONFIRMATION BIAS

Selection for experience should be understood to include either (or both) the selection of features and the selection of objects. In Payne's pliers/gun experiment, the prime pretty clearly isn't affecting whether one sees the pliers, but it may be influencing how the pliers look when you see them, by influencing which features one attends to. But in general, prior states such as beliefs, desires, or fears could influence not only which features one experiences, but also which objects one experiences.

Taken by itself, selection for experience is unremarkable. But when selection for experience operates as a means of confirmation bias, so that one ends up experiencing (roughly speaking) what one antecedently fears or wants or believes, an epistemological problem arises. For instance, suppose you want it to be true that in general, red things are square. Call this the red-square desire. Unbeknownst to you, your red-square desire affects how you distribute your visual attention when looking at a crowded display, to the extent that you don't see other red shapes. Although the display includes shapes of many colors, among the red shapes your attention is drawn only to the squares. You see green triangles and orange circles, but overlook the red shapes that aren't square. A similar selection effect could also hypothetically be induced by a prior belief in the generalization that red things are square. The core of the problem is to assess what kind of rational support the red-square experiences provide for the generalization that all the red things in the display are square.

How might such distributions of a subject's perceptual experience come about? Other red shapes besides squares could be unconsciously perceived, without representations of them ever becoming conscious. Alternatively, a prior state such as a red-square desire or the belief that in general red things are square might direct attention only to those places where red squares are likely to be found, even if the subject isn't aware that these places are likely to contain red squares.

These examples are unrealistic, but the unrealistic part is limited to the content of the states involved. A more realistic type of example involves stereotypes or prejudices that operate implicitly, and impact behavior in evaluative contexts. For instance, often evaluators assessing candidates for hire harbor implicit negative views of "outgroups," and focus disproportionately on the negative features of outgroup members when evaluating them. The negative features are not always (or even very often) hallucinated, but a fuller picture of the candidate would show that those features are counterbalanced by positive features. One of the psychological structures that could underlie this behavior is the type of selection effect described here, where the positive features are simply not consciously taken in to begin with.¹³ This manifes-

¹³ Another psychological structure that could underlie the same behavior is that the positive features are perceived, but given less weight in assessment (Steinpries 1999, Goldin and Rouse

tation of prejudice is more realistic than our red-square example, but our example lets us isolate one of the epistemic problems raised by this type of phenomenon.

What kind of rational support do the red-square experiences provide for the generalization that all the red things in the display are square? On the one hand, you see the red squares in the display perfectly well, without any illusion or hallucination—just as a hiring committee may register the negative features of outgroup candidates without any distortion. On the other hand, the idea that these experiences rationally support the red-square generalization can seem suspicious—just as suspicious as the hiring committee’s judgment might seem that all of the outgroup candidates are relatively weak, when this judgment is made on the basis of their selective focus on those candidates’s weaknesses. This case raises the possibility that facts about what experiences you don’t have can influence the rational role of experiences that you do have.

The problem can be sharpened by considering a series of experiences. Imagine a psychological experiment in which you’re given a visually crowded display, such as the kind found in the game *Where’s Waldo?*. The experimenter gives you a series of three search tasks to find all the targets belonging to a specified type within a challenging time limit.¹⁴ First the experimenter tells you to find all the sad faces in the display as quickly as you can, and to let her know when you think you’ve found them all, giving you a short time limit. After finding several sad faces, the time in between searching and finding sad faces lengthens, until you cease to find any more. When you reach the time limit, you are not completely confident that you have found them all, but you report your results to the experimenter. She says: “Well done! You found them all.”

Your second task is the same, except with happy faces. Encouraged by your success with the sad faces, you search for the happy faces, and once you stop finding any more, you report your result before the time limit is up, feeling more confident this time. Once again the experimenter tells you that you found all the targets.

Your third task is to find the red items in the display. Unbeknownst to you, a prior hope influences how your attention is distributed—the hope that in general, red things are square. Buoyed by your success in the previous search tasks, you report confidently that you have found all the red things. Because of how you distribute your attention, you miss all the red things that weren’t

2000). We might see this phenomenon as a temporary shift in the criterion for what counts as a strong candidate. A similar hypothesis about Payne’s misclassification result would be that when shown the Black prime, subjects jump to the conclusion that the metallic shiny object is a gun, effectively lowering their criterion at the level of judgment for what will count as a gun. Standard shifting could also occur in the transition from subpersonal perceptual processing to experience, if subjects formed a gun-experience in response to the combination of initial sensory inputs and the Black prime. That would be a case of cognitive penetration.

¹⁴ Thanks to Eric Mandelbaum for inventing this experiment.

square. You noticed that all the red things that you found in the display were square, and formed the belief that in this display, *all* the red things are square.

We can call the scenario in the third stage of the experiment the perverse scenario. In the perverse scenario, the prior red-square hope prevents you from consciously perceiving non-square red things. The perverse scenario can be contrasted with a bad-luck scenario. In a bad-luck scenario, you don't antecedently hope that in general, red things are square and you aren't subject to any other systematic red-square related bias in your dispositions to attend, but just by chance, you fail to notice the non-square red things in the third stage of the experiment.

The sharpened question concerns the rational support provided by the subject's experiences for the universal generalization that all the red things in the display are square:

Distribution Question: Do your red-square experiences provide just as much rational support for believing that all the red things in the display are square, in the perverse and the bad-luck scenarios?

If the answer is Yes, then the fact that the red-square experiences in the perverse scenario came about via the red-square hope does not affect how much rational support those experiences provide, and there is no other relevant difference between the two scenarios.

Some theories of rational belief may ultimately embrace the Yes answer, on the grounds that the subject is unaware of the role of the red-square hope in selecting red squares for experience. For instance, many internalist theories of justification deem etiological features of experience irrelevant to its rational role, so long as those features are forgotten, or were never within the subject's awareness, or are otherwise inaccessible to the subject at the time of the experience.¹⁵ From this observation, one might conclude that only externalist theories of justification, such as reliabilism, can ground the No answer. But this conclusion misses an internalist option that grounds the No answer in the relations between the experiences and the psychological states that give rise to it—even if the subject is unaware of those relations, and regardless of any considerations about reliability. This option locates some of the features that determine the rational status of experiences in places that are external to the subject's awareness, but still internal to the subject's mind.

To bring the internalist option into focus, we need to distinguish between two ways in which single-square experiences could provide rational support for the universal generalization that all the red things in the display are square.

Incremental Support: Each red-square experience, at the time that you have it, provides incremental rational support for the universal generalization that all the red things in the display are square.

¹⁵ See the discussion of the problem of forgotten evidence in Feldman and Conee 2001.

Overall Support: The red-square experiences, once you have all of them that you're going to have, provide rational support for the universal generalization that all the red things in the display are square.

From now on I'll call the generalization "the universal generalization" for short. When the Distribution question asks whether the red-square experiences provide rational support for the universal generalization, we could interpret it as asking about either incremental support or overall support.

There is also a different dimension along which the Distribution question can be made more specific. The original question asks about the rational role of the red-square experiences in providing rational support for the universal generalization, without specifying whether the experiences do this all by themselves, or together with other factors that are present in both scenarios. Given your success on the search task (for faces) that precede the search for red things, such a factor might seem to be present in the form of a well-founded belief, based on the experimenter's feedback, that you are good at these search tasks, or that you are not systematically excluding some of the search targets.

