

© 2017

Georgina Jayne Gardiner

ALL RIGHTS RESERVED

IN PURSUIT OF UNDERSTANDING

by

GEORGINA JAYNE GARDINER

A dissertation submitted to the  
School of Graduate Studies  
Rutgers, The State University of New Jersey  
in partial fulfillment of the requirements  
for the degree of  
Doctor of Philosophy  
Graduate Program in Philosophy  
written under the direction of  
Ernest Sosa  
and approved by

---

---

---

---

---

New Brunswick, New Jersey

October 2017

## **ABSTRACT OF THE DISSERTATION**

**In Pursuit of Understanding**

**By GEORGI GARDINER**

**Dissertation Director:  
Ernest Sosa**

Understanding consists in integration and coherence amongst beliefs, the individual's grasping of these connections, and the explanatory power of the individual's representations of the world. Understanding is under-theorised in contemporary epistemology. Most epistemological research focuses on knowledge, and related epistemic kinds such as knowledge-relevant justification. This dissertation begins by motivating the value of thinking about understanding in epistemology. This is the aim of section A, 'Understanding and Value'.

I begin, in 'Understanding, Integration, and Epistemic Value', with a debate about the ultimate bearers of epistemic value. Veritism holds that attaining true belief (and avoiding false belief) is the sole ultimate epistemic good. All epistemic value is ultimately epistemically valuable in virtue of its relation to these dual aims. Value pluralists deny this, and posit plural sources of epistemic value. I suggest this debate has reached an impasse, and I appeal to the nature and value of understanding to make progress. I argue that veritism is committed to a supervenience thesis stating that if two sets of beliefs are identical with regard to true and false beliefs and propensity to gain further true beliefs and avoid false beliefs, then they have

equivalent levels of epistemic value. I argue that comparing mere true belief with understanding indicates the supervenience thesis is false. Understanding contributes value that does not reduce to the value of true belief.

I motivate this claim by arguing that the integration and coherence-making relations characteristic of understanding have epistemic value, and this value does not reduce to the value of true belief. I argue, further, these relations amongst beliefs do not themselves reduce to further true beliefs; the structure of beliefs that constitutes understanding is not simply a matter of aggregating further true beliefs.

This paper thereby motivates the claim that thinking about understanding can help advance and inform debates in epistemology. It also motivates a second—more important—claim: understanding has a distinctive nature and value. When we understand, our beliefs and knowledge cohere; structural, organising relations are formed and maintained. This epistemic kind is not simply a matter of possessing more knowledge. (Or, if understanding is an agglomeration of knowledge, it is a distinctive and important kind of knowledge, one characterised by structural relations amongst beliefs.) I thus hope to motivate the importance of theorising about, and pursuing, understanding.

The second essay, ‘Evidentialism and Moral Encroachment’, considers a recent debate in social epistemology. Like in the first essay, I employ an existing debate to motivate the importance of focusing on understanding in our epistemological theorising.



Evidentialist views hold that the epistemic justification of a doxastic attitude depends solely on evidential considerations, such as whether the attitude fits the available evidence. Evidentialism is an ‘intellectualist’ view. Intellectualism in epistemology holds that the whether a belief is epistemically justified depends exclusively on truth-related factors. Since these truth-related factors can include whether the belief was reliably formed, whether it manifests intellectual virtue, and so on, intellectualism is (usually understood as) more permissive than evidentialism, which restricts the relevant features to evidential considerations, such as possessed evidence. It is worth emphasising at this juncture that evidentialism can be understood in various ways, depending on what evidential considerations are relevant to epistemic justification, whether features other than epistemic justification bear on what you ought believe, and so on.

Intellectualist positions, including evidentialist views, hold that practical and moral considerations do not bear on whether a belief is epistemically justified. Recently advocates of moral encroachment, by contrast, argue that moral considerations affect whether a belief is epistemically justified. Moral encroachment holds that if the belief can morally wrong a person or group then more evidence is required to epistemically justify that belief. Advocates of moral encroachment employ examples like the following:

**Administrative Assistant.** A consultant visits an office. He knows that few people visit the office who are not employees of the firm and that almost every woman employee is an administrative assistant. The consultant sees a woman walking down the corridor and forms the belief ‘she is an administrative assistant’.

Advocates of moral encroachment argue that the consultant’s belief is morally wrong, and the moral features of the case render the belief epistemically unjustified. Some theorists

focus on the fact that the consultant does not rule out relevant alternatives—such as that the woman is not an administrative assistant—and argue that this alternative is rendered salient by moral factors. Other theorists focus on the idea that the belief can wrong a person and so the moral stakes are raised and the thresholds for justified belief is thereby higher. Either way, moral encroachment maintains that evidence that would normally suffice for belief is rendered epistemically insufficient by moral features.

In ‘Evidentialism and Moral Encroachment’ I defend intellectualist views against this argument for moral encroachment. I argue firstly that the beliefs and epistemic practices used to motivate moral encroachment, such as the consultant’s belief described above, are best understood as unjustified according to orthodox, intellectualist epistemology. The beliefs exhibit myriad epistemic errors, including failing to reflect the available evidence. I explain that if the putatively morally wrong belief is also epistemically unjustified according to orthodox epistemology, the moral wrong does not impugn evidentialism: moral error and misfit with evidence align. This defence of evidentialism against moral encroachment accords with orthodox views of racism, sexism, and similar prejudices, which hold that orthodox epistemic error is central to the nature of fault.

I then argue that the moral status of a belief depends on the understanding in which the belief is embedded. It is a mistake to evaluate the beliefs in these vignettes without further information about the person’s broader understanding. Whether the consultant’s belief is sexist depends on broader features of how he thinks about women and employment. This argument places understanding front and centre as a paramount feature of normal thought and philosophical theorising: when morally evaluating a belief we must consider the

understanding in which the belief is situated. This stems from the centrality of understanding in thought. (Although I do not explore the topic in the following essays, this kind of perspective could motivate an ‘understanding first’ approach to epistemology, according to which understanding is the primary, or central, epistemic kind or end, and all other epistemic kinds and ends should be understood with reference to their relationship to understanding.)

In section B, ‘Teleologies and Understanding’, I aim to illuminate teleological approaches in epistemology, and apply teleological theorising to understanding. Teleological approaches to understanding an epistemic kind, concept, or attribution proceed by asking questions such as ‘what is the *purpose* of the epistemic kind?’, ‘What *role* has it played in the past?’, or ‘If we imagine a society without it, why would they feel the *need* to invent it?’ The idea behind the teleological approach is that examining the function of the epistemic kind illuminates the contours of the kind itself.

In ‘Teleologies and the Methodology of Epistemology’ I first contrast teleological approaches with some other methods in philosophy. I then construct a three-way taxonomy of teleological approaches and illustrate each with an example. The first kind of teleological approach is practical explication, which aims to identify roles that the epistemic kind currently actually satisfies. Klemens Kappel and Chris Kelp claim, for example, that knowledge attributions mark when to cease inquiry.

The second kind is historical genealogy, which aims to uncover the roles played by the epistemic kind in the past, and to trace its evolution. I describe the example of epistemic attributions employed in scientific practice during the seventeenth-century Enlightenment.

According to Steven Shapin and Martin Kusch, seventeenth-century scientists faced the challenge of identifying whose testimony they could trust in order to advance the collective pool of scientific knowledge. These early scientists determined they could trust ‘sources of acknowledged integrity and disinterestedness’, by which they meant those who had financial freedom and high social status, and so would not be motivated to dissemble for financial and social gain. They meant, that is, landowning gentlemen. According to Shapin and Kusch, epistemic language and concepts were used to exclude from the pool of potential informers those who lacked these qualifications. I explore what this historical genealogy might indicate about the contours of contemporary knowledge attributions.

The third kind of teleological approach is hypothetical genealogy. The approach is genealogical as it explores the beginning and history of the concept, and hypothetical because rather than tracking the actual history of the concept, it posits a fictional narrative. Edward Craig developed an influential hypothetical genealogy for the concept of knowledge and the practice of knowledge attributions. He imagines an ‘epistemic state of nature’, in which people are language using and minimally co-operative, but lack epistemic concepts. Craig hypothesises that such a society, in virtue of lacking knowledge attributions, would need a way to tag good informants. He argues that knowledge attributions served the function of tagging good informants, and he articulates how knowledge attributions might evolve over time and what this indicates about the contours of knowledge attributions.

I then explore ways that teleological theorising might be fruitfully integrated within broader epistemological theorising. Finally, I evaluate an objection to the Craigian teleological approach, pressed by Hilary Kornblith. Kornblith criticises the teleological approach for

relying on the false assumption that humans create concepts, such as that of knowledge, to satisfy needs. He contends knowledge is a natural kind that fulfils the roles it does because of its nature; it does not have its nature because of the roles it fulfils. Although sympathetic to Kornblith's concerns, I defend some kinds of teleological approach against Kornblith's criticisms. Throughout 'Teleologies and the Methodology of Epistemology', I aim to shed light on teleological approaches in epistemology.

In 'Understanding and Emulation' I tie together my aim of illuminating and developing teleological approaches in epistemology and my aim of theorising about understanding. In particular I advance a tentative teleological account of understanding attributions.

I first articulate several features of understanding and understanding attributions. The subject matters that are understood, such as mechanics or languages, tend to exhibit complexity, such as unifying patterns or causal relations, to grasp. The subject matter of knowledge, by contrast, can be simple. Understanding is not easily transmitted by testimony. Understanding attributions are often connected to appreciation and emotion, such as in aesthetics or interpersonal relationships. Understanding, even more than knowledge, can underwrite the ability to manipulate and predict outcomes. It allows us to apply insights to novel cases. I argue that understanding is more stable than knowledge: knowledge can be quickly gained and lost, whereas understanding characteristically takes a long time to develop and cannot be quickly undermined. Understanding plausibly has a distinctive epistemic value.

These features of understanding, I suggest, accords with the proposal that understanding attributions tag those worthy of intellectual emulation. I explain what this proposal means,

and what it indicates about the nature and value of knowledge. I situate the proposal as a development and extension of Craig's hypothesis that knowledge attributions serve to tag good informants. The need Craig identifies is increasing our stock of true beliefs, whilst avoiding false beliefs. Craig hypothesises that knowledge attributions serve those needs. But acquisition of true beliefs and avoidance of false beliefs are far from our only epistemic needs. We also need to improve ourselves as inquirers. We want to improve the way we think, evaluate, and explore. We would benefit from being more creative hypothesis-formers, more effective problem-solvers, and better retainers of information. We want to be better at navigating complex ideas, comprehending challenging materials, and adjudicating between conflicting claims. As good inquirers we would be well-positioned to detect misleading evidence, misinformation, and deceit. I thus identify a second need: improving ourselves cognitively.

One effective way to improve ourselves cognitively, I argue, is to identify those with desirable and worthwhile epistemic traits and practices, and emulate them. I suggest that this is a characteristic function of understanding attributions, and I articulate what this proposal indicates about the nature and value of understanding and understanding attributions.

Finally, I evaluate competing proposals. These include proposing that attributions of a competing epistemic kind—such as knowledge, know how, or wisdom—serve to tag those worthy of intellectual emulation, and that understanding attributions serve to fulfil a different function, such as tagging experts or those who can provide explanations. I articulate weaknesses of these proposals, and I conclude by suggesting my favoured proposal can capture what is correct about the alternative proposals.

In section C I turn to the role of normalcy in thought. I begin with a recent debate in meta-philosophy. A common method in analytic philosophy challenges a target claim with a putative counterexample. An author describes a case and poses a question about it, such as whether an action described is permissible or a character depicted possesses knowledge. Reflecting on the vignette is meant to provide evidence for or against a theory or claim, or otherwise illuminate a philosophical subject matter. A well-known example of this is invoking a Gettier case to oppose the claim that knowledge is justified true belief. Timothy Williamson raises the question, what is the content of the operative judgement we have in response to counterexamples, such that it legitimately challenges the target theory. Given the role of this judgement in theorising, the judgement should be true, psychologically plausible and supported by the vignette. An account of this judgement should illuminate how and why vignettes can contribute to philosophical understanding.

In ‘Normalcy and the Contents of Philosophical Judgements’ I evaluate several competing proposals about the content of the judgement. Examples include ‘necessarily, anyone who stands to p in the relation described by the vignette has a justified true belief that p but does not know p’, ‘it is possible that someone stands to p as described by the vignette, and she has a justified true belief that p but lacks knowledge that p’, and ‘counterfactually, were someone to stand to p as described in the vignette, then she would have justified true belief that p but not know p’. I articulate weaknesses with each proposal, and I argue that these problems point to a promising alternative.

I propose that upon hearing a vignette, we think of what constitutes a normal, typical, generic, or natural instance of the vignette. If the normal versions of the vignette exhibit justified true belief without knowledge, then the vignette constitutes a challenge to the target theory. We judge something like, ‘Normally, whenever a person is related to a proposition such that the vignette obtains, then that person will have justified true belief without knowledge.’ I argue that the normalcy proposal is superior to the alternatives. This argument includes the claim that the normalcy proposal offers a unified account of other roles that vignettes play. Vignettes can motivate claims and principles, illustrate distinctions, communicate ideas, help identify the phenomena of interest, and frame research questions. And they provide evidential weight against claims that are not, or do not entail, necessity claims; these claims are not subject to straightforward refutation by counterexample. An account of the judgement is better if it unifies these roles. I argue that the normalcy account, unlike rival accounts, can explain the range of roles that vignettes play in philosophical theorising.

I conclude ‘Normalcy and the Contents of Philosophical Judgements’ by suggesting various ways one might develop the normalcy proposal. At present the proposal remains schematic, and ultimately its viability will be revealed only through further theorising.

The final essay, ‘The Burden of Proof and Statistical Evidence’, examines the epistemic requirements of legal standards of proof. Legal standards of proof, such as ‘beyond reasonable doubt’ or ‘preponderance of evidence’, are commonly interpreted as statistical likelihoods. In this essay I explain how ‘proof paradoxes’ challenge this interpretation of standards of proof. Proof paradoxes are pairs of cases. In one case the verdict of liability is



supported by highly probabilifying statistical evidence. In the second case the statistical evidence is weaker, but liability is supported by individualised, non-statistical evidence, such as eyewitness testimony. The statistical likelihood of guilt given the evidence is higher in the first case than in the second case. And yet, the paradox continues, the legal burden of proof is plausibly unsatisfied in the first case but satisfied in the second case. I explain how this suggests that the standard of proof is not simply a statistical likelihood.

I then evaluate three competing proposals for what the legal burden of proof requires. The first proposal, advanced by Judith Jarvis Thomson, holds that legal burdens of proof require individualised evidence with an appropriate causal relation to the person's liability. On Thomson's view individualised evidence, unlike merely statistical evidence, guarantees the claim is true.

The second proposal, advanced by Enoch, Spectre, and Fisher, holds that the burden of proof requires that the supporting evidence be sensitive to the liability of the accused. They argue that the burden of proof is satisfied only if (roughly speaking) were the claim false, the evidence would likely not obtain. This is because, Enoch, Spectre, and Fisher argue, relying on merely statistical evidence undermines incentives to obey the law; requiring that the standard of proof be reached through sensitive evidence generates the right incentive structure for citizens to obey the law.

Martin Smith argues that a legal burden of proof requires that the evidence normically supports liability. Evidence normically supports a conclusion when, roughly speaking, given the evidence,  $p$  would normally be true. If the evidence obtains yet  $p$  is false, some

abnormality or malfunction has occurred. The error demands explanation. Statistical evidence does not normically support a conclusion because, on Smith's view, normalcy does not reduce to statistical frequency.

In 'The Burden of Proof and Statistical Evidence' I critically evaluate the three proposals. I conclude that the normalcy proposal—or at least kindred proposal—is most promising for explaining the epistemic demands of legal burdens of proof. In conclusion I articulate what the normalcy proposal indicates about legal burdens of proof, and explain how it can capture what is correct about the other two proposals.

## Acknowledgements

Firstly, thanks to my advisor Ernest Sosa. Ernie is the most wonderful advisor anyone could wish for, and I am very fortunate to have worked with him during my doctorate. Thanks also to my committee members, Susanna Schellenberg, Jonathan Schaffer, and Liz Camp, who have provided invaluable comments on my research and invaluable support throughout my studies at Rutgers. Thirdly, thanks to my external advisor, Catherine Elgin, who inspired my focus on understanding as an undergraduate and continues to be an inspiration.

Special thanks also to Duncan Pritchard, Alvin Goldman, John Greco, and Stephen Grimm for their continued support. I am very grateful.

Many thanks to Ernest Sosa's dissertation group at Rutgers. The group provided valuable feedback on almost every essay I have written in the last four years. The members of the group during my dissertation were Austin Baker, Bob Beddor, David Black, Laura Callahan, Marilie Coetsee, Megan Feeney, Will Fleisher, and Lisa Miracchi.

Thanks also to Alex Antony, Boris Babic, Rima Basu, Peter Baumann, Renee Bolinger, Rodrigo Borges, Campbell Brown, Evan Butts, J. Adam Carter, Eddy Chen, Matthew Chrisman, Marcello Di Bello, Sandy Goldberg, Simon Goldstein, Emma Gordon, Michael Hannon, David Henderson, Doug Husak, Jesper Kallestrup, Jennifer Lackey, Ting An Lin, Peter Milne, Andrew Moon, Jessie Munton, Ted Poston, Wayne Riggs, Pamela Robinson, Daniel Rubio, Martin Smith, Kurt Sylvan, Chris Tweedt, Beth Valentine, Peter van Elswyk,

Alec Walen, Tedla Woldeyohannes, and anonymous reviewers for providing valuable comments on earlier drafts of these essays.

