

7. Inferential Abilities And Common Epistemic Goods

Abrol Fairweather & Carlos Montemayor

While the situationist challenge has been prominent in philosophical literature in ethics for over a decade, only recently has it been extended to virtue epistemology¹. Alfano argues that *virtue epistemology* is shown to be empirically inadequate in light of a wide range of results in social psychology, essentially succumbing to the same argument as virtue ethics. We argue that this meeting of the twain between virtue epistemology and social psychology in no way signals the end of virtue epistemology, but is rather a boon to *naturalized virtue epistemology*. We use Gerd Gigerenzer's models for bounded rationality (2011) to present a persuasive line of defense for virtue epistemology, and consider prospects for a naturalized virtue epistemology that is supported by current research in psychology.

Keywords: inference, reliabilism, bounded rationality, naturalized virtue epistemology, situationism

1. *From Moral to Epistemic Situationism*

A wide range of psychological research on trait attribution and rationality has chipped away at what appeared to be a solid empirical footing for virtue ethics, thereby challenging the adequacy of virtue ethics on the very point that appeared to be a primary strength². Philosophers such as Gilbert Harman (2000) have been led to question the very existence of character traits, and others like John Doris (2003) have denied their robustness and explanatory value. Character trait attributions enjoin predictive and explanatory commitments that simply fail too often to meet norms of epistemic success that require the manifestation of epistemic virtues. Doris(2002) argues that traditional forms of virtue ethics cannot be empirically adequate and normatively adequate at the same time. Only recently has the situationist challenge been applied to virtue epistemology. In two recent papers, Mark Alfano (2012) defends the first thorough application of situationism to responsibilist virtue epistemology, and (2013, this volume) develops the first thorough challenge to virtue reliabilism. Alfano's challenge to virtue reliabilism is based a diverse range of empirical results in social psychology on rationality, inferential abilities and trait attribution.

Alfano nicely frames the psychological challenge to virtue epistemology as an inconsistent triad:

¹ See also forthcoming work on this issue by Doris and Olin, Heather Battaly, Christian Miller in *Naturalizing Epistemic Virtue*. (Fairweather & Flanagan eds.)

² The literature on relevant research is enormous, for some comprehensive treatments see Doris (2002), Miller (2008), Alfano (2011, 2012, and forthcoming)

a) *inferential non-skepticism*: most people know quite a bit inferentially

b) *inferential reliabilism*: inferential knowledge is true belief acquired and retained through inferential reliabilist intellectual virtue

c) *inferential cognitive situationism*: people acquire and retain most of their inferential beliefs through heuristics rather than intellectual virtues.

The dilemma for the virtue epistemologist is that empirical adequacy will require accommodating the empirical work presented by situationists and thus will have to accept (c). But if this empirical research shows that all too rarely will an agent meet virtue theoretic standards for epistemic success, we are now unable to account for the robustness of knowledge affirmed in (a). Failing to meet the non-skepticism desiderata would be a *normative inadequacy* in virtue epistemology because any such theory will be unable to assign positive epistemic standings to actual beliefs in a way that keeps pace with the actual frequency of human knowledge. Alternatively, if a virtue epistemologist crafts norms that assure meeting (a), such an account will now fall prey to *empirical inadequacy* because no such account will be supported by research in psychology presented in support of (c). Alfano argues that virtue epistemology must be the discarded commitment. Call this the challenge of *epistemic situationism*.

This essay will challenge Alfano's argument on many points and will defend a naturalistic account of reliable inferential abilities that will not only meet Alfano's challenge virtue reliabilism but will also illuminate the nature and norms of inference, rationality and assertion. Key to this defense will be the Bounded Rationality (BR) research program started by Herbert Simon and recently developed by Gerd Gigerenzer according to which fast and frugal heuristic reasoning often outperforms optimizing rationality for bounded rather rational agents. Gigerenzer (2011) argues that norms of rationality cannot be reduced to assessments of ideal-approximation, or how closely an agent approximates an ideal rational outcome. If Gigerenzer is right here, situationists are applying the wrong kind of norms in their interpretation of the research on rationality.

We take a closer look at Alfano's triad in Section 2 and then get right to Gigerenzer's account of "ecological rationality" in Section 3. Section 4 defends a number of reliable inferential abilities that are supported by relevant empirical work on knowledge of syntax, communication, assertion and directed memory and Section 5 examines Ernest Sosa's virtue reliabilism in light of the situationist challenge and concludes with a novel proposal for rigidifying the normal conditions for epistemic assessment to the (psychologically normal) conditions for assertion. We argue that an empirically grounded account of normal conditions for epistemic assessment can be provided by work examining work on the psychological and semantic processes involved in assertion.

2. A Closer Look at Inference

Since Alfano restricts his argument to inferential virtue reliabilism and not perceptual virtue reliabilism, it will be important to get clear on what is meant by inference. Alfano should be able to show that the same considerations that count against extending his argument to perceptual virtue reliabilism do not also count against extending his argument to perception like forms of inference, and, independent of the specifics of Alfano's argument, it will be necessary to get as clear as possible on the nature of inference to assess the merits of any argument against inferential reliabilism.