To take account of the fact that the red-square experiences might provide rational support of either kind for the universal generalization, in combination with some additional factor X , we can reformulate the distribution question to focus on whether the two scenarios differ in whether they contain any values for X , allowing that it might be a null factor. In case the additional factor might be different for incremental support and overall support we can use different subscripts for X in the two questions:

Incremental support question: Is there a factor X_I present in the perverse and bad-luck scenarios, such that the combination of each red-square experience with X_I provides just as much incremental support for the universal generalization in both scenarios?

Overall support question: Is there a factor X_O present in the perverse and bad-luck scenarios, such that the combination of the red-square experiences with X_O provides just as much overall support for the universal generalization in both scenarios?

The incremental support question faces a controversy over whether, if a proposition $\sim e$ is evidence for a hypothesis $\sim h$, then e is evidence for h . For instance, let h be the universal generalization, and let e be the proposition that *location L does not contain a red non-square*. Given that the proposition ($\sim e$) *location L in the display contains a red non-square* provides evidence for the proposition ($\sim h$) *it is not the case that all the red things in the display are square*, would the proposition (e) that *location L does not contain a red non-square* provide evidence for the universal generalization? One way for e to be true is if location L contains a red square. Given that the discovery of red non-squares would disconfirm the universal generalization, does the discovery of red squares add rational support to it? Kaplan (1996) argues that it does, Popper (1959) argues that it doesn't. But even if Kaplan's side is right, and

each red-square experience incrementally supports the generalization in equal amounts in both scenarios, that support might be insignificantly small. If you weren't in New York City last night when a murder occurred there, then you didn't commit the murder. But suppose you were in New York City last night. On Kaplan's view, the fact that you were there provides some evidence for the hypothesis that you committed the murder, but (as he emphasizes) the amount of evidence it provides is so little that it would need to be supplemented with lots of other evidence to provide any significant support for this hypothesis. Just as what's relevant in the murder example is a bigger collection of evidence, so too in the red square case, the version of the distribution question we should be asking is a version about a significant amount of rational support, not just any amount. Since if any red-square experiences provide any significant amount of rational support for the universal generalization, all of the red-square experiences taken together do, I'll interpret the distribution question as asking about overall support from now on. In this context, "rational support from red-square experiences for the generalization" is overall support from those experiences for the generalization. The official Distribution question is the question about overall support.

Official Distribution question = Overall support question: Is there a factor X_O present in the perverse and bad-luck scenarios, such that the combination of the red-square experiences with X_O provides just as much overall support for the universal generalization in both scenarios?

My strategy for introducing the internalist grounding the No answer to this question is to argue first that the No answer is compatible with *evidentialism*, the thesis that the rational beliefs are exactly those that fit the evidence, where a subject's evidence supervenes on her mental states.¹⁶ Evidentialism may seem to be at odds with the No answer. But it isn't. By seeing how evidentialism can be combined with answering No to this question, an internalist grounding of the No answers comes to light.

Why Evidentialism Seems Forced Into Answering Yes

If evidentialism is true, then when we ask whether you have rational support for P, what's relevant is just the evidence for P you have. The fact that there's some other evidence that bears on P, but which you don't have, isn't relevant.¹⁷ There may be norms governing the search for evidence—how to do it, where to look, when it is okay to stop, etc. But according to standard versions of evidentialism, these norms are different from norms governing what is rational to believe. How much evidence you have, and the process by which you got it, doesn't affect what's reasonable to believe in light of that evidence.

¹⁶ Feldman and Conee 2001.

¹⁷ For discussion and defense of this idea, see Foley and Fumerton 1982, Kelly 2002, Feldman 2003.

The same idea applies straightforwardly to experiences, on the assumption that experiences are or provide evidence. If the visual search process had gone differently, for instance because there was no selection effect by the red-square hope of the sort described, then your visual search probably wouldn't have terminated without your finding some other red shapes. But according to the evidentialist idea as applied to experiences, when we ask whether the red-square experiences provide incremental rational support for the distribution propositions, we're asking whether those propositions are supported by the experiences you have. The fact that you could have had some other experiences bearing on them is not relevant.

How would the No answers to the Distribution Question come to be at odds with evidentialism? Evidentialism entails that two subjects with exactly the same evidence have the same amount of rational support for the same propositions. The conflict arises when we combine No and evidentialism with an assumption about the evidence present in both scenarios, resulting in the following inconsistent triad:

- (a) In each scenario, the amount of rational support from red-square experiences you have for the generalization is different.
(= *the 'No' answer*)
- (b) If your evidence for the generalization is the same both times, then your rational support for the generalization is the same both times.
(= *consequence of Evidentialism*)
- (c) Your evidence for the generalization is the same in both scenarios.
(= *hypothesis about experiences and search process*)

What grounds are there for (c)—that your evidence is the same in both scenarios? One might argue: the evidence for the generalization in both cases consists in (i) your experiences, plus (ii) your evidence about the search process, and nothing else. These two factors seem evidentially on par in both scenarios:

Same experience → same evidence: Since the red-square experiences in the two scenarios are the same, they contribute the same evidence to the generalization belief in the two scenarios.

Rational search termination: The subject has no reason to think that the visual search was terminated too soon, given her success on the previous tasks.

Since evidentialism is at odds with answering No only if claim (c) is true, evidentialism can be reconciled with No by denying (c).¹⁸ But on what grounds can (c) plausibly be denied?

¹⁸ An intermediate response splits the difference between Yes and No, by rejecting thesis (a), but holding that nonetheless, beliefs in the distribution propositions based on the red-square experiences would be epistemically ill-founded. According to the split-the-difference position, in the perverse scenario, the red-square experiences provide propositional justification for the generalization, but do not support the formation of doxastically justified belief in the generalization.

2. HOW EVIDENTIALISM CAN ANSWER NO

Writ large, what seems to force evidentialism into answering Yes is that the red-square experiences seem to provide a rational route to the generalization, when combined with the (apparently justified) belief that you have looked long enough to reach an answer to the question defining the task. In the red square example, the task is to find all the red things, and the supporting belief is based on the experimenter telling you that you succeeded on the previous, similar search tasks. In the hiring case, there is also a search task—to find features of the candidates that are relevant to deciding whether they are a strong candidate. And committee members might have an analogous, apparently justified supporting belief, based on their past success in identifying strong candidates, that they have looked long enough to answer the question.

To see why evidentialism is not forced into answering Yes, we need to look more closely at the structure of the red-square experiences that figure in this reasoning. I'm going to make a case for the existence of a temporally extended, *compound experience*. With the compound experience in the picture, the reasoning that seems to push the evidentialism toward Yes can be defused in two steps. First, the compound experience is the only one that could reasonably be thought to combine with the background beliefs just described, in order to support answering Yes. But second, it is open to the evidentialist to hold that this compound experience is downgraded by the selection effect. This position gives them a way (and apparently the most plausible way) to deny claim (c), even on the assumption that the background beliefs that one has looked long enough are justified.

Introducing the evidentialist position involves examining compound experiences (2.1), and then using that notion to defuse the reasoning that seems to force evidentialism to answer Yes (2.2). In section 3, I sketch an argument for the claim that the selection effects epistemically downgrade the compound experiences in each case. Since the argument provides general reason to think the compound experience is downgraded, a fortiori it provides reason for evidentialists to reject claim (c) and answer No to the Distribution question.

2.1. *Individual vs Compound Experiences*

When you look for red things in the display, you have an experience of searching. When you find the red squares, and then keep looking for red things but cease to find any other red shapes, that too is an experience—an experience of searching and not finding any more targets.