Thanks also to Louise Antony, Natalie Ashton, Yuval Avnir, Heather Battaly, Ben Bronner, Lara Buchak, Sam Carter, James Chase, David Coady, Andy Egan, Frances Egan, Branden Fitelson, Amy Flowerree, Miranda Fricker, Scott Friedman, Alex Geddes, Kristina Gehrman, Thony Gillies, E.J. Green, Alex Guerrero, Chris Hauser, Chris Kelp, Hilary Kornblith, Martin Kusch, Jon Kvanvig, Sarah-Jane Leslie, Clayton Littlejohn, Michael Lynch, Sarah Moss, Bernhard Nickel, Alasdair Richmond, Luis Rosa, David Rose, Cat Saint-Croix, Mary Salvaggio, Jennifer Saul, Mark Schroeder, Eli Shupe, Una Stojnić, Barry Stroud, Emily Sullivan, Josh Watson, Christopher Willard-Kyle for helpful discussions and insights.

Thanks also to the active members of the Facebook group ‘Board Certified Epistemologists’. The group is extremely helpful for locating relevant reading materials and for fruitful epistemological discussion. I am grateful to the group and to Clayton Littlejohn for founding and maintaining the group.

I am grateful to the Philosophy Departments at the University of Edinburgh, Northwestern University, Saint Louis University, and Rutgers University for their support. Many thanks to the ‘Varieties of Understanding’ and ‘Wisdom and Adversity’ Projects at Fordham University, the ‘Intellectual Humility’ Project at Saint Louis University, the Philosophy Department at the University of Tennessee, and the Templeton Foundation for the educational and research opportunities they made available during my graduate studies.

Parts of this work was presented at the *Episteme* Conference on Cognitive Ecology, Long Island Philosophical Society, the ‘Knowledge, Understanding and Wisdom Conference’, Bled, Slovenia, and the Pacific Division Meeting of the American Philosophical Association. Parts were also presented at Saint Louis University, Swarthmore College, the University of Houston, Underwood International College at Yonsei University, the Central European University, Northwestern University, Oxford University, the Pontifical Catholic University of Rio Grande do Sul, and at various seminars and meetings at Rutgers University and the University of Edinburgh. I am very grateful to those audiences for their feedback and insights.

Finally, thanks to Jon Garthoff for his immeasurable support.

### Acknowledgement of Previous Publication

Essay I is previously published as ‘Understanding, Integration, and Epistemic Value’ (2012) *Acta Analytica* 27(2): 163–181. Essay II is forthcoming as ‘Evidentialism and Moral Encroachment’ *Believing in Accordance with the Evidence: New Essays on Evidentialism*, ed. Kevin McCain, Springer. Parts of essay III were previously published as ‘Teleologies and the Methodology of Epistemology’ (2015) *Epistemic Evaluation: Purposeful Epistemology*, eds. John Greco and David Henderson, Oxford University Press, pp. 31–45, and parts of the essay are drawn from *A Defence of Robust Virtue Epistemology* (2011), which was my Master’s Thesis at the University of Edinburgh. Essay V was previously published as ‘Normalcy and the Contents of Philosophical Judgements’ (2015) *Inquiry* 58(7): 700–740. And essay VI is forthcoming as ‘The Burden of Proof and Statistical Evidence’ *Routledge Handbook of Applied Epistemology*, eds. David Coady and James Chase, Routledge.

## Table of Contents

Abstract .....	ii
Acknowledgements .....	xiv
Acknowledgement of Previous Publications .....	xvii
Table of Contents .....	xviii
 Introduction .....	 1
 A. Understanding and Value .....	 4
I. Understanding, Integration, and Epistemic Value .....	5
Abstract .....	5
1. Introduction .....	5
2. Veritism .....	9
3. The Value of Cognitive Integration .....	16
4. Veritist Responses .....	22
5. The Inadequacy of Veritism .....	26
6. Recapitulating .....	35
Acknowledgements .....	38
Bibliography .....	38
 II. Evidentialism and Moral Encroachment .....	 40
Abstract .....	40
1. Friendship and Evidence .....	40
2. Recent Challenges to Evidentialism .....	43
3. The Challenge from Moral Encroachment .....	46
4. The Inadequacy of Merely Statistical Evidence ...	52
5. Objections to Moral Encroachment .....	56
6. Resisting Moral Encroachment .....	63

7.	Understanding .....	72
	Acknowledgements .....	78
	Bibliography .....	78
	B. Teleologies and Understanding .....	82
	III. Teleologies and the Methodology of Epistemology .....	83
	Abstract .....	83
1.	Introduction .....	83
2.	Orthodox Approaches .....	85
	2.1 Extension-first approach .....	85
	2.2 Intension-first Approach .....	87
	2.3 Linguistic Approach .....	88
3.	Teleological Approaches .....	90
	3.1 Practical Explication .....	93
	3.2 Historical Genealogical Teleology .....	95
	3.3 Hypothetical Genealogical Teleology ....	97
4.	Roles for Teleologies .....	99
5.	Kornblith's Contention .....	106
6.	Conclusion .....	110
	Acknowledgements .....	110
	Bibliography .....	111
	IV. Understanding and Emulation .....	113
	Abstract .....	113
1.	Teleological Approaches .....	113
2.	Characteristics of Understanding .....	117
3.	Understanding and Emulation .....	123
4.	Other Potentially Emulation-Worthy Kinds ....	132
5.	Other Potential Functions .....	138
	Acknowledgements .....	139
	Bibliography .....	140



C. The Role of Normalcy in Thought .....	142
V. Normalcy and the Contents of Philosophical Judgements	143
Abstract .....	143
1. Introduction .....	143
2. Two Initial Proposals .....	145
3. The Necessity Proposal .....	150
4. The Possibility Proposal .....	157
5. Williamson's Counterfactual Proposal .....	167
6. The Emendation Approach .....	179
7. The Normalcy Proposal .....	183
7.1 Virtues of the Normalcy Proposal .....	188
7.2 A Unified Account .....	193
7.3 Resistance to the Proposal .....	196
7.4 Decision Points .....	200
Acknowledgements .....	206
Bibliography .....	206
VI. The Burden of Proof and Statistical Evidence .....	208
Abstract .....	208
1. The Golden Thread .....	208
2. Statistical Evidence .....	210
3. Causal Relations and Guarantees .....	213
4. Sensitivity and Incentives .....	222
5. Normic Support .....	229
Acknowledgements .....	235
Bibliography .....	236
Conclusion .....	240
Bibliography .....	248

## **In Pursuit of Understanding**

### **Introduction**

This series of essays has six main themes. The first is the nature and value of understanding, which is principally explored in the essays ‘Understanding, Integration, and Epistemic Value’, ‘Evidentialism and Moral Encroachment’, and ‘Understanding and Emulation’. The second theme is philosophical methodology, which is the topic of the essays ‘Normalcy and the Contents of Philosophical Judgements’ and ‘Teleologies and the Methodology of Epistemology’. Philosophical methodology also features substantially in ‘Understanding and Emulation’.

The themes of philosophical methodology and the nature of understanding are interrelated since focusing on understanding—rather than knowledge—as a central pursuit of philosophy ameliorates several putative meta-philosophical problems. Widespread, persistent disagreement is often thought to undermine knowledge, for example, and such disagreement is ubiquitous within philosophy. If philosophy principally pursues knowledge, this disagreement threatens the success of the enterprise. But this threat is ameliorated if we instead conceive of understanding as the central pursuit of philosophy. Understanding is more compatible with widespread, persistent disagreement, and so the existence of this disagreement does not threaten the viability of pursuing understanding in philosophy.

Within philosophical methodology, one of my aims is to explore and illuminate teleological approaches in epistemology. This third theme is the focus of section B, ‘Teleologies and

Understanding’, which comprises the essays ‘Teleologies and the Methodology of Epistemology’ and ‘Understanding and Emulation’.

The fourth theme is how to understand inference from merely statistical or inadequate evidence. This is the topic of the essays ‘Evidentialism and Moral Encroachment’ and ‘The Burden of Proof and Statistical Evidence’. I also discuss this topic in ‘Normalcy and the Contents of Philosophical Judgements’. In the dissertation’s conclusion I suggest ways that focusing on understanding as a central epistemic end, rather than knowledge, can illuminate how we should respond to weak or otherwise inadequate evidence.

A fifth theme is the role of normalcy in thought, which I explore in the last two essays ‘Normalcy and the Contents of Philosophical Judgements’ and ‘The Burden of Proof and Statistical Evidence’. In the dissertation’s conclusion I sketch some relationships between normalcy and teleology. In future research I shall examine further the relationship between understanding and the role of normalcy in thought.

Finally, many of these essays are situated within social epistemology. ‘Evidentialism and Moral Encroachment’ explores the ethics of belief and what we should believe about others based on their social groups. Many teleological approaches in epistemology are deeply social, since they focus on the social roles of epistemic kinds and attributions. This is the topic of the essays ‘Teleologies and the Methodology of Epistemology’ and ‘Understanding and Emulation’. Legal epistemology, which is a branch of social epistemology, is explored in ‘The Burden of Proof and Statistical Evidence’. Thus four of the six essays are within social epistemology.

Although my central aims in this dissertation are to illuminate the nature and value of understanding and to shed light on topics in meta-philosophy, the essays form a kind of constellation: different connections and themes appear from different perspectives. The essays can be read independently as stand-alone essays.

## A. Understanding and Value

## I. Understanding, Integration, and Epistemic Value

**Abstract.** Understanding enjoys a special kind of value, one not held by lesser epistemic states such as knowledge and true belief. I explain the value of understanding via a seemingly unrelated topic—the implausibility of veritism. Veritism holds that true belief is the sole ultimate epistemic good and all other epistemic goods derive their value from the epistemic value of true belief. Veritism entails that if you have a true belief that *p*, you have all the epistemic good qua *p*. Veritism is a plausible and widely held view; I argue that it is untenable. I argue that integration among beliefs possesses an epistemic value independent from the good of true belief, and so has value veritism cannot account for. I argue further that this integration among beliefs comprises the distinctive epistemic value of understanding.

### 1. Introduction

This paper has two themes. The first is to argue against veritism, the thesis that true belief is the sole ultimate epistemic good<sup>1</sup>; the second is to illuminate the nature and value of understanding. In particular I motivate the thought that reflecting on the nature of understanding should cast doubt on veritism. This is because understanding involves a kind of integration and structure among beliefs that is epistemically valuable independently from the value of true belief.

---

<sup>1</sup> Veritism is also known as epistemic value T-monism. See Pritchard (2011).

Understanding is enjoying a resurgence of interest in epistemology, and deservedly so, for it seems to have a distinctive value, one not held by other epistemic standings such as true belief and knowledge. Arguably understanding is the aim of scientific, historical, and other inquiry; understanding puts us in special contact with reality, perhaps more powerfully than merely knowing isolated facts. If we understand something we grasp how it works, or why it is thus. In this paper I examine the value of understanding. I do not offer a full account of the nature of understanding, but I begin by briefly sketching two paradigmatic kinds of understanding of interest to epistemology in order to highlight some key features of understanding.<sup>2</sup> One kind of understanding is when we understand an object such as a system, structure, or pattern. This could be, for instance, a historical period, a machine, a person or a field of study. This kind of understanding is usually attributed by the verb followed directly by a noun, with no ‘that’ clause:

Jill understands American history.

Tom understands his wife.

Lorna understands car engines.

Steve understands calculus.

Following Kvanvig we can call this objectual understanding.<sup>3</sup>

---

<sup>2</sup> For similar accounts of understanding see Grimm (2012), Kvanvig (2003; 2009: chapter 8), Elgin (2006; 2009), and Riggs (2003; 2004).

<sup>3</sup> See Kvanvig (2003; 2009). Kvanvig contrasts objectual understanding with propositional understanding. The latter is understanding that, understanding when, understanding who etc. He holds that understanding of explanations reduces to understanding that. However this is a controversial view (see Grimm (2012)). An alternative view holds that understanding of explanations reduces to objectual understanding, as causal structures or explanations are objects that we can comprehend in a way isomorphic to our

The other kind of understanding is when we understand why something happens. This is understanding of explanations:

Sally understands why the Democrats lost the election.

Joan understands why her house burnt down.

Typically when an agent possesses objectual understanding she knows facts about the object, and often enjoys a kind of know-how associated with the object (it is hard to see how Laura understands car engines if she does not know how to fix or build any, and if Steve can never obtain correct results using calculus we would doubt that he understands it). When someone understands an explanation she knows facts about the causes and antecedent conditions, and can often reason counterfactually (if Sally understands why the Democrats lost the election then she can reason about the outcomes where antecedent conditions were different, say, if they had more money). Understanding is more resilient to false belief: we can understand something even if some, or a great deal, of the beliefs we hold about it are false or merely approximately true.<sup>4</sup> But these features do not fully capture the nature of understanding; perhaps they do not even capture the essential nature or defining property of understanding. When we understand an object we also grasp relations among the various parts of that object; we see its structure. When we understand an explanation we see how different elements of the causal or explanatory web hang together and how various facts

---

comprehension of other objects such as transport systems and political systems. Whether and how kinds of understandings reduce to each other is an important question, however it falls outwith the focus of this paper.

<sup>4</sup> See Kvanvig (2003), Riggs (2009), and Elgin (2006; 2009).



interact. We understand which features are important and which are peripheral. It seems that in general when we think about the nature of understanding, as distinct from true belief, knowledge or other epistemic standings, what springs to mind is coherence among beliefs and a grasping of the relations between parts.<sup>5</sup> This structure and integration, I suggest, is a central feature of understanding, and it is the nature of this coherence that I hope to illuminate.<sup>6</sup>

The route I shall take, however, is unusual. I hope to shed light on the nature and value of this coherence via insights about the ultimate bearers of epistemic value. This debate divides veritists—those who hold that true belief is the sole ultimate bearer of epistemic value, and so anything that possesses epistemic value does so in relation to the value of true belief—and those who deny this claim, either because they hold that some other epistemic good explains all other epistemic value,<sup>7</sup> or because they are pluralist about the sources of ultimate epistemic value. In this paper I motivate pluralism: I argue that integration among beliefs, such as that characteristic of understanding, has an epistemic value that is not reducible to the value of true belief.

In section two I explain veritism and explore the claims that veritists are committed to. I then present two cases in section three which suggest a value independent from the value of true belief. In section four I examine the stances a veritist could take in response to these

---

<sup>5</sup> Kvanvig (2009: 96).

<sup>6</sup> For more on the integration of beliefs being characteristic of understanding, see Kvanvig (2003; 2009: chapter 8). See also Elgin (2006; 2009), Grimm (2012), and Riggs (2003; 2004).

<sup>7</sup> Williamson, for instance, would probably endorse knowledge monism. See Pritchard (2011). See also Williamson (2002).

cases. In section five I explain why these responses are not promising strategies. In section six I draw together the strands of my argument to show how the two themes of the paper have developed.

## 2. Veritism

When we think something has value we can always ask the further question, in virtue of what is it valuable? Perhaps it is valuable that I donate to Oxfam, for instance. We can ask, in virtue of what is this donating valuable? Suppose its value lies in its enabling Oxfam to provide resources to those in need. We can then ask, ‘in virtue of what is providing such resources valuable?’ A plausible response is: because it contributes to human flourishing. We can then ask the question again, ‘in virtue of what does human flourishing have value?’ At this stage the explanation may bottom out. Human flourishing is perhaps not valuable in virtue of another valuable thing. Following Pritchard, I call this ultimate value.

**Ultimate value.** Something whose value is not derivative from the value of any other good.

In the example flourishing is an ultimate good, it doesn’t derive its value from the value of anything else. We can contrast this with non-ultimate, or dependent, value. In the example giving to Oxfam is a non-ultimate good—it is valuable, but only because flourishing is valuable. The value of donating is dependent on the value of flourishing. In this case the dependency relation is instrumentality—the donation is instrumental in causing flourishing—but dependency relations are not exhausted by instrumentality relations. Something can be (non-ultimately) valuable in virtue of being constitutive of an (ultimate) good, or an enabler of an ultimate good, or a necessary condition, or a part of an ultimately

good whole, and so on.<sup>8</sup> The key question, therefore, is not whether the phenomenon is valuable because it *causes* an ultimately valuable thing, but whether its value derives *in some way* from the value of something else.

Suppose there are different categories of value, such as epistemic value, moral value, and aesthetic value, then we can ask a variant of the above question: Of something that has epistemic value, in virtue of what does it have epistemic value? Telescopes have an epistemic value, for instance, and we can ask ‘in virtue of what do they have epistemic value?’ Suppose they are epistemically valuable because they allow us to see distant objects. In virtue of what does seeing distant objects have epistemic value? Suppose seeing distant objects has epistemic value because we can thereby gain true belief and avoid false belief. Perhaps this has epistemic value because true belief (and avoiding false belief) in general has epistemic value. And perhaps, like flourishing in the moral case, this is where our chain of questioning bottoms out: There is no other epistemic good that explains the epistemic value of true belief; true belief is of ultimate epistemic value. For simplicity I shall henceforth refer to ‘true belief’, and ‘gaining true belief’, but I mean this to include avoiding false belief as well.