Graham on association and inference

In "Psychological Capacity and Positive Epistemic Standing" (2013), Peter Graham distinguishes a number of distinct abilities and capacities which underlie different kinds of positive epistemic standings. These include critical reasoning, propositional thinking, perceptual representation and sensory registration. Graham claims that only the latter is a genuinely inferential capacity, though certain forms of association can be easily confused with reasoning and inferring. On the distinction between associating and inferring, he says:

"Associating is one 'intelligent' capacity for learning about and navigating one's environment. And it is widely thought that no matter how much representation actually goes on in animals when associating, associating isn't reasoning or relying on inference. Logical reasoning is a different kind of psycho-logical capacity." (2103, pg. 154)

Of particular interest is Graham's discussion of research by Premack and Call on inference in apes and chimpanzees. In full view of chimpanzees, researchers took two boxes and placed an apple in one and a banana in the other, and then proceeded to eat the banana out of one of the boxes, again in full view. When given the opportunity to pick from the boxes, the chimps went right to the other box containing the apple. Graham explains that the researchers concluded

"That the chimps reasoned something like this: there is an apple in box *A* and a banana in box *B*. But there is no longer a banana in *B*, so there's just an apple in *A*. That's why they went right for *A*. Animals that don't reason like this, but presented with the same information, might still look for a banana in box *B*, or might only slowly make their way to box *A*."

To show that the chimps were inferring and not just associating, Joseph Call's

experiments with apes involved putting two opaque cups in full view, one full of food, the other empty, and shook both in front of the subjects. If the cup with food was shaken, the apes went right for it. If the cup without food was shaken, they went right for the other one. Call concluded that apes were reasoning something like: 'when there's no noise, there's no food in the shaking cup, so grab the other one'.

This is particularly interesting because similar research with dogs showed that they rely on associative intelligence rather than "logical guidance", reasoning or inference. With dogs searching for a ball placed behind one of three screens, the speed of their search would slow down as they went from 1 to 2, and again from 2 to 3, whereas with children performing a similar task the speed of their search will increase. This explained by the fact that inference in the child shows the failure to find the target in the first attempt as making it more likely that it is behind 2 or 3, and if not 2 then definitely 3. When the dog responds to the failure to find the ball as an 'extinction trial' (Graham, *ibid*) that signals it is less rather than more likely to be found in screens 2 and 3. Children exhibit the kind of inferential intelligence attributed to apes and chimps rather than the associative intelligence of the dogs.

The best explanation of these findings is that very robust, stable inferential abilities exist and provide the basis for distinguishing inferential intelligence from merely associative intelligence. Graham distinguishes the inferential reasoning of chimps, apes and children from the full blown "critical reasoning" exhibited by most adult humans, but the existence of basic inferential abilities is sufficient for our purposes here. We argue that the stability of inferential abilities for many epistemic tasks is confirmed by extensive psychological findings as well as theoretical considerations.

In "Epistemic Virtues and Cognitive Dispositions", Henderson and Horgan (2009) distinguish between *classically inferential processes* and *inferential processes broadly construed*. The former are much more restricted kinds of processes and, they argue, have been the dominant focus of epistemology since the modern period. Inferential processes broadly construed:

"are simply those cognitive processes in which beliefs are formed or maintained on the basis of the information. Being based on information is a causal notion, pointing to arrays of counterfactual dependencies and to dispositions. This is the broadest and most tolerant notion of an inferential process"

Classically inferential processes are restricted by two additional things, "the information figuring in the inference is *explicitly represented* in the cognitive system that is the agent....Second, the causal processes whereby beliefs are fixed (formed, revised, or retained) must be *occurrently isomorphic* with the deductive and inductive support relations obtaining between the information that the agent

possesses.” (ibid.) Cases where content that is occurrently not represented (and thus not classically inferential) but is nonetheless causally salient in belief formation will be broadly but not classically inferential, and beliefs formed on the basis of perception rather than on the basis of other beliefs may be a large and interesting example of such cases³.

Henderson and Horgan argue that properly recognizing the distinction between classical and broad inferential processes supports epistemic virtue theory. In particular, they argue that epistemic virtues have the right dispositional structure to support a theory of broad and classical inference, and thus “ a superior epistemological perspective will give significant attention to virtues—to epistemically good dispositions.” Graham’s distinction between associative and inferential intelligence is not essentially about what content is or is not represented as in Henderson and Horgan’s distinction between classical and broad inferential processes, but this just further demonstrates the heterogeneity of inference and the range of human cognitive activity where stable and reliable inferential dispositions show up. Their distinction between classical inferential processes and inferential processes broadly construed will also have some resonance with Gigerenzer’s distinction between optimizing rationality and ecological rationality discussed at length below.

