CHAPTER 1

Introduction
Naturalized virtue epistemology

Abrol Fairweather and Owen Flanagan

1 Virtue epistemology: metaphysical and normative

This volume aims to launch a powerful and largely unexplored position in epistemology: naturalized virtue epistemology. Most debates in virtue epistemology have been decidedly axiological and aim to clarify the goals, values, and ends constitutive of epistemic evaluation. Value-driven inquiry has now become quite complex in the large literature on the value problem (and the related Meno problem), which examines whether the value of knowledge can be reduced to the value of any proper subset of its parts (Zagzebski 1996; Kvanvig 2003; Pritchard 2007). Normative epistemic inquiry has also been useful in meeting more traditional problems in epistemology, such as Gettier problems (Turri 2011) and problems of epistemic luck more generally, as well as the structure of knowledge (as etiological rather than foundational or coherentist), and Chisholm's "problem of the criterion" (Riggs 2007). Virtue epistemology has opened many new areas of inquiry in contemporary epistemology including: epistemic agency (Greco 1999; Zagzebski 2001; Sosa 2007), the role of motivations and emotions in epistemology (Fairweather 2001; Hookway 2003), the nature of abilities (Millar 2008; Greco 2010; Pritchard 2012), skills (Greco 1993; Bloomfield 2000), and competences (Sosa 2007), the value of understanding (Kvanvig 2003; Grimm 2006; Riggs 2009), wisdom (Ryan 1999; Zagzebski 2013), curiosity (Whitcomb 2010; Inan 2012) and even education policy and practice (Baehr 2011). The virtue turn in epistemology that started with the early work of Sosa (1991) and Zagzebski (1996) has now produced a large and mature literature in normative epistemology.

While the growth and impact of virtue epistemology has been impressive and important, it has come with insufficient attention to the empirical grounding of these normative theories, and thus runs the risk of endorsing free-floating epistemic norms cut loose from the real-world phenomenon they must evaluate. To this end, virtue epistemologists should heed

the exhortation given by Anscombe in "Modern Moral Philosophy" (1958) to constrain normative theorizing in ethics with an empirically adequate moral psychology, and might even do so optimistically since Anscombe (and Foot, later Geach, and still later MacIntyre) was led to endorse virtue theory precisely because it appeared more psychologically plausible than deontology or consequentialism. The same cautionary (and perhaps optimistic) point holds for epistemic psychology and normative epistemology.

Greater concern for empirical adequacy can be seen in a few empirically focused works on virtue ethics (Russell 2009; Snow 2010; Slingerland 2011; Miller 2013), but these laudable efforts have come largely in response to 'situationist challenges' which question the existence and/or explanatory salience of character traits on the basis of research in social psychology. Flanagan (1991) was the first to raise worries about situationism for virtue theory, recommending a certain caution with reference to overconfidence about character traits. But swiftly an extremist position arose and philosophers were asked to take seriously a debate that was already long dead in psychology about whether there were any such things as character traits (Flanagan 2009). Nonetheless, the long tradition of empirical challenges to character-based psychology has been largely ignored by virtue epistemologists until very recently (see Axtell 2010; Alfano 2012). Moreover, it cannot be assumed that any satisfactory empirical ground for virtue ethics will offer the same grounding for virtue epistemology. For example, the successful CAPS (Cognitive-Affective Personality System) research program used by Snow and Russell to ground virtue ethics shifts the focus of evaluation from "objective stimulus" to "construed stimulus," but arguably skips over the main epistemic question in the process, which is precisely whether the subjective stimulus tracks or accurately represents the objective stimulus. Virtue epistemology will have to earn its own response to situationism. If the demonstrated success on the normative side can be paired with equal success on the empirical side, an extremely powerful perspective in epistemology is well in hand, arguably one of most promising general epistemologies on offer today. The current volume is the first collection of essays on naturalized virtue epistemology.

The most important motivation for naturalized virtue epistemology does not come from naturalism per se, but rather from the insight that virtue theory is heavily metaphysical and empirical, as well as normative. Any attribution of a virtue to an agent will assume some taxonomy of agent-level dispositions that are both explanatory and praiseworthy. The truth-makers for virtue attributions will thus be of a highly nuanced metaphysical kind, and this must be some form of disposition in any theory

considered truly virtue theoretic. Any truly virtue-theoretic account of a phenomenon will need some form of these metaphysical commitments, otherwise it becomes unclear how virtue-theoretic normativity is distinct from a range of other sources of normativity (e.g. responsibility, rule-consequentialism, deontology). Virtue-theoretic normativity is distinct in part because of this dispositional commitment.