There may additionally be a *practical* value to telescopes and the true beliefs they engender: learning about distant objects for instance has practical value because we can see the enemy approaching, spot land whilst navigating at sea, or identify fruitful hunting ground from afar. True belief may have a moral or aesthetic value too. We may be morally obliged to not remain wilfully ignorant, for instance. It may even be the case that if there were no such

---

<sup>8</sup> My account of ultimate value is largely based on Pritchard (2011). Pritchard however, does not discuss dependency relations other than instrumental relations (cf. Pritchard (2011: 2)).

moral or practical value to true belief, then there would be no epistemic value to it either. That is to say that all epistemic value is, at base, dependent on non-epistemic value. All epistemic value, according to this suggestion, is ultimately derivative on non-epistemic goods. However we should note that this suggestion is consistent with there being ultimate epistemic values, epistemic goods that explain the epistemic value of all other epistemic goods. In other words, if something has epistemic value we can ask our question ‘in virtue of what epistemic good does it have epistemic value’, and this question is coherent and important even if all epistemic value is ultimately derivative from non-epistemic value.

Our chain of questioning about the epistemic value of telescopes has led to the suggestion that true belief is of ultimate epistemic value. That is, perhaps there is no further epistemic value that grounds the epistemic value of true belief. This view makes sense of much of our epistemic behaviour: once we judge a belief to be true we cease inquiry into the matter, conversely we find it hard to be satisfied epistemically if we don’t think we have reached the truth. It seems that belief aims at truth, furthermore, and in some sense belief is proper if and only if true. When we ask someone to testify, we want her to tell the truth. All this suggests that true belief plays a keystone role in our epistemic lives; perhaps it is what gives value to all other epistemic practices and values?

That true belief is of ultimate epistemic value is a widely held and plausible view. It is not one that I shall dispute here. But there is a further question, whether true belief (and avoiding false belief) is the *sole* ultimate epistemic good, or are there others in addition to true belief. This question asks whether *all* epistemic goods are valuable in virtue of their

relation to true belief, or if some have epistemic value in virtue of other epistemic goods. Those who hold that true belief is the sole ultimate epistemic good are known as veritists.

**Veritism.** Having true belief (and avoiding false belief) is the sole ultimate epistemic good.<sup>9</sup>

Those who deny veritism can join one of three families: they can retain monism about the ultimate source of epistemic value, but hold that something other than true belief, such as knowledge, understanding or justification, grounds the epistemic value of all other epistemic goods. Alternatively they may be pluralist and hold that true belief is not one of the ultimate epistemic values, or they may be pluralist, but argue that whilst true belief may be an ultimate epistemic value, it is not the sole ultimate epistemic value. Pluralists may hold, for instance, that gaining true belief and avoiding false belief does not exhaust the sources of epistemic value because having beliefs that are well grounded is a second, independent source of epistemic value. Some support for this view is that even if a belief is true, if it is held for bad reasons, such as superstition, guessing or coin tossing, then something valuable is missing from the epistemic point of view. Such pluralists will argue that since the irrationally-held belief is nonetheless true, the fact that something of value is missing suggests that not all epistemic value bottoms out in truth.

This debate between veritists and their detractors about the ultimate source of epistemic value is interesting in its own right, and may have practical ramifications for how we conduct

---

<sup>9</sup> Note that this account slightly differs from other articulations of veritism in the literature, such as Pritchard (2011). These differences do not concern us here, and my arguments against veritism, insofar as they are successful also challenge veritism as Pritchard renders it, in virtue of undermining the ‘ultimate value’ aspect of his account.

inquiry (because when we are pursuing epistemic value, our epistemic enterprises should be shaped by what has epistemic value). This debate also has consequences for other areas of epistemology: veritism plays a central role in formulations of the Meno Problem, for instance.<sup>10</sup> I believe that arguments for value pluralism are fruitful in another area of epistemology: they illuminate the distinctive nature of understanding. Progress on the question of what has ultimate epistemic value is notoriously difficult, however, with intractable stalemates as both sides brutally intuit what possesses ultimate epistemic value. A methodological question arises, how can we isolate epistemic value, and how can we infer what value derives from which goods. In addition to illuminating the value of understanding, I hope to make some progress on the veritism/pluralism debate, by suggesting how the dialectic might proceed, and offering some metaphilosophical points about the debate.

We can put pressure on veritism by first exploring what the position is committed to. We are then better positioned to see what a case against veritism may look like. Veritism holds that true belief (and avoiding false belief) is the sole ultimate epistemic good. This means that all epistemic goods are valuable insofar as they tend to result in true belief. But notice this makes all other epistemic goods dependently valuable on the value of true belief. To see why, recall that donating to Oxfam was only valuable because it was instrumental in providing much needed resources, which itself was only valuable because it would increase flourishing. But now we can ask: what if Oxfam and similar aid organisations become redundant? What if there were no more people in need of resources, indeed everyone is self-sufficient, maximally happy and free from suffering? Now it seems that there is no longer

---

<sup>10</sup> The Meno Problem is the problem of explaining why knowledge is more valuable than mere true belief. See, for instance, Riggs (2009). For how veritism relates to the Meno Problem see Pritchard (2011).

value in giving to Oxfam. The value of donating was dependent on the value of that which it could produce (resources); if sufficient resources are already present, donating no longer has value.<sup>11</sup> This suggests a very general truth about dependent, as opposed to ultimate, goods. If something, A, is valuable in virtue of something else, B, and B is already present, then A is no longer valuable. This is a widely held thesis about value:

**General axiological thesis.** If the value of a property is only derivatively valuable relative to producing a further good, and that good is already present, then the property confers no additional value.<sup>12</sup>

When we combine veritism and the general axiological thesis, we can see that veritism entails that epistemic goods such as justified beliefs, reliable belief forming methods, competent testifiers and the like are epistemically valuable, but only insofar as they lead to more true beliefs. Once we have obtained the true belief, then the belief is not more valuable in virtue of being justified, reliably formed etc. The value of these other properties is swamped by the fact that the belief is true. Thus veritism is committed to the swamping hypothesis:

**Swamping hypothesis.** If a belief is true then it cannot be any more valuable in virtue of having other properties (such as being formed well, being formed by a reliable method etc). The value of the other properties is swamped by the fact that the belief is true.<sup>13</sup>

The swamping hypothesis entails that if you have a true belief that  $p$ , you have all the epistemic good qua  $p$  (and if you have every possible true belief you are maximally epistemic

---

<sup>11</sup> This example is reminiscent of Zagzebski's (1996) coffee machine analogy.

<sup>12</sup> See also Pritchard (2011: 5), from where I adapted this formulation of the general axiological principle.

<sup>13</sup> For discussions of the swamping hypothesis see, Kvanvig: (2003 chapters three and four), Pritchard (2011), Riggs (2002), and Zagzebski (2003).

good). We certainly see veritists explicitly endorsing these claims, but even where they don't explicitly endorse it, the swamping hypothesis is entailed by veritism combined with the general axiological thesis. From the swamping hypothesis we obtain the following supervenience thesis:

**Supervenience thesis.** If two sets of beliefs are identical with regard to the value of true and false beliefs, and propensity to gain more true beliefs and avoid false beliefs, then they have equivalent levels of epistemic value.

By making explicit the commitments of veritism in this way, it becomes clearer how we can make a case against veritism.<sup>14</sup> Firstly, we could put pressure on veritism by challenging the swamping hypothesis. This would occur if we find an agent who is epistemically deficient even with maximally true beliefs. The deficiency points to epistemic values not reducible to the epistemic value of true belief. The first case in section three is of this kind.

We could also put pressure on veritism by challenging the supervenience thesis. If we find two belief sets that differ in epistemic value, even though there is no difference with regard to truth (they contain the same amount of true belief, and propensity to gain further true belief), then this suggests veritism is false. Again, this is because the difference in epistemic value points to a source of epistemic value independent from the epistemic value of true belief, and so falls outwith the bounds of veritism. This is the structure of the second case introduced in section three.

---

<sup>14</sup> Anecdotally, the way I got into this project, was wondering what a case against veritism would look like, and how could we make progress on this seemingly intractable debate.



The epistemic deficiency of the agent in the first kind of case and the difference in value between the two belief sets in the second kind of case cannot be explained by reference to true beliefs. This is because in both cases the volume of true beliefs is not at stake. Something else must explain the value judgements. This second thing is either a second ultimate epistemic good or else it is a non-ultimate epistemic good whose value derives from an ultimate epistemic good other than true belief. With either of these disjuncts there is at least one epistemic end independent from true belief. Note that the method of using the cases to elicit the thought that something is missing from the epistemic point of view cannot distinguish between the disjuncts (that the missing property is an ultimate epistemic good itself or it derives its value from an ultimate epistemic good independent from true belief), but either disjunct is problematic for veritism as both are incompatible with the claim that true belief is the sole ultimate epistemic good.

The dialectic stands: veritism cannot account for an epistemic value that isn't merely instrumental value towards gaining true beliefs and avoiding false beliefs, so if we can locate such a value, veritism should be considered false. The cases I introduce in section three motivate the view that structure and integration among beliefs has such a value.

### **3. The Value of Cognitive Integration**

Take the following case. Markus believes wildly irresponsibly. He does not care about believing truly or with evidential justification. He randomly believes propositions, and does not enquire about their truth. Unbeknownst to Markus an omnipotent being alters the world

ensuring that everything Markus believes is true. Markus cannot have a false belief. His belief forming method is reliable and modally stable.<sup>15</sup>

Intuitively, however, from the epistemic point of view he could be epistemically better off. Markus' beliefs lack integration: as he never reflects on his beliefs, none of his beliefs serve as reasons in favour of any of his other beliefs and he never sees connections among his beliefs. In Kvanvig's terms, he does not grasp the 'coherence-making relationships' among his beliefs.<sup>16</sup> By integration I mean both have basing relations among beliefs (that is believing *p* based on believing *q*), and also grasping relations such as this belief is an instance of that general rule; this belief is similar to that belief; this belief is about the same thing as that belief. Markus does not reflect on his beliefs, so he does not grasp those connections; his beliefs do not form a web of beliefs. It is precisely these relations that lead to coherence and understanding.

Recall that I need to establish that i.) something about Markus is epistemically missing, and ii.) what is missing is epistemically valuable. If I establish these two things, given that Markus is not lacking anything qua the truth of his beliefs, we should consider veritism false. I have argued the first bit—that cognitive integration is missing. There is reason to believe that integrated true beliefs are more valuable than non-integrated true beliefs. Kvanvig argues

---

<sup>15</sup> For those familiar with the Temp Case introduced by Greco to motivate the evolution from simple reliabilism to agent reliabilism, it is like that case, except that it is all the agent's beliefs, not just the ones about temperature. See Greco (1999).

<sup>16</sup> Kvanvig (2003: 192).

that finding such relationships organises and systemises our thinking beyond mere addition of more (justified) true beliefs.<sup>17</sup>

A veritist might reply that organised, systemised thinking is only valuable epistemically insofar as it is instrumental to attaining further true belief. It seems however that holding fixed non-epistemic ends and volume of true beliefs it is preferable to have organised, integrated true beliefs. Given our cognitive limitations, systemised and organised beliefs are better as they allow a richer picture epistemically (allowing us to perceive patterns, connections, which truths are general principles, which facts are more important, how true beliefs link, and so on). These things are epistemically valuable.

An organised, systematised belief set allows us to reason from one belief to another and draw conclusions for ourselves. An agent who has coherent, connected beliefs doesn't merely believe isolated facts, but can really grasp for herself why they are true. Kvanvig notes that this is a more satisfying closure to inquiry; if we can draw inferences ourselves, and see connections among facts, we thus sense that we fully grasp the subject matter.<sup>18</sup> This distinction is illustrated when we consider the difference between on the one hand reading a book where we virtuously accept each individual proposition in isolation, and on the other hand actively grasping how the arguments in the book hang together, believing the propositions because we see how they follow from previous facts. With the former mode of reading we may believe of each proposition that it seems plausible, and that we trust the author, and so affirm the proposition with justification. But with the latter mode of reading

---

<sup>17</sup> Kvanvig (2003: 202).

<sup>18</sup> Kvanvig (2003: 202).

we actively grasp the relationships among these facts and how the conclusions must follow. We gain the subjective experience of the subject matter making sense and hanging together. Kvanvig's point about inquiry *feeling* complete once we have grasped the coherence making properties of a body of facts is a psychological one, but maybe it is not merely a psychological fact: perhaps inquiry *is* complete once we have actively and successfully grasped these relations among facts. If this is right then Markus never reaches the end of inquiry. He believes truly, but does not fully understand the subject matter, because he does not grasp how it all hangs together.

Kvanvig also cites William James' view that in our epistemic lives we are strongly averse to the feeling that we are being duped or are missing something.<sup>19</sup> Kvanvig offers the intriguing suggestion that grasping this structure among beliefs provides some sense of reassurance that we are not missing something. We will feel that we have grasped the whole picture. And again to 'de-psychologise' the point, then it will be some protection from *in fact* missing something: in virtue of having grasped the relations between facts in a body of information, rather than merely believing isolated facts, then missing items will be noticed. Similarly I suspect that well integrated beliefs protect epistemic agents against being duped, because when our beliefs cohere we feel that we grasp the reasons for ourselves, and in fact it is harder to be fooled.

Of course, the veritist might object here. She may hold that all of these properties that Markus lacks are either true beliefs (such as the true beliefs about relative importance of other facts and about which beliefs ground other beliefs) or are valuable only insofar as they

---

<sup>19</sup> See Kvanvig (2003: 202–3) and James (1897).

lead to further true beliefs (such as the capacity to reason from known beliefs to other truths or to infer that you are not missing something). Such a response would safeguard veritism, because I would have failed to show an epistemic value that veritism cannot countenance. I return later to the value of integration, whether such value merely reduces to the value of true belief, and whether grasping connections among beliefs merely reduces to having more true beliefs, however first let's look at the second case against veritism. Here I show how two belief sets can be identical with respect to true belief yet nonetheless have differing epistemic value. This puts pressure on the supervenience thesis, and accordingly on veritism.

This second case doesn't use human epistemic agents, which allows me to control for more factors, thus sharpening our intuitions about the case. It edits out the 'noise'. We may be tempted to respond to the Markus case, for instance, by thinking that he is such a terrible epistemic agent, or such an exotic case, that questions regarding his epistemic status have little bearing on the real world. Consider two data collection robots, Alpha and Beta.<sup>20</sup> Alpha has a set of truths recorded as a list (fact 1, fact 2, fact 3...). Beta has a set of truths, identical in propositional content, but Beta also records relations among truths.

As I see it Beta has something akin to hyperlinks found in websites such as Wikipedia. They send you from one page to another, and this creates a web of relevance and relations among facts. Beta also utilises tools such as subheadings to convey non-propositionally what is important and how facts relate to each other. Beta's beliefs are structured. Alpha's beliefs, by

---

<sup>20</sup> By using data collection robots—robots whose task is up-taking data and storing it—I hope to isolate the epistemic domain. The robots do not process information for practical ends or navigate around the world. They merely seek to represent the world in their memory system. Isolating the epistemic domain in this way may help us to have sharper intuitions about purely epistemic value.

contrast, is merely stored as a list.

Beta is better. There is more epistemic value held within Beta's memory system. Were we to use these robots to gain understanding about the world using the data they stored, Beta gives us a head start. Using Alpha we would need to grasp for ourselves the relations among the truths listed by Alpha, so would have to do more work. Beta stores them for us. Beta's system of facts is more akin to having understanding (and understanding, presumably, is epistemically better than merely truly believing). The difference between Alpha and Beta is precisely those connections that form basing relations among beliefs, allow beliefs to be justified by other beliefs, and link beliefs together in our thinking. They are coherence-making properties. Given that the difference between Alpha and Beta is not a difference in true propositions stored, the additional value of Beta is not supplied by the volume of truth.<sup>21</sup> Thus comparing Alpha and Beta puts pressure on the supervenience thesis:

**Supervenience thesis.** If two sets of beliefs are identical with regard to number of true beliefs and propensity to gain more true beliefs, then they have equivalent levels of epistemic value.

And so it accordingly puts pressure on veritism.

---

<sup>21</sup> To pre-empt an objection at this stage: A reader may argue that I am allowing Beta to have more truths than Alpha, because Beta is storing 'extra' truths such as the relations between facts. In response, however, for any such relation that Beta has (such as 'Fact 2 is an instance of the general law expressed in Fact 1') we should allow that Alpha stores that same fact in its list. If Beta stores a connection, Alpha can write about the connection on its list. Alpha may even have *more* propositions as a result. I shall later argue that in spite of holding these true beliefs about the relations, Alpha's list fails to store the relations in the valuable way that Beta does.

#### 4. Veritist Responses

At this juncture we should pause and take stock. I have introduced two cases I think are problematic for veritism. The first is a case where the agent seems to have an epistemic deficiency even though all of his beliefs are true and whenever he forms a new belief it is guaranteed true. This threatens veritism by challenging the swamping hypothesis. Secondly we have the robots case, where two sets of stored propositions that are identical with regard to truth nonetheless have different epistemic values. This puts pressure on the supervenience thesis.

There are four responses to these cases available to the veritist. The second and third are specific to the robots case, the first and fourth apply to both cases. The first, which I mentioned above, holds that the structure and integration possessed by Beta, but lacking in Alpha and Markus, is merely instrumentally valuable for attaining further true beliefs (or in some other way derivatively valuable from the value of true belief). The second response holds that there is something amiss about the Alpha/Beta case, perhaps because the two robots are in fact identical in their epistemic properties, or one of the robots is not possible. Either way the approach challenges the cogency of the example. The third response holds that Alpha and Beta actually have the same amount of epistemic value. The fourth admits that whilst the structure and integration possessed by Beta, and lacking in Alpha and Markus, is valuable the *nature* of this structure can be accounted for in terms of true beliefs. This response holds that the structure among beliefs can be reduced to further beliefs—and so can be countenanced by veritism (albeit in a slightly indirect way). This is because a veritist can maintain that everything of epistemic value is valuable derivatively from the value of true belief, *or its nature can be reduced to true belief*. My focus is largely on the fourth possible

response, but first I look at the other three responses, because I think these responses reveal interesting contours of the veritism/pluralism debate.