3. A Closer Look At The Triad

Returning to Alfano, the above discussion shows that inferential virtue reliabilism can rely on the existence of cognitive capacities between mere association and full blown critical reasoning (Graham), or on cognitive dispositions that are broad and non-representational rather than classical (Henderson & Horgan). In both cases, we have empirical support not only for distinguishing kinds of inferential abilities, but also for the claim that some of these are clearly stable, robust and reliable in human beings. Below, we take a much closer look at the three principles in Alfano’s triad before moving to our defense in the following section.

A. *Non-Skepticism*: “People know quite a bit through inference”. Alfano follows Moore in advancing an optimistic intuition about the frequency of knowledge in human cognition. While we are in broad agreement that “people know quite a bit through inference”, (NS) contains an important ambiguity. Clearly enough, an adequate theory of knowledge must ‘get it roughly right’ about how often and when to attribute knowledge and other valuable epistemic achievements to actual human believers. However, it is important to note that a theory can fail to satisfy (NS) in at

³ Henderson and Horgan distinguish a third form called “argumentative inference” to cover cases where the isomorphism between logical relations and causal relations within the information represented by the agent fails to hold.

least two different ways; a theory might *over-attribute failure* or *under-attribute success*. We will see that, depending on how read the normal conditions constraint on abilities and virtues, some beliefs will be neither epistemic successes nor epistemic failures, and this might be used to show that virtue epistemology does not violate (NS) even if we grant Alfano's *inferential reliabilism* and *inferential cognitive situationism*. This will turn out to be an important distinction when we examine Ernest Sosa's virtue reliabilism and will ultimately require looking more closely at 'normal conditions' requirements in virtue theory to determine how to treat these cases.

Local skepticism: While (NS) is an anti-skeptical intuition, it is fully compatible with "local skepticism". It may be true that people know quite a bit through inference all told, but nonetheless perform very poorly within certain domains of inquiry. Cognitively limited creatures using fallible methods of inquiry will be expected to have certain dark areas in their full set of beliefs, even when they reliable in their actual inferential practices all told. If the research he relies on only supports local virtue theoretic failures then situationism will not be enough to push virtue epistemology to skepticism. As shown below, some cognitive failures actually entail broader cognitive success. The response available here to virtue epistemologists is to broadly individuate epistemic abilities being evaluated in virtue epistemology, and narrowly individuate the failures shown in the empirical research.

Alfano presents (NS) as a widely shared intuition about knowledge. True enough, but, just as clearly, it is true that our untutored intuitions about knowledge might be slightly off the mark in any number of ways. In particular, our intuitions about the frequency of doxastic success may turn out to far outstrip the actual frequency of doxastic success. So, we add that if given sufficient reason, (NS) can be revised down. Here is one reason to revise (NS) downward. If the best interpretation of the empirical data implies that *any plausible theory of knowledge* will violate (NS), then there is good reason to revise (NS) downward. If any plausible epistemic theory will fall victim to (NS) in the face of the situationist's empirical results, then this cannot be a special problem for virtue epistemology in particular. Thus, virtue epistemology must be shown to run afoul of the non-skepticism principle in ways that other plausible epistemic theories do not.

B. Virtue Epistemology: Alfano argues that both responsibilist and reliabilist virtue epistemology are threatened by empirical findings (Alfano 2011, Alfano *forthcoming*). However, it will be difficult to establish both. Consider this: If we grant that Alfano's empirical findings show epistemic irresponsibility and thus succeed against virtue responsibilism, we must assume that the subjects accurately represented the stimulus in the cases they did not manifest the virtue-relevant outcome. This is because, if *S incorrectly* represents a given stimulus that is actually *P* as being *P**, and subsequently fails to achieve a virtue relevant outcome (*O*) by performing action (*A*) on the basis of *P**, *S can still act responsibly* by *A-ing* so long as

the virtue in question would require S to A when P* obtains. This is actually very common – When we say “I see why you would have thought that”, and while we might express disagreement with a conclusion we can also grant cognitive responsibility as intended above. Responsibility also appears consistent with misrepresentation in “new evil demon” cases. Thus, if the failures Alfano cites are to count against the attributability of responsibilist virtues, these very agents must be accurately representing the stimulus conditions. But, it then appears that *virtue reliabilism must be true if virtue responsibilism is shown false in the way Alfano proposes*. Going the other way, if we grant Alfano that virtue reliabilism is shown false, then his argument against virtue responsibilism cannot succeed because we cannot assume that the agents are correctly representing the stimulus.

Perception like inferential abilities: Alfano appears to be aware of this, and thus only targets reliabilism for ‘inferential’ rather than perceptual knowledge. This appears to avoid the dilemma above, since the abilities assumed in representing the stimulus accurately seem to be perceptual. However, this is also problematic. Representing an epistemic environment or cognitive task does not easily reduce to reliably perceiving one’s environment, and perception itself must likely involve very inference-like cognitive actions and abilities. Also, a wide range of research in bounded rationality, language acquisition and assertion shows that there are perception-like forms of inference that are very stable and reliable. Collectively, these research programs make it extremely likely that human beings have basic inferential abilities that are stable and reliable across an impressive range of situations and environments.