Many of the chapters in this volume examine empirical findings on the nature of cognitive dispositions and personality traits (Alfano, Battaly, Miller, Pritchard), and this is clearly one direction for naturalized virtue epistemology to take. However, some naturalistically minded virtue theorists will prefer traditional accounts from Aristotle, Plato, the Stoics, Aquinas, and others that appeal to some form of teleology, function, or essence rather than the most recent results in psychology. For Aristotle, virtue is not any admirable or praiseworthy quality, it is a quality that makes you good at performing your function, and thus must draw on a substantive account of human nature (see Korsgaard 2008). We say that naturalized virtue epistemology can be pursued metaphysically or empirically, and possibly both at the same time. Contemporary empirical work in psychology on person-level dispositions like 'generosity', 'conscientiousness', and 'narrow-mindedness', rationality and inference is thus one fertile source for naturalized virtue epistemology, as is metaphysical work on the essences, natures, and kinds relevant to sustaining the human form of life. One metaphysical debate long underway within virtue epistemology is whether epistemic virtues are best seen as faculties (see Sosa and Greco), skills (Greco, Bloomfield) or character traits (Baehr, Zagzebski). But these issues have largely been decided on normative grounds, with different philosophers arguing that different disposition types support the set of normative standings necessary to make the full range of evaluations that matter in epistemology. Naturalized virtue epistemology (scientific or metaphysical) shifts to an empirical basis for understanding the nature and explanatory power of faculties, skills, and character traits, and thus continues to ask the same questions about which disposition types to countenance and how to account for the causal-explanatory properties of the virtues. As we will see below, this empirical turn will have important (and largely positive) normative implications for virtue epistemology.

Virtue attributions also engender predictive commitments such that to attribute a virtue (V) to an agent (S) is to assert that (S) will satisfy a set of conditionals linking situations (C) and virtue-relevant behaviors (B) in roughly the following way: if an agent S is V, then most of the relevant following conditionals will be true of S: $(C_1 \rightarrow B_1)$, $(C_2 \rightarrow B_2)$, $(C_3 \rightarrow$

B_3) … ($C_n \rightarrow B_n$). Given these behavioral-predictive commitments, virtue attributions make a real claim about the world and the truth value of such attributions should be empirically testable. Whether an agent does or does not possess virtue-constituting dispositions will be an empirically testable proposition. If we know the actual cognitive dispositions that constitute a given virtue, empirical research should inform us of the causal mechanisms that underlie their manifestations across a range of relevant circumstances. We argue below that this is very important information for the normative aspect of virtue epistemology. The most plausible and promising way of handling these empirical commitments is to do it naturalistically. In the remainder of this Introduction, we examine two significant worries for a would-be naturalized virtue epistemology and then summarize the essays here collected.

One problem a naturalistic turn might create for virtue epistemology is the persistent worry about normativity in naturalistic theories. Any difficulties on this score will be particularly worrying for a value-driven epistemology. A second worry is that the relevant results from the sciences will signal bad news for virtue epistemology. It may turn out that, according to our best psychological theories, people either do not possess or rarely manifest the kinds of dispositions countenanced by virtue epistemology.

2 Worries about normativity: Quine and Moore

Any call to naturalize an epistemic theory immediately brings us to W. V. O. Quine and his famous paper "Epistemology Naturalized" (1969). Quine is often seen as defending the extreme position that a purely descriptive empirical psychology should outright replace traditional epistemology. Since the latter is inherently normative (see Kim 1988), Quine has been accused of being unable to account for the normativity of epistemology and "changing the subject," and this makes naturalized epistemology and normative epistemology appear to be very strange bedfellows. How can the approach to knowledge that aims to remove normativity from epistemology (naturalized epistemology) be united with the most overtly normative approach to epistemology (virtue epistemology)? One reply is that naturalization does not eliminate the normative so much as it tames and explains it (Flanagan 2006). Another way of thinking about naturalism and normativity that goes back to G. E. Moore (1993), rather than Quine. The Moorean approach to naturalism and normativity will be just as important in the present context because virtue epistemology has always borrowed heavily from virtue ethics, and Moore had moral

properties in mind when he raised the Open Question Argument against naturalistic accounts of goodness. Let us begin with Moore, and then work our way back to Quine.

Moore's challenge to naturalism in ethics is that we can easily see that no predicate expressing a natural property can have the same meaning as the predicate 'is good'. For any candidate natural property (e.g. pleasantness), a competent speaker can, without contradiction, make the following remark: "I see that *x* is pleasant, but is it good?" The goodness of *x*, even given that it is pleasant, is an 'open question', unlike the goodness of *x* in "I see that *x* is good, but is it good?", which is clearly a closed question. Moore argued that any natural property proposed to analyze a moral property will be subject to open question arguments, and thus commits the naturalistic fallacy. Two points interest us here. First, it is now very clear that Moore is making the wrong kind of semantic demand, namely that any adequate analysandum must entail or have the same (conceptual) meaning as the analysans. Second, epistemic normative properties might fare much better than normative properties in ethics in meeting the Moorean demand. On the first point, we now know that sameness of conceptual meaning is the wrong thing to demand between a concept and its proper analysis. With advances in semantics from Putnam, Kripke, and others, we now say that statements like 'water is H_2O' involve terms that differ in (conceptual) meaning, but pick out the same property. This is a very successful naturalistic analysis of water, even though it would fail to meet Moore's demand for sameness of meaning. While the specific demand is wrong, Moore is certainly right to think that some constraint must be imposed on naturalistic accounts of normative properties. But if not sameness of meaning as Moore demands, then what? A modified Moorean demand might require that some statements couched in naturalistic terms will constrain prescriptions (perhaps through an 'ought implies can' principle), provide a supervenience base for normative facts, or provide truth conditions for normative claims.

Epistemology may have an easier time meeting these modified Moorean demands than ethics. Alvin Goldman (1994) argues that normative epistemic facts about when an agent is justified or has knowledge supervene on natural facts about the reliability of the processes used in forming the relevant beliefs. This might be plausible as a general thesis in epistemology, but most virtue epistemologists insist that reliabilism cannot adequately explain the value of knowledge. But if reliabilism cannot account for the additional value of knowledge over true belief, it cannot give the right supervenience basis for normative epistemic properties. One additional

normative standing that facts about reliability will need to support involve 'credit' for holding a true belief, which is often identified by virtue epistemologists as the value that knowledge has but a mere true belief might lack. Moreover, the presumably natural facts about when a given true belief is 'due to', 'caused by', or 'from' an ability in the agent are very plausibly taken to be the supervenience base of normative facts about epistemic credit. While this points to a seemingly clear way in which virtue-reliabilism meets modified Moorean demands on normative epistemology, the 'because of' relation has proved notoriously difficult to properly formulate, and some virtue epistemologists have given up any such attempt. This should not be taken to show that virtue epistemology fails to meet Moorean demands after all, but rather that we need a more naturalized virtue epistemology that furthers our understanding of the relevant causal-etiological relations to show the way.

We now return to Quine. Jagewon Kim (1988) raised the classic normativity challenge to Quinean naturalized epistemology, but his objection has now received many responses which mute the perceived force of his original objection. Naturalized virtue epistemology may nonetheless face a special normativity worry because virtue epistemology has consistently been championed as the most overtly normative way to do epistemology. It is one thing to show that an epistemic theory (e.g. process reliabilism) can eke out some normative content once naturalized, but quite another to bill your epistemic theory as robustly and overtly normative and then attempt a naturalization project.

Fortunately, the situation vis-à-vis normativity is not all that bad. Naturalized virtue theory has a number of sources for maintaining the normatively focused inquiry that has distinguished it from other perspectives in epistemology. Research in evolutionary theory, meta-cognition, and bounded rationality all strongly support the existence of robust cognitive dispositions. Remembering that the metaphysical side of virtue epistemology requires a taxonomy of cognitive dispositions and an account of their causal explanatory role in producing epistemically assessable agent-outcomes, empirical facts about these dispositions and their manifestations will provide normative content for virtue epistemology. Normative facts about what an agent with virtue (V) in situation (C) ought to do are identical to facts about how the relevant disposition will actually manifest when in (C). If we are talking about a real psychological disposition, the way it ought to manifest in various situations is an empirical issue and is best understood by consulting work in cognitive psychology, social psychology, evolutionary biology, and other relevant sciences. Given that

disposition (D) is a virtue and that (D) is attributable to an agent (S), we can now say that S normatively-ought to do whatever S dispositionally-ought to do. A dispositional-ought just is a normative-ought here, and the former is an empirical issue. It will be true of all meaningful dispositions (and certainly of natural dispositions) that the bearer of the disposition ought to manifest certain behaviors in certain conditions: the fragile glass ought to break, the sugar ought to dissolve. Call this a merely dispositional-ought. Considered just as such, dispositional-oughts have no prescriptive force. But with just the premise that D is a virtue, we can now collapse what S normatively-ought to do into what S dispositionally-ought to do. Since what S dispositionally-ought to do is an empirical issue, so too is what S normatively-ought to do. A number of additional important normative statements for virtue epistemology will be inferred using principles of disposition theory relating to masking, mimicking, and finking, but this takes us beyond the scope of the present inquiry.

It may rightly be objected that, as stated, we have simply assumed that a given disposition (D) is a virtue to make the point above, and the main normative work of virtue theory is precisely to determine exactly that. It may be that naturalized virtue epistemology leaves this particular question to an autonomous virtue theory to work out, but it is clear that merely identifying which dispositions qualify as virtues does nothing to provide prescriptive content for specific situations; that depends on how the virtue constituting dispositions will actually manifest across the relevant range of circumstances. So empirical work on dispositions is still doing important normative work here. A more robust naturalized virtue epistemology will aim to settle the question of which dispositions will be virtues on empirical grounds as well. It is not our intent here to distinguish either of these projects from the other, but just to note that naturalized virtue epistemology can take either modest or robust forms, and that the metaphysics of dispositions might perform essential work for generating the prescriptive content of virtue theory.

3 Metaphysical worries: situationism and virtue epistemology

If a naturalized virtue epistemology can provide empirically informed and falsifiable disposition attributions that also ground prescriptive epistemic norms, this would be an attractive general epistemology. However, some philosophers claim that work in social psychology supports skepticism about the existence of the right kind of person-level dispositions,

virtues and the like, as well as their causal powers. This is the situationist challenge. Owen Flanagan (1991) first formulated the situationist challenge in his brief for psychological realism in ethics, and suggested a measured response to the research, a response which was already available in Aristotle, and that had been sensitively explored in contemporary literature by Martha Nussbaum (1986), among others. There are persons, characters, traits, virtues and vices, but persons, traits, virtues, and the like, are fragile, of variable consistency, resiliency, and robustness. Quickly, an extreme, occasionally eliminativist, response to the situationist challenge emerged in ethics. There are no such things as character traits or virtues, or if there are such things they are effervescent, ephemeral, short-lived, merely linguistic vapors (Harman 1999, Doris 2002). The current consensus favors the moderate view, one that was absorbed inside psychology thirty-five years before philosophers got all fussed-up, to the effect that there are genuine morally relevant dispositions that occur in complex, dynamic causal fields and that need to be carefully individuated; for example, honesty-at-work versus *honesty tout court* (Flanagan 1991, 2009; Merritt 2000, Sreenivasan 2002, Miller 2003, Kamtekar 2004, Russell 2009, Snow 2010, Slingerland 2011). Dispositional traits of variable scope, range, depth, texture, and kind (virtues, vices, thinking styles, emotional patterns, temperamental traits, skills, habits, etc.) are real, and they partake in complex causal relations with all the other inner and outer features of the complex ecologies that make for human life as we know it. The causal powers and actual role of any particular disposition depend on it, its kind, how superficial or deep it is, and on how it interacts with the rest of a person and her world at any given time.

Still, some philosophers continue to ring the bell for the extremist view arguing that situations out-predict and out-explain character traits or other dispositional features (Alfano 2012, 2013). There are two mistakes here: first, even if we restrict ourselves to only situations, roughly, current stimuli, and the traits we antecedently ascribe to some person (she is honest), it is false empirically that situations out-predict and out-explain character traits (Flanagan 2009). But second, conceiving prediction and explanation in terms of only two variables, situations and traits, is impoverished, a sophist's trick. We commonly use knowledge of history, culture, politics, and economics in explanation and prediction in addition to knowledge about some individual and the situation she is in. Zip code, sadly, is more predictive of educational achievement than almost any other variable. Zip code does not name a situation, nor does it name a character trait. It designates a world. The key point is that the predictive power of situations

fades to relative insignificance the more we know about history, culture, economics, race, gender, personality, and traits.

Whereas ethics generally, and virtue ethics in particular, has had twenty-five years to work out its response to the hyperbole about situationism, virtue epistemology needs to make its own peace with this research.

A number of contributions in this volume (Alfano, Battaly, Miller, Pritchard) examine situationist-style arguments applied to virtue epistemology rather than virtue ethics. Mark Alfano (2012, 2013) has tried to fill this dialectical niche by applying situationism to both responsibilist and reliabilist virtue epistemology, arguing that both fall prey to an inconsistent triad along the following lines: (1) Non-Skepticism: most people know quite a bit; (2) Virtue Epistemology: knowledge is virtuously formed true belief; (3) Epistemic Situationism: most behavior is explained by often trivial features of situations, rather than by personal characteristics or traits. This issue will be examined in a number of chapters in this book. The consensus appears to be similar to the one in moral psychology: namely, that (3) Epistemic Situationism is false.

If virtue epistemology is held accountable to empirical results, it will have to be to the best and most diverse range of results, but research in meta-cognition, the big-five personality factors, and bounded rationality presents much better prospects for naturalized virtue epistemology. As these research programs mature, it is highly likely that the basic commitments of virtue epistemology will be vindicated.

Below we introduce the chapters that appear in this volume. They address a wide range of issues relevant to the project of developing a naturalized virtue epistemology.

"Warrant, functions, history" by Peter J. Graham

Peter Graham defends a proper functionalist form of virtue-reliabilism which takes proper functioning as its central norm. Epistemic warrant consists in the normal functioning of belief-forming and -sustaining psychological processes. When the process has forming and sustaining true beliefs reliably as a function, especially an etiological function, a token belief that results from a properly functioning system quite plausibly has a form of warrant. While Graham shapes his understanding of proper functioning around evolutionary theory and philosophy of biology, he extends his account of etiological functions to include artifacts and social facts. Graham provides a compelling epistemic teleology with significant connections to evolutionary theory and the social sciences.

"The epistemic 'ought'" by Ram Neta

Many naturalistic epistemologists take empirical evidence concerning our cognitive powers to cast doubt on the claim that our beliefs ought to be consistent, or that our degrees of belief ought to be probabilistically coherent. Neta argues that, in doing so, they are implicitly confusing cognitive competences with epistemic virtues. In this chapter, Neta distinguishes competences from virtues, explains why epistemic "oughts" imply that we have the relevant cognitive competences, but not the related epistemic virtues, and finally describes the constitutive relationship between epistemic virtues and epistemic "oughts." Neta defends a virtue-theoretic account of epistemic oughts in terms of normal epistemic functioning.

"Naturalism and norms of inference" by Carrie Ichikawa Jenkins

Jenkins defends a normative and naturalistic account of basic inference. While reliabilist virtue epistemology has focused on abilities like perception and memory, and responsibility, virtue epistemology has focused on traits like open-mindedness and conscientiousness, much less attention has been given specifically to virtues of the act of inferring. Suppose that a subject S makes a token inference from [P & P→Q] to Q. To focus ideas, let's take this to be a matter of her transitioning from a belief in the premise to the adoption of a new belief in the conclusion which she now holds on the basis of her belief in the premise (which is retained throughout). This inference (as distinct from S's belief in the premise and/or conclusion) is available for distinctively epistemic normative assessment. The inference itself might be described as epistemically good or bad, warranted or unwarranted, justified or unjustified, appropriate or inappropriate, correct or incorrect, permissible or impermissible, and so on.

"Indirect epistemic teleology explained and defended" by David Copp

David Copp examines connections between meta-ethics and "meta-epistemics" and defends a pluralist epistemic teleology that provides a new motivation for epistemic instrumentalism. J. L. Mackie proposed that morality as a "device" needed to solve a "problem" faced by humans because of "certain contingent features of the human condition." Copp proposes that the standards of normative epistemology are similarly