So to return to the first response, the veritist can acknowledge that the integration among beliefs that Beta stores (and Alpha and Markus both lack) has epistemic value. But they hold that this value is derivative from the value of true belief. If Markus's belief were cognitively integrated, for instance, then he would form more true beliefs in the future. And Beta's structure is valuable because it allows us to form more true beliefs than Alpha's list of facts does. Thus, they maintain, the value of grasping connections among beliefs, having a coherent belief set, and basing beliefs on others can be explained by veritism: it's instrumental to future true beliefs.

And the veritist can give a story to bolster the response: when considering the Markus case we are asked to imagine that the truth of his future beliefs is secured by the demon, so that the truth of his beliefs aren't dependent on having integrated beliefs in the way that a typical epistemic agent's are. But this aspect of the thought experiment doesn't stop us, in practice, from judging that the missing integration among beliefs is valuable. This is because we are used to thinking about the real world, where integration is instrumentally valuable. We base our value judgements on what we are accustomed to valuing, and we can't be moved to amend our value judgements so easily. Thus, the veritist would argue, the thought experiment is inadequate—being asked to imagine that Markus' beliefs are guaranteed true is



not sufficient to enable us to abstract away from the prosaic value of integration when evaluating his somewhat more exotic epistemic life.<sup>22</sup>

But bear in mind that my aim is to put pressure on veritism, by making the position look less attractive, and here does seem to be an area where holding the veritist thesis in this way looks like foot-stomping or question begging. By simply denying our intuitions about the cases, their position is in danger of reducing to brute assertion that the integration is valuable solely because it leads to more true belief, and that the ‘richer picture’ integration enables just means possessing more true beliefs. It therefore makes veritism a less attractive position.

The second available response denies the cogency of the example, perhaps by saying that Alpha and Beta are actually identical or one of them is not possible. Without hearing the particular form of the objection it is hard to formulate a reply. But as a pre-emptive response, it certainly seems as though Alpha and Beta are different beings, and that they are intuitive and possible beings. Furthermore cognitive science models cognition in ways that distinguish them, and render them cogent and useful models of different kinds of information storage.<sup>23</sup> The burden of proof seems to be on the veritist.

---

<sup>22</sup> This response is similar to Goldman’s explanation of why we do not judge a reliable clairvoyant to know the results of her reliably clairvoyanced beliefs. This case is problematic for Goldman because according to his account of knowledge a clairvoyant knows, yet we typically intuit that she does not know. He holds that intuitive judgements about the case exhibit a failure of our imaginative capacities (Goldman 1993: 279).

<sup>23</sup> My preliminary investigations with cognitive scientists are a cause for optimism: they like the comparison between the two robots, and think it illustrates something important about cognition and computing. A future project will be to model Alpha and Beta using the tools of cognitive science, to better understand how the kinds of integration Beta stores map onto human cognitive resources and processes.

The third response holds that the two robots actually have the same epistemic value. They could thus maintain, as seems natural to veritism, that a list of true beliefs encompasses all epistemic value, and that further epistemic properties, such as the ‘Wikipedia style’ links, contribute no further value. They would deny that robot Beta is epistemically more sophisticated or further along the road towards understanding. I find that implausible. Certainly if I was buying one I would choose robot Beta. But I find this third response of interest because my investigation grew from asking what a counterexample to veritism would look like, whether we could ever truly argue against the thesis, and more generally how we could make progress in this debate. If these two robots have the same value, this casts doubt on the distinction between pluralism and veritism. I would start to suspect that something fishy is going on and that the debate is specious. This is because we need to be able to illustrate a substantial difference between two sides of a debate. If Alpha and Beta have the same amount of value, I would wonder what the distinctive claims of veritism are (it seems as though it can account for anything).

The fourth response is the one I talk about in greatest depth. The response concedes that Beta has more value, but argues that everything Beta has can be stored as propositions (as a list of facts), so everything Beta possesses can be stored by Alpha. If this is true my example doesn’t threaten at least some forms of veritism. This is because the veritist can grant that, structure and integration among beliefs has value, and the value isn’t merely derivative from the value of true belief (such as being instrumental in gaining more true beliefs in the future), but then contend that the nature of the structure can be reduced to true belief. This response protects the veritist notion that when it comes to epistemic value, it all boils down to true belief.

One way to view the dialectic is that Alpha is the showcase of veritism: Veritism must provide a positive account of Alpha, where Alpha has the capacity to reach maximum epistemic value. This is because veritism holds that true belief is the ultimate epistemic good and all other epistemic goods are merely derivatively epistemically valuable. Alpha's cognitive state is similar to a set of true beliefs, and so it shouldn't be missing anything of epistemic value. If Alpha is missing anything, this indicates some epistemic good independent from the value of true belief. To defend Alpha veritists thus have to argue for at least one of the four responses.

### **5. The Inadequacy of Veritism**

We have now articulated a promising route for veritism. Whilst perhaps a slight variation on a paradigmatically veritist position, it is certainly within the spirit of veritism, and so should be amenable to adherents of the position. Veritists can allow that structure and integration among beliefs has epistemic value, and this epistemic value is not merely instrumentally valuable for attaining further true beliefs in the future. But they can defend the veritist position by arguing that such integration and structure is actually propositional in nature or can be reduced to further true beliefs. This would mean that everything that Beta stores can be stored by Alpha. Alpha, the poster child of veritism, is not inferior to Beta after all.

But this response won't work. Relations among beliefs, including both basing relations (believing *p* based on believing *q*) and structural relations (such as 'this truth is an instance of this general rule, and 'this truth is similar to that truth') among a coherent web of beliefs do not reduce to further beliefs. They have a non-propositional element. To explain why I draw

on Stroud's essay 'Inference, Belief and Understanding'.<sup>24</sup> Stroud argues that structural dependencies among beliefs can never be fully accounted for by adding more beliefs. In other words, the way beliefs are connected in the cognition of an agent can never be captured propositionally.

Stroud introduces Carroll's story 'What Achilles Said to the Tortoise'.<sup>25</sup> In the tale the tortoise refuses to make inferences 'required by logic'. He'll believe the proposition A, 'things that are equal to the same are equal to each other', and the proposition B, 'the two sides of this triangle are things that are equal to the same'. But the tortoise refuses to make the inference and thereby believe on the basis of A and B the further proposition Z 'the two sides of this triangle are equal to each other'.

The tortoise will accept a supplementary hypothetical proposition C 'if A and B are true then Z is true',<sup>26</sup> but he still will not make the inference and thereby believe Z. He will accept the further hypothetical D 'if A, B, and C are true, then Z is true', but accepting these four propositions still will not suffice for him to believe Z. He will not put the beliefs together and thereby infer the conclusion. In addition to the original premises, the tortoise is willing to accept further iterative hypothetical propositions, each of which 'strengthens' the argument by stating that the conjunction of the original premises implies the conclusion (and each of which he asks Achilles to write down because whatever logic teaches him is

---

<sup>24</sup> Stroud (1979).

<sup>25</sup> Carroll (1895).

<sup>26</sup> Following Stroud I have altered the modality of the conditionals. In Carroll's original paper the conditional C is 'If A and B are true Z must be true'. See Stroud (1979: 179).

important enough to write down).<sup>27</sup> A regress quickly develops. The point that philosophers usually take Carroll to be making is a point about logic: logical dependencies among logical relations cannot be captured within the content of the logical relations themselves. It is something external to that content.

Stroud dissents from the standard interpretation. He argues that the ‘jumps’ or inferences the tortoise refuses to make do not reveal anything about logical consequence and truth. As Stroud says, ‘the conjunction of A and B does imply Z, whatever the tortoise happens to think about it’.<sup>28</sup> Insofar as we consider the logical relations themselves, ‘strengthening’ the argument is a futile project, because the logical relations were always there, and always implied the conclusion and will not get any better at implying that conclusion by adding premises that express the relations among the premises and the conclusion. Stroud notes, however, that when we think about *belief* in the propositions, things are different. Sometimes ‘strengthening’ in the epistemic sense of making explicit the relations among propositions (such as relation C) that were always ‘present’ but that we had not realised, can bring us to accept Z. Such epistemic strengthening is not futile: We might tell a friend ‘if you believe A, and you believe B, then you must believe Z’, and they thereby come to infer Z. And so whilst what a proposition logically entails does not change, what an agent infers from it can change. But as Stroud notes, telling our friend the relation between two propositions that entail each other does not offer a further reason to believe Z, but rather it comments on the

---

<sup>27</sup> Thomson calls such an argument ‘strengthened’ when it includes all the original premises, and a supplementary premise which states that the conjunction of the original premises imply the conclusion. Thomson (1960: 97).

<sup>28</sup> Stroud (1979: 180).

original reasons that the agent possessed.<sup>29</sup> And Stroud explains how the fact that these supplementary premises (C, D, E etc.) are not additional reasons can reveal an important point about inference and belief. In his words, it ‘raises a quite general problem about one’s beliefs being inferred from or being based on another, whether the one belief implies the other or not.’<sup>30</sup>

To see why what it means to believe P based on Q does not involve another reason, such as reason R connecting P and Q, imagine that we are trying to represent someone’s beliefs, including which beliefs are based on others. In the same way that Achilles wrote the logical propositions in his notebook, our ‘epistemologist’s notebook’ aims to capture in sentences someone’s cognitive state.

We begin with a simple case: the agent believes P based on believing Q. We write down P and Q. But we have not yet captured the relationship in the agent’s mind between P and Q. Merely recording that she believes P and Q does not record that she believes Q based on P. Even if P logically entails Q, it might be that she has never made the connection between P and Q, and so believes Q for some other reason. We could write down the further fact that she believes Q based on P but as Stroud notes, this would be unsatisfactory because it leaves unexplained what it means to base one belief on another. It simply states that beliefs are so based, but does not explain what this means.<sup>31</sup> So we are tempted to write down that the agent sees some relation R between P and Q, and thus assume that we have recorded what it means to base a belief in Q on P. R could be the proposition that ‘P implies Q’, or ‘P is a

---

<sup>29</sup> See Stroud (1979: 182–4). See also Thomson (1960: 100–2).

<sup>30</sup> Stroud (1979: 184–85).

<sup>31</sup> Stroud (1979: 186).

reason to believe Q'. Then it seems we have recorded in propositional form P, Q and that the agent sees the relationship between them. But herein lies the rub. The agent might believe that P, P implies Q, and Q, but yet have never put the beliefs together. She may believe Q for an entirely different reason. But notice that if we write down S, a proposition that expresses the relationship between P, R, and Q (which may state for instance 'the conjunction of P and R implies Q'), we still haven't recorded what it is for the agent to have put those beliefs together in her mind because, of course, she may never have put the propositions P, R and S together. She may believe Q for an entirely unrelated reason. We haven't recorded that she sees the relation among the parts, has put them together in her mind, and thereby believed Q. In Stroud's words:

A list of everything that a person believes, accepts or acknowledges must leave it indeterminate whether any of those beliefs are based on others [...] Even if God himself looked into our heads and inspected all the members of our beliefs set, He could not thereby determine whether any of our beliefs are based on others. That is not a question which can be settled by any facts, however complex, about what we know, believe or accept.<sup>32</sup>

Carroll's story thus illustrates that if we want to record that the agent believes one thing based on another, we must add a non-propositional element. Structure among beliefs cannot be reduced to beliefs.<sup>33</sup> This structure must be something else, such as a competence, disposition or practical capacity. Whatever it is, Stroud argues, it is something captured by an account of belief, but not the content of beliefs.<sup>34</sup>

---

<sup>32</sup> Stroud (1979: 187–8).

<sup>33</sup> Stroud (1979: 189).

<sup>34</sup> Stroud (1979: 194).

This structure and integration among beliefs is how we justify beliefs and form new beliefs. It is paramount in cognition. The structure and integration is how everything ‘fits together’; they are the coherence making properties that are constitutive of understanding.<sup>35</sup> Stroud doesn’t talk about value in his paper, but in section three we saw some reason to think that these structural and coherence making properties have value. And now we have reason to think they cannot be reduced to mere truths, as the veritist’s fourth response claims. Where does this leave the veritist? To return to the two cases above: what Markus lacks cannot be expressed in terms of propositions, or even *believed* propositions. Because he doesn’t make connections among beliefs himself he is missing something that cannot be added as further isolated beliefs. We could allow Markus all the true beliefs in the world, but he would still be missing something of epistemic value. We also now have reason to think that with just a list

---

<sup>35</sup> When Stroud talks about the grasping of relations among beliefs being constitutive of understanding, the paradigmatic kind of understanding he has in mind is linguistic understanding—what it means to understand a sentence (Stroud (1979: 192)). He does not restrict himself to linguistic understanding, (Stroud (1979: 192, n. 2)) however, and it may be that focusing on these relations forges a connection between linguistic understanding and epistemic understanding (the latter being the kind that epistemologists study, where paradigmatic examples include understanding of a system or subject and understanding of explanations). The suggestion would be that linguistic understanding of a proposition requires the ability to make at least some of the connections and inferences from the proposition, (and so if an agent cannot do this she doesn’t linguistically understand the sentence and so can not even believe it, let alone epistemically understand it), but epistemic understanding requires seeing *more* of these connections, and fitting the proposition into a wider web so that it coheres with other things you believe. We may say that we (linguistically) understand that ‘hydrogen and oxygen combine to make water’ once we can grasp at least some inferences entailed by it, such as ‘this water contains some oxygen’. But we have epistemic understanding only once we are positioned to make many more connections between this truth and other facts, such as the workings of the Periodic table, the principles of chemistry, and the behaviour of water. It may be that both linguistic understanding and epistemic understanding come in degrees, depending on how many such connections you can make, and that at their boundaries they shade into each other: a really good linguistic understanding of a sentence allows you to see so many connections to other facts that it is akin to weak epistemic understanding of the subject matter.



of propositions Alpha cannot store structure among facts. Beta has an epistemic advantage because what Beta has—structure among facts—Alpha cannot store.

Stroud's reasoning is reminiscent of regress arguments employed against various articulations of Plato's Doctrine of the Forms and against various theses of Nominalism combined with the doctrine of transcendental universals.<sup>36</sup> The regress arguments I have in mind are Bradley's Regress and the Relational Regress (Armstrong attributes the latter to Ryle). They are both descendants of Plato's Third Man Argument.<sup>37</sup> The details of the how the arguments apply in metaphysics and the differences between the different regresses do not concern us here, and so we can safely leave those aside, but I shall sketch the basic structure of the regress, and explain how it relates to the Alpha and Beta robots case, and thus how it applies to the question of veritism.

Suppose we have a metaphysics that comprises objects such as chairs and tables, which are particulars, and properties such as redness and bachelorhood, which are universals.<sup>38</sup> We can then ask what it means for a particular to instantiate a universal. We have a red table for instance, but what is the relationship between the universal property of redness and the particular object the table? We can posit a relation R—the relation of instantiation—between the particular (the table) and the universal (redness): the table instantiates redness. But what

---

<sup>36</sup> For a useful summary of these regress arguments see Armstrong (1978: esp. chapters 7 and 10).

<sup>37</sup> See Bradley (1897: chapter 3) and Ryle (1939: 137–8) for the attribution to Ryle see Armstrong (1978: 70). Plato's Third Man Argument, unlike the regress arguments discussed here, use principles of self-predication and one-over-many to get going.

<sup>38</sup> Bradley describes these particulars as having a basic substrate, on which the properties can be 'bound'. The instantiation relation may also be known as 'participation'—the table being red can be expressed as 'the object (table) participates in the universal (redness)'.

does it mean for the relation of instantiation to obtain in this case? It seems that to explain relation R we need to posit another relation, the relation of instantiation between the particular and relation R. We can call this second relation R'. But relation R' remains unexplained. If R' is the same relation as R, then we have explained the (first-order) relation of instantiation using the (first-order) relation of instantiation, and it is a circular analysis.<sup>39</sup> If R' is different from R, we will need to posit a third relation, R'', which holds between R' and R. But what does it mean for R'' to be instantiated in this case? The regress begins. This regress suggests that what it means for a particular to instantiate a universal cannot be analysed in terms of a relation between the particular and the universal, where a relation is taken to be something which itself must be instantiated.<sup>40</sup>

How does this regress apply to the veritist debate? Recall Alpha and Beta. Beta stores data using non-propositional structures, such as 'hyperlinks' that record connections among facts. Alpha, by contrast, stores facts solely as a list. To safeguard veritism, veritists must establish that Alpha has the capacity to store anything of epistemic value that Beta can, and thereby maintain that everything of epistemic value can be reduced to true propositions. Suppose that Beta stores the following facts,

Fact 1. Ice is lighter than water (general law).

Fact 2. This piece of ice is lighter than this piece of water (instance).

---

<sup>39</sup> See Armstrong (1978: 72).

<sup>40</sup> See Armstrong (1978: 107).

Notice that fact 1 is a general law of nature and fact 2 is an instance of that general law. Beta can record this relation between those two facts.

Fact 3. Fact 2 is an instance of the general law articulated in fact 1.