C. Inferential Situationism: Alfano reports that psychological research shows that the inferences people actually make employ heuristics rather than optimizing methods of formal logic and probability theory, citing a wide range of studies from Kahneman and Tversky (1973, 2011). However, heuristics as studied in bounded rationality research present a more optimistic story. Simon, Gigerenzer and others take seriously the fact that rationality theory studies a cognitively limited creature and have flourishing research programs that suggest heuristic use is often *more reliable* for a cognitively limited agent than using an optimizing rule. Alfano states that empirical results show that inferential beliefs are typically formed by heuristics rather than intellectual virtue. At a minimum, it must also be shown that heuristics cannot be virtues. In the current context, this will be a question of whether they are reliable. Properly understood, we argue that heuristics can be sources of relevant epistemic success when properly selected in the right environments. Because of some slothfulness involved in the process, it might be difficult to argue that heuristic reasoning is a form of responsibilist epistemic virtue, although there are prospects for a ‘heuristic responsibilism’ in Gigerenzer’s recent work. However, we only aim to defend virtue reliabilism against Alfano, so this will be an independent issue. While the appropriate conditions for using heuristics are very narrow and they can lead to mistakes in reasoning, we suggest that human beings can manifest a certain kind of epistemic virtue through the appropriate use of heuristics. These will be *frugal*

virtues .

3. Gigerenzer's Ecological Rationality, bounded agents and epistemic norms

Initiated by the pioneering work of Herbert Simon (1972), research on bounded rationality takes seriously that the subjects of epistemic evaluations are *cognitively limited*, and that heuristics often play an important role in successful human reasoning. Gerd Gigerenzer (2011) has now developed Simon's early insights into a well developed naturalistic epistemic perspective he calls *ecological rationality*, and has recently presented his research to mainstream epistemologists. Our interest here is to see how Gigerenzer's work provides the basis for an empirically grounded inferential virtue reliabilism that can meet the challenge from epistemic situationism.

Less is more, sometimes

Perhaps the most essential point in the bounded rationality research is that limited cognitive agents will often perform *less reliably* when using an ideal or optimizing epistemic rule than when properly employing fast and frugal heuristics. Gigerenzer illustrates this with the example of an outfielder tracking a fly ball who could potentially mathematically calculate the trajectory of the ball or apply some formal method to determine its future location and a strategy for catching it. Or, they could just keep the ball held fixed at the center of their visual field and keep running. The latter is a far more reliable way to succeed in catching the ball, even though the former would yield more accurate information if allowed to run to completion. In such cases, rational agents should not do what ideal epistemic rules prescribe. This is a very important point. Optimizing rules are not always epistemically normative for limited cognitive agents and cannot fully prescribe what a limited cognitive agent ought to do. We get problematic results about rationality when we lose sight of this, but the results are not surprising when this is kept in mind.

Reinterpreting Linda

Consider Gigerenzer's (2011) interpretation of the Linda case from Kahneman, the main example examined in Alfano (2013). To quickly review the case: Infamously, when asked whether, given a character description of Linda, it is more probable that she is (a) a bank teller or (b) a bank teller and active in the feminist movement, 85% of the subjects answered (b), clearly committing the "conjunction fallacy" and violating basic theorems of probability calculus. Gigerenzer notes that subjects are required to use syntactic, content blind rules of reasoning where the values of the variables are not relevant to getting the answer right, agents do not have additional cues from context or a specific rule to use. However, if one asks 'how many' instead of 'how probable', research shows better results. When asked *how many* out of 100 people that satisfy Linda's description would be bank tellers and *how many* would be bank tellers and active in the feminist movement, subjects' performance significantly improves and they do not commit the conjunction fallacy. This shows

that different framing of logically equivalent information gets very different results, and the framing is thus playing a big role here. Gigerenzer shows that when the same information is presented relative to certain frames, people answer quite rationally. The poor performances that worry Alfano may thus be very *local*, and we agreed above that local failures are consistent with Non-skepticism.

The significance of cognitive limitations can be easily missed. The point is not just that our threshold standards for *approximating ideal epistemic rules* should be informed by facts about cognitive limitations. This is a reasonable enough view, but the stronger implication is that we need to use an entirely different kind of norm. This would be a major shift in normative epistemology away from ‘ideal-approximation assessments’. Adam Morton (2013) puts this point very well: *from the fact that we have an ideal epistemic rule, it does not follow that non-ideal epistemic agents should be evaluated in terms of how closely they approximate the ideal epistemic rule*. Morton argues that ideal-approximation norms are not sufficient instruments for evaluating limited cognitive agents. But what will this other kind of norm look like? How is it different from an optimizing norm? Will it be anything like a virtue?

Heuristics: What are they? Why do we need them?

In “Bounded Rationality: Models for some fast and frugal heuristics”, Arlo-Costa (2013) develops results in formal epistemology that nicely illustrate the structure of heuristic reasoning, with both good and bad results for epistemic normativity. Drawing on Simon’s famous image of “the two great scissor blades of rationality”, *heuristic reasoning implements threshold evaluations for selected criteria that exploit reliable features of task environments rather than performing computations on sets of evidence*. That is not to denigrate optimizing rationality, but rather to emphasize that optimizing rationality and ecological rationality are two distinct and equally legitimate forms of rational response, two equally good and importantly different scissor blades. Traditional thinking about rationality sees it as normatively governed by optimizing norms alone, and is thus monistic in this sense.

Ecological reasoning transitions from threshold and criteria assessments to search and stopping rules, and Arlo-Costa shows that this can be formalized as a reliable type of reasoning. He also suggests a weak and strong reading of ecological rationality, and correctly locates Gigerenzer as advocating the stronger form.

Weak Ecological Rationality: heuristic reasoning can be and often is near optimal when used in appropriate circumstances. Since optimizing norms are nearly approximated, no deep revision in epistemic norms is necessary, we just expand the rational strategies for satisfying them to include heuristic inferential processes and abilities.

Strong Ecological Rationality: It is rational to use heuristics even when doing so goes against the dictates of optimizing rationality. This is normatively

revisionist compared to traditional conceptions of rationality. This is Gigerenzer's stated view.

If strong ER forces significant revision to epistemic norms, some virtue epistemologists will approach bounded rationality cautiously⁴, while others might embrace Gigerenzer as a fellow epistemic revisionary. The issue raised here is over the tenability of a certain kind of epistemic value monism. In epistemology, value monism typically refers to something like Goldman's "t-value monism"⁵, which claims that truth and the reliable means to it will be the sole values in normative epistemology. Many have argued for dualist or pluralist accounts of epistemic value in the literature on the "value problem", but the question above is over the scope of optimizing norms rather than whether truth is the sole epistemic value. Are optimizing norms sufficient to provide evaluation and/or guidance for the full range of cognitive tasks and achievements relevant to epistemology?

Weak heuristic rationality does not give up the truth goal, but constrains norms for success around the *cognitive limits* imposed on real world decision-making. Yet, even weak ER is at odds with optimizing rationality in the *content of guidance norms*, since rational agents will be instructed to do very different things when they are being ecologically rational than when they are being optimally rational. Weak ER thus preserves a fundamental commitment to truth, but will have a different prescriptive content than many traditional epistemic theories.

Ecological virtues?

Gigerenzer (2011) proposes a novel 'ecological definition of terms' according to which elements of an epistemic theory will actually represent complex relations between an organism and its environment, rather than properties seated entirely in either. This is a significant move regarding the ontology of epistemology. The agent is now just part of a broader epistemic ecology, and this ecology is the fundamental unit of analysis for evaluating human rationality. However, *this move potentially conflicts with virtue epistemology's emphasis on the agent*, and thus may be metaphysically (rather than normatively) revisionist. Virtue epistemology is defined as being 'agent based' rather than 'belief based', but it is not clear that Gigerenzian 'ecological virtues' would still be agent-based enough and in the right way to properly constitute what virtue epistemologists have in mind⁶. Many virtue epistemologists rely on some form of *agent-based credit for success* to both answer the value problem and respond to a range of problems related to epistemic luck, including Gettier Problems (Greco, Sosa, Zagzebski, Turri). Person-level abilities are

⁴ Greco and Pritchard both clearly endorse the "traditional epistemic project", while others like Zagzebski and Roberts and Woods (2007) and Axtell are more revisionist in how they see the epistemic project, but not in how they see they epistemic virtues.

⁵ See Goldman 2001, Pritchard 2007.

⁶ See discussions of "the direction of analysis" in Greco (2009), Blackburn (2001)

an important arrow in the quiver and this will have to be worked out in any virtue theoretic formulation of ecological rationality.

On the other hand, perhaps virtue theory has always been understood ecologically. Dispositions have the very ecological structure Gigerenzer refers to because even the most robust dispositions like fragility and solubility will only manifest with the help of “reciprocal causal partners”¹. Abilities and dispositions are also sensitive to environmental cues through “normal conditions” requirements that account for relevant forms of masking and mimicking, only some of which will imply agent culpability for lack of success. If disposition theory can bring virtue epistemology and ecological rationality together in a single account of reliable inferential abilities, virtue epistemologists will have at least the basis of a powerful empirical response to the situationist’s empirical challenge. The social dimensions involved in cultivating virtues are also ecological in the relevant sense. Since virtues most likely have an ecological structure to begin with, accommodating Gigerenzer’s research will not require metaphysical revisionism in virtue epistemology.

4. Knowledge of Syntax, Directed Memory and basic inferential abilities

Knowledge of Syntax : Heuristics are not the only inferential ability supported by research in psychology. Research on generative grammar, language acquisition and communication all show that human beings have very stable and robust inferential abilities, though perhaps these are *basic inferential abilities* compared to the higher order calculations that Alfano appears to be concerned with⁷.

Basic inferential abilities are critically involved in acquiring the lexicon and generative rules of a language. Knowledge of syntax requires the manipulation of information according to strictly formal rules. Children have epistemic skills that allow them to learn any language based on these rules and their modal robustness is extraordinary. A vast amount of research in neuroscience and linguistics aims at explaining this robustness⁸. Specifically, scientists have tried to understand how it is possible for infants to learn a language given the incredibly diverse contexts they are in, the impoverished stimuli they are exposed to, the complexity of the grammatical rules etc. Despite there being many open questions, it is clear that some kind of inferential abilities are essential to language acquisition and, like perceptual skills involved, are remarkably stable across different situation types and

⁷ See Bach, Montemayor and Proust for accounts of basic action that may be amenable to a theory of basic abilities. If there are basic actions, there are very likely abilities to cause the actions. These abilities might themselves be inferential even if the basic action is not itself an inference.

⁸ See Jackendoff (2003), Chomsky (1986 and 1987) and Hornstein (1984).

individual differences. Although knowledge of syntax is highly formal, humans manifest such knowledge at a very early age, and they do so reliably and without conscious effort or monitoring. Infants do not need classes of universal grammar and rules of syntax in order to distinguish the syntactic components of (in many cases poorly constructed) utterances of a language. They are certainly not introspecting on these rules, or accessing evidence that could justify them to parse an utterance in terms of subject and predicate. What the infant is doing is highly complex, but the infant performs this incredible epistemic task in a *perception-like fashion*. Widely accepted results in linguistics and cognitive science also show that there is something inferential going on.

Notice that Gigerenzer makes it much easier to see this case as inferential but not optimizing or computational. Without claiming too much, we can at least say that the kind of inferential ability the child is manifesting will be something like the inferential abilities manifested in Gigerenzer's account of ecological reasoning. The inferential abilities involved in language acquisition may well be among the most basic capacities that heuristics feed on. But it will be enough for our purposes here if knowledge of syntax requires a distinct type of inferential ability, because this will still be problematic for Alfano independently of this convergence with Gigerenzer. The situationist may insist that even the most robust epistemic dispositions can be easily disturbed by very easy manipulations of the stimuli, perhaps the framing effect in the Linda case and other studies on the effect of font size show precisely how fragile these abilities are. Even here, they only function well when seemingly irrelevant environmental variables are not present. This in turn might threaten the anti-luck and safety intuitions endorsed by many virtue epistemologists.

In response, we would like to provide an illustration of why although information processing may always be disturbed under laboratory settings, this by no means threatens the stability of epistemic dispositions. For instance, in the Stroop task, the interference between inclinations (the automatic inclination to read a word vs. identifying a color) does not entail that the capacities involved are unreliable because of alleged context sensitivity. The capacities to read and detect color are incredibly reliable across subjects in many conditions. Interference only shows that having two inclinations affects processing. Any virtue conceived as a stable disposition will be disturbed or "masked" under some conditions. But being disturbed in non-standard situations is just part and parcel of *being a disposition*, but this is a point we will return to below.

It is worth noting that this response to Alfano requires widely individuated abilities for color recognition. It was conceded that in some cases the very narrow disposition to 'identify color R in *disturbing conditions C*' may not be reliable, so it is the broad recognitional ability "identifying color R" that allows reliability and

susceptibility to disturbing conditions⁹. This appeal to broad abilities is not an *ad hoc* move just to defeat the situationist, but rather is the most natural way of understanding epistemic dispositions in light of the most recent evidence in psychology and linguistics. This also seems supported by the evolved basic capacities Gigerenzer cites as the life blood of ecological rationality. If Alfano insists that inference must be rule-based, formal and regimented, one can hardly think of a type of inferential process that satisfies these constraints better than knowledge of syntax. One constantly uses the rules of syntax to parse words, identify their meanings, and translate from one language to the other. Knowledge of syntax is necessary to understand and know the meanings of any expression. So it is not trivial that these robust, widespread and stable epistemic capacities are performed in a perception like fashion.

Communication and cognitive success: Mellor

The inferential abilities shown above in knowledge of syntax and language acquisition point to more complex inferential abilities involved in communication. What do people know when they communicate and how are inferential abilities involved? D.H. Mellor's (1990) theory of communication is based on widely shared knowledge of 'utility conditions' or what a person needs to know in order to effectively communicate. Mellor's account is inferential and shows more inferential abilities involved in language acquisition above. These will not be limited to examples of children, since we increasingly all communicate all the time.

Mellor's argues that communication involves a form of indirect inferential knowledge that is analogous to indirect observational knowledge. Communication is "the production in the audience of beliefs about what the speaker believes he believes". We what we assert when we communicate is that 'S believes that he believes p'. Mellor argues roughly that x gets the belief that p from what Y believes, but not directly from what Y says, but from what Y believes he believes. The cognitive abilities involved in communication are stable, reliable and inferential.

Thus, an vast array of inferential capacities build on top of another, allowing humans to engage in a complex network of epistemic exchanges, with clear implications for social epistemology. Communication, based on knowledge of syntax and language, allows for shared forms of reliable true belief production, for instance by testimony or by collective evidence gathering. The principles of effective and epistemically virtuous communication are deeply linked to the norm of assertion as knowledge: one must only assert what one knows. We expand on this issue below.

⁹ For a nice account of recognitional abilities see Millar 2008.

Meta-Cognition and Epistemic Feelings

A recent line of research on meta-cognition suggests that there are specific brain activities that monitor and control its own cognitive operations. At a minimum, metacognitive control involves (a) self prediction (b) post-evaluation and in many cases (c) intermediate evaluation. An interesting development in understanding meta-cognition comes from work on “epistemic feelings” – the feeling of knowing, certainty, doubt. Rather than being highly reflective, computational and costly, epistemic feelings are efficient ways of achieving meta-cognitive control. This ‘non-cognitivist’ account pushes us closer to Gigerenzer’s gut reactions than Bayesian calculation. Lepock (2013, this volume) argues that meta-cognition does not require meta-representation, just an adequate model. Epistemic feelings are implicit assessments of our cognitive operations, and these are used in the process of meta-cognitive regulation very effectively. Proust argues that meta-cognition regulates “mental actions” in a way that is analogous to the regulation of bodily actions by the motor system. Meta-cognitive dispositions are the motor system of the mind.

For our purposes, we need to consider whether any of this is inferential, or shows reliable inferential abilities. Borrowing from Proust (2001), performing mental actions like judging, deciding, solving, active attending, looking, listening involve *self-prediction*. This involves modeling available strategies for likelihood of achieving some cognitive goal and searching these for salient features. *Strategy selection on the basis of self-prediction certainly seems inferential*. If epistemic feelings are involved in this process in the way suggested above, this will be a less costly, fast and frugal process, and it will still be inferential.

‘Post evaluation assessments’ of mental actions evaluate how successful the selected operations were for the task at hand. This is a rule based judgment that also looks very inference like. Intermediate assessments are more controversial, but they present interesting assessments of likelihood of success of the selected strategy *while the strategy is being implemented*. Monitoring the selected strategy occurs largely by recognizing and responding to epistemic feelings that provide implicit assessments of how well things are going in a current cognitive operation. Even if this does not sound inferential, selecting a complimentary strategy (deciding to search your address book rather than your memory) certainly does. Complimentary strategies involve an inference that another strategy will be more effective in this particular task situation than the one currently being performed. That is clearly an inference. This kind of inference is either employed or available to an agent in most rational actions, and the ability to select complimentary strategies is a reliable and robust basic inferential ability.

Knowledge of logic

Now consider the basis for any type of formal rule of inference: knowledge of logic. We have the capacity to reason according to modus ponens and this capacity is part of a set of stable dispositions to draw deductive inferences that are truth-preserving. One may actually say that these dispositions constitute what we *mean* by deductive inference¹⁰. If this is the case, then one could not know the meaning of what a deductive inference is without having such stable epistemic dispositions. It is a truism that basic deductive reasoning (for example an application of modus ponens) can be achieved without explicit understanding of such rule and that these dispositions, like those underlying knowledge of syntax, are remarkably stable. Demanding an explicit *understanding* of the rules for deductive reasoning increases cognitive demands, and although we can be trained to have such explicit understanding, this is not a necessary condition to have the stable dispositions that are implicit in our capacity to identify these rules. More importantly, requiring such explicit understanding is open to traditional objections against accessibilism and deontological accounts. Thus, it seems that the best strategy is to characterize these fundamental rules for deductive reasoning in accordance with our perception-like model¹¹.

The situationist seems to face a new dilemma. Either we possess stable epistemic dispositions that allow us to identify valid rules for deductive inference or we don't. If we do, then situationism is false. If we don't, it is not clear how we are able to understand what we mean when we talk about, for instance, modus ponens. For it is not clear that highly unstable and easily disturbed capacities would help us succeed in specifying what we mean *in every situation* by the fundamental rules (modus ponens, modus tollens, etc.). Thus, it would not be entirely clear that we mean the *same* fundamental rules when we characterize a piece of deductive reasoning as modus ponens or something else. The situationist needs to explain why the psychological evidence would have such a dramatic result and this strongly suggests that situationism is in trouble. Obviously, the easy way out of this dilemma is to affirm that situationism is false, which is what we propose¹².

¹⁰ See Boghossian, 2000.

¹¹ Notice that this is quite different from having a conscious-intellectual "seeming," which is one way of defining intuitions.

¹² This is a concrete way of making a point suggested to us by Lauren Olin in conversation, which is that relativism is much more *troubling* in the epistemic case, as compared to the moral case. If we are right, situationism is also a lot more *implausible* in the epistemic case.

5. Sosa, assertion and normal conditions

Consider how Sosa's view would handle Alfano's challenge. Sosa requires that a belief must be AAA in order to count as knowledge (2007). It must be accurate (true), adroit (skillful) and apt (true because skillful). In Alfano's Linda case, Ash Paradigm and others do we have a failure of accuracy, adroitness or aptness? Clearly a failure of accuracy, whatever the right answer is, it cannot be the conjunction in the Linda case. It may also be a failure of adroitness in inductive reasoning, or a failure of accuracy because of a failure of adroitness. In either case, an agent would not know on Sosa's account, but since Alfano's challenge is to reliabilism, showing a general failure to meet the accuracy condition will be enough. If a belief is not accurate, it is not apt and is thus not knowledge. If we have too many of these kinds of failures, then we have face Alfano's worry about respecting non-skepticism.

Sosa can respond here by saying that the experimental conditions take an agent outside of normal conditions because of the presence of *interfering conditions*. Sosa says that an archer may fail to hit the mark if drugged or in a tornado, but in these cases there is no relevant sense in which the archer has failed because success is only expected in normal conditions for the exercise of the relevant competence. This is a widely accepted point about dispositions – they are only expected to manifest when in normal conditions, that is when not masked or finked. Failures outside of normal conditions are not relevant failures. If the epistemic situations that agents are placed in are not a normal for the cognitive competence being tested and they get a false belief, the agent does not *fail* any more than the archer does when in Hurricane conditions. Since Sosa requires normal conditions for relevant failures, he can say that there is no (virtue theoretic) epistemic failure shown in the research. Thus, even if there are many such cases, that does not imply that there are many epistemic failures, and thus does not push the virtue epistemologist to violating the non-skepticism principle.

Alfano might respond that these are clearly “normatively irrelevant” features of the environment, so they cannot/should not be built into the stimulus conditions. Building in too many normal conditions threatens to give a vacuous account that a virtue will manifest except when anything is preventing it from manifesting. We might avoid this through with a turn to disposition theory and distinguish between *culpable and non-culpable masks*. Alfano will need to show that environmental elements introduced in the research are environmental conditions that we should expect agents with a given ability to perform well in, even when presented with masks or finks. Alfano would need to show that these are not cases of non-culpable masking or finking. Without claiming to answer this important questions, it does not appear that Alfano has made good on this additional and necessary premise.

Assertion requires the same constraints, and situationism has extremely strong implications with respect to assertion. If the situationist is right, it would be

quite difficult to satisfy the norm of assertion, threatening not only inference, but also testimony, public assessment of evidence and, ultimately, basic communication. We believe that having these consequences is a form of *reductio ad absurdum* for the situations it challenge in epistemology.

References

- Alfano, M. (2011). Expanding the situationist challenge to responsibilist virtue epistemology. *Philosophical Quarterly*. Vol. 62 Issue 247. 223-249.
- Alfano, M. (2012). Expanding the situationist challenge to reliabilist virtue epistemology. *Forthcoming*.
- Boghossian, P. (2000) "Knowledge of Logic" in P. Boghossian and C. Peacocke, *New Essays on the A Priori*, Oxford University Press, pp. 229-254.
- Chomsky, N. (1987) *Generative Grammar: Its Basis, Development and Prospects*, Studies in English Literature and Linguistics, Kyoto University.
- Chomsky, N. (1986) *Knowledge of Language*, Praeger.
- Gigerenzer, G. (2008) *Rationality for Mortals: How People Cope with Uncertainty*, Oxford University Press.
- Graham, P. (2011) 'Epistemic Entitlement' *Nous*, online January 2011.
- Greco, J. (2010) *Achieving Knowledge: A Virtue-Theoretic Account*, Cambridge University Press.
- Harman, G. (2000). The nonexistence of character traits. *Proceedings of the Aristotelian Society*, 100: 223-226.
- Harman, G. (1999). Moral philosophy meets social psychology: Virtue ethics and the fundamental attribution error. *Proceedings of the Aristotelian Society*, New Series 119, 316-331.
- Henderson and Horgan (2011) *The Epistemological Spectrum*. Oxford University Press.
- Hornstein, N. (1984) *Logic as Grammar*, MIT Press.
- Jackendoff, R. (2003) *Foundations of Language: Brain, Meaning, Grammar, Evolution*, Oxford University Press

Kahneman, D. (2011) *Thinking Fast and Slow*, Farrar, Straus and Giroux.

Kahneman, D. & Tversky, A. (1973). On the psychology of prediction. *Psychological Review*, 80, 237-251.

Simon, H. A. (1972). Theories of bounded rationality. *Decision and organization*, 1, 161-176.

Sosa, E. (2011). *Knowing Full Well*. Princeton: Princeton University Press.

Sosa, E. (2007). *A Virtue Epistemology*. Oxford: Oxford University Press.

Zagzebski, L. (2010). Exemplarist virtue theory. In Battaly (ed.), *Virtue And Vice: Moral and Epistemic*, pp. 39-55. West Sussex: Wiley-Blackwell