This option is at odds with a plausible principle linking propositional justification to doxastic justification. According to the linking principle, if a factor X provides propositional justification for a proposition, then if you formed a belief on the basis of X, the belief would not thereby be doxastically *unjustified*. Without this linking principle, the resulting notion of propositional justification would be watered down a lot, leaving it obscure in what sense X provided S with propositional justification for P in the first place, given that she can't even use X to form a well-founded belief in P.

These experiences differ phenomenally from the experiences of a red square that punctuates the experience of searching. Those individual experiences can occur in isolation, without occurring in the course of a search for red things. We can therefore distinguish individual red-square experiences (experiences of single red squares) from more complex, temporally extended experiences that involve a sequence of individual red-square experiences.

For any putative temporally extended experience, we can ask what makes it stand out as a natural subexperience in the subject's stream of consciousness—one that is both temporally extended, yet stops short of the entire stream of consciousness, or the entire "ray" of experience that extends forward in time given a temporal starting point. There are doubtless multiple temporally extended experiences that are salient aspects of the subject's mental life. Call the *compound experience* the sequence of the single-square experiences, together with the experience of having looked for red things in the display after finding the red squares, and not finding any more red shapes. By hypothesis, the compound experience is temporally extended, starting with the sequence of red-square experiences, and continuing through the experience of looking for red things and not finding any. We can define the compound experience further by answering specific questions about its duration, its relation to memory, its content, and its phenomenal character.

Regarding duration: when does the compound experience end? It is natural to think that it ends when the subject stops looking at the display. But what if by the time the subject stops looking at the display, she has forgotten that she saw some (or even all) of the red squares she saw when she had the single-square experiences? For some purposes it may be illuminating to focus on a stretch of a subject's experience, regardless of what she remembers. But here, it is useful to focus on experiences that the subject could explicitly endorse as bases for believing all of the contents of the experience. For that purpose, it is useful to build into the nature of compound experience that the subject retains in memory its earlier parts by the time she has the later parts.

Regarding memory: compound experiences satisfy the following assumption. According to the *memory assumption*, at the end of the compound experience, the subject remembers experiencing each red square that she saw. (The memory assumption is compatible with the subject forgetting about some of the red squares over the course of seeing them, so long as she remembers them at the end).

Regarding content: what contents does the compound experience have? Given the memory assumption, it is natural to think that the content of the compound experience includes a conjunction of the contents of all of the individual red-square experiences. Conjoining these contents yields a content of the form:

Conjunctive content: this₁ red thing is square and this₂ red thing is square ... and this this_N red thing is square

... where N is the number of red things the subject sees. (I am ignoring the differences between “this red thing is square” and “this is a red square,” and assuming that the red things in the display look red to the subject). Factoring in the experiences of searching and not finding any more red things plausibly adds another conjunct:

Conjunctive content + : this₁ red thing is square and this₂ red thing is square ... and this this_N red thing is square and no other red thing in the display is square.

If a subject has an experience with this conjunctive content, then her experience conveys something to the effect that the red things she has seen in the display are all the red things there are. And if it does that, then it seems to include in its contents the universal generalization:

Universal generalization: All the red things in the display are square.

The considerations suggest that a subject in the red-square experiment could have a compound experience with universally quantified contents.

Regarding phenomenal character: one might reasonably wonder whether experiences could have universally quantified contents. Could such content answer to any phenomenal aspect of the experience? If not, then either no experience ever has such contents, or if it does, its link to the subject’s mental life is tenuous.

One might think that our experiences can only ever take a stand on how many F ’s there are, by representing that there are at least that many F ’s. Consider the experience of seeing three pens on a cluttered table. Your experience might represent that there are three pens on the table, but not take a stand on whether there are any other pens among the clutter. If there turned out to be more pens underneath the clutter, the experience would not thereby be inaccurate. If experiences of quantity were like this, then experiences could never represent that the F ’s in the relevant domain are all the F ’s there are in that domain. And if they can’t do that, then they can’t represent that all the F ’s in that domain are G .

But some experiences provide natural examples of experiences that represent that the F ’s in the relevant domain are all the F ’s there are in that domain. For example, suppose you saw three eggs lined up in an egg carton, and had an experience that represents not only that three eggs are in the carton, but also that the other spaces in the carton are empty. If there were (by some amazing illusion) four eggs in the egg carton your experience would be inaccurate. It seems straightforwardly part of the phenomenal character that the eggs in the egg carton are exactly three. Going with that, the experience represents that three eggs are all the eggs in the carton. And if it can do that, then it can represent that all the eggs in the carton share some feature (they are brown, or egg-shaped, or are eggs, for instance).¹⁹ I’ll call the phenomenal feature that goes with universally quantified contents ‘phenomenal that’s-allness’.

¹⁹ I’ve spoken as if experiences could represent properties such as being an egg or being a pen, but the same points would hold for egg-shaped or pen-shaped volumes.

By following on the experience of searching and not finding any more red things, the compound experience has phenomenal that's-allness. In this respect, it is more like the egg-carton experience than the cluttered-table experience.

In a context where one has been searching for a target (such as eggs), an egg-cartony experience has a salience structure that is stable with respect to its target. A salience structure for a subject is a structure in which some objects and properties are experienced by S and others, though present, are not experienced by that subject.²⁰ In a stable salience structure, which items are targets is no longer evolving in the experience. Once the salience structure is stable, you may suddenly notice another target. But when that happens, then (assuming the additional target is experienced as such) the phenomenal that's-allness ceases. When the salience structure is still evolving in the context of a search, so is the structure of experienced affordances—roughly, experienced possibilities of action. If you are in the midst of looking for red things in the display, then as yet unsearched areas of the display are presented to you as places where red things might be found. If the salience structure of your experience of seeing the display was still in flux, you would either still be looking around at the display, or else you experience parts of the display you haven't seen yet as possibly containing red things.²¹ Once the phenomenal that's-allness emerges, these affordances are gone.

Putting these observations together, we could say that with respect to the target, the transition into phenomenal that's-allness can be a transition from a kind of experiential flux, in which the salience structure of targets is still evolving, and a kind of experiential fixity, in which the salience structure of targets is stable. This type of experiential flux and fixity could be seen as a way for an experience to treat the question of which experienced objects or features are targets as open or closed.

Treating a question as open or closed can have obvious ethical implications as well as epistemological ones. Abstracting from perceptual experience, it is intuitively possible to prematurely write someone off as unsuitable for a task, just as it is possible to treat it as an open question whether someone is suitable for a task, even when there is already evidence to settle the question one way or another. Social contexts like these can strengthen our understanding of what it is to treat a question as open or closed, by illustrating how these options can make some social relationships into live possibilities while making others socially inaccessible. In this way, in social contexts, transitions between flux and fixity can be mechanisms of social inclusion or exclusion. It is natural to think that these mechanisms sometimes operate in ethically bad

²⁰ Watzl 2010 and 2011 discusses structures of salience, though he defines it in terms of the structure of attention to objects and properties, rather than the structure in which objects and properties are experienced at all.

²¹ More exactly, you would experience the display as containing parts that are so far unseen by you and that possibly contain more red things.

(or good) ways by operating in epistemically bad (good) ways. In section 3, I argue that partly by virtue of its transition from flux to fixity, the compound experience is afflicted by an epistemic shortcoming.

2.2. *Why Evidentialists Don't Have to Answer Yes*

With compound experiences on the table, we can return to the idea that the subject of the experiment seems to be able to combine her knowledge that she succeeded on the first two search tasks with her red-square experience, to get significant rational support for the universal generalization. If she can do that in both the bad-luck and perverse scenario, then by evidentialist lights, claim (c) in the inconsistent triad seems true (the subject has the same evidence in both scenarios for the universal generalization), and so does the answer Yes to the Distribution question (Yes, the subjects in the two scenarios have the same amount of rational support from their experiences for the universal generalization.)