But Beta can also store the relation expressed by fact 3 non-propositionally, as ‘hyperlink’ of some kind between fact 1 and fact 2. The connection itself (the structure) between fact 1 and fact 2 is made available in Beta’s system. Now we turn to Alpha. Alpha can also store fact 1, fact 2, and fact 3. But suppose Alpha wants to store the structure of the three facts, namely that fact 3 is expressing a relation between fact 1 and fact 2. Whilst the isolated facts are recorded, the actual connection between 1 and 2 is not yet made. We may ourselves realise the relation as we read the facts—we may put them together in our minds using our cognitive capacities and dispositions—but Alpha has not itself stored the relation. Alpha can record another fact, fact 4, which expresses that fact 3 relates fact 1 with fact 2. But, as should be no surprise, this further fact does not establish a connection between the isolated propositions. The facts can be listed, but their relation is not registered. This regress, which is isomorphic to the regress above, suggests that a relation such as instantiation between two facts is not stored in the facts themselves, and cannot be stored by a further fact, but instead must be stored by something non-propositional. Alpha can store facts about connections and relations, but the regress argument suggests that the actual structural relation—the connection itself—is something that Alpha cannot store. The implication for veritism isn’t just the implausibly large number of beliefs that it needs to posit to gain structure among beliefs, as every connection needs to be explicitly recorded—although that alone is a troublesome consequence for veritism—but the real problem for veritism is that such

connections cannot in principle be forged with the addition of further mere beliefs. Structure requires some non-propositional element.

So whatever it is for two propositions to be connected, it seems it is something that Alpha cannot grasp. And if, as Kvanvig and I suggest, such relations have value which is not merely instrumental value for gaining more true beliefs, then this is a problem for veritism.

## **6. Recapitulating**

To recap, Stroud's argument suggests that structure among beliefs—how beliefs cohere or are 'put together' in the mind—cannot be reduced to the content of beliefs themselves. The regress argument sketched above similarly suggests that structure among facts (such as instantiation relations) cannot be reduced to relations where those relations are expressible by facts, as this leads to a regress. Both conclusions can fruitfully be summarised as suggesting that structure among (believed) propositions does not reduce to propositions.

In section three we examined reasons for thinking that structure among beliefs in cognitive systems has epistemic value, and that this value is not merely reducible to the value of gaining more true belief. To this end we considered Markus, who seemed to be lacking something epistemic even though his beliefs were all true. Part of what he lacked stems from the isolated nature of each belief; his beliefs are not well integrated and he doesn't grasp the coherence making features of his belief set. He lacked any understanding of the subject matter. We also considered Beta, who seems a more powerful cognitive tool than Alpha, even though both Alpha and Beta store large numbers of truths (and in fact every truth stored by Beta can be stored by Alpha, it is just the relations among those truths that Alpha

cannot store). The difference in value between Alpha and Beta stems from the structure among the 'beliefs' stored by Beta.

When we put these two threads together veritism is not as tenable as it initially seems. Integration among beliefs has epistemic value not reducible to the epistemic value of true belief, and the nature of this integration cannot be reduced to further true beliefs. This suggests that true belief is not the sole ultimate epistemic good. I present the argument here.

P1.) Qua propositions stored, Alpha and Beta are identical (by hypothesis).

P2.) Beta is epistemically better than Alpha.

P3.) What Beta has cannot be captured by a list of stored propositions (by Stroud and regress argument).

C1.) Therefore, something of epistemic value (structure) can not be captured by stored propositions (by P2, P3).

P4.) That thing (structure) is not valuable merely as an instrument to further true beliefs (Kvanvig's point; Markus' beliefs are guaranteed true; no difference in truth between Alpha and Beta, understanding has a distinctive value which seems to lie in the way beliefs cohere).

C2.) Something (structure) has epistemic value that is not merely instrumentally valuable to true belief, and that thing is not reducible in nature to stored propositions.

Conclusion two puts pressure on veritism because if something (structure among propositions) has epistemic value that is not merely good in virtue of the value of true belief, then this suggests that traditional veritism, as outlined in part two, is false. If, in addition, that valuable thing is not reducible in nature to true belief, as Stroud and the regress argument suggests, then this puts pressure on the 'veritist retreat' I described above. The

veritist retreat allows that some epistemic properties (such as structure among beliefs) have epistemic value that isn't merely in virtue of the epistemic value of true belief, but holds that the nature of these properties can be reduced to true belief. They can thereby maintain the veritist notion that everything of epistemic value in some way 'comes down to' the value of true belief. But if structure among beliefs cannot be reduced to further true beliefs, this retreat fails.

But in addition to putting some pressure on veritism, I hope also to have established something about the nature of understanding. Recall that coherence among beliefs is a characteristic feature of understanding and when understanding obtains the various parts of an explanation or subject matter 'hang together' in the agent's mind. In other words, when an agent understands, the agent's beliefs are organised and systematised. As Kvanvig writes,

When the question is whether one has understanding, the issues that are foremost in our minds are issues about the extent of our grasp of the structural relationships (e.g., logical, probabilistic, and explanatory relationships) between the central items of information regarding which the question of understanding arises.<sup>41</sup>

The nature of this cognitive structure remains underexplored in philosophical literature, but it does seem to account for at least some of the distinctive value of understanding.<sup>42</sup> I hope to have motivated the thought that the organisation and systematisation characteristic of understanding comprise the same kind of relations that are not reducible to the content of beliefs. This suggests the special epistemic value of understanding cannot be captured in

---

<sup>41</sup> Kvanvig (2009: 3).

<sup>42</sup> In fact I suspect that the cognitive integration characteristic of understanding is the *sole* property that explains the special epistemic value of understanding over and above other epistemic standings such as knowledge and true belief, but this claim lies outwith the arguments in this paper.

terms of propositions believed. Instead it is about the way beliefs are put together within the agent and this is non-propositional. That the special value of understanding lies in something non-propositional, and is therefore divorced from issues of truth, may shed light on why understanding can be so valuable even when some of the believed propositions are false, models, or mere approximations. It explains why when we consider the nature of understanding, truth takes a back seat—non-propositional connections come to the fore, and the question shifts to whether the structural relationships are grasped.<sup>43</sup>

### Acknowledgements

Thanks to Bob Beddor, David Black, Campbell Brown, Liz Camp, Andy Egan, Will Fleisher, Scott Friedman, Jon Garthoff, Emma Gordon, David Henderson, Jon Kvanvig, Peter Milne, Duncan Pritchard, Ernest Sosa, and Barry Stroud for comments on earlier drafts of this work and useful suggestions. Thanks also to audiences at the University of Edinburgh, Northwestern University, and at the ‘Knowledge, Understanding and Wisdom Conference’, Bled, Slovenia for helpful questions.

### Bibliography

- Alston, William (2005) *Beyond Justification* Cornell University Press.
- Armstrong, David (1978) *Nominalism and Realism, Volume One* Cambridge University Press.
- Carroll, Lewis (1895) ‘What the Tortoise said to Achilles’ *Mind* 4(14): 278–80
- Elgin, Catherine (2009) ‘Is Understanding Factive?’ *Epistemic Value* Adrian Haddock, Alan Millar, and Duncan Pritchard (eds.) Oxford University Press, pp. 322–30.
- \_\_\_\_\_ (2006) ‘From Knowledge to Understanding’ *Epistemology Futures* Stephen Hetherington (ed.) Oxford: Clarendon Press, pp. 199–215.
- Gardiner, Georgi (ms) ‘Understanding and Emulation’.
- \_\_\_\_\_ (2012) ‘Understanding, Integration, and Epistemic Value’ *Acta Analytica* 27(2): 163–81.
- Goldman, Alvin (1993) ‘Epistemic Folkways and Scientific Epistemology’ *Philosophical Issues* 3: 271–85.
- Greco, John (2010) *Achieving Knowledge* Oxford University Press.
- \_\_\_\_\_ (2009) ‘The Value Problem’ *Epistemic Value* Adrian Haddock, Alan Millar and Duncan Pritchard (eds.) Oxford University Press, pp. 313–21
- \_\_\_\_\_ (1999) ‘Agent Reliabilism’ *Philosophical Perspectives* 13: 273–96.
- Grimm, Stephen (2012) ‘The Value of Understanding’ *Philosophy Compass* 7(2): 103–17.

---

<sup>43</sup> This essay is previously published as Gardiner (2012).

- Hannon, Michael (ms) 'What's the Point of Understanding?'
- Hills, Alison (2016) 'Understanding Why' *Noûs* 50(4): 661–88.
- James, William (1987) 'The Will to Believe' *The Will to Believe and Other Essays* NY: Longmans.
- Korsgaard, Christine (1983) 'Two Distinctions in Goodness' *Philosophical Review* 92(2): 169–95.
- Kvanvig, Jonathan (2009) 'The Value of Understanding' *Epistemic Value* Adrian Haddock, Alan Millar, and Duncan Pritchard (eds.) Oxford University Press, pp. 95–112.
- \_\_\_\_\_ (2008) 'Pointless Truth' *Midwest Studies In Philosophy* 32(1): 199–212.
- \_\_\_\_\_ (2003) *The Value of Knowledge and the Pursuit of Understanding* Cambridge University Press.
- Nagel, Jennifer (2007) 'Epistemic Intuitions' *Philosophy Compass* 2(6): 792–819.
- Pritchard, Duncan (2011) 'What is the Swamping Problem?' *Reasons for Belief* Andrew Reisner and Asbjørn Steglich-Petersen (eds.) Cambridge University Press, pp. 244–59.
- \_\_\_\_\_ (2009) 'Knowledge' *Central Issues of Philosophy* John Shand (ed.) Wiley-Blackwell.
- \_\_\_\_\_ (2007) 'Recent Work on Epistemic Value' *American Philosophical Quarterly* 44: 85–110.
- Riggs, Wayne (2009) 'Knowledge, Understanding, and the Meno Requirement' *Epistemic Value* Adrian Haddock, Alan Millar, and Duncan Pritchard (eds.) Oxford University Press, pp. 331–8.
- \_\_\_\_\_ (2004) Review of 'The Value of Knowledge and the Pursuit of Understanding' by Jonathan Kvanvig, *Notre Dame Philosophical Reviews*.
- \_\_\_\_\_ (2003) 'Understanding "Virtue" and the Virtue of Understanding' *Intellectual Virtue: Perspectives from Ethics and Epistemology* Michael DePaul and Linda Zagzebski (eds.) Oxford University Press, pp. 203–226.
- \_\_\_\_\_ (2002) 'Reliability and the Value of Knowledge' *Philosophy and Phenomenological Research* 64: 79–96.
- Stroud, Barry (1979) 'Inference Belief and Understanding' *Mind* 86: 179–196.
- Williamson, Timothy (2000) *Knowledge and its Limits* Oxford University Press.
- Zagzebski, Linda (1996) *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge* Cambridge University Press.



## II. Evidentialism and Moral Encroachment

**Abstract.** Moral encroachment holds that the epistemic justification of a belief can be affected by moral factors. If the belief might wrong a person or group more evidence is required to justify the belief. Moral encroachment thereby opposes evidentialism, and kindred views, which holds that epistemic justification is determined solely by factors pertaining to evidence and truth. In this essay I explain how beliefs such as ‘that woman is probably an administrative assistant’—based on the evidence that most women employees at the firm are administrative assistants—motivate moral encroachment. I then describe weaknesses of moral encroachment as a critique of these beliefs. Finally I explain how we can countenance the moral properties of such beliefs without endorsing moral encroachment, and I argue that the moral status of such beliefs cannot be evaluated independently from the understanding in which they are embedded.

### 1. Friendship and Evidence

In her essay ‘Epistemic Partiality in Friendship’ Sarah Stroud argues that sometimes friendship requires that our beliefs not fit the evidence. Being a good friend, Stroud argues, can require epistemic partiality.<sup>1</sup> On hearing a disturbing anecdote about our friend, Stroud suggests, friendship demands we sometimes resist believing what the available evidence

---

<sup>1</sup> See also Keller (2004), Hazlett (2013; 2016), and Piller (2016). For discussion see Ryan (2015), Kawall (2013), and Crawford (forthcoming). I do not think Stroud (2006) establishes that the norms of friendship conflict with orthodox epistemic norms, but I do not evaluate this claim in this paper.

indicates. Instead we should disbelieve the story or re-interpret it to reflect less poorly on our friend. We should think well of friends and give them the benefit of the doubt. This duty to our friends does not extend to non-friends. In short, Stroud argues, we ought to be epistemically biased towards our friends. Stroud writes,<sup>2</sup>

Friendship positively demands epistemic bias, understood as an epistemically unjustified departure from epistemic objectivity. Doxastic dispositions which violate the standards promulgated by mainstream epistemological theories are a constitutive feature of friendship. Or, to put the point as succinctly—and brutally—as possible, friendship requires epistemic irrationality.

If epistemic norms of impartiality—norms demanding that doxastic attitudes reflect available evidence—genuinely conflict with the demands of friendship, what ought we do? Stroud articulates three broad options:<sup>3</sup> Perhaps, given the indispensability of friendship for a good life, when epistemic norms conflict with the requirements of friendship, ‘so much the worse for epistemic rationality’.<sup>4</sup> According to this first option, we have most reason to be epistemically irrational; demands of friendship override the demands of epistemic rationality. The second option maintains that when the norms conflict there is no overriding ought—there is nothing that all-things-considered you should do. There are simply two conflicting norms: what you should do as a friend and what you epistemically should do.

The third option Stroud considers holds that the tension between epistemic demands of friendship and the orthodox view of epistemic normativity indicates the received

---

<sup>2</sup> Stroud (2006: 518).

<sup>3</sup> See also Hazlett (2013), Heil (1983), Aikin (2006), Preston-Roedder (2013), and Enoch (2016) for discussion of how to understand conflict between epistemic norms and the requirements of friendship or morality.

<sup>4</sup> Stroud (2006: 519).

understanding of epistemic norms is inadequate. We epistemically should be partial to our friends; we are not committing an epistemic error when we believe against the evidence in favour of friends. According to the third option, epistemic norms ought to reflect the distinctive, partial epistemic demands of friendship. Stroud writes,

If standard epistemological theories condemn as irrational something which is indispensable for a good life—so that we have compelling reason *not* to comply with the demands of those theories—then perhaps we should question whether those theories offer an adequate account of epistemic rationality after all. Why accept a conception of epistemic rationality on which it is something which we have very strong reasons to avoid. It might be better to rethink the assumption that epistemic rationality requires the kind of epistemic objectivity or impartiality from which friendship seems necessarily to depart [...] [R]ather than concluding that friendship is epistemically irrational, we could instead conclude that our previous ideas of epistemic rationality were too narrow.<sup>5</sup>

The third option holds that if something is indispensable to the good life, epistemic norms must answer to the epistemic demands of that domain. In some cases it is *epistemically* permissible or required to not proportion belief to the evidence, and instead believe in a way that promotes flourishing, friendship, or some other ideal.

Stroud doubts the third option is viable, and dubs it ‘unattractive’ and ‘dubiously available’.<sup>6</sup> Epistemic norms, Stroud holds, seem to answer to attaining the truth and avoiding falsehood, reflecting evidential considerations, and aiming at knowledge and understanding. The epistemic domain is independent from other pursuits, such as friendship or happiness.<sup>7</sup>

---

<sup>5</sup> Stroud (2006: 522, emphasis in original).

<sup>6</sup> Stroud (2006: 519–22).

<sup>7</sup> See Adler (2002), Shah (2006), Kelly (2002), and Chignell (2010). This view is widely regarded as orthodoxy, although see Grimm (2011).

Although Stroud doubts the viability of the third option—that epistemic norms answer to the demands of domains such as friendship—she notes a virtue of the option. If epistemic norms genuinely conflict with the demands of friendship, then—given the indispensability of friendship—epistemic norms relinquish their claim to overriding authority about what we ought to do and believe. Sometimes we ought not be epistemically rational. Forgoing the priority of epistemic norms represents a substantial cost, Stroud notes, since epistemic norms are usually taken to be authoritative. The third option preserves the overridingness of epistemic norms.

Committed evidentialists might at this juncture emphasise the availability of a fourth option, mirroring the first: if there is a genuine conflict between the norms of epistemic rationality and the epistemic demands of friendship, well, so much the worse for friendship. Perhaps friendship, like frenemies and nemeses, are things that we overall ought not cultivate. This fourth option, whilst unappealing, retains the overridingness of epistemic norms.

## **2. Recent Challenges to Evidentialism**

Although Stroud was skeptical about its prospects, recently several theorists have endorsed the third option. These theorists re-interpret epistemic norms to reflect perceived normative demands from other domains. If friendship, morality, or agency require particular doxastic attitudes, these attitudes are epistemically permitted or required. There is nothing *epistemically* improper about other considerations influencing belief. This flood of views opposes evidentialism, which holds that epistemic justification depends solely on the available

evidence, and kindred ‘intellectualist’ positions that maintain epistemic justification depends solely on truth-related factors.<sup>8</sup>

Berislav Marušić (2015), for example, advocates the following principle,

If we should sincerely promise or resolve to  $\phi$ , it is rational to believe that (we will  $\phi$  if we sincerely promise or resolve to  $\phi$ ).

Marušić claims a person should believe she will successfully fulfil her promises and resolutions despite evidence indicating she will fail. A person should believe she will stop smoking, for example, even if the available evidence indicates she will relapse. On Marušić’s view such beliefs are not beholden to evidential considerations; evidential considerations are the wrong standards for evaluating beliefs about one’s own promises or resolutions.<sup>9</sup>

Clayton Littlejohn (2012) argues a special class of normative beliefs cannot be both justified and false. This means some beliefs—such as beliefs about what one morally ought do—cannot be epistemically justified if they are morally mistaken. Since a non-moral belief that was evidentially supported to the same degree would be epistemically justified, Littlejohn’s view opposes the evidentialist principle that whether a belief is epistemically justified

---

<sup>8</sup> I do not argue for evidentialism; instead I defend evidentialism against arguments pressed by advocates of moral encroachment.

<sup>9</sup> See James (1956/1897) and Aikin (2008) for discussion of a related Jamesian idea: that antecedent beliefs concerning a prospective friendship might be necessary conditions for the success of the friendship, before the evidence supports those beliefs. The (evidentially unsupported) beliefs are thus necessary for their own (future) truth.

depends solely on whether the belief fits the evidence. Moral considerations bleed into epistemic normativity.<sup>10</sup>

Rima Basu (ms b) argues there is a moral error with treating people as subjects to be studied and predicted. Basu invokes Sherlock Holmes as exemplifying this error. Holmes makes observations, inferences, and predictions about others with a scientific or disinterested perspective. Basu argues this is a moral mistake, even when the resulting belief is neutral or positive, such as inferring what the person ate for breakfast based on arcane clues or predicting the person has likely read *The New Jim Crow* because she is an African American scholar. Basu holds this moral mistake bears on the epistemic rationality of such beliefs.<sup>11</sup>

Mark Alfano (2013: chapter four) suggests attributing virtues to others in the absence of evidence can be epistemically permissible because such attributions can be self-fulfilling. The attribution causes the person to conduct themselves in ways consonant with the virtue possession, and so contributes to its own truth. Crucially for Alfano's opposition to evidentialism, attributions of vice do not share this permissibility: if the epistemic permissibility stems wholly from evidence concerning self-fulfilling prophecies, and vice attributions were also self-fulfilling, attributions of vice would also be epistemically permitted. Alfano's view opposes evidentialism because he claims moral facts influence what

---

<sup>10</sup> Thanks to Clayton Littlejohn for helpful discussions on this topic.

<sup>11</sup> Thanks to Rima Basu for helpful discussion of these issues. Armour (1994: 795) suggests 'race-based predictions of an individual's behaviour insufficiently recognize individual autonomy by reducing people to predictable objects rather than treating them as autonomous entities' but, unlike Basu, Armour does not claim this is a distinctively epistemic error.

one epistemically ought believe. In Alfano's words, one ought 'to speak and to think what ought to be'.<sup>12</sup>

Mark Schroeder (forthcoming) agrees with Stroud that how we should evaluate evidence concerning loved ones depends on whether the evidence reflects well or poorly on them, and argues we should interpret their behaviour partially. But unlike Stroud, Schroeder argues this is required by *epistemic* normativity. Given the high stakes of such beliefs, Schroeder argues, it is an epistemic error to form beliefs about loved ones impartially. The importance of our loved ones in our lives provides epistemic reason to withhold belief and interpret evidence in a partial manner.<sup>13</sup>

### 3. The Challenge from Moral Encroachment

For the remainder of this essay I focus on one family of recent opposition to evidentialism, namely the challenge from moral encroachment. In sections three and four I articulate the putative conflicting demands that motivate moral encroachment. I then, in section five, survey some problems with moral encroachment, which provide motivation to deny the view. In sections six and seven I explore how evidentialism, and kindred views, can explain the apparently conflicting normative demands without eschewing evidentialism.

---

<sup>12</sup> Alfano (2013: 108).

<sup>13</sup> For further examples of recent theorists arguing that factors deemed non-epistemic by orthodox epistemology bear on the epistemic status of a belief, see Rinard (2015; 2017), McCormick (2015), Pace (2011), Dotson (2008; 2014), Ross and Schroeder (2014), Stanley (2005; 2015, especially chapter six), Fantl and McGrath (2002), Guerrero (2007), and Buchak (2014). These theorists either argue that epistemic norms answer to norms in other domains, or deny there are distinctively epistemic norms. For further discussion, see also Hazlett (2016), Fritz (2017), Natalie Ashton (2015), Ashton and McKenna (ms), and Arpaly (2003: chapter 3).

Several theorists have recently argued that in some cases if a claim concerns a morally significant subject matter we epistemically ought to be more inclined to suspend judgement. If a belief might wrong a person or group, the threshold for justified belief is higher than for a belief that is morally neutral. More evidence is required to justify the belief. These theorists advocate *moral encroachment*: moral features of a belief can affect whether the belief is epistemically justified.<sup>14</sup>

**Moral encroachment.** What it is epistemically rational for a person to believe can, in at least some cases, be affected by moral factors.

Advocates of moral encroachment deny that epistemic reasons for or against belief are exhausted by considerations pertaining to evidence and truth.<sup>15</sup> Moral considerations can bear on epistemic justification.

To illuminate the position, consider the following three examples:

**The Cosmos Club.** Historian John Hope Franklin hosts a party at his Washington D.C. social club, The Cosmos Club. As Franklin reports, ‘It was during our stroll through the club that a white woman called me out, presented me with her coat check, and ordered me to bring her coat. I patiently told her that if she would present her coat to a uniformed attendant, “and all of the club attendants were in uniform,” perhaps she could get her coat’. Almost every attendant at the Cosmos Club is black and few members of

---

<sup>14</sup> Advocates of moral encroachment include Basu (ms a, ms b, ms c, ms d), Schroeder (forthcoming), Basu and Schroeder (forthcoming), Moss (forthcoming), Bolinger (ms), and Pace (2011). See also Munton (ms), Fritz (2017), and Enoch (2016) for discussion. See also Arpaly (2003: chapter 3) for related discussion. Note that Arpaly’s discussion concerns the normativity of *false* moral beliefs, and in Arpaly’s view morally wrong beliefs also exhibit an epistemic error.

<sup>15</sup> Schroeder (forthcoming) specifies that on his view there are only non-evidential epistemic reasons *against* belief; there are no non-evidential epistemic reasons *for* belief.



the club are black. This demographic distribution almost certainly led to the woman's false belief that Franklin is an attendant.<sup>16</sup>

**Administrative Assistant.** A consultant visits an office. He knows that few people visit the office who are not employees of the firm and that almost every woman employee is an administrative assistant. The consultant sees a woman walking down the corridor and forms the belief 'she is an administrative assistant'.<sup>17</sup>

**Tipping Prediction.** Spencer works as wait staff at a restaurant. He sensed that white diners tipped more than black diners. He researched the trend online, and read about a well-documented social trend that black diners tip on average substantially lower than white diners. Spencer weighs the evidence before reaching his belief about the social trends. A black diner, Jamal, enters Spencer's restaurant and dines in a booth outside of Spencer's area. Spencer predicts Jamal will tip lower than average for the restaurant, and later discovers his prediction was correct.<sup>18</sup>

Advocates of moral encroachment argue these beliefs are morally wrong despite being based on evidence that renders the claim likely true and, in the third vignette, being true. But, they argue, this does not exemplify a tension between moral requirements and epistemic permissibility. Since the relevant belief or evidence is a kind that can morally wrong, it is either the wrong kind of evidence to support belief or the evidence fails to justify the belief because of the high stakes. The belief based on merely demographic, statistical, or weak evidence is epistemically faulty, and this is because of the moral significance of the belief.

---

<sup>16</sup> See Franklin (2005: 4; 340) and Gendler (2011). Gendler invokes this example to illustrate a putative tension between the demands of morality and the demands of epistemic normativity. Basu (ms a), Schroeder (forthcoming), Basu and Schroeder (forthcoming), and Bolinger (ms) have since invoked Franklin's experience to motivate moral encroachment. See also the similar 'Mexican restaurant' case discussed by Schroeder and Basu (forthcoming) and Basu (ms b: 5). This kind of error is ubiquitous. As Obama observes in Westfall (2014), 'there's no black male [his] age, who's a professional, who hasn't come out of a restaurant and is waiting for their car and somebody didn't hand them their car keys.'

<sup>17</sup> Adapted from Moss (forthcoming).

<sup>18</sup> Adapted from Basu (ms a: 3).

Just as there are several variants of pragmatic encroachment, there are also several varieties of moral encroachment. Some theorists maintain the belief is epistemically wrong because it fails to eliminate a salient relevant alternative, such as the woman's not being an administrative assistant or that Jamal will leave a large tip. The relevant alternative is rendered salient by its moral import. Since the person fails to eliminate relevant alternatives the belief is not epistemically justified.<sup>19</sup> Some theorists maintain the belief has high stakes. The costs of being wrong—or the accumulated costs of error when many people commit the same error—contribute to systemic marginalisation. Given the high stakes, more evidence is required.<sup>20</sup>

Basu (ms b) maintains one should refrain from believing based on facts that are due to racism, and these cases exemplify this error. Basu writes,

[The woman in the Cosmos Club vignette] ignores a relevant moral feature of her environment: the fact that she relies on—the South's racism—makes her ignorant to the way in which she wrongs by forming beliefs about individuals on the basis of facts that are due to racism. Whereas facts may not themselves be racist, they can be the result of racism and racist institutions and policies, thus when forming beliefs on the basis of them it seems appropriate to ask for more moral care. (p. 12)

and,

It is *the history of racism at the Cosmos Club* that makes relying on race, despite it being the best indicator and the strongest evidence that someone is a staff member (in the context of the Cosmos Club), problematic. That is the moral stake in question that an epistemically responsible agent must be sensitive to. If the best evidence that someone is

---

<sup>19</sup> See for example, Moss (forthcoming) and Bolinger (ms). For the role of relevant alternatives in epistemology, see Lewis (1996) and Dretske (1970).

<sup>20</sup> Basu (ms a, ms c), Schroeder (forthcoming), Basu and Schroeder (forthcoming), and Fritz (2017). Bolinger (ms) also emphasises the epistemic significance of the harms of error, including the aggregate harms of many people committing the same errors based on demographic evidence. For more on the role of stakes in pragmatic encroachment, see Fantl and McGrath (2002) and Stanley (2005).

a staff member is a consequence of an unjust and racist policy, then you still need to look for more evidence. (p. 14, emphasis in original)

Renee Bolinger and Sarah Moss focus on the epistemic wrong of forming beliefs about people based on purely statistical evidence, and argue that moral factors render such beliefs epistemically flawed. Schroeder and Basu indict beliefs on ‘insufficient’ evidence, regardless of whether that evidence is merely statistical or otherwise inadequate.

Details aside, the key to the criticism of evidentialism and kindred views is that evidence that would normally suffice for belief is rendered epistemically insufficient by moral features. These claims are in tension with evidentialist claims that what one epistemically ought believe is solely a function of evidential considerations, and that epistemic justification supervenes on strength of available evidence. The claims oppose any ‘intellectualist’ position that holds epistemic justification depends solely on truth-relevant factors. In what follows I focus mainly on the moral encroachment view advocated by Basu and Schroeder, but draw on ideas advanced by Bolinger and Moss.

This recent tide of anti-evidentialist thought takes as its starting point the indisputable fact that society is structured by racist institutions. Given this, advocates of moral encroachment argue, some of our evidence will be racist or will support racist conclusions. If we believe according to this evidence, as evidentialism and other orthodox epistemological views require, our beliefs will thereby be racially biased. The same applies *mutatis mutandis* for sexist,

homophobic, transphobic, and other such prejudiced beliefs.<sup>21</sup> Schroeder articulates a challenge: He writes,

Gendler argues that in cases like [the Cosmos Club] there is a conflict between epistemic rationality and avoiding implicit bias—given underlying statistical regularities in the world, many of which are directly or indirectly caused by past injustice, perfect respect for the evidence will require sometimes forming beliefs like the woman in the club. But the belief that the woman forms is racist. I hold out hope that epistemic rationality does not require racism. If it does not, then the costs of [the woman’s] belief must play a role in explaining why the evidential standards are higher, for believing that a black man at a club in Washington, D.C. is staff. And I believe that they are—a false belief that a black man is staff not only diminishes him, but diminishes him in a way that aggravates an accumulated store of past injustice. (Schroeder forthcoming, p. 15)

The challenge Schroeder articulates is to explain how—despite widespread inequality and oppression in society—epistemic practices can rationally respond to evidence and yet not thereby be morally amiss. If epistemic normativity is not affected by moral considerations, how can one countenance the normativity of the above vignettes? Schroeder (forthcoming) and Basu and Schroeder (forthcoming) argue that endorsing moral encroachment on belief satisfies this challenge, and in the above quote Schroeder avers that evidentialism cannot satisfy the challenge.<sup>22</sup>

---

<sup>21</sup> See especially Basu (ms a; ms b; ms c), Schroeder and Basu (forthcoming), Bolinger (ms), and Gendler (2011) for statements of this view.

<sup>22</sup> Strictly speaking Schroeder suggests that evidentialism should embrace evidentialism by allowing that, even though only evidence can justify a belief, what qualifies as a sufficient evidence for justification can vary depending on the moral stakes. I will not evaluate whether the resulting view can qualify as a species of evidentialism, but it certainly differs from how evidentialists have hitherto understood the view.

#### 4. The Inadequacy of Merely Statistical Evidence

The vignettes in section three describe outright beliefs about a person based on statistical demographic evidence. Spencer outright believes ‘Jamal will tip less than average’, rather than the qualified belief ‘Jamal will *probably* tip less than average’. One response to Schroeder’s challenge holds these vignettes thereby exhibit an epistemic error. By concluding a fact from evidence that merely probabilifies the fact, the person has gone beyond the evidence. Evidentialism decrees we should apportion belief to the available evidence; the beliefs violate the decree. If there is an orthodox epistemic fault the accompanying moral fault does not impugn evidentialism: moral error and misfit with evidence align. This defence of evidentialism against the moral encroachment challenge accords with orthodox views of racism, sexism, and similar prejudices, which hold that orthodox epistemic error is central to the nature of the fault.<sup>23</sup>

Two considerations support this response. Firstly, we can compare the vignette’s beliefs with morally neutral beliefs. Plausibly in morally neutral cases similar kinds of evidence do not support outright belief; the evidence only supports credences or beliefs about what is likely. Suppose, for example, you know 95% of the birds in the aviary are yellow, and one bird has just died. This evidence typically licenses the qualified belief that ‘*probably* a yellow bird died’. But it does not license the outright belief that ‘a yellow bird died’.<sup>24</sup> (Or perhaps the evidence licenses a weak and easily unseated species of outright belief. I return to this point in section six.)

---

<sup>23</sup> See, for example, Mills (1997), Gordon (1995, 2000), Arpaly and Schroeder (2014), Arpaly (2003), and the discussion of ‘restricted accounts’ in Basu (ms a). See also Munton (ms), who describes an underappreciated epistemic error commonly infecting racist beliefs about statistics.

<sup>24</sup> This example is inspired by Moss (forthcoming).

Secondly, comparison with other kinds of evidence arguably also indicates the epistemic (and accompanying moral) fault is basing an outright belief on merely statistical evidence. In the original vignettes the beliefs are based on highly-probabilifying statistical evidence. The beliefs are not based on non-statistical individualised evidence. Consider a revised vignette, in which the statistical evidence is considerably weaker and not playing a significant epistemic role. The beliefs are instead based on non-statistical individualised evidence. The individualised evidence is less probabilifying, so that in the revised vignette the overall evidence is more likely to lead to a false belief.

Suppose, for instance, that the racial demographics at the Cosmos Club are more equitable, and the woman instead bases her belief on weak testimony. Someone told her Franklin was staff, but the woman later realises the testifier seemed ignorant about the club in general or did not check carefully who he was pointing towards. The consultant—who in this revised case has no particular sense of the demographic distribution within the office—was expecting an administrative assistant to approach around that time, as arranged, and assumed the person approaching was the appointed person. Spencer's belief that Jamal will tip less than average is based wholly on snippets of misheard and misinterpreted conversation. Jamal was charismatically explaining that his teacher used to rail against high tipping rates and 'tip inflation', and Spencer thought Jamal was voicing his own views.

Each of these three revised cases is under-described. There are many additional epistemically significant, evidentially-relevant factors, such as the office layout and how frequently people walk the office corridors. Please fill in the details so the vignettes exhibiting non-statistical,

individualised evidence are more likely to lead to false beliefs than the original vignettes, in which the beliefs are supported by highly-probabilifying statistical evidence. The reasoning in the revised vignettes, although slightly hasty, is not particularly irresponsible or unusual. I contend that in these revised cases the moral error seems less significant, even though the chances of the beliefs being false are higher than when they were based on highly-probabilifying merely statistical evidence.

This suggests an error exhibited by the original vignettes is that outright beliefs were based on *statistical* evidence, not that the beliefs were based on *insufficient* evidence. The beliefs supported by ‘less probabilifying’, individualised evidence (that is, evidence that is less likely to lead to accurate beliefs and that supports lower credences) are not as improper as the beliefs supported by highly probabilifying statistical evidence (that is, evidence that supports higher credences). If correct this suggests the original examples, rather than supporting moral encroachment, instead exemplify the proof paradox: the cases do not illustrate that the higher stakes mean a higher degree of evidential support is required for justified belief. Instead the cases indicate that merely statistical evidence does not typically support an outright, unqualified belief.<sup>25</sup>

---

<sup>25</sup> For background on the proof paradox and the inadequacy of merely statistical evidence, see Thomson (1996), Gardiner (forthcoming b), Bolinger (ms), Buchak (2014), and Smith (2010). See also the related lottery paradox (Kyburg (1961), Nelkin (2000), and Hawthorne (2004)). If belief aims at knowledge and beliefs based on merely statistical evidence fail to be knowledge, this might explain the fault of outright belief based on merely statistical evidence: even if correct the belief cannot be knowledge on that kind of evidential basis. Beliefs with faults such as poor testimony, misidentifying an anticipated greeter, or misinterpreting anecdotes do not share this flaw. See also Moss (forthcoming: 166). Note the beliefs might well be flawed in more than one way. A belief might be faulty because based on statistical evidence and also faulty because the evidence is insufficient given the high stakes. Thanks to Sarah Moss for emphasising this point.

Perhaps extremely-probabilifying statistical evidence can support outright belief. Perhaps, for instance, believing your ticket did not win the national lottery is epistemically justified. But purely statistical demographic evidence on this order does not typically arise. Cases about gender, race, sexuality, and so on with this kind of extreme statistical evidence are rare, and I am not sure we have good intuitions about these cases. Normal cases have much weaker and more complicated demographic evidence. I return to this in section six.

Basu's argument that moral requirements affect the demands of epistemic rationality—that is, her case against evidentialism—requires a 'rational racist'. A rational racist is someone whose beliefs align with the evidence, yet whose corresponding belief is racist.<sup>26</sup> Basu holds that Spencer qualifies. I have argued that Spencer's belief is epistemically flawed in virtue of going beyond the available evidence.

Thus we can reconsider the original three vignettes, but replace the unqualified belief with a corresponding belief about what is likely: Spencer believes that Jamal will *likely* tip lower than average. The consultant believes the woman is *probably* an administrative assistant. The woman at the Club believes Franklin is *probably* staff. Perhaps in these cases the person simply believes according to the evidence. (In section six I cast doubt on the claim that the beliefs about what is probable are supported by the available evidence.)

Some theorists maintain that even beliefs about what is likely, if based on merely statistical evidence, can be morally wrong. Correspondingly, they hold, such beliefs are thereby

---

<sup>26</sup> See also Schroeder and Basu (forthcoming).



epistemically impermissible.<sup>27</sup> Beliefs about likelihoods can judge an individual on demographic data, and pigeonhole people, even if also allowing that the person might diverge from the relevant statistical regularities. Moss (forthcoming) suggests such beliefs violate a moral demand that we bear in mind that a person might differ from arbitrary members of their relevant reference classes.

In section five I articulate some reasons to resist the conclusion that moral considerations affect epistemic justification in such cases. In sections six and seven I articulate some strategies evidentialists can employ to meet Schroeder's challenge and so explain how believing in accordance with the evidence can be morally appropriate despite widespread inequality and oppression. Some of my comments apply to both outright beliefs about a person and qualified beliefs about what is likely, others apply only to the latter.

## 5. Objections to Moral Encroachment

Moral encroachment suffers from many of the same weaknesses that afflict other versions of pragmatic encroachment, namely the counterintuitive consequences of holding that

---

<sup>27</sup> See Moss (forthcoming, especially section 10.4), Armour (1994), and Basu (ms d). Note Moss discusses the normativity of belief in a probabilistic content (that is, a set of probability spaces), rather than beliefs concerning likelihoods given certain contextually determined information. Some of Schroeder and Basu's motivations for moral encroachment extend to moral encroachment about beliefs representing what is likely. Basu argues, for example, that believing someone shoplifted based on statistical evidence is wrong because it hurts (Basu, ms a: 11). But similarly believing someone probably shoplifted on this evidence also hurts. Basu and Schroeder (forthcoming) argue that you should not believe on weak evidence that your spouse has fallen off the wagon, given the high stakes, even if the same evidence would license belief about a stranger's drinking. But presumably similar considerations apply to the belief that your spouse *probably* fell off the wagon. Bolinger (ms) and Schroeder (forthcoming) discuss, but do not endorse, moral encroachment on credences.

considerations that do not bear on the truth a belief can affect its epistemic justification.<sup>28</sup>

Below I articulate some worries that apply to moral encroachment.

One reason to resist moral encroachment is the risk of tensions amongst the epistemic statuses of related beliefs. Advocates of moral encroachment typically hold it is morally and epistemically permissible to believe something about a person based on statistical evidence if the moral stakes are low. Moss (forthcoming: 233) suggests that believing that ‘someone probably has brown eyes, on the grounds that most people have brown eyes’ is typically morally neutral and so is epistemically justified. On this view we can believe of a person selected randomly from the world population that they probably have brown eyes. But now consider the world’s incarcerated population. Given systemic racism, brown-eyed people are overrepresented in prison populations. If believing of a randomly selected person that they probably have brown eyes is licensed by the evidence, surely believing of a randomly selected inmate that they probably have brown eyes is also licensed, since it is better supported by the same kind of evidence. But being incarcerated is a morally significant property. The moral stakes are raised. Moral encroachment suggests we should not believe of a randomly selected prison inmate that they probably have brown eyes, since this belief has high moral stakes. This example illustrates two problems for moral encroachment. Firstly, moral encroachment renders unjustified the better-supported belief, whilst endorsing the less supported belief. Secondly, there seems to be a tension amongst believing that an arbitrarily selected person

---

<sup>28</sup> See for example Eaton and Pickavance (2015), Ichikawa, Jarvis, and Rubin (2012), Worsnip (2015), and Munton (ms: 28–9). Schroeder emphasises that his version of moral encroachment posits beliefs that are stable over time, and holds that he thereby avoids many objections that encroachment views typically confront.

probably has brown eyes, that brown-eyed people are overrepresented in prisons, and not believing that an arbitrarily selected prisoner probably has brown eyes.

Most advocates of moral encroachment hold that more evidence is required if the belief contributes to, or accords with, the disadvantage of socially disadvantaged groups. The stakes are lower if the target belief asperses historically advantaged groups or commends members of disadvantaged groups.<sup>29</sup> But this asymmetry might also vindicate tensions amongst beliefs. To illustrate, suppose the evidence Spencer marshals justifies race-based beliefs about how specific customers will likely tip. (I articulate some doubts about this in section six.) And suppose in accordance with moral encroachment, Spencer believes on this evidence that non-black diners will tip higher than average, yet believes of no diners that they will tip lower than average. (He believes his evidence indicates the patrons will tip less well than average, but he refrains from this belief.<sup>30</sup>) In this case Spencer's beliefs seem epistemically amiss: A teacher who believes of half his class that they will perform better than average while withholding concerning the other half might be 'kind' or 'sweet', but they are not exhibiting epistemic rationality. Perhaps Spencer and the teacher are being laudable in some way, but they are not conforming to an epistemic ideal. This objection highlights that moral encroachment endorses such doxastic attitudes as an *epistemic* ideal.

---

<sup>29</sup> Basu (ms a; ms c), Schroeder (forthcoming), and Basu and Schroeder (forthcoming). Bolinger (ms) also emphasises that moral considerations arise particularly when we contribute to overall patterns of oppression. Idiosyncratic beliefs about an individual based on statistical evidence, such as that a black person cannot draw well based on their race, are less harmful than stereotypical beliefs, such as that black people consume more narcotics.

<sup>30</sup> Reflecting on this case also raises the concern that the edicts of moral encroachment are not psychologically possible, since it is not possible to suspend judgement despite compelling evidence. I will not evaluate the psychological availability of suspending belief despite evidence in this essay, in part because I think the evidence in these cases is weaker than usually appreciated.

If moral encroachment were true, one might gain evidence for a claim, but thereby learn the claim has morally high stakes, and so be less justified in believing the claim. Gaining the new evidence undermines one's epistemic justification for the claim.<sup>31</sup> Illustrations of this idea are a little difficult to articulate, since whether the illustration succeeds depends on details of the particular version of moral encroachment. But nonetheless an example might help convey the structure of the worry.

Bolinger holds that beliefs about individuals based on statistical inference are permitted if there is a 'permissible signal' underwriting the relevant reference class. Permissible signals include features such as attendant's uniforms, but do not include features such as race. Permissible signals, on Bolinger's view, affect the epistemic justification of a belief without being truth-relevant; believing on non-permissible signals, such as race, raises the stakes. The epistemic significance of permissible signals is thus a non-evidentialist feature of Bolinger's view.

Suppose you learn a gang distributes drugs in a particular area. You see someone who looks like he might be a gang member selling drugs, and base your belief that he is a gang member on 'permissible signals' such as clothing and behaviour. You do not have negative attitudes towards drug selling or gang membership, and your belief seems fairly well supported by evidence. Suppose you learn the gang is Asian, and all gang members are Asian. The person you see is Asian. Assuming that the base rate of Asian people in the area is not very high, the

---

<sup>31</sup> See also Eaton and Pickavance (2015).

person's race is plausibly further evidence for your belief.<sup>32</sup> Yet this evidence might—depending on particular details of the moral encroachment view—render the belief morally high stakes, since it is now a belief partially based on race. Thus gaining further supporting evidence for the belief can alter its status from epistemically justified to unjustified.

In many cases it is not straightforward whether a belief has moral valence, or whether the valence is positive or negative. Consider claims such as gay men are more likely to be promiscuous than gay women. Most white people with dreadlocks have attended a drum circle. Lesbian women are often less 'ladylike' than straight women. Most women are paid less than most men. The moral significance of these kinds of claims is controversial. If whether a belief is justified depends on the moral properties of the belief, this uncertainty and complexity bleeds into whether the belief is epistemically justified. It can underwrite contextualism about epistemic justification: perhaps homophobic people require more evidence before endorsing statistical inferences about sexual orientation and behaviour, for example, whereas non-homophobic people require less evidence.

The moral significance of a belief can also depend on who the belief is about. And so according to moral encroachment if a community is marginalised one might require more evidence to justify beliefs about their behaviours than for non-marginalised groups, even with the same behaviours and same evidence. But this seems implausible. It seems counterintuitive that the belief that 'unemployed people smoke more cannabis on average than employed people' requires less evidence than, for example, the belief that 'wealthy

---

<sup>32</sup> Thanks to Renee Bolinger for pointing out that the base rate of Asian people in the area bears on the epistemic significance of race in this example.

youths smoke more cannabis on average than less wealthy youths'. Plausibly, given their similarity, these beliefs require the same amount of evidence to justify. Moral encroachment risks making epistemic justification contingent on myriad complex social factors that are intuitively irrelevant to epistemic justification.

Further reasons to resist moral encroachment stem from considerations of social justice. Black people are overrepresented in the US prison population, and acknowledging this fact matters for social justice. An important feature of this claim about demographic distribution is how it affects particular individuals. A person's skin colour makes it more likely they—the individual—will be incarcerated. If we ought to acknowledge that a person is disproportionately likely to be imprisoned if they are black, we also ought to acknowledge that a randomly selected black person is more likely to be incarcerated than a randomly selected white person. The injustice is not simply systemic injustices concerning overrepresentation; central to the injustice is the effect on individuals' life chances. Particular *individuals* are more likely to be imprisoned.

When a particular person is incarcerated, underemployed, participating in crime, and so on, one potential source of injustice is that their race, gender, or other social category means the outcome was more likely.<sup>33</sup> And these are social facts we ought to acknowledge.

In some cases acknowledging base rate facts about someone can help frame their accomplishments. It is relatively rare, for instance, for a first-generation college student to

---

<sup>33</sup> The effect of social group on likelihoods is often indirect. A person's race might affect their likely economic circumstances, for example, which can affect the probability they are incarcerated.

become a professor in America. If Ali is a first-generation college student who became a professor then plausibly she merits particular praise, since there is a higher chance she overcame distinctive obstacles. If so, the reason is not simply that first-generation students are underrepresented amongst the professoriate. This does not explain the particular accomplishment of Ali as an individual. The relevant fact is that it was less likely for Ali to become a professor (relative to her colleagues), given she was a first-generation college student.

Recognising how base rates bear on individuals can help interpret behaviour. Some choices might seem unwise or difficult to understand, but the choices can be illuminated by how they fit into background base rates.

Perhaps the central reason to resist moral encroachment is that epistemic normativity answers solely to considerations pertaining to evidence, truth, reliability, comprehension, and so on. This reason is perhaps both the most and least compelling reason. It is the most compelling reason since the idea that epistemic justification depends solely on how a person responds to evidential and other truth-relevant considerations is a central motivation for people who deny encroachment. It is the least compelling since this is precisely what advocates of encroachment deny. Plausibly, though, there is a strong default in favour of the view that epistemic justification depends on considerations pertaining to evidence, truth, and so on. The burden of proof falls squarely on those who argue that the moral stakes influence epistemic justification.<sup>34</sup>

---

<sup>34</sup> Kim (2017: 7) and Piller (2016). As Ichikawa, Jarvis, and Rubin (2012) comment, “The most widely discussed argument to date against pragmatic encroachment is that it is

I do not deny that common epistemic practices are morally faulty, and that widespread epistemic practices contribute to systemic inequality, disadvantage, and oppression. What I hope to resist is that ideal epistemic practices should be influenced by moral or political considerations.<sup>35</sup> Instead my hope is that impartial epistemic practices, including impartial evaluation of the available evidence, are morally permissible. There is no tension between epistemic and moral norms because tracking the truth accurately cannot be morally wrong.

## 6. Resisting Moral Encroachment

In this section I articulate some strategies for defending evidentialism and kindred views against moral encroachment.

Many real life beliefs are morally problematic. Sexism, racism, and other prejudice are widespread. But these real life beliefs also exhibit myriad epistemic errors.<sup>36</sup> People are poor at statistical reasoning. They overestimate patterns, extrapolate too readily from limited and biased sources of information, and engage in motivated reasoning. Confirmation and availability biases contribute to the epistemic faults of such beliefs. If the morally wrong belief is also epistemically unjustified according to orthodox epistemology, the moral wrong does not impugn evidentialism. Arguments for moral encroachment need to abstract away

---

counterintuitive.’ Although see Grimm (2011) and Marušić (2015) for nuanced discussions of the burden of proof concerning pragmatic encroachment.

<sup>35</sup> Right wingers would delight in the idea that left wingers resisted impartial evaluation of the evidence when adjudicating facts about individuals based on race and other social categories, and that they did not aim to maximise true belief concerning crime, education level and so on.

<sup>36</sup> Kahneman (2011), Kunda (1990), Arpaly (2003, especially chapter three), Munton (ms), and Gendler (2011).



from the myriad, ubiquitous flaws of real life beliefs and insist that a belief with no epistemic flaw of this kind is also immoral; my contention is that advocates of moral encroachment have failed to do this.

It would be impossible to articulate here all the ways that such beliefs commonly err epistemically. Below I sketch some ways most relevant to the examples used to motivate moral encroachment.

Crime data provide common examples of the putative tension between epistemic and moral demands.<sup>37</sup> But the differences in base rates among social groups for the relevant kinds of social facts are typically low, and the overall percentage of people who actively commit crime is very small. People overestimate these differences and overestimate overall rates. Very few people commit robbery, for example. So even if commission of robbery is higher amongst black men than white men, this says almost nothing about the chances concerning any particular black man. Given the tiny proportion of people who commit robbery, and the small differences in rates amongst races, any association between a person and robbery based on base rates is a flagrant epistemic error. Even *if* an arbitrary black person is more likely than an arbitrary white person to commit robbery, they are still extremely unlikely, and the difference is minute. Any association forged between a particular person and crime risk is based on racial prejudice and irrational fear.<sup>38</sup>

---

<sup>37</sup> See, for example, Munton (ms), Basu (ms d), Gendler (2011), and Armour (1994).

<sup>38</sup> See also Armour (1994: 792–3).

To further illustrate the trouble with everyday statistical reasoning concerning crime: infamously when some white people see a black person nearby they worry about crime. (Consider, for instance, the phenomenon of women pulling their purses closer.) But most crime is committed by people of the victim's race. This statistic indicates white people should be more suspect of other white people. But, then, this statistic is largely underwritten by the pattern that people commit crime near where they live, and American housing is not very integrated, so one 'should' correct for that... The 'reasoning' could continue. My point is not to estimate which demographics one should associate with crime risk. My point is instead that almost every association between an individual and behaviour such as crime based on a social category such as race commits basic epistemic mistakes.

Secondly, we overestimate the epistemic significance of race, gender, and similar social categories when we estimate likelihoods. The inappropriate salience of race as a reference class is exhibited in the Cosmos Club case, where the customer should have instead relied on the more probative reference class of whether Franklin was wearing a uniform or dinner attire. In the administrative assistant vignette the consultant knows that most women in the office are administrative assistants. But this belief does not license the judgement that a particular woman is likely an administrative assistant. This belief is legitimate only if her being a woman is the canonical reference class from which to extrapolate.<sup>39</sup> The consultant might instead judge the likelihood of her being an administrator based on other reference classes she belongs to: the fact she is an older woman, an older person, a person in a business suit, a woman exiting a private office, or a person walking down the hallway talking

---

<sup>39</sup> See Bolinger (ms, especially the appendix), Leslie (forthcoming), Moss (forthcoming), Hájek (2007), Venn (1866), Reichenbach (1949), and Munton (ms). See also Armour (1994: 791; 809–14).

into a mobile phone. These different reference classes alter the probability that the person is an administrative assistant. But people tend to focus on gender and race as salient reference classes, even if they are less probative than alternative reference classes. Perhaps, in other words, the consultant's all too human focus on gender led him to neglect the fact that the woman was wearing a power suit, hiring someone via mobile phone, and asking her assistant to bring coffee. Or perhaps the consultant neglected his belief that the administrators are almost all young, and this person is older, or she exits a door labelled 'laboratory' and is wearing a lab jacket.

Advocates of moral encroachment compare morally significant beliefs like 'the woman is likely an administrative assistant' with morally neutral beliefs like 'a yellow bird has likely died'. They argue that, given their similarities, any epistemic difference between the beliefs must arise from moral differences. But the reference class problem indicates an important difference between these cases. When we learn that most birds in the aviary are yellow and one has died, it is very likely that we draw on all available evidence when we conclude that likely a yellow bird died. (Some ornithologists might have relevant beliefs about avian life expectancy to draw on.) But in the consultant's case it is extremely unusual that the information described in the original vignette exhausts the consultant's information. He would have substantial supplementary evidence about mannerism, bearing, clothing, actions, and so on. And he would likely, given widespread cognitive biases, overestimate the epistemic significance of gender.

In real life cases our evidence is shifting, complex, nuanced, and varied. And new evidence is in many cases readily available if we inquire. These details cast doubt on whether the person

draws on all available evidence in forming their belief. The Cosmos Case, often used as a central motivating case for moral encroachment, exemplifies this: the woman ignored the counterevidence of clothing, and any ‘mannerism’ evidence that would have likely been available given Franklin’s evening plans. Since he was hosting friends at his club, he would likely have been acting differently from staff. And Franklin was *eighty* years old at the time, so he would have appeared extremely old for a club attendant.<sup>40</sup> Thus it is unlikely that the total available evidence supported the woman’s belief. If the belief exhibits epistemic errors according to orthodox epistemology—such as failing to respond to the available evidence—the cases do not impugn evidentialism. The evidentialist can explain the moral error without revising epistemic normativity. The kind of epistemic error committed—narrowly focusing on features such as race and gender and failing to countenance other individuating features of the person—plausibly underwrite the kind of moral error the person commits.<sup>41</sup>

Jessie Munton (ms) addresses the sense of moral unease we can feel when considering social statistics such as ‘Black Americans commit disproportionately more violent crime than white Americans’. These beliefs can be true and well supported by evidence, yet generate moral discomfort. Munton does not endorse moral encroachment, and so does not hold that moral considerations provide an epistemic reason to withhold belief in these cases. (Munton (ms) also does not examine applying general social statistics to individuals.)

---

<sup>40</sup> Franklin (1995: 4). Of course similar errors happened when Franklin was younger. But my point is to illustrate that there is usually counterevidence in these real life case that are not represented in artificial, oversimplifying vignettes. This counterevidence contributes to the affront. If the person were not prejudicially associating ‘Black’ with ‘staff’ she would likely heed the counterevidence.

<sup>41</sup> There may also be a moral and an epistemic flaw in persistent attention to particular facts. This flaw might also be exhibited in the cases advocates of moral encroachment describe. I owe this suggestion to Jessie Munton and Dan Greco.

Instead Munton highlights an underappreciated epistemic error that often accompanies beliefs about true social statistics. People can believe the statistic but fail to accurately understand the appropriate reference class. They will thus misinterpret the counterfactual properties of the statistic and misapply the statistic to novel cases. People might falsely believe that the statistic indicates that black people are more criminally inclined by nature, for example, rather than appreciating that the statistic indicates that social marginalisation and oppression leads to increased crime rates. Munton notes that although the epistemic error might be more typical and troubling concerning social statistics, the error can also arise concerning morally neutral beliefs, such as statistical claims concerning tree heights. This error, Munton argues, might underwrite the sense of moral unease the social statistics generate.

Munton emphasises that an effective way to correctly identify the relevant reference class is to *understand* what *explains* the statistic. Munton writes,

One way of ensuring that that the domain of the statistical belief is appropriately circumscribed is to hold a set of associated beliefs that offer an explanation of the regularity in question [...] I am arguing that even a simple statistical belief may draw on a rich web of further belief and behavior. [...] But an important upshot of this account is that the most naturally reported description of [the avowed statistical belief] is really the tip of an iceberg, in the sense that it is a small part of a network of beliefs which provide additional implicit, sometimes explanatory, content. The epistemic good-standing of a belief depends on what is going on ‘under the water’, that is, on the broader belief structure. (p. 14)

In section seven I return to the importance of the understanding in which the beliefs are embedded for illuminating the normativity of the vignettes that advocates of moral encroachment use to motivate their view.<sup>42</sup>

Many of these examples exhibit, or readily bring to mind, other wrongs in addition to epistemic errors. The woman at the Cosmos Club behaves rudely. Spencer seems to disapprove of or resent poor tippers.<sup>43</sup> Describing the consultant's belief about administrative assistants and gender, without any context for why he focuses on this, suggests he might disdain administrators. Or perhaps we simply project perceived *normal* opinions onto Spencer and the consultant.<sup>44</sup> Relevant real life cases will typically include similar moral flaws. The anti-evidentialist strategy pursued by Basu (ms a, c) and Basu and Schroeder (forthcoming) relies on a person whose beliefs and epistemic character impeccably follow the evidence, and whose moral behaviour is faultless, and yet who morally wrongs another in virtue of his beliefs. But if the examples exhibit—or conjure images of—other wrongs this complicates the anti-evidentialist strategy. Perhaps the sense of wrong can be (partially) explained by these other wrongs.

The ubiquity of (flawed) beliefs about people based on weak or merely demographic evidence might generate the sense that respecting the evidence is morally problematic. And

---

<sup>42</sup> Munton and I developed our ideas independently. I think there are fruitful connections between the ideas.

<sup>43</sup> Spencer's noticing the trend, his keenness to find evidence, and his applying the generalised belief to Jamal might be evidence of prejudice. See Arpaly (2003) for related discussion.

<sup>44</sup> See Gardiner (2015) for more about how we interpret vignettes by applying our understanding of how they would normally be fleshed out. See also Ichikawa and Jarvis (2009).

so it might generate a sense that we ought to revise epistemic normativity in light of this ubiquitous wrong. But if these ubiquitous beliefs also always include epistemic errors—errors countenanced by orthodox epistemology—this undermines the threat to orthodox epistemology.

Another evidentialist strategy for responding to Schroeder's challenge emphasises that the beliefs licensed by demographic evidence are easily unseated. The person should readily revise the belief in light of new evidence. Suppose the consultant's total evidence supports the belief that the woman is, or probably is, an administrative assistant. Perhaps the consultant sees a woman's name on an employee roster, for example, and so possesses no additional individualising evidence. Evidentialism and kindred views should emphasise the belief licensed by the evidence is a working hypothesis, or tentative belief, one that could be easily dislodged. Advocates of moral encroachment, by contrast, tend to emphasise the 'settled' nature of belief. Schroeder writes,

[S]ince forming a belief is taking on an ongoing commitment into the future, it will be rational to form the belief that *p* up front only if the strategy of counting on *p* in reasoning is one that is expected to bear good fruits *over time*. In the simplest case, in deciding whether to believe that *p*, you are deciding whether to *always* be disposed to count on *p* in reasoning. (Schroeder, forthcoming, p. 11, emphasis in original)

Schroeder thus emphasises a conception of belief as stable and not easily unseated. Similarly Basu talks of 'settling on a belief' (ms a; ms c). Bolinger (ms) writes,

[T]o accept that *p* is to add *p* to the stock of propositions that you are ready to act on without further consideration. When an agent accepts *p*, she dismisses the possibility that *p* is false from consideration, and takes *p* as a premise in her practical reasoning [...] Accepting that *p* involves deciding to move from an epistemic partition including at least

some  $\sim p$  spaces with  $> 0$  probability to one without. Deciding to accept  $p$  is deciding to give  $\sim p$  no cognitive space in future deliberation. (p. 2, 9)

Plausibly there is something morally wrong with firmly settling on a judgement about a person on weak or merely statistical evidence. But this is consistent the moral permissibility of forming a readily-revised belief about a person based on this evidence.<sup>45</sup> Evidentialists can emphasise the role of this kind of belief in responding to weak or statistical evidence in general.

Another potential response emphasises that belief is often inappropriate when more evidence is readily available. Plausibly one should not form a belief on weak or merely statistical evidence when stronger or individualised evidence is easily obtainable, and this feature of epistemic normativity underwrites the epistemic error in many cases. I am sympathetic to this idea. It is not, however, compatible with some stricter forms of evidentialism. It is plausibly compatible with the evidentialist idea that whether a doxastic attitude is epistemically justified depends solely on evidential considerations. But it is

---

<sup>45</sup> I was writing this book chapter in a coffee shop when two young men approached and asked how my homework was coming along. I explained that I was writing a book chapter, so it wasn't homework exactly, but that I was enjoying thinking about the topic. I do not think they wronged me by assuming I was doing homework. Perhaps most people in a cafe who look relatively young, wear informal attire, and make notes in books and papers are doing homework; not many are writing book chapters. The base rates favour their initial belief. Plausibly their belief simply accorded with the evidence. But my interlocutors couldn't shake their initial belief. They assumed they misheard me ('You are writing *about* a book chapter, you say?') They acted extremely surprised that I was writing a book chapter, and it took a number of rounds of questioning before they revised their belief, such as skeptically asking for the book title. Plausibly being committed to their initial belief, and reluctant to revise this belief in light of new evidence, was morally poor treatment. But evidentialism can countenance this thought. One of the young men, who was about to enrol at a local community college, offered me some writing advice: 'Use examples' he suggested, 'to explain your points'. I hope, dear reader, you appreciate the example.



incompatible with stricter evidentialist claims such as whether a doxastic attitude is epistemically justified depends solely on the strength of one's currently available evidence.<sup>46</sup>

## 7. Understanding

Consider the following example,<sup>47</sup>

Joan notices four young men together in an alley in Baltimore. It appears they are using their bodies to shield their activity from people in the street. It looks like they are exchanging money and packages.

This evidence is inconclusive, but it suggests the people are engaged in the sale of controlled substances or contraband. Suppose Joan forms the belief that 'those people are selling illegal drugs' or 'those people are probably selling illegal drugs'. (In my view in almost every such case only the latter belief is epistemically warranted, but perhaps I am unusually diffident.) What can we say about Joan's belief, morally?

My contention is that we cannot yet tell; we lack sufficient information. It depends on what else Joan believes and how she integrates her judgement with existing beliefs.

Joan's observation might remind her of her background beliefs that drug traffickers on the street are selfish, ruin communities, are violent, often carry guns, and endanger law abiding citizens. Or she might start reflecting on her beliefs about the economic inequality that leads

---

<sup>46</sup> For discussion of the epistemic significance of readily available evidence, and how this relates to evidentialism, see McCain (2004), Conee and Feldman (2011; 2004), and DeRose (2011).

<sup>47</sup> For versions of moral encroachment focusing on the distinctive wrong of forming a belief about a person based on statistical evidence, instead consider a relevant 'base rate' example.

people to sell drugs, the social pressures to participate in the activity, and the way that young people in poorer areas have more financial responsibilities than wealthier peers. She might consider these social pressures whilst bearing in mind the individual choices and agency involved. She might integrate her observational belief with her recollection of a newspaper article articulating how members of upper socioeconomic groups exchange drugs in privately owned, secluded places whereas members of lower socioeconomic groups tend to do so in public, exposed places. Seeing the group might make her worry about how the activity will affect future social prospects of the participants, and she might connect this to her beliefs about racism in the criminal justice system. She might be angry and unsympathetic, since she views drug dealers as preying on poor marginalised individuals. Or she might be relieved, since she is looking to purchase drugs. In short, the moral character of Joan's belief depends on the broader understanding in which it is embedded. This understanding comprises Joan's beliefs and the connections between them.<sup>48</sup>

Similar approaches apply to other beliefs discussed. Consider the consultant's belief that the woman is probably an administrator, which is based on demographic base rates. His belief might be embedded in an understanding according to which it is appropriate that women occupy lower status jobs, since they ought be servile and pursue less ambitious careers. Or it might be embedded in ideas about women's oppression and the systemic challenges that women face in the workplace. Or the consultant might view administrators as the true experts in how to improve a company since they have the clearest perspective on the weaknesses and strengths of the organisation. Plausibly it is the understanding the belief is

---

<sup>48</sup> Joan's understanding might also include (connections to) her relevant emotional reactions.

embedded in—or that we take the belief to be embedded in—that explains much of the perceived wrong in the beliefs described in the vignettes. (As noted in section six, many of these beliefs plausibly exemplify other wrongs, such as not responding properly to evidence.)

In Franklin's anecdote about the Cosmos Club the woman behaves poorly. This rude behaviour complicates the probative value of the vignette as a motivation for moral encroachment, since the poor behaviour (and poor response to evidence, given Franklin's clothing and advanced age) might explain the moral error. The poor behaviour also indicates her understanding was unenlightened. She seems to lack an anti-racist understanding. This understanding taints the moral value of her belief that Franklin is an attendant.

If the understanding the beliefs are embedded in explains the moral fault, one can explain the wrongness of such beliefs without appeal to moral encroachment.

Perhaps some beliefs are evaluable as morally wrong regardless of the understanding they are embedded in. These beliefs might include 'women should be subjugated by men', 'black people are all bad at their jobs', and so on. But these beliefs are manifestly not supported by evidence, and so they do not threaten evidentialism and kindred views. Part of the moral wrong, moreover, includes the deplorable understanding these beliefs are embedded in.

Holding that the moral valence of the belief depends on the person's broader understanding is compatible with also holding that moral encroachment explains an additional moral and epistemic error. The explanations are compatible, and advocates of moral encroachment might well endorse my emphasis on understanding. My argument is that if focusing on

understanding can illuminate moral faults of evidentially-supported beliefs, this undercuts the motivation to endorse moral encroachment. Moral encroachment is not needed.<sup>49</sup>

Note too this explanation can illuminate the potential moral wrongs of beliefs that are exceedingly well supported by evidence. If the error of Spencer's belief is he lacks sufficient evidence given the high moral stakes, as Basu holds, then the error should disappear if Spencer possesses sufficient evidence. But if the moral evaluation of Spencer's belief depends on whether his overall understanding disdains black people for (putatively) tipping less on average, the wrong will remain regardless of how well supported the belief is. If a person's overall understanding is racist, a problem remains despite the evidence they collect in support of individual beliefs.

This difference is important since some of the putatively problematic beliefs advocates of moral encroachment discuss are true beliefs. If the belief is true, then typically the belief will be well-supported by further evidence. If someone is racist or sexist, but they collect further evidence for their true beliefs, this does not abate the moral error. The way to abate the

---

<sup>49</sup> Thinking about understanding can play a further role in accounting for the epistemic normativity of these kinds of beliefs. When we form beliefs there is a chance to gain true belief, which is valuable, and a risk of false belief, which is disvaluable. One question moral encroachment seeks to answer is how to weigh these competing considerations. Moral encroachment replies that the relative weight depends on the moral stakes. If the moral stakes are high, we should be risk averse in belief, and so seek more evidence. (Although see Worsnip (2015) for an objection to encroachment as a response to weighing the relative risks of error.) Wayne Riggs (2003) instead proposes that the relative values of attaining truth and avoiding error can be weighed by how beliefs contribute to, or impede, understanding. If Riggs's proposal is fruitful it provides a second way that theorising about understanding undermines a motivation for moral encroachment.

moral error is to alter one's understanding—towards a more accurate understanding—so that it is no longer sexist, racist, or otherwise morally wrong.

To illustrate consider the following fictional circumstance. Suppose that girls' scores on standardised maths tests are on average lower than boys' scores.<sup>50</sup> We could learn this fact and embed it in a non-sexist understanding: the difference in test scores indicates girls receive inferior educational opportunities, or girls' mathematical acumen is not well-measured by current testing methods. We might connect the result to our understanding of the pressures of gendered cultural expectations. We embed the fact in a framework of beliefs and attitudes that does not denigrate girls, even if they perform less well on average on maths test. Suppose we later learn that girls are simply less good on average than boys at maths. (Remember, this example is fictional.) We should then embed this new information in a non-sexist understanding: since women have equal moral status to men, this result means mathematical acumen is irrelevant to moral status. We might think about strategies to help support girls in maths education and we might recalibrate how we credit people for individual accomplishments in maths. That girls are less good than boys at maths on average would be as morally irrelevant as that women are on average shorter or that men are on average more susceptible to disease and early death. In short, there is no difference between genders, races, sexualities, and so on, that affects moral status.

Moral encroachers and I share the view that if there is a moral mistake then there is also an epistemic mistake. According to moral encroachment the moral error grounds the epistemic

---

<sup>50</sup> Recall the relationship between group averages and individual averages is often misunderstood and misinterpreted.

error. On my view—which I think accords with the orthodox view of both epistemic normativity and the nature of racist belief—it is the epistemic error that gives rise to the moral error.

To forestall a potential confusion: One can appropriately believe something about a person that reflects a moral problem with that person. I do not doubt this. In some cases it is true that the person has a significant flaw. One might believe truly of Fred that he is a domestic abuser, and this belief might be morally and evidentially appropriate. My view holds that there cannot be a justified belief about someone that reflects a moral problem with that person, *where that belief is based on demographic information*. Any such belief is either making an epistemic error, such as those discussed in sections six and seven, or the property of the person that is taken to be bad is not in fact bad. The latter is explored in the illustration concerning maths scores. The latter error is commonly exemplified when, for instance, sexist people take the fact that women are physically weaker on average to show that women are inferior to men, or when racist people take the double negation of Black American English to show there is something wrong with the dialect. They misunderstand the moral significance of the property.

Advocates of moral encroachment aim to describe a person whose beliefs are epistemically impeccable—well supported by the evidence and conscientiously considered—yet morally wrong because racist. My contention is that no such belief can exist. If a belief is morally wrong then there is some corresponding prior epistemic error. The belief is not well supported by the evidence and/or it is not interpreted through a morally appropriate understanding, and that understanding is not epistemically well supported. If a belief is

epistemically well supported it cannot be racist since no true fact is genuinely racist. With the right background understanding we see that since everyone is equal, any differences based on gender, race, and so on are morally insignificant.<sup>51</sup>

### Acknowledgements

Many thanks to Rima Basu, Renee Bolinger, Jon Garthoff, Kevin McCain, Jessie Munton, Ted Poston, and Ernest Sosa for valuable comments on an earlier draft of this essay. Thanks also to Natalie Ashton, Sarah Moss, Clayton Littlejohn, Cat Saint-Croix, and Mark Schroeder for helpful discussion about these ideas. Finally, thanks to Jennifer Saul and several active members of the Facebook group Board Certified Epistemologists for drawing my attention to relevant literature.

### Bibliography

- Adler, Jonathan (2002) *Belief's Own Ethics* Cambridge, MA: Bradford/MIT.
- Aikin, Scott (2008) 'Evidentialism and James's Argument from Friendship' *Southwest Philosophy Review* 24(1): 173–80.
- \_\_\_\_\_ (2006) 'