

Cognitive Processes for Epistemologists

Jack Lyons

University of Arkansas

Abstract:

Individuating cognitive processes is important epistemological work. This is most obvious in connection with the well known generality problem for reliabilism. I don't believe that solving the problem of process individuation constitutes a solution to the generality problem, although I do think it goes a substantial way toward solving that problem. I also think that this solution solves a number of other problems, and not only for reliabilists. In this paper, I develop a psychological criterion of process individuation. On the basis of some plausible, general, empirical assumptions, I argue that if we turn the matter of process type individuation over to cognitive psychology, what we get back is pretty much exactly what we were looking for: a principled way of assigning a unique process type to each belief token, which appears to get the cases intuitively right as well. Toward the end, I apply this theory to the case of hallucination, where it has interesting results, leading us to a halfway position between the traditional view and a disjunctivist view.

Individuating cognitive processes is an important task for epistemology. This is most obviously the case for reliabilism, which faces the well known generality problem. In this paper, I want to develop a psychological criterion of process individuation. As will become clear below, I don't believe that a solution to the problem of process individuation constitutes a solution to the generality problem, although I do think it goes a long way toward solving that problem. I do think, however, that it solves a number of problems, and not just for reliabilists. After explaining how I think processes should be individuated, I apply this theory of processes to the case of hallucination, where it has interesting results. The basic idea of a psychological criterion for process individuation is not a new one, although I think the motivations, details, and implications here are new.

Although process individuation has loomed large in discussions of the generality problem, I don't think that what I offer here constitutes a solution to the generality problem. I try to clarify the relationship between a theory of processes and the generality problem in section I. In section II, I explain the sense in which processes are of interest to other epistemologies besides reliabilism. In section III, I lay out the general view about how I think processes should be individuated. The theory of process individuation I favor is similar to some extant solutions to the generality problem, but I develop things in importantly different directions. Section VI makes the argument that this way of individuating processes gives us the desired epistemological results, and Section V discusses the distinction between processes as understood here and the kinds of rules for reasoning encountered in a logic or critical thinking class; these two are too often conflated in the literature. In section VI, I explore the ways in which the psychological criterion has implications for hallucination. In the end, I think it opens up a halfway position between the traditional view and a disjunctivist view.

I. Processes and the Generality Problem

I want to insist right up front that I'm not offering this as a full solution to the generality problem. A theory for individuating process types only gives us a partial solution to the generality problem, because the degree of reliability is determined not just by the type of process used but by the environment in which it is used. Suppose, for example, that the process I'm using to form a belief right now is reliable (has a high true/false ratio) in a 1 mile radius from here but unreliable in a 10 mile radius; it is reliable again in a 1000 mile radius. Is the process reliable or not?¹ Or suppose an ordinary person has her brain removed and put in a vat, ala the standard skeptical scenario, for five minutes, and then replaced. Do her processes count as reliable during those five minutes? During the short period immediately afterward?²

Discussions of the generality problem have tended to focus on the individuation of cognitive processes to the neglect of the individuation of the relevant environment. I imagine that this is because those pressing the issue against reliabilism (primarily Feldman and Conee [Feldman 1985; Conee and Feldman 1998]) have been treating processes synecdochally to illustrate more vividly the larger problem. To solve the generality problem, we will need not just a way of individuating processes, but a way of individuating environments as well, and I fear that the latter will be substantially harder than the former. This is partly because I can't foresee anything for individuating environments that is at all analogous to the way I will defend for individuating processes.

Still, there is much to be gained by having a better grasp on processes. First, there are a good many discussions of reliabilism where the whole train of argument simply gets derailed because we lack a good sense of what sorts of things might count as a cognitive process, and this should bring clarity to these discussions. Second, I think that the difficulty of individuating environments is a real problem and not merely an artifact of having a solution to the other problem first; I think I am not merely moving the lump from one part of the carpet to another by having a theory of process individuation without having a theory of environment individuation. It is more like pressing out one of two lumps while leaving the other unaffected. Finally, as I want to argue next, a theory of process individuation should be of interest to other epistemologists, and not just reliabilists.

II. Beyond Reliabilism

Reliabilism is a view about what makes the good processes good; one needn't accept that to agree with reliabilism about (a) the truth of a general process theory of justification, and (b) which processes are good.

¹ Pollock 1986 makes this point; he claims that color vision is reliable on earth but not in the universe more generally.

² For further difficulties along these lines, see Lyons 2012.

A “state theory” is one that holds that justification is determined by the current state of the organism, where the current state is specified in nonfactive, nonnormative, and ahistorical terms. A “process theory” is one that holds that the causal processes responsible for a belief can affect the justificatory status of that belief. A very weak version of this claim restricts such influence to basing relations; pretty much everyone holds that basing is (at least sometimes, at least partially) causal, and that what one’s belief is based on can affect the justification of that belief. Since this much is nearly uncontroversial, I will restrict my attention (and the use of the term) to those process theories that claim that the nature of the belief-forming and/or -sustaining³ processes can affect justification in a way that goes beyond the influence involved in basing.

The obvious place to look for this additional influence is where beliefs receive their justification from something other than beliefs. The standard candidates are beliefs that result from introspection, perception, memory (what Sellars calls IPM judgments), and perhaps a priori intuition.

Some epistemologists (e.g., Lyons 2009) hold that these beliefs typically aren’t *based on* anything; these are epistemologically basic beliefs for which the agent doesn’t have (and/or doesn’t require) grounds, or evidence, or reasons. This isn’t to say that the beliefs are unjustified, only that for these beliefs, such evidentialist concepts as grounds, reasons, and basing a belief on them, are simply inapplicable. Surely *something*, however, must distinguish the justified beliefs from the unjustified, even within these categories, and the nature of the belief-forming process seems to fit the bill.

Other epistemologists hold that these beliefs are based on nondoxastic states, typically conscious experiential states, but allow that the causal history of the experiential state itself can affect the justification of the belief that is based on it. My current belief that *p* is based on a nondoxastic state of seeming-to-remember that *p*, but if that seeming-to-remember is the result of an *unjustified* earlier belief that *p*, then the experience cannot serve as a justifying ground for my current belief, even if this belief is properly based on the experience. (Senor 1993, Goldberg 2010). Similarly, my perceptual experience might be largely the result of wishful thinking; if so, this should negatively affect the epistemic status of the consequent perceptual belief (Lyons 2011, Siegel 2011). The standard variant of this idea holds only that the causal history of these experiences affects the epistemic status of the beliefs based on them, but one interesting way of elaborating on this basic idea is to claim that the etiology of the experience actually gives it a negative epistemic status (similar to a belief’s being unjustified), which renders it incapable of conferring justification on a belief (Siegel forthcoming, Sosa 2007).

The process theory is embraced most obviously by reliabilists,⁴ but one needn’t be a reliabilist to hold that the process by which a belief is formed is relevant to its justification. Though many of the epistemologists just cited are reliabilists of one sort or another, the expressed views about memory and wishful thinking could easily be embraced by nonreliabilists.

³ If the process that sustains a belief is different from the one that originally produced it, the former will often contribute as much or more to the belief’s justification as the latter. To simplify matters, I will use “belief-forming processes” as an abbreviation for belief-forming and belief-sustaining processes.

⁴ Process reliabilists, that is. Indicator reliabilism (e.g., Alston 1988) is an ahistorical view, neglecting the importance of causal factors beyond those that determine basing. For some problems that result from this, see Goldman 2009, 2011, Lyons forthcoming.