A first attempt to challenge this reasoning would be to deny that the subject's success on the first two search tasks (for faces) gives her reason to think she has succeeded on the third search (for red things). But the subject has no inkling that her red-square hope is impairing her performance on the third tasks, and all three tasks are superficially similar. And the red-square hope has no influence on this belief—its basis is the experimenter's feedback, together with the experience of looking for targets and eventually not finding any more. So this attempt to challenge the reasoning seems weak. I am going to grant to the proponents of Yes that it fails.

A different attempt to challenge the reasoning is to hold that due to the selection effect by the red-square hope, the individual red-square experiences in the perverse scenario do not provide any evidence at all.²² This option entails that the individual red-square experiences would not even rationally support the single-square propositions such as those that could be expressed by *this₁ is a red square*, *this₂ is a red square*, etc., or the *plurality proposition* that the display contains some red squares. That position is arguably too strong, since on the face of it, the subject is excellently positioned to form justified beliefs with these contents: she veridically sees the red squares, and her experiences result from normal processes. We can even suppose that the process of detecting red squares generates few false positives and few false negatives. Typically, when visual experiences with these features are endorsed, the resulting beliefs are justified.

But if individual red-square experiences provide rational support for single red-square propositions, then why can't these individual experiences be combined with the background knowledge about one's success on the previous

²² I set aside the question whether this position rules out that self-ascriptions of red-square experiences made in the usual distinctively first-personal way can be rational.

search tasks to provide reason for the universal generalization? Here we come to the first step in defusing the reasoning. Suppose you had only individual red-square experiences, without the compound experience. More exactly, suppose you experienced several red squares, but did not have the experience of looking for more red squares in the display and not finding any more. The allotted time ran out before you felt you had looked long enough to complete the task. Merely having a series of individual red-square experiences leaves open that you might experience the display as having large areas in which you hadn't yet looked. If you had an experience like that, then your background knowledge that you succeeded on the previous search tasks would not give you much reason, if any, to think you have succeeded in finding all the red things. In the previous tasks, you had analogous compound experiences as well—experiences of looking for targets and eventually finding no more. If you lacked that experience on the third task, your background knowledge of previous successes would not go very far in giving you reason to believe the universal generalization.

These considerations suggest that the type of experience that the background knowledge needs to combine with to support the universal generalization is something like a compound experience, rather than individual red-square experiences. And this brings us to the second step in defusing the reasoning that seems to force the evidentialist's hand. While it seems implausible that individual red-square experiences are epistemically downgraded, a case can be made that the compound experience is epistemically downgraded.

3. WHY THINK COMPOUND EXPERIENCES ARE DOWNGRADED?

Why think the selection effect downgrades the compound experience in perverse scenario? Reliabilists might hope to appeal to reliability considerations.²³ A potential difficulty for reliabilist treatment of the red-square example lies in identifying which processes leading from red-square experiences to the belief in the universal generalization have to be reliable, for the belief to count as justified. This is a version of the Generality Problem (Conee and Feldman 1998). In the perverse scenario, the process generating red-square experiences is a highly reliable process for forming beliefs about where the red *squares* are, a slightly less reliable process for forming beliefs about where the colored shapes are, and an unreliable process for forming beliefs about where the red *shapes* are. Which of these processes (if any) determines whether the belief that the display contains red squares, red shapes, and colored shapes is justified, and what principled way is there to identify the processes that matter?

This problem can be avoided by internalist explanations of the downgrade that start from the idea that an experience or belief can be a conduit of ill-

²³ See Lyons 2011.

foundedness. An experience or belief is a conduit of ill-foundedness, if any subsequent belief that is formed primarily on its basis would thereby be ill-founded. To show that an experience or belief is a conduit of ill-foundedness, it is not necessary to specify which type of unreliable process (if any) makes it ill-founded, and in this way the Generality Problem is avoided.

If an experience is a conduit of ill-foundedness, is it epistemically downgraded? It is plausible to think that being a conduit of ill-foundedness and being epistemically downgraded are linked in this way. If any belief formed on the basis of the experience would thereby be ill-founded, then it is hard to see what is left of the idea that the experience provides a significant amount of justification for it.²⁴ If so, then we can argue that the compound experience is downgraded, by arguing that it is a conduit of ill-foundedness.

Let's call a dependence relation whereby a hope, fear, or prejudice that P leads to a belief or experience with content P an 'irrational dependence relation', if it would make the belief ill-founded. In the case of belief, we can distinguish between the relation of dependence on a hope that makes it ill-founded, and the features that explain why the belief is made ill-founded by that dependence relation. And then we can ask: which features of belief explain why depending on a hope in that way makes beliefs conduct ill-foundedness? If a belief's dependence on a hope makes it conduct ill-foundedness because the belief has feature X, and experience E with the same dependence relation on a hope has feature X, then that's a reason to think that dependence relation makes E conduct ill-foundedness. That is my argumentative strategy for defending the thesis that the compound experience is downgraded, I first compare the compound experience with a belief, B1, that has the same content and depends on a hope in the same way. Then I argue that this dependence relation makes B1 conduct ill-foundedness, by virtue of a feature that beliefs in general share with experiences.

To execute this strategy, we need an example of B1. Consider a route to belief via wishful (or fearful) thinking. Here you don't experience any red things at all, but you form the belief that all the red things in the display are square because you hope or fear that this generalization is true. A different kind of wishful (or fearful) thinking operates by overgeneralization. You see some red squares (as such), and the red-square hope (or fear) makes you jump to the conclusion that all the red things in the display are square.

In both of these wishful (or fearful) thinking cases, the red-square hope (or fear) causes the subject to form B1. And in both cases, because of this influence, B1 is a conduit of ill-foundedness, for subsequent beliefs formed primarily on its basis. (Maybe the ill-foundedness of B1 would wash out if the basis of a subsequent belief B2 was big and complex enough and included B1. But it

²⁴ The underlying principle here is the same as the one discussed earlier in rejecting the "split the difference" option for evidentialism (note 18). For more on the link between being a conduit of ill-foundedness and epistemic downgrade, see Siegel 2013a, and McGrath 2013b.

doesn't always wash out. For instance if B2 = thing1 is a red square, if you formed B2 on the basis of B1, B2 would thereby be ill-founded.)

In these cases, are the features that make B1 conduct ill-foundedness shared by the compound experience? Let's call C the universal generalization that all the red things in the display are square. We have already noted a first causal feature that applies to both the direct and the overgeneralization routes to B1: the red-square hope (or fear) causes there to be a state with content C, in a manner characteristic of wishful thinking. The compound experience ends up with content C via the selection effect imposed by the red-square hope, whereas B1 ends up with it via wishful (or fearful) thinking, either directly or by overgeneralization.

A second causal feature is shared by the compound experience and by B1 when B1 is formed by overgeneralization. This feature is a transition, caused by the red-square hope (or fear), from a state of flux to a state of fixity. We've seen that in the case of the compound experience, the red-square hope causes a transition from experiential flux (when the salience structure is still evolving, and the subject is still looking around the display for red things), to experiential fixity (when the salience structure stabilizes into phenomenal that's-allness with respect to red things). A similar transition can be found in the overgeneralization route to B1. At the level of belief, the transition between flux and fixity can be thought of as a transition from suspending judgment on whether a proposition P is true, to settling on an answer to that question (an answer other than "I don't know"). This transition is normatively assessable. In some cases, if one starts out in a state of suspended judgment on what the answer is to the question one is investigating, sometimes it is epistemically appropriate to continue suspending judgment in the face of incoming information, rather than to settle on an answer to that question.

When B1 is formed via overgeneralization, in the course of a search for red things in the display, there is a similar transition from flux with respect to which things in the display are red (while you are still looking around the display) to fixity on this question. Just as in the compound experience case, the transition from experiential flux to experiential fixity is caused by the red-square hope, in the case of B1, the analogous transition is caused by the red-square hope. In the compound experience, the red-square hope causes an experience (part of the compound experience) of looking around the display for red things and ceasing to find any more of them. In the case of B1, the red-square hope causes the overgeneralization from the individual red-square experiences.

Given the assumption that the compound experience and B1 (in the overgeneralization case) depend on a red-square hope in the same way, we can ask why depending on a red-square hope in that way makes beliefs conduct ill-foundedness. Are the features of B1 that explain why the overgeneralization we've described leads to ill-foundedness shared with experiences?

A first set of reasons to think they are shared is that the features that distinguish belief from experiences don't explain why irrational dependence relations make beliefs conduct ill-foundedness.

The major difference between beliefs and experience is that we don't subject them to the same normative assessment. We consider some beliefs to be unjustified or irrational, and we subject all beliefs to normative assessment. But we rarely, if ever, criticize someone on the grounds that they have an irrational or unjustified experience. We don't regularly subject experiences to normative assessment. And even if some experiences are downgraded—which could be seen as a form of epistemic normative assessment—it sounds odd to criticize someone for having the experience they have, even if it is downgraded.

But on closer examination, these disanalogies are not relevant to finding features that explain why irrational dependence relations ill-found beliefs. When B1 is ill-founded by irrational dependence relations such as wishful thinking, the explanation of what makes it ill-founded is not the fact that B1 is irrational. Rather, the explanation goes around the other way. B1 is irrational, because of the way it depends on wishes or fears.

An opponent might reply that beliefs can be ill-founded by as wishful or fearful thinking, because they need to be caused by rational processes to provide justification, and wishful or fearful thinking are not rational processes. In reply, it doesn't seem true in general that beliefs have to be caused by rational processes to provide justification. For instance, when developmental psychologists invoke "core cognition" to explain behavior of infants, they are naturally interpreted as positing beliefs, such as beliefs that give shape to core concepts of mechanical causation and agency, that constitute our intuitive physics, and that underlie our primitive number systems.²⁵ These beliefs were not ontogenetically acquired at all, a fortiori they are not caused by rational processes. For all that, they may be rational beliefs to hold, and presumably they provide rational support for subsequent beliefs in these domains. If so, then it is not in general true that beliefs need to be caused by rational processes to provide justification.

What features of belief explain why they can be ill-founded by irrational dependence relations on wishes, fears, or prejudice? My answer is that beliefs are made to conduct ill-foundedness by irrational dependence relations, by virtue of the fact that belief is a mode of endorsing content. Beliefs are states whose contents subjects endorse, and thereby give shape to their outlook on the world. Whatever else beliefs are, they are endorsements of contents.²⁶

Why think that being an endorsement of content explains why irrational dependence relations make beliefs conduct ill-foundedness? Consider the relationship between ill-founding and well-founding. Schematically, the

²⁵ Carey 2009.

²⁶ What about other belief-like states, such as low credences, the cognitive underpinnings of implicit bias, delusions, or specially circumscribed situations in which subjects act as if P were true? Are these all endorsements of content as well? Low credences can be set aside because the argument requires only that there is a class of beliefs that are endorsements of content. (I am assuming without argument that reductionist approaches that would dismiss our usual notion of belief or replace it with credences cannot be right.) Whether or not the other states are beliefs, they are belief-like to the extent that they involve endorsements of contents.

reasoning can be put like this: B1 ill-founds B2 only if B1 has feature F. B1 has feature F only if it is an endorsement of content. So B1 ill-founds B2 only if it is an endorsement of content. Less schematically, when a belief B1 ill-founds another belief B2, an epistemically good-making feature of B1 is removed. That feature is: providing rational support for B2. For instance, if B1 were well-founded, and it combined with other states in the right way, then it could provide rational support for B2. But in at least some cases, for B1 to provide rational support for B2, B1 has to be an endorsement of content. For instance, having a very low credence or weak belief in P could not make it reasonable to endorse the disjunction P-or-Q. In general, the relations underlying coherence requirements on belief need a steady attitude all the way through. Putting these points together, in at least some cases, for B1 to ill-found B2, B1 has to be an endorsement of content. So being an endorsement of content explains why irrational dependence relations make beliefs conduct ill-foundedness. If being an endorsement of content explains why beliefs are ill-founded by irrational dependence relations, then this is explained by a feature that is shared by experiences. Experiences are also states whose contents subjects endorse, and thereby give shape to their outlook on the world. By virtue of being endorsements of contents, beliefs and experiences have similar psychological roles. Like beliefs, experiences provide input to the reasoning (including action plans) that we actually go through—whether that reasoning conforms to epistemic and practical norms or not. Experiences and beliefs play similar roles in reasoning (good and poor, epistemic and practical) because are both are modes of endorsing content.

Of course sometimes, epistemic rationality dictates that we quarantine experiential endorsement, so that it does not interact in the usual ways with our other beliefs and with our behavior. Other times, rationality dictates that we don't quarantine the experience—but we do that anyway, refusing to believe our eyes when we should. In both cases, the same is true of belief. Sometimes epistemic rationality dictates that we give up a belief; other times it dictates that we hold on to it.

An opponent might object to the hypothesis that being an endorsement of contents explains what makes irrational dependence relations ill-founded beliefs, on the grounds that if true, it shows too much. If being an endorsement of content explains why beliefs are made to conduct ill-foundedness by irrational wishful or fearful thinking, then it is natural to think it also explains why beliefs can be made to conduct well-foundedness by good reasoning. But then by parallel reasoning about experience, experiences can also be made to conduct well-foundedness by good reasoning. That result may seem to go too far. The idea that experiences can be justified by person-level, psychological antecedents is at odds with our practices of epistemic normative assessment. Either that part of our practices should be revised, or else the reasoning is too powerful.

Two points can be made in reply. First, suppose it followed from the fact that experiences can conduct well-foundedness by being caused by rational

processes, that experiences had to be caused by a rational process in order to provide justification. That consequence might be worrisome, since it seems false, and seems to rob experience of the potential to stop regresses of justification. But the consequence does not follow. So the result can respect the plausible idea that for experiences to provide justification, they do not need to be the result of good reasoning, or any other rational process operating over person-level psychological states.

Second, our conception of experiences as a-rational should be revised in any case, to account for the rational status of a kind of dependence relationship that can hold between experiences, or more exactly between subexperiences (parts of the same overall experiences)—even if these processes do not guarantee that the dependent subexperience conducts well-foundedness. For instance, suppose you experience x as edible, as a result of experiencing it as orange and jutting out from the tree trunk, together with your background belief that mushrooms with those features are edible. Here your experience of the mushroom as edible could be seen as a rational elaboration of the rest of your experience (x is orange and jutting out) combined with your background belief. To take another example, acrophobes standing on high balconies tend to overestimate their distance from the ground, compared with people standing on the same balcony who are not afraid of heights.²⁷ An acrophobe's fear might lead her to experience a balcony as being at a distance that is dangerous to fall from, and that subexperience might explain why she experiences the distance to the ground as magnitude $D +$. Here too, there is something like good reasoning that unfolds within the experience: if a balcony is a dangerous height to fall from, then it is at least $D +$ from the ground. The starting point of this reasoning is fear-induced subexperience of danger, but from there it unfolds into another subexperience of distance in a way that seems rational, not in the normative sense, but in the way that rationalizations do. If experiences can stand in rationalizing relations to each other, then the part of our epistemic normative practices on which we prescind from assessing experiences from having rationally assessable causes should be revised.

I've argued that the compound experience is downgraded by the role of the red-square hope in selecting only square red things for experience, and in producing an experience with the universal generalization as its content. If this conclusion is correct, it challenges the presumption that the etiology of experience is always an a-rational process. The causal history of an experience, and more generally the dependence relations it stands in to hopes, fears, or prejudice, can be rationally assessable, and more specifically they can be irrational. And when it is, it reduces or eliminates the rational support that the experience provides for believing its contents.

I now turn to a different kind of selection effect: anti-selection of experiences for uptake into belief.

²⁷ Stefanucci and Proffitt (2009).

4. ANTI-SELECTION FOR UPTAKE THAT LEADS TO CONFIRMATION BIAS

Here is an ordinary phenomenon. You open the fridge to see whether there is any mustard inside. There is a jar of mustard in the fridge, and you see it when you open the door. But instead of forming the belief that there's mustard in the fridge, you close the door and find yourself as unsure as you were at the start as to whether the fridge contains mustard. You didn't come to believe and then forget that the mustard is there; instead, you never formed any mustard-belief in the first place, because your mustard-experience—the experience you have when you see the mustard, which we can suppose does not involve any illusion or hallucination—never fed into any belief-formation process at all.

Whereas the mustard case as stated doesn't specify the factors that led your mustard-experience to get passed over in belief-formation, in the kind of anti-selection for uptake that leads to confirmation bias, a prior psychological state does this, such as a prior fear, hope, or belief. For instance, if you strongly hope that there is mustard in the fridge, and that hope "freezes" your mustard-experience when you see the nearly empty jar, leading you to form the belief that the mustard jar is full, then that would be the kind of anti-selection for uptake of interest here.

To help us discuss this kind of anti-selection for uptake, some terminology is useful. First, let's call a belief B_p a *bypassing belief*, just in case B_p is formed or maintained in part by a process that bypasses an experience that would uncontroversially rationally bear on p , if it were selected for uptake—holding constant the rest of what the subject believes. By definition, a subject will either initially form a bypassing belief upon having a bypassed experience, or will form the belief prior to having that experience and maintain it afterward, in which case her state of believing p is a bypassing belief only in the period after her experience is bypassed. I'll call experience E *pertinent* to a bypassing belief B_p , just in case E would rationally bear on p , if it were not bypassed. Since the definition of pertinence involves a counterfactual, to assess whether an experience is pertinent to a belief, we need a way of individuating experiences across worlds. There are different ways to individuate experiences, but the way I'll be individuating them holds constant phenomenal character and level of attentiveness.

A simple example of pertinence can be found in typical perceptual demonstrative beliefs. Consider beliefs that would be natural to express by uttering a sentence of the form *that is an F*, where the speaker is perceiving what she's intending to refer to in using the demonstrative.²⁸ In tasks that ask subjects to classify certain targets as belonging to a class A or class B, subjects are supposed to form demonstrative beliefs, of each target they see, that it is either (an) A or (a) B. For instance, if asked to classify what they see as a

²⁸ In addition to perceiving what she intends to refer to, the intention and the perception have to be related in the right way. For discussion, see Siegel 2002.

tool or a gun, or as red or green, subjects are asked to form demonstrative beliefs, of the targets, that they are a tool (or a gun), or red (or green). Here, the experiences of the things demonstrative referred to are pertinent to the demonstrative beliefs.

Anti-selecting one's experience for uptake differs from discounting it. Discounting an experience is a way of acknowledging a defeater for the experience.²⁹ For instance, suppose you open the fridge and have a mustard-experience, but you know that you are liable to hallucinate mustard, so you don't form the belief that the fridge contains mustard. Here your mustard-experience would be feeding into the process of belief-formation, and discounted in that process. In contrast, if your mustard-experience is anti-selected for uptake, it does not feed into beliefs about what the fridge contains at all.

The contrast with discounting can be brought out further by comparing different types of irrational perceptual belief. While some cases of discounting give experiences only as much evidential weight as they deserve, other routes to discounting experience seem doomed to irrationality. For instance, consider an anxious subject who has a pervasive, ominous gut feeling of uncertainty that makes her hold all perceptual beliefs less firmly, across the board. The fact that this discounting originates in anxiety seems to make the process of forming beliefs irrational. Or consider a subject who blankly refuses to draw the natural conclusions from her experience, even though she has no background belief that appearance are misleading, nor any (prior) gut feeling of uncertainty. She sees the jar of mustard in the fridge, it looks the way she expected the mustard to look, but she doesn't form the belief that there is mustard in the fridge. Like the anxious subject's beliefs that there probably isn't any mustard in the fridge (or her low credence in the proposition that the fridge contains mustard), this subject's agnosticism about whether the fridge contains mustard seems irrational.

The anxious subject and the resolutely agnostic subject are both crazy. But they are crazy because of the way they respond doxastically to experience, when the experience is pertinent to the belief. In contrast, when experiences are anti-selected for uptake, they have no more than a causal impact on the subject's forming a belief to which the experience is pertinent.

Anti-selection for uptake is a more likely explanation of certain dynamics of belief than discounting. At time 1, you believe that the mustard jar is in the fridge, mostly full. Shortly afterward, at time 2, you open the fridge, and there's the mustard, just where you knew it would be, but it is visibly empty and looks empty to you. You see the mustard jar, it looks empty, but when you shut the fridge at time 3, you come away believing exactly what you started out believing: that the mustard jar is in the fridge, mostly full. A likely

²⁹ Defeat of experience is introduced by Pollock in his 1974. See also the discussion of undercutting defeat in Pollock 1986.

explanation for your behavior—more likely than crazy discounting—is that you failed to integrate the information about the mustard you got from seeing it with your prior beliefs.

A similar explanation could apply to an experiment reported by Triesch et al. (2003), who were investigating change-blindness in the context of a sorting task. Participants who are fixating a block while moving it (in a virtual environment) sometimes fail to report changes in height or color, even when the change is relevant to their motor task, such as moving the tall blue blocks to one place and short yellow ones to another. The changes are *less* frequently missed when the change *is* relevant to the task, and that link between reporting changes and task-relevance is the central result of the study. But when subjects fail to report the change, it is as if, after the block was first seen, but before it changed color or height, the channel that normally allows information from experience to guide belief and action is shut off. If that's what happens, then the stretch of one's perceptual experience after the block changes color is bypassed in the process of maintaining a prior belief. After the un-noticed change in color or height, the experience plays a central role in making it the case that the subject maintains a belief about the block.³⁰

5. CAN ANTI-SELECTED EXPERIENCES RETAIN THEIR RATIONAL ROLE?

Anti-selection of experience for uptake can generate a type of fragmentation in the mind, in which an experience that could potentially impact which beliefs one forms is prevented from having that kind of impact. Beliefs are sometimes compartmentalized in a way that has this same effect. You might believe you made an appointment with X alone at noon, while also believing that you made an appointment with Y alone at noon, without realizing that you made conflicting appointments. Here you fail to believe the straightforward consequence that you made conflicting appointments, as well as failing to believe the contradiction that you both will and will not meet with X alone at noon. In a similar example offered by Lewis (1982), a resident of New Jersey believes that Nassau Street in Princeton runs north-south and is parallel to the train tracks, and also believes that the train tracks run east-west and are parallel to Nassau Street. Here too, the subject fails to believe a straightforward consequence of their beliefs, since they fail to believe the contradiction that Nassau both does and does not run east-west. The metaphor of mental compartments (or mental fragments) is apt for describing all of these cases, because a type of inferential integration among one's mental states is missing.

The compartmentalization among beliefs that was just described differs from anti-selection for uptake in several ways. First, your belief that you

³⁰ The Triesch experiment itself does not establish that the behavior it elicits is best explained by anti-selection for uptake, but this hypothesis is an open possibility.

have an appointment with X at noon may be accessed by some strands of reasoning while not being accessed by other strands. For instance, the belief that you'll see X at noon may ground your belief that you'll soon see X and can give her back her umbrella, while being unaccessed by a strand of reasoning that would lead you to discover your conflicting appointments. In contrast, anti-selected experiences may be unaccessed by *any* strand of reasoning. Each of your appointment beliefs belongs to a separate compartment with its own internal rational complexity, whereas an anti-selected experience may belong to a compartment with no internal rational complexity to it.

Second, in the Nassau Street case, the contents of the resident's beliefs are inconsistent. The appointments case may involve contradictory contents as well, such as *I will meet with X alone at noon* and *I will meet with Y alone at noon*. Whether the belief *states* in these cases are inconsistent is a substantive question to which we'll return shortly. In contrast, the content of an anti-selected experience and its bypassing belief need not be inconsistent.

Third, the subjects in the belief cases seem to fail to believe straightforward logical consequences of other things that they believe, and this failure seems due to the compartmentalization. In contrast, although the subject of experiences that are anti-selected for uptake will not believe the contents of that experience, or any of their consequences on the basis of that experience (since she does not believe anything on its basis), she may believe those contents on other grounds.

Despite these differences, anti-selection of experience and compartmentalization of belief raise the same epistemological question: what rational bearing do the unintegrated states have, if any, on the rational status of beliefs that they cannot impact, so long as they remain compartmentalized?

We can apply this general question to our cases of anti-selection of experiences for uptake by asking: When an anti-selected experience is pertinent to a bypassing belief, what impact does the anti-selection have, if any, on the rational status of the belief? There are two main options: no impact or some impact. Either the pertinent experience is rationally inert vis-à-vis the bypassing belief (the inertness option), or it isn't (the impact option). The inertness option licenses combinations of experience and beliefs which, in a unified mind, would be irrational. In contrast, on the impact option, bypassed experiences retain their rational force with respect to all beliefs to which they are pertinent. Here, if a combination of experience and belief would be irrational if the experience was selected for uptake, it remains irrational, even when the experience is bypassed. On a global version of the impact option, all bypassed experiences retain whatever rational force that they would have if they were selected for uptake, and on a global version of the inertness option, anti-selection for uptake always removes any rational force that the bypassed experience would have, were it not bypassed.

The impact and inertness options interact in a straightforward way with evidentialism about rational belief. In an evidentialist framework, the difference between the impact and inertness options directly translates into two

positions on whether perceptual bypass is a means of ignoring evidence.³¹ The inertness option holds that anti-selecting a pertinent experience cannot make any corresponding bypassing beliefs irrational. In its evidentialist form, the inertness option entails that experiences provide evidence only if they are selected for uptake (not necessarily uptake into belief, but uptake into even the most minimal forms of inferential processes). In contrast, the impact option holds that anti-selection of a pertinent experience can make corresponding bypassing beliefs irrational. In its evidentialist form, the impact option entails that anti-selection of experience can be a means of ignoring evidence.

I think a *prima-facie* case can be made for a version of the impact option, independently of the evidentialist framework. One could try to defend the global impact option, by appealing to a putative requirement of ideal rationality. According to the putative rational requirement, an ideally rational subject would always respond appropriately to her experiences, and perceptual impact thus leaves subjects one step closer to that ideal than anti-selection, where the subject does not have the opportunity to respond doxastically to her experience at all, either rationally or irrationally. According to this line of thought, anti-selection of experience that is pertinent to what the subject believes is a departure from ideal rationality.

We can distinguish the line of thought just sketched from the idea that ideal rationality requires you to consider all your evidence in forming doxastic attitudes. What's at issue between the global impact option and its opponents is whether bypassed experiences always have the status of evidence. The strategy sketched above would need to provide grounds for the idea that anti-selected experiences always have this status.

In contrast to the global impact and inertness options, in principle anti-selected experiences might differ in whether they are rationally inert, depending on the etiology of the anti-selection. Here, certain routes of anti-selection for uptake make a case for at least a moderate impact thesis. In particular, when experiences are anti-selected for uptake in a way that leads to confirmation bias, it is plausible that the bypassed experience retains its rational force *vis a vis* the beliefs to which it is pertinent.

Consider a subject *S* who comes to believe that the mustard jar in the fridge was full on the basis of *Q*'s testimony, and *Q*'s testimony gives *S* good grounds for this belief. But suppose that *S* overvalues *Q*'s opinion about nearly everything, and tends to give *Q*'s opinions more weight than they deserve, because in general, *S* wants *Q* to be right. *S*'s dispositions to overvalue *Q*'s opinion cause *S* (or causes *S*'s belief-formation system) to anti-select her experience when she opens the fridge and sees the nearly empty jar of mustard inside, with the result that after looking in the fridge, *S* keeps on believing what she believed before she looked in the fridge—that the mustard jar is full. Here anti-

³¹ At least, it can be so translated, given the assumption that experiences typically provide evidence for external-world beliefs.

selection happens in the service of a type of confirmation bias, and a type that perceptual experience, in general, is well-suited to correct. In this type of case, S's post-fridge mustard-belief seems epistemically less well grounded than the pre-fridge belief was, even though both beliefs were based on Q's good testimony. A straightforward explanation of why it is less well-grounded is that S fails to heed the excellent reason, provided by her bypassed experience, to believe that the mustard jar is nearly empty. If her bypassed experience provides her with this type of reason for belief, then the impact option is correct.

A second consideration targets the inertness option. In response to the cases of compartmentalized beliefs discussed earlier, one might hold that those beliefs are rationally inert with respect to one another, by virtue of the fact that they are never activated at the same time. For instance, Lewis's reason for thinking that the New Jersey resident does not believe a contradiction is that the east-west belief and the north-south belief never "come into action" at the same time.³² And if it is part of believing a proposition that one is disposed to act in ways that would satisfy what one desires if the proposition is true, then it is hard to see what ways of acting those would be, if the proposition has the form $P \& \text{not-}P$. (Here we return to the question set aside earlier about whether the belief states in the Nassau Street case are inconsistent).

It seems compatible with Lewis's own description of the Nassau Street case that even when the belief that p is dormant, it rationally impacts the activated belief that $\text{not-}p$. For instance, if the belief that p is well-grounded, whereas the belief that $\text{not-}p$ is not, then the belief that p would make it less rational to believe $\text{not-}p$. That position would give beliefs rational impact on one another, even when they are compartmentalized. So Lewis himself seems uncommitted on the question whether compartmentalization per se neutralized rational impact that the belief in p would otherwise have with respect to a belief in $\text{not-}p$. But someone might hold that compartmentalized beliefs are always rationally inert with respect to one another, due to the fact that they never both "come in action" at the same time.

What exactly is it for a belief to come into action? One option is that a belief comes into action, only while the subject is behaving as if the belief is true, relative to her desires (or to some of her desires). Another option is that a belief comes into action, only while the subject is (right then!) drawing inferences from it. We can set aside this important question, though, because on either of these ways of elaborating what it is for a belief to come into action, anti-selected experiences do not seem to come into action selectively, the way compartmentalized beliefs do. In our original mustard example, for example, the anti-selected mustard-experience fails to guide behavior at all. So if the fact that compartmentalized beliefs only selectively come into action is

³² "Different fragments came into action in different situations, and the whole system of beliefs never manifested itself all at once." (Lewis, p. 436).

supposed to underwrite their rational inertness with respect to beliefs in other compartments, then this feature is not available to underwrite the inertness option for bypassed experiences.

In a Triesch-style example of anti-selected experiences, the experience may feed into some beliefs but not into others. For instance, the experience may feed into beliefs about where the block is on its path from the moving belt it was plucked from to the pile, without feeding into beliefs about what color it is. This description of the case assumes that there is a single experience representing both color and movement, rather than ontologically separate experiences of motion and color, but it is convenient to use the labels “color-experience” and “motion-experience” for two aspects of an overall experience that represents the block changing its color in mid-path. We could imagine a case in which, due to limited cognitive resources, after the color change, the motion-experience feeds into a belief about its trajectory, only if the color-experience fails to feed into beliefs about its color, resulting in a color-belief that is not properly updated. In this type of case, the bypassed experience and the bypassing belief both guide behavior, just as both fragmented beliefs do. But unlike the fragmented beliefs, the bypassed experience guides behavior at the same time as the belief that bypasses it. So the feature of fragmented beliefs that might be taken to license rational inertness—the idea that they do not both “come into action” at the same time—does not apply to the case of anti-selected experiences.

CONCLUSION

Can selection effects on experience influence their rational role? I’ve suggested that some forms of selecting objects for experience can epistemically downgrade those experiences—for example when the selection effect is explained by its role in confirmation bias. For these cases, the proposed answer to the title question is Yes. I also suggested that when experiences are anti-selected for uptake, they can retain their rational force. For these cases, the proposed answer to the title question is No.

REFERENCES

- Balci, E. and Dunning, D. 2006. “See What You Want to See: Motivational Influences on Visual Perception” *Journal of Personality and Social Psychology* 91(4), 612–25.
- Barrick, C. 2002. “Color Sensitivity and Mood Disorders: Biology or Metaphor?” *Journal of Affective Disorders* 68, 67–71.
- Byrne, A. 2009. “Experience and Content” *Philosophical Quarterly* 59(236), 429–451.
- Carey, S. 2009. *The Origin of Concepts*. New York: Oxford University Press.
- Chalmers, D. 2006. “Perception and the Fall from Eden”. Reprinted in *The Components of Consciousness*. New York: Oxford University Press, 2012.

- Churchland, P. 1988. "Perceptual Plasticity and Theoretical Neutrality" *Philosophy of Science* 55, 167–87.
- Conee, E. and Feldman, R. 2001. "Internalism Defended". Reprinted in Feldman and Conee 2004, 53–82.
- 1998. "The Generality Problem for Reliabilism" *Philosophical Studies* 89(1), 1–29.
- Feldman, R. and Conee, E. 1985. "Evidentialism." Reprinted in Feldman and Conee 2004, 83–107.
- 2004. *Evidentialism: Essays in Epistemology*. New York: Oxford University Press.
- Fodor, J. 1983. *Modularity of Mind*. Cambridge: MIT Press.
- Foley, R. and Fumerton, R. 1982. "Epistemic Indolence" *Mind* 91, 361, 38–56
- Glüer, K. 2009. "In Defence of a Doxastic Account of Experience" *Mind and Language* 24(3), 297–327.
- Goldin, C. and Rouse, C. 2000. "Orchestrating Impartiality: The Impact of 'Blind' Auditions on Female Musicians" *The American Economic Review* 90, 715–41.
- Hansen, T. et al. 2008. "Color Scaling of Discs and Natural Objects at Different Luminance Levels". *Visual Neuroscience* 23, 603–10.
- Huemer, M. 2013. "Epistemological Asymmetries between Belief and Experience" *Philosophical Studies* 162(3), 741–748.
- Kaplan, M. 1996. *Decision Theory as Philosophy*. Cambridge: Cambridge University Press.
- Kelly, T. 2002. "The Rationality of Belief and some other Propositional Attitudes" *Philosophical Studies* 110, 165–96.
- Levin, R. and Banaji, M. 2006. "Distortions in the Perceived Lightness of Faces: The Role of Race Categories" *Journal of Experimental Psychology* 135(4), 501–12.
- Lewis, D. 1982. "Logic for Equivocators" *Noûs* 16(3), 431–41.
- Lyons, J. 2011. "Circularity, Reliability, and the Cognitive Penetrability of Perception" *Philosophical Issues* 21(1), 289–311.
- McGrath, M. 2013a. "Cognitive Penetration and Bad Basis Counterexamples" in Tucker, ed. (2013). *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*.
- McGrath, M. 2013b. "Siegel and the Impact for Epistemological Internalism" in *Philosophical Studies*, 162(3), 723–32.
- Martin, M. G. F. 2002. "The Transparency of Experience" *Mind and Language* 4(4), 376–425.
- Payne, B. K. 2001. "Prejudice and Perception: The Role of Automatic and Controlled Processes in Misperceiving a Weapon" *Journal of Personality and Social Psychology* 81(2), 181–192.
- Pollock, J. 1974. *Knowledge and Justification*. Ithaca: Cornell University Press.
- 1986. *Contemporary Theories of Knowledge*. Lanham, MD: Rowman and Littlefield.
- Popper, K. 1959. *The Logic of Scientific Discovery*. London: Hutchinson.
- Proffitt, D. 2006. "Embodied Perception and the Economy of Action" *Perspectives on Psychological Science* 1, 110–22.
- Raftopolous, A. 2009. *Cognition and Perception: How do Psychology and Neural Science Inform Philosophy?* Cambridge, MA: MIT Press.

- Siegel, S. 2002. "The Role of Perception in Demonstrative Reference" *Philosophers' Imprint*, 2(1), 1–21.
- 2011. "Cognitive Penetrability and Perceptual Justification" *Nous* 46(2), 201–222.
- 2013a. "The Epistemic Impact of the Etiology on Experience" *Philosophical Studies*, 162(3), 697–722.
- and Silins, N. 2013. "The Epistemology of Perception." In *The Oxford Handbook of the Philosophy of Perception*. Edited Mohan Matthen. Oxford: Oxford University Press.
- Stefanucci, J. K. and Proffitt, D. R. 2009. "The Roles of Altitude and Fear in the Perception of Height" *Journal of Experimental Psychology, Human Perception and Performance* 35(2), 424–38.
- Steinpreis, R., et al. 1999. "The Impact of Gender on the Review of Curriculum Vitae of Job Applicants and Tenure Candidates: A National Empirical Study" *Sex Roles* 41, 509–28.
- Sturgeon, S. 2000. *Matters of Mind: Consciousness, Reason and Nature*. London: Routledge.
- Triesch, J. et al. 2003. "What You See is What You Need" *Journal of Vision* 3, 86–94.
- Tucker, C., 2013. *Seemings and Justification: New Essays in Dogmatism and Phenomenal Conservatism*. Oxford: Oxford University Press.
- Watzl, S. 2010. *The Significance of Attention*. Columbia University PhD Thesis.
- 2011. *Attention as Structuring of the Stream of Consciousness*. In *Attention: Philosophical and Psychological Essays*. Edited by C. Mole, D. Smithies, and W. Wu. New York: Oxford University Press.