The two main branches of virtue epistemology are virtue reliabilism and virtue responsibilism (Baehr 2011). Virtue epistemologists of both sorts hold that why one holds a belief has a lot to do with the epistemic status of that belief. A virtue responsibilist, for instance, might hold that a belief that results in part from openmindedness will have a better epistemic status than one that results from partisanship. This is presumably to be understood in terms of differences in belief-forming processes.

Evidentialism holds that the justification of a belief is determined by the quality of the evidence on which it is based, but it allows that the quality of evidence is in turn determined by other factors, and these could include belief-forming processes. This is worth elaborating in some detail. Feldman and Conee (2004), for example, endorse Mentalism, which is the claim that any two mentally identical agents are justificationaly identical. Since they also endorse evidentialism, they are committed to claiming that the evidential status of a piece of evidence supervenes on the mental features of the cognizer. It could do so trivially, if, for instance, evidential relations hold necessarily (in which case, they'd supervene on whatever you like), but their view is compatible with a more substantive supervenience as well. Surely they mean for 'mental' to be restricted to nonfactive mental states and properties, but there is no reason to think they intend it only to include synchronic mental factors. The fact---which I cannot now remember---that I decided four years ago to believe *p* for no reason at all, might, consistent with their mentalist evidentialism, undermine the evidential status of my current seeming to remember that *p*. My having had a certain learning history might make a given experiential state evidence for me that there's a two-year old pileated woodpecker nearby, where that very same experience would have a much more generic evidential import otherwise (Lyons 2012).

Even going beyond IPM beliefs, to cases where the relevant etiology of a belief is a matter of basing on other beliefs, there is a distinction to be drawn between current basing and historical basing. In their discussion of inference, Pollock and Cruz (1999) distinguish between the "genetic argument" and the "dynamic argument". Any fairly complicated piece of reasoning is going to require more steps than can fit into working memory; some arguments rely on reasons that haven't been consciously entertained for several years.⁵ The genetic argument is this protracted set of historical reasons for a belief, while the dynamic argument is a snapshot of what the agent is directly basing a belief on at any given time, usually the most recent few steps of the genetic argument, along with a state of seeming to remember that provides reassurance of the earlier steps. Pollock and Cruz's own view is that one's justification is determined by the dynamic, rather than the genetic argument,⁶ but we need not follow them in this view. In fact, I suspect that few epistemologists would.

Consider someone reasoning though a fairly long argument. Keep in mind that we can only hold a small number of independent items in working memory at any given time (estimates range from the famous 7 +/- 2 down to the more recent 4 [Miller 1956, Cowan 2001]). At an

⁵ I will follow Pollock and Cruz in supposing that the contents of working memory at *t* are coextensive with what one is consciously thinking about at *t*. The supposition simplifies the exposition quite a bit.

⁶ The only role for the genetic argument is that if *S* has a defeater for her dynamic argument if she has a reason to think that there was something wrong with her genetic argument (with 'genetic argument' here presumably appearing de dicto, rather than de re). The actual genetic argument seems therefore to be irrelevant to them.

early stage in the proof, the reasoner affirms the consequent, but a few seconds later, she is no longer consciously thinking about these earlier steps. What she's consciously thinking of now are the more recent stages, while she is also in some sense aware of the fact that she seems to remember doing the earlier parts correctly. To forestall certain irrelevant complications, let's assume the agent doesn't believe affirming the consequent is valid, has no reason to believe so, was unjustified at the time in the belief that resulted from it, etc. Here's another important assumption: let's suppose---as I think is entirely natural---that she is no more consciously attending to the fact that her earlier reasoning mnemonically seems to have been fine than she is on the fact that p and $q \supset p$ were her original reasons for believing q .

I admit that there is some temptation to think that in cases where I am consciously and deliberately attending to my apparently memory of having done everything right, that this could potentially override the fact that I actually reasoned incorrectly and render the final conclusion justified. (I admit to *feeling* the temptation, not to succumbing to it!) But I myself don't feel any temptation to think that the poorly reasoning agent is justified in the case where the reassuring memory appearance is no more accessible to her than the poor reasoning itself, where both are not in, but merely readily available to, working memory.

This sort of case seems to be quite typical. When we engage in long reasoning, we are neither forming explicit beliefs about genetic arguments nor explicitly attending to mnemonic seemings regarding them. When we make mistakes in inference, we are typically unaware of doing so. Surely these mistakes in these circumstances lead to unjustified beliefs, and surely compounding these mistakes by further inference doesn't change this fact. The preservationist about memory thinks it implausible that one could create justification simply by holding an unjustified belief long enough. My argument here is analogous to, but does not depend on, this preservationist argument concerning memory: it is implausible that one could create inferential justification out of unjustified beliefs simply by keeping very poor track of her reasons for believing things.

Once more, it seems that the nature of the belief-forming process is relevant to the justification of the final belief. Note that in none of these cases---memory, wishful thinking/perception, openmindedness, and inference---did reliability considerations play a role. The argument for a process epistemology is independent of the argument for the reliabilist version of that epistemology.

Thus, I think that a concern for processes is not limited to reliabilist epistemologies. There are reasons to prefer a process theory to a state theory, and these reasons should appeal to various kinds of internalists as well as to reliabilists. It is not obvious that the different epistemologies will all converge on the same understanding of processes, but I want to put the various epistemologies aside for now and try to figure out, in a general way, how we should individuate belief-forming processes. I believe this can be done in a way that is neutral regarding the overarching epistemology, that most everyone will agree on what the relevant belief-forming

process type is---i.e., the type that determines the epistemic status of the belief---without having to agree on much else.⁷

III. Process Individuation

We need a principled way of individuating processes, and it should at least roughly capture our pretheoretic intuitions about which process types are relevant. It seems obvious---and I don't want to deny---that vision in bad lighting conditions, or of objects at long distances, does not give us as much justification as does perception in good lighting conditions and at close distances. A state theory is going to hold that the kind of percept typical of vision in bad lighting conditions confers less justification than does the kind of percept typical of vision in good lighting conditions. If we can make *vision in bad lighting conditions*, or something very much like it, one of our relevant process types, then our process theory can more or less parallel---and therefore, more or less subsume the successes of---the state theory. The parallel will be approximate, for there will be cases where, e.g., a "good" percept is formed in bad lighting conditions, and the two theories will offer conflicting verdicts about justification. The disagreement about these odd cases will separate the process theorists from the state theorists, but the existence of pervasive agreement about the ordinary cases is reassurance that there's a real dispute here and not just a misunderstanding.

My proposal is that we identify the epistemically relevant process with the (narrowest) *psychological* process that formed the belief, by which I mean the narrowest process that cognitive psychology holds to be responsible for the belief.⁸ To get this proposal off the ground, I will need to say a bit more about what I mean by 'cognitive psychology', to give some sketch of the general outlines of what processes might look like under this proposal, and to provide some reason to think there might be just one such process per belief.

By 'psychology' (and 'psychological') I mean true, finished psychology, which I presume will look a lot like current, mainstream psychology. I won't try to argue this here, but I predict that, although psychology will incorporate varying amounts of insight from such movements as behaviorism, Gibsonianism, embodied cognition, and the like, the basic theoretical framework will be the representationalist, information processing theory that has been dominant for the last half century or so.

It is important to make such predictions explicit, because the framework determines what types of variables the cognitive psychological theory will and will not include. A psychological

⁷ I am assuming a pretty strong version of a process theory here, as a process theory as such only requires that the process can *affect* the justification of the belief, while I'll be assuming that the process *determines* the degree of justification. The only real consequence this should have for what follows is that it will require me to individuate processes more finely---as finely as we want our epistemic distinctions to be---than a weaker process theory might have to.

⁸ This bears obvious similarities to solutions to the generality problem proposed by Goldman 1979, 1986, Alston 1995, and Beebe 2004, especially to Goldman's suggestion that "the critical type is the narrowest type that is causally operative in producing the belief token in question" (1986, p. 50). In a longer version of this paper, I would enumerate some of the more important differences between my account of cognitive processes and their solutions to the generality problem. I am, at the least, elaborating on and adding to what these authors say.

process is a sequence or set of transitions among psychological states or events, but different theoretical frameworks would have different views about what sorts of states or events count as psychological. Mainstream information processing psychology, like common sense, draws a sharp contrast between internal, psychological states, and external, environmental states. Thus, our psychological theory will not include in its description of psychological processes environmental variables in the way that an enactivist or other revisionary psychology might; instead it will include representations of the environmental conditions that the Gibsonian or behaviorist theory would invoke, depending, of course, on the degree to which the Gibsonian or behaviorist theory contained insights worth subsuming and incorporating.⁹ Behaviorist, enactive (e.g., Noë 2004), and certain kinds of Gibsonian psychologies would be much more willing to think of psychological states as reaching partly out into the environment. Disjunctivists view certain mental states as relational, and a psychology that was friendly to disjunctivism would not construe the mental/psychological as purely internal in the way I am suggesting.

This prediction imposes considerable constraints on process individuation. For example, we can't really count *vision in such-and-such lighting conditions* as a psychological type, because our information-processing psychology doesn't recognize lighting conditions as a *psychological* variable. It's not how much light is out there illuminating the distal stimulus; it's how much more or less information about that stimulus becomes available to the organism as the causal consequence of the lighting (or what kinds of information---reductions in light intensity might result in coarser color discriminations before affecting shape identification, etc.).

Which variables psychology countenances will depend on (i) the methodological and general substantive presumptions of the field, as just sketched (hence no genuinely environmental variables, but only internal reflections or analogues of them), (ii) which general factors the theory takes to be causally relevant (so internal analogues of lighting conditions, but not of days of the week), and (iii) the level of detail at which the theory finds differences to be less important than similarities (there must be *some* differences between visually recognizing something to be a candle and visually recognizing something to be a ukulele, since the outcomes are different, but the theory presumably treats these as different applications of a single process).

I have already elaborated on (i), and (ii) is pretty self-explanatory; *use of modus ponens on Tuesday* won't count as a distinct process from *use of modus ponens on Thursday*, not only because day of the week is external, but because day of the week---along with its internal analogue---is causally irrelevant to the psychological mechanisms.

I should say more about (iii). A familiar theme from the philosophy of psychology is that it is often as important for a theory to classify things as being the same as it is to catalogue differences. Science is in the business of stating what informative generalizations there are to state, and this requires that the science categorize items as falling under common types (Fodor 1974, Pylyshyn 1984). The *raison d'être* of the special sciences is to notice and state

⁹ This is not, of course, to deny that the dependent and independent variables are sometimes---indeed typically---external, public items. The point is that mainstream orthodoxy sees psychological processes as being entirely internal, rather than reaching out into the world. This is a weaker assumption than methodological solipsism, which requires one to reject either a representational theory of mind or an externalist theory of the contents of representations. Psychological processes might be built up out of transformations from inner states to inner states, even if the identity conditions for some or all of these inner states depend constitutively on relations to external states or properties (Burge 2012).

commonalities that are invisible from the lower, implementational levels. Part of what makes psychology distinct from, say, neuroscience, is that the former but not the latter relies on an intentional taxonomy. What makes a particular psychological theory the psychological theory it is, is in part that it invokes the particular intentional and causal categories it does, that it treats certain processes as the same and others as different.

The current restriction to psychological processes is important here, for if we processes to be individuated at the neural level, the taxonomy would be far too fine-grained.¹⁰ But the second point, about a psychological theory refusing to divide its intentional and causal categories more finely than it does, is at least as important. It is an empirical though not very surprising fact that the process we use to visually recognize ukuleles is the same as the process we use to visually recognize candles. It is an empirical question whether visual recognition of biological entities involves the same process as does visual recognition of artifacts; there is reason to think that a parts-based decompositional algorithm works well for certain items (including, mostly, artifacts) and a holistic algorithm works better for other (including, mostly, natural) items (Biederman 1987, Strat & Fischler 1991, Gerlach 2009).

If identifying artifacts involves a different algorithm than identifying natural objects does, then these are distinct identification processes. Difference of algorithm is sufficient for difference of psychological process. This allows *deduction about social contracts* to be a different process from *deduction about abstract matters* if, say, Cosmides and Tooby (1992) are right about Darwinian algorithms, but allows these to be the same process if some other psychological theory turns out to be true. In essence, individuating by algorithm makes the content of the computation irrelevant, except where that content somehow results in the system's using a different algorithm/process than it would have.

But now the challenge is to show that we can distinguish everything we want distinguished. What about vision in good light vs. vision in bad light? How are these different algorithms/processes, rather than different uses of the same process with different values for the variables? Similarly, we need to find a way to treat the kinds of things we talk about in critical thinking classes. There's no problem distinguishing modus ponens from affirming the consequent; these are clearly different algorithms, but how does argument ad hominem differ from other processes? What makes something an argument ad hominem seem to be a difference in content. What about false dilemma? Hasty generalization?

Different algorithms are sufficient for different processes but not necessary. *Vision in low light* (or rather, the internalized version of it: *vision that starts with low intensities and contrasts*) counts as a different psychological process from *vision in bright light* because the difference in lighting conditions makes a psychological difference; it affects processing. *Vision on Tuesday* doesn't count as a different process from *vision on Thursday* because the day of the week doesn't affect processing. *Vision in the presence of a convincing hologram* isn't a different process from *vision in the presence of real objects* because again, the difference doesn't affect processing.

¹⁰ There would be many more types to choose from, in part due to the multiple realizability of psychological properties. Some common psychological type might end up having an idiosyncratic realization in a single individual who is highly unreliable for entirely environmental reasons. This is a kind of complication a reliabilist would be better off just avoiding.

And now we need to avoid going too far in the other direction. The visual recognition of dogs presumably doesn't involve a different process from visual recognition of cats. In some sense, of course, the stimulus' being a dog does make a difference to processing: it results in dog identifications rather than cat identifications. The same algorithm is used for both, but that can't be all there is to sameness of processing, or else vision in low lighting conditions would count as the same as vision in good lighting conditions.¹¹ If we allow just any old difference in content to yield different processes, then the whole project is trivialized: using a given algorithm to add 3 and 4 would count as a different process from using that same algorithm to add 3 and 6.

The difference, I think, is that there is a *systematic* difference in processing across the high/low light conditions but not across the cat/dog conditions. The idea this: the dog/cat difference is one of *mere content*; the only difference between processing dogs and processing cats is that the one is about dogs and the other about cats. The difference in lighting conditions, on the other hand, constitutes a difference in general parameters of processing. It affects processing across a wide range of subject matters; dogs, cats, oak trees, tables, colors, textures are all harder to recognize in low light. The point is not that lighting conditions affect the *reliability* of the process---that's not an allowable criterion here, as reliability is not a *psychological* variable---it is that that lighting condition alters the input-output functions in a systematic way, i.e., across many different input and output values.¹²

I'm not sure how to make this much more precise, but I have little doubt that any psychologist would claim that light levels make a difference to processing while the distal stimulus' being a dog or cat does not. I won't try to define what counts as a "wide range of subject matters", since I think that some conception of same/different processes is already part of the science, and my suggestion here is that we epistemologists rely on whatever same/different process conception is to be found in the science. I am not engaged in conceptual analysis here but merely trying to clarify what I take the scientific conception to be. Of course, it is not an accidental feature of psychology that it refuses to treat dog perception as involving a different process than cat perception, or more generally, that it refuses to take non-systematic content differences as making a difference to processing, or indicating different processes. As mentioned earlier, the science deliberately dismisses certain differences as insignificant, as part of a larger goal of systematizing and unifying phenomena, as seeing distinct events as exhibiting common patterns and falling under general laws.

Again, what we are looking for are the narrowest process types. Of course, psychology is free to employ more inclusive concepts of process types, so that, say, visual object recognition might count as a psychological process at some coarse level of description---so long as there are

¹¹ This probably isn't the best example. There may be computational steps that are taken in low light conditions that are not taken in high light conditions; certain filling-in or contrast enhancement operations may be required in low light but not high light, thus making for a genuine difference of algorithm across the two conditions. The point, however, is that I want certain values of the variables to count as making for different process types, even in cases where the same algorithm is in fact used.

¹² This is one major difference between the current proposal and that of Beebe 2004, which invokes reliability considerations to make process distinctions finer than the algorithmic level will allow (though he does so for very different reasons than I do here). I focus on differences in the input-output function, but there may be other ways in which the values of certain variables have systematic effects, e.g., by influencing processing speed, etc.

interesting, perhaps lawlike, generalizations that can be made about visual object recognition per se---even if there are two or more distinct algorithms for visually recognizing objects. So the proposal is this:

x and y are tokens of the same (narrowest) process type iff x and y execute the same algorithm, where the values of the variables x and y range over, if different, don't result in any systematic differences in processing.¹³

This is somewhat vague, as the required systematicity is explicitly the sort of thing that could come in degrees. But this imprecision shouldn't impugn the whole project here, for it is merely the imprecision endemic to and inherited from the scientific conception of psychological processes. Sciences don't normally explicitly define their kinds in any significant way, and a certain amount of vagueness is, I think, to be expected of scientific kinds (witness *species*, *planet*). The present case seems rather typical in this regard; it is certainly typical of the kinds involved in the cognitive sciences. The fact that it is easy to determine---without taking any epistemological considerations into view---whether the required systematicity obtains in a given case shows that the conception is not unworkably vague. This may not be math, but it's not poetry either. Our desideratum was that we find a principled way of individuating processes, not that we find a mathematically precise way. If this degree of precision is good enough for science, it should be good enough for philosophy.

IV. Assessing the Psychological Criterion

I have offered a psychological criterion for the individuation of cognitive processes for epistemological purposes. But why believe that this is the---or even a---good way to individuate processes? I think there are several reasons.

1. It is principled. Given a pretty modest assumption of psychological realism, the psychological criterion (PC) gives us a single process type for any given belief. Because the typing is done for other purposes, without any epistemological ramifications in mind, it is clear that we can't be cherry-picking processes to suit our intuitive epistemological judgments. We might do as much in (mis)applying PC, but the typing itself is independent of our epistemic concerns.
2. It is relatively clear. I conceded above that there is a certain amount of vagueness, but I contend that it is only the standard vagueness one finds in scientific taxonomies more generally.
3. It integrates our epistemology with our science. This wasn't one of the original desiderata, and perhaps only the most scientific among us would have demanded it ab initio. But it's certainly a pleasant bonus whenever a theory we are defending meshing nicely with other things we have some reason to believe are true. One might even argue that an epistemology that needs to make up its own process taxonomy therefore loses out in respect to simplicity to one that can adopt a ready-made taxonomy, but for now, I'll take this merely as a nice perk.

¹³ Distinctions among cognitive processes have little to do with distinctions among cognitive modules. A single process may span several modules; a single module may use several distinct processes. The individuation criteria are of course quite different for processes and for modules.

4. It seems to deliver epistemological verdicts that are at least more or less intuitively correct; it makes pretty much the epistemic distinctions we want made. This, of course, needs extended defense, and the rest of this section will try to make this case.

Note first of all that PC should be acceptable to internalists and externalists alike. Given the substantive psychological commitments assumed here, processes are going to have to be internally individuated. Reliabilists, of course, will claim that the difference between the good and bad processes hinges on something external, but I am proposing PC for process theories more generally, and an internal criterion for process individuation makes this proposal one that internalists can accept.

It is common these days to distinguish between two main varieties of internalism. Access internalism holds that justifying factors must be internal to the agent in the sense of being in some special way epistemically accessible, typically available to introspection or reflection. Mentalism claims that justifying factors must be internal in the sense of supervening on the nonfactive mental states of the agent. When I claim that PC is friendly to internalism, I mean that it is friendly to mentalism, not to access internalism. Whatever access internalists take access to be, it's unlikely that we will always have access to which process produced a given belief. This doesn't really have much to do with PC, however; access internalism doesn't sit well with a process theory, no matter how we individuate processes. Our introspective access to the causes of our beliefs is notoriously poor, and if these causes are epistemically relevant, as a process theory insists, then access internalism was probably off the table anyhow.

Although PC offers an internalist-friendly criterion, there is nothing here that is at odds with externalism. Externalism, again, is primarily a theory about process *goodness*, not process individuation. Reliabilists, including myself, have typically been happy to individuate processes internalistically---they have frequently insisted on it (Goldman 1979, 1986, Alston 1995). There is, however, one kind of externalism that will be left out of the fold. Not all disjunctivists are externalists, but some are, and an externalist disjunctivist will not want to individuate processes internalistically. I discuss this more in section VI below. I won't try to placate disjunctivists, but I don't know that it would have mattered much; as I will argue below, (epistemological) disjunctivists are committed to a state theory anyhow, so how we process theorists individuate processes should be of little concern to them in any event.

Second, PC allows, but does not require, fairly fine-grained epistemological distinctions. If some factor makes a systematic difference to processing, it makes for a different process, so we will have a great number of processes to support a large number of epistemic distinctions. If a visual identification process uses shape information, that affects the process type; if it uses color information instead, or in addition, then these too affect the process type. If lighting conditions influence the psychological processing, then they affect the process type. If factors internal to the process, such as degree of match between stored template and target representation, affect processing, then they affect process type (Goldman 1986, p. 50). If drunkenness or distraction

effect processing, then they affect the process type.¹⁴ Within limits, more distinctions is better, especially since we aren't obligated to use them all.

But I have argued that there are limits, that differences of 'mere content' do not affect process type. This means that, everything else being equal, my typical perceptual identification of dogs uses the same process as my typical perceptual identification of cats. If the justification of a belief is fixed entirely by the nature of the process that produced it, then these beliefs will have to have the same degree of justification (again, assuming all else is equal). We have to say this, even if my environment is such that all my perceptual beliefs about dogs are true and all my perceptual beliefs about cats are false. The prevalence of feline impostors will affect the reliability of the process, but across the board and not only for cats, and by an amount determined by the proportion of the uses of that process that are cat-involving. Some reliabilists might not like this result, but I find it fairly intuitive. If I am in fake barn country, my perceptual beliefs about barns are, intuitively, justified but Gettiered. PC explains this.¹⁵

Third, PC seems to handle the standard cases discussed in connection with the generality problem. Consider a well-known discussion from Conee and Feldman:

The token event sequence in our example of seeing the maple tree is an instance of the following types, among others: visually initiated belief-forming process, process of a retinal image of such-and-such specific characteristics leading to a belief that there is a maple tree nearby, process of relying on a leaf shape to form a tree-classifying judgment, perceptual process of classifying by species a tree located behind a solid obstruction, etc. The number of types is unlimited (1998, p. 2).

It is clear that none of these types is likely to count by PC. The first is too broad. The others are all too narrow, adverting to a particular output content, though sometimes too broad as well, as there are multiple visual processes that could identify trees. The last two come close to getting at the right process: it is a process of shape-based (natural?) object recognition, within certain important parameters; the presence of an occluding object is a systematically relevant factor (although if the interposing object is transparent, then it doesn't count as an *occluding* object), as is the fact that the judgment is being made at a basic level of generality, rather than a subordinate or superordinate level.¹⁶ I am not in a position to guess how many parameters might systematically influence processing, though surely there is a fact of the matter and the number won't be unlimited.

¹⁴ Drunkenness, unlike distraction, is not obviously a *psychological* variable and if not, cannot directly affect process type (in this way, it is like lighting conditions). Clearly, however, blood alcohol content has general psychological consequences (reduced inhibition, reduction of attention span, inhibition of memory consolidation, etc.) which can and do systematically influence processing.

¹⁵ I have claimed (Lyons forthcoming) that indicator reliabilism gets the wrong result in these cases, classifying them as unjustified, since barn experiences are poor indicators of the presence of barns, but that process reliabilism gets the right result, classifying these beliefs as justified. I wasn't able there to argue for this latter clause, but PC provides this missing argument.

¹⁶ Concepts are hierarchically organized into basic, or entry, level concepts (e.g., *horse, dog, car*), superordinate (e.g., *animal, vehicle, furniture*), and subordinate (e.g., *German shepherd, 1962 Chevy Impala*) (Rosch 1978). Which concepts count as entry level for a given subject depend on his or her level of expertise (Tanaka & Taylor 1991). I am treating the current example as involving someone with enough expertise to have *maple* as an entry level concept.

Similarly, Feldman (1985) asks whether we should take the relevant process type to be “the perceptual process, the visual process, processes that occur on Wednesday, processes that lead to true beliefs, etc” (p. 159). Again, it is clear that none of these is a *psychological* process.

Kripke offers this counterexample to reliabilism: *S* blindly believes everything the pope says, and unbeknownst to us all, the pope really is infallible. *S* is then using an infallibly reliable process, even though her beliefs are unjustified. Obviously, this trades on treating *believing the pope* as a process, which it won't be on PC, since matters of content do not generally result in difference of process. There are two interesting points I want to make about this. First, Goldman (1979) rules this out by stipulating that contents don't enter into the individuation of processes, but this stipulation amounts to an ad hoc amendment, not really motivated by the rest of his theory. By contrast, it is a natural, indeed inevitable, result of PC. Second, as has been repeated noticed above, the role of contents is a complicated one on PC. If *S* believes everything the pope says because she believes he is reliable, or because she admires him greatly and implicitly trusts anyone she admires so much, then it is easy to see what the genuinely psychological process is, but it is one that is considerably more general than *believing the pope*. It is *possible*, though empirically unlikely, that the pope is such a special stimulus for *S* that his proclamations trigger the use of a distinct process that doesn't have any other applications, perhaps in something like an extreme version of the way language or faces are special classes of stimulus for normal humans. For such a cognizer *believing the pope* might count as a distinct process.¹⁷

V. Psychology and Organon

When we think commonsensically about belief forming processes, especially in epistemological contexts, there are two main approaches, or sets of guidelines, that come to mind for individuating process types. Sometimes we are inclined to type processes by invoking an internal, psychological vocabulary: ‘seeing an object’, ‘remembering’, ‘deductively reasoning’, ‘guessing’, etc. Other times we are inclined to invoke the sort of vocabulary used in a logic and critical thinking class: ‘constructive dilemma’, ‘argument ad hominem’, ‘hasty generalization’, ‘proper statistical sampling’, etc. The contrast is, to put it a bit crudely, between a *psychology* and an *organon*; it is between a mechanism of belief-formation and a set of advisory rules for pursuing inquiry. It is not obvious that these two approaches can be reconciled, as the latter seems to be concerned often with external features, often with particular contents. In neither case can the “methods” recommended or forbidden by the organon be viewed as psychological processes. *Arguing ad hominem* is no more a psychological process than is *seeing a tree*. A

¹⁷ Whether this is a problem for reliabilism depends on the details of the reliabilist theory in question. On my own view (Lyons 2009), the agent would also need what amounts to a “pope module” for the belief to be justified in the absence of inferential support from other beliefs and would need to be in an environment where popes really are reliable. This does little to show that my real-life dogmatic cousin is highly justified if it turns out that the pope really is, unbeknownst to us all, infallible. For my cousin's beliefs are surely influenced by other psychological factors; if he didn't have such an emotional investment in Catholicism and depend on it so heavily for his sense of identity, he wouldn't believe some of the pope's more implausible utterances. Other versions of reliabilism assess reliability not in the world where the process is used but in the actual world or in “normal” worlds. Whether these views would have trouble here depends on whether the pope is infallible in the actual or normal worlds.

standard account of what makes something an argument *ad hominem* is that it contains certain premises (about the rottenness of one's dialectical opponent), but this is a difference of mere content. *Proper statistical sampling* doesn't count, not only because 'proper' is a normative term, but because 'statistical sampling' does not specify psychological state transitions. The fact that reliabilists have often flitted back and forth between the psychological and organonic vocabularies without apparent awareness of doing so has probably contributed a lot to the sense that processes are being individuated in an *ad hoc*, willy-nilly manner.

I have been developing a taxonomy of cognitive processes along the former, psychological, line in defending PC. The organonic vocabulary does not pick out cognitive processes, by my own account of cognitive processes. Yet epistemologists frequently use this organonic vocabulary to pick out belief-forming processes, to explain why a certain belief is or isn't justified.¹⁸ How can the organon, and the use to which we like to put it, be reconciled with PC?

Consider deduction. Some psychologists believe that we engage in deductive reasoning by using a kind of "inner logic", whereby we reason from premises in a formal, syntactic, rule-governed way, irrespective of the contents of the premises and conclusions (Braine & O'Brien 1998). Most psychologists, however, don't believe this; they think we reason deductively by forming and transforming mental models (Johnson-Laird 2001), or by drawing analogical inferences from known cases (Shastri & Ajjanagadde 1993), or some other heuristic method.¹⁹ Evidence for the various theories consists mainly of the observation that we are better at some deductive problems than at others, which is measured primarily in terms of pervasive and predictable errors.

These non-syntactic theories, which explain our capacity for deductive reasoning in terms of heuristic processes, would be paradoxical if 'deduction' had to pick out a particular kind of cognitive process, since deduction is paradigmatically non-heuristic. Rather, 'deduction' is ambiguous among problem, function, and algorithm: it might mean a certain posed problem that needs to be solved; or it might mean a certain input-output mapping that constitutes a correct solution to the problem; or it might mean a method of effecting that mapping, of deriving certain conclusions from certain premises. The study of deductive reasoning in psychology normally takes 'deduction' in one of the former two senses, as---we might say---a *task level* phenomenon, rather than an algorithmic level phenomenon, and the different theories of deductive reasoning propose different algorithms.²⁰

¹⁸ See, e.g., Goldman's original (1979) discussion: "processes whose belief-outputs would be classed as unjustified [include] . . . confused reasoning, wishful thinking, reliance on emotional attachment, mere hunch or guesswork, and hasty generalization. . . . By contrast, . . . processes [that] are intuitively justification-conferring . . . include standard perceptual processes, remembering, good reasoning, and introspection." (p. xx)

¹⁹ A heuristic method is one that trades efficiency for accuracy; it uses "rules of thumb" to increase problem-solving power and speed, while getting the right answers often enough. Even the syntactic views take the inner logic to differ in important respects from the natural deduction systems codified in logic books.

²⁰ For the most part, the distinction between problem and function doesn't matter much, so long as we keep in mind that our actual confrontations with deductive problems doesn't normally result in our computing deductive functions, but only approximations thereto, given the robust and recurring errors we make. The function is the mapping actually performed, while the problem determines the mapping that in some sense ought to be performed. For simplicity's sake I lump these two together using the appropriately ambiguous term 'task'.

Like other process theorists, I think that the actual way by which the conclusion is reached matters quite a bit to the epistemic status of the belief. Suppose we have a quick-and-dirty, “System 1” (Schneider & Shiffrin 1977, Kahneman & Frederick 2002, Kahneman 2011) process for solving deductive problems that trades some accuracy for speed by using heuristic processes and whose inner workings are a mystery to us as cognizers. In addition, we have taken and taught logic class and have a slow but highly reliable “System 2” process, whereby we deliberately and knowingly apply the rules of logic to the problem at hand. It seems obvious that which of these two processes we are using on a given occasion should affect the justification of the belief. In both cases, we are engaged in deductive reasoning in the task-level sense, although only System 2 uses a deductive algorithm. The *task* is always deduction, but the *belief-forming process* is deductive only when we use System 2.

So here is a first hint of how the organon matches up with the psychology. We often say someone is engaged in deductive reasoning, or make claims about what counts as proper deductive reasoning, without really intending any claims about the actual psychological processes the subject is or should be undergoing. Rather, we are specifying the problem the subject faces and specifying the appropriate input-output mappings. Logic gives us an organon: a set of rules for reasoning properly, but these rules don’t provide a mechanism for computing a function; they only characterize the function to be computed. They tell us what counts as correct reasoning but not how to go about reasoning correctly. The organon doesn’t care *how* you achieve deduction, or some close approximation to it, so long as you *do*.

The organonic rules are thereby framed at a much higher level of abstraction than the cognitive processes. An organon advocates broad classes of cognitive processes that share similar input-output mappings. In this way, the organon is not in competition with the specification of cognitive processes, for the latter attempts to explain how a cognizer does something, and the organon is saying what in a general way it is they’re doing. At the same time, and more importantly, the organon marks some classes of processes as good and some as bad. So the “methods” mentioned in the organon do not count as belief-forming processes according to my psychological criterion. But that’s okay; they’re not supposed to. The organon doesn’t *name* belief-forming processes; it *sorts* them.

Though processes and organonic principles have often been conflated, my claim here is consistent with the use to which the organonic principles are typically put. When we say, for example, “What *S* did just now was epistemically wrong, because it was an instance of the argument ad hominem,” we mean only that whatever actual belief-forming process *S* just used *amounted to* her arguing ad hominem. *Committing the ad hominem fallacy* is a task-level phenomenon, not a process-level phenomenon. That doesn’t mean people don’t do it; it means that they do it by way of doing something else. If arguing ad hominem is epistemically bad, then a number of process-level phenomena are therefore bad: *downgrading or discounting someone’s argument because of a personal dislike for the speaker* is a process-level entity, or at least it’s getting closer. More accurately, there is a certain standard suite of processes for evaluating argumentation, and emotional factors, like dislike for the speaker (or admiration, for that matter), can have systematic effects on these processes. When the emotion affects processing, it plays a role here similar to the role played by occluding objects or poor lighting conditions (rather, their

internal analogues) in the perceptual case; it is a factor that systematically affects processing, so its influence here makes for a distinct process.

Similarly for other rules of critical thinking. The rules tell us to attribute sample characteristics to the larger population only when the sample is representative of the population. Such a rule doesn't pick out any psychological process. Still, we have intuitive, automatic processes for generalizing from a sample, at least some of which processes are not sensitive to the representativeness of the sample. Those of us who have studied inductive generalization have at our disposal a quite different process, the one that we engage when we consciously, deliberately follow the rules for taking into account various kinds of evidence for and against representativeness. The intuitive, System 1 inductive process is clearly psychologically distinct from the deliberative, System 2 inductive process, and the organon tells us that the latter process is epistemically better than the former process.

And similarly for epistemic virtues. Openmindedness is not a psychological process, but certain classes of cognitive processes exemplify openmindedness and others do not. As with ad hominem argument, it is a matter of which factors are causally influencing reasoning. When current inferential processes are affected by the desire to have one's previous commitments stay unrefuted, this desire serves as a systematically causally relevant parameter, again, in roughly the way that the internal analogue of lighting conditions does. Just as a certain visual recognition algorithm with one set of lighting conditions counts as a different process from the same algorithm with different lighting conditions, so too a single inferential algorithm might yield up different processes corresponding to different strengths of influence from the desire to have been right.

I have only addressed a few examples, but I think this is enough to illustrate the general strategy. It has been a fairly serious mistake on the part of reliabilists to confuse the elements of an organon with psychological processes, and it is really the latter that we have in mind when we embrace a process theory. Nevertheless, there is a close relationship between the organon and the processes, i.e., between the task level and the algorithmic level. The organon groups together various algorithmic level processes and classifies them as epistemically good or bad.

VI. Hallucination and Veridical Perception

I want to draw out one interesting consequence of the psychological criterion of process individuation in conjunction with the more general process theory. Besides the intrinsic interest of the consequence, the discussion illustrates the psychological criterion and serves to show that it really does impose substantive constraints. We can't just say anything we want about which process was operative.

Disjunctivists hold that veridical perception and hallucination are two different kinds of state. This is at least sometimes supposed to have epistemological implications; a disjunctivist metaphysics might, for instance, be used as part of an argument for "epistemological disjunctivism" (Byrne & Logue 2008), the view that one's epistemic position is better in a case

of veridical perception than it is in a case of hallucination.²¹ Disjunctivists tend to prefer to talk about knowledge, rather than justification, perhaps because the view is more plausible in this context. Veridical perception is clearly and uncontroversially epistemically superior to hallucination---even *veridical* hallucination---if ‘epistemically’ here means ‘having to do with knowledge’. It is far from obvious, however, that veridical perception is *justificationally* superior to hallucination. This is how I will understand epistemological disjunctivism: not only is the belief false or deviantly caused and therefore at best Gettiered, but the hallucinator doesn’t even have the same evidence that the veridical perceiver has. The hallucinator is therefore not even justified, and that’s why she doesn’t know.²²

This approach to hallucination runs contrary to a dominant and pervasive strand not only in contemporary epistemology, but in the tradition going back at least to Descartes. The traditional view holds not only that there is a common experiential factor in veridical and hallucinatory cases, but that the external causal history of this experience---that which makes it either a veridical or hallucinatory experience---is irrelevant to the justificational status of the belief. This traditional view has some intuitive plausibility, and it is closely connected to the intuitive plausibility of an influential objection to reliabilism, which holds that since the victims of a Cartesian demon are obviously justified in their beliefs, which are nearly all false and thus unreliably caused, reliability and justification come apart. The epistemological disjunctivist of the sort I am envisioning here has to make the same counterintuitive claims about the epistemic status of demonworlders as does the steadfast reliabilist.

The psychological criterion for process individuation suggests a third alternative. It is consistent with the disjunctivist’s claim that hallucination (even though phenomenologically indistinguishable from the real thing) at least sometimes interferes with the justificational status of beliefs (and not merely their status as knowledge). However, I will hold that this is because of the *process* of hallucination, not the *state* of hallucination. Like the traditionalist, and contra metaphysical disjunctivism, I assume that the experiential *state* of hallucination does not differ from a veridical experiential state; however, the *process* of hallucination does, sometimes, differ from the typical veridical experiential process. And this is where the epistemic difference lies.²³ The traditionalist is correct in holding that the external causal history of the experience is epistemically irrelevant---the external causal history cannot affect process type, according to PC. However, the *internal* causal history of a belief is highly relevant to its epistemic status, and this sometime differs between cases of hallucination and veridical perception.

Veridical perception is, of course, driven by the bombardment of the sense organs by various types of energy and is at least in large measure a *data driven* affair. There are active debates, of course, about how much, if any, perceptual processing is top-down, or conceptually driven (Raftopoulos 2005), but no one denies that the activation of the primary sensory cortices,

²¹ I don’t want to try to explicate disjunctivism in any detail here. See Haddock & Macpherson 2008 for more on the view. The contrast is with the view that, vaguely put, there is something directly present to the mind in the case of veridical perception, and this same thing is directly present to the mind in hallucination as well. I’ll leave illusion out for now.

²² McDowell 1998 seems to hold this view, though he’s less explicit than one would like.

²³ I am leaving behind concerns about knowledge, so when I use ‘epistemic’ I will mean ‘justificational’.

and corresponding tokenings of low level perceptual representations, have a large role to play in ordinary, veridical perception.

As philosophers understand hallucination, however, it is at least sometimes an *highly* top-down affair; Macbeth's dagger has a conceptual origin, rather than an origin in low-level perceptual processing.²⁴ This means that, even though the distal cause is irrelevant to process individuation, according to PC, there will still be a difference of process between hallucination and veridical perception---not because of the veridicality per se, but because veridical perception tends to use a different (data driven) algorithm than the (top-down) algorithm involved in hallucination.

Not everything we might want to call hallucination has this sort of conceptual origin. One could experience a hallucination as the result of a data driven process exactly like that of normal perception. In fact, I think this is how we are to understand the standard skeptical scenarios: the mad neuroscientists are tricking us by stimulating our *primary* sensory cortices; we aren't tricking ourselves in the way that Macbeth is. I don't know how dream imagery works (and know even less about how it would work were it to be experientially indistinguishable from normal perception), but I suspect that much of the activation we see of the primary sensory areas (Hobson & Pace-Schott 2002) is top-down and thus involves a different process from normal perception.

The epistemological implications of this will depend on what kind of process theory one endorses. A process theory per se is neutral with respect to which processes deliver what degree of justification. Depending on what it is about a process that determines its associated degree of justification, a process theory is not committed to finding an epistemic difference every time there is a process difference, only the other way around. A process theorist could thus admit that hallucination typically involves a different belief-forming process from normal perception, while attributing the same degree of justification to the respective beliefs.

Reliabilism, however, is constrained to treat processes epistemically differently if they have different reliabilities, and since the kind of top-down process involved in hallucination is presumably a lot less reliable than the bottom-up process involved in normal perception, the former beliefs will be a lot less justified than the latter. Whether this result transfers to the brain-in-the-vat case will depend on how the reliabilist individuates environments. If being envatted makes for a new environment, then the beliefs will be unjustified; if not, they will retain their original degree of justification. Dreams will probably go the way of hallucination; the kind of process involved is probably pretty much restricted to dreaming, and if dreams produce beliefs, these processes will be unreliable and the beliefs unjustified. The reliabilist may or may not take these implications to be problematic for her theory. Some people find epistemological disjunctivism to be an intuitively plausible view; I doubt they would be bothered by the implications just discussed. Some reliabilists are not bothered by the fact that their view implies that demonworld victims are unjustified in their beliefs (Lyons 2012); they may not be bothered by these implications either. If demonworlders are unjustified, why make a special exception for dreamers?

²⁴ I am not aware of any relevant science on this sort of thing, since the kind of hallucination that philosophers are concerned with---where things are sense-experientially exactly as they would be if veridical---never happens in real life.

Epistemological disjunctivism distinguishes hallucination from normal perception on external grounds: a veridical perception is simply one where I'm standing in a particular relation to a distal object, while hallucination is an introspectively indistinguishable state where I'm not related to any such object. By dividing the territory this way, disjunctivism fails to distinguish relevantly different processes; Macbeth-style top-down hallucinations have the same status as hallucinations caused by the direct stimulation of the primary sensory cortices, because they're both hallucinations, and hallucinations have a special epistemic status. Depending on the details of the theory (and these really do have to do with the generality problem---see section II above) the reliabilist may be able to claim that the latter beliefs are justified.

How implausible is it that reliabilism would be committed to claiming that hallucination (of a certain sort) yields unjustified beliefs? Is this something that only a disjunctivist could swallow? I don't think so. I have no sympathy for disjunctivism but quite a lot of sympathy for the view that some hallucinatory beliefs are therefore unjustified.

Macbeth-style hallucination is endogenous: it starts with certain desires or beliefs and produces sense-experiences on their basis. Brain-in-a-vat hallucination (which is so different that we tend not to call it hallucination) is exogenous: it starts with stimulation of the primary sensory areas, or with the sense organs. For all we care, it could start with holograms. Let us put aside the question of exogenous hallucination and ask about endogenous hallucination. How absurd is it to hold that the resulting beliefs are unjustified?

Let's approach from afar. The belief bias (Evans et al 1983) is the phenomenon whereby your believing the conclusion of some argument makes you more likely to judge the argument to be valid. More generally, the cogency of an argument is decided in part by prior beliefs about the truth of the conclusion. Some epistemologists (e.g., Huemer 2007) hold that we make such judgments on the basis of intellectual seemings, or appearances, nondoxastic states that play the same epistemic role as nondoxastic sensory seemings, or appearances. The belief bias seems clearly fallacious. It is very plausible---though not uncontroversial---that being influenced in this way is epistemically illegitimate, that the assessment of validity is unjustified if the apparent cogency results from prior belief in this way. If this assessment is correct, then it is not merely the nature of the intellectual seeming that determines whether it confers justification, but the source of the seeming as well. In particular, those seemings that arise endogenously out of antecedent cognitive biases do not carry the same epistemic weight as do seemings that do not arise thus.

Now turn to the perceptual case. Suppose that extreme zealotry and wishful thinking influence my perceptual experiences in such a way that what would have otherwise been random meaningless patterns in clouds or the surface of a grilled cheese sandwich now take on for me the shape of the Virgin Mary; I really want to see the Virgin Mary, and I believe wholeheartedly---though for no good reason---that she's trying to reveal herself to me, and this belief and desire conspire to influence my visual experience, which I take as evidence for the belief that the Virgin Mary is revealing herself to me. Intuitively, this belief is unjustified, and it doesn't serve as confirmation for the belief that she wanted to reveal herself to me, especially since this belief was part of the cause of the experience in the first place (Siegel 2011, Lyons 2011). This is illusion, rather than hallucination, but it is a case of endogenous effects on

perceptual experience, and a case where that endogenous influence intuitively renders the belief based on that experience unjustified.

But if wishful thinking, in conjunction with unjustified expectations, can render an illusory experience epistemically inert (i.e., incapable of conferring justification), why couldn't, say, terror do the same to a hallucinatory experience? That is, if we have already agreed that the causal history of an experience can affect that experience's ability to justify beliefs, then there shouldn't be any automatic resistance to the idea that some hallucinatory beliefs might be unjustified, especially when we remember that it is not the hallucinatoriness per se that is blocking justification, but the fact that the experience is endogenously generated. No, the hallucinator doesn't realize she's hallucinating, but the wishful thinker generally doesn't realize that that's what she's doing either.

So process theories can draw epistemic distinctions between (some kinds of) hallucination and veridical perception. Reliabilist versions of the process theory *must* draw these distinctions, at least if they endorse the psychological criterion of process individuation developed here. This yields an interesting, and I think plausible, alternative that splits some of the difference between disjunctivism and the traditional view.

VIII. Conclusion

I have argued for a psychological criterion for individuating belief-forming processes for epistemological purposes. That is, we turn over to empirical psychology the question of which processes are of the same type. General theoretical considerations about the generalization-stating aims of science indicate that the psychology will not individuate the types too narrowly, and I have argued that if we consider relevant parameters in addition to identity of algorithm, the psychology will not individuate them too coarsely. In addition, I think it is a virtue of the theory that it integrates our epistemology with our science. Although I insist that a theory of processes does not offer a solution to the generality problem, it does solve part of that problem and could potentially bring light to a number of discussions of process theories, especially reliabilism. An example of such a discussion is the case of hallucination and veridical perception. If we agree that processes matter to epistemology and we accept the psychological criterion for process individuation, then the options regarding what we can or must say about the various kinds of hallucination, illusion and veridical perception, are clarified.

Much of what I say about processes is based on a contentious and unargued assumption about the overall contours of empirical psychology. I assume, for instance, that neoGibsonian movements are dead ends I would be delighted to make that argument elsewhere and would be thrilled if more epistemologists would also take more interest in such things. But for here, it is enough to show that (a) how psychology turns out might have important consequences for epistemology, and (b) at least one plausible view about that matter yields epistemological consequences that are independently acceptable.

References:

- Alston, W. P. (1988). An internalist externalism. *Synthese* 74, 265-283.
- (1995). How to think about reliability. *Philosophical Topics* 23, 1–29.
- Baehr, J. (2011). *The Inquiring Mind*. Oxford: Oxford University Press.
- Beebe, J. (2004). The generality problem, statistical relevance and the tri-level hypothesis. *Nous* 38, 177-95.
- Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review* 94, 115-147.
- Braine, M. D. S, & O'Brien, D. P., eds. (1998). *Mental Logic*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Burge, T. (2010). *Origins of Objectivity*. Oxford: Oxford University Press.
- Byrne, A. & H. Logue (2008). Either/or. In A. Haddock & F Macpherson, eds. *Disjunctivism: Perception, Action, Knowledge*. Oxford: Oxford University Press. pp. 57-94.
- Conee, E. & R. Feldman (1998). The generality problem for reliabilism. *Philosophical Studies* 89, 1–29.
- (2004). *Evidentialism: Essays in epistemology*. Oxford: Oxford University Press.
- Cosmides, L. & J. Tooby (1992). Cognitive adaptations for social exchange. In J. Barkow, L. Cosmides, & J. Tooby, eds. *The Adapted Mind*. Oxford: Oxford University Press.
- Cowan, N. (2001). The magical number 4 in short-term memory: a reconsideration of mental storage capacity. *Behavioral and Brain Sciences* 24, 87-114.
- Evans, J., J. L. Barston & P. Pollard (1983). On the conflict between logic and belief in syllogistic reasoning. *Memory and Cognition* 11, 285-306.
- Feldman, R. (1985). Reliability and justification. *The Monist* 68, 159-174.
- Fodor, J. A. (1974). Special sciences (Or: the disunity of science as a working hypothesis). *Synthese* 28, 97-115.
- Gerlach, C. (2009). Category-specificity in visual object recognition. *Cognition* 111, 281-301.
- Goldberg, S. (2010). *Relying on Others*. Oxford: Oxford University Press.

- Goldman, A. I. (1979). What is justified belief? In George Pappas, ed. *Justification and Knowledge*. Dordrecht: Reidel.
- Goldman, A. I. (1986). *Epistemology and Cognition*. Cambridge: Harvard UP.
- Goldman. (2009). Williamson on knowledge and evidence. In D. Pritchard and P. Greenough, eds., *Williamson on Knowledge*, pp. 73-91. Oxford: Oxford University Press.
- (2011). Commentary on Jack Lyons's *Perception and Basic Beliefs*. *Philosophical Studies* 153, 457--66.
- Haddock, A. & F. Macpherson, eds. (2008). *Disjunctivism: Perception, Action, Knowledge*. Oxford: Oxford University Press.
- Hobson, J. A. & E. F. Pace-Schott (2002). The cognitive neuroscience of sleep: Neuronal systems, consciousness and learning. *Nature Reviews Neuroscience* 3, 679-93.
- Huemer, M. (2007). Compassionate phenomenal conservatism. *Philosophy and Phenomenological Research* 74, 30-55.
- Johnson-Laird, P. N. (2001). Mental models and deduction. *Trends in Cognitive Science* 5, 434-42.
- Kahneman, D. (2011). *Thinking, Fast and Slow*. New York, NY: Farrar, Straus & Giroux.
- Kahneman D. & S. Frederick (2002). "Representativeness revisited: Attribute substitution in intuitive judgement". In T. Gilovich, D. Griffin, D. Kahneman, eds., *Heuristics and Biases: The Psychology of Intuitive Judgment*, pp. 49–81. Cambridge, UK: Cambridge University Press.
- Lyons, J. C. (2009). *Perception and basic beliefs*. New York: Oxford University Press.
- (2011). "Circularity, reliability, and the cognitive penetrability of perception". *Philosophical Issues* 21, 289-311.
- (2012). Should reliabilists be worried about demon worlds? *Philosophy and Phenomenological Research*.
- (forthcoming). Goldman on evidence and reliability. In H. Kornblith and B. McLaughlin, eds., *Goldman and His Critics*. Malden, MA: Blackwell.

- McDowell, J. (1998). Criteria, defeasibility, and knowledge. In J. McDowell, *Meaning, Knowledge, and Reality*. Cambridge, MA: Harvard University Press.
- Miller, G. A. (1956). The magical number seven plus or minus two: some limits on our capacity for processing information. *Psychological Review* 63, 81–97.
- Noë, A. (2004). *Action in Perception*. Cambridge, MA: MIT Press.
- Pollock, J. (1986). *Contemporary Theories of Knowledge*. Savage, MD: Rowman & Littlefield.
- Pollock, J. L., & J. L. Cruz (1999). *Contemporary Theories of Knowledge*, 2nd ed. Lanham, MD: Rowman and Littlefield.
- Pylyshyn, Z. W. (1984). *Computation and Cognition*. Cambridge, MA: MIT Press.
- Raftopoulos, A. (2005). *Cognitive Penetrability of Perception: Attention, Action, Strategies, and Bottom-Up Constraints*. New York: Nova Science.
- Rosch, E. (1978). Principles of categorization. In E. Rosch & B.B. Lloyd, eds., *Cognition and Categorization*, pp. 27–48. Hillsdale, NY: Lawrence Erlbaum Associates.
- Schneider, W. & R. M. Shiffrin (1977). Controlled and automatic human information processing I: Detection, search and attention. *Psychological Review* 84, 1–66.
- Senor, T. D. (1993). Internalistic foundationalism and the justification of memory belief. *Synthese* 94, 353-76.
- Siegel, S. (2011). Cognitive penetrability and perceptual justification. *Nous* 46, 201-22.
- Siegel (forthcoming). The epistemic impact of the etiology of experience. *Philosophical Studies*.
- Shastri, L. & V. Ajjanagadde (1993). From simple associations to systematic reasoning: A connectionist representation of rules, variables and dynamic bindings using temporal synchrony. *Behavioral and Brain Sciences* 16, 417-494.
- Sosa, E. (2007). *A Virtue Epistemology*. Oxford: Oxford University Press.
- Strat, T. M. & M. A. Fischler (1991). Natural object recognition: A theoretical framework and its implementation. *Proceedings of the International Joint Conference on Artificial Intelligence*.
- Tanaka, J. W. & M. Taylor (1991). Object categories and expertise: Is the basic level in the eye of the beholder? *Cognitive Psychology* 23, 457–82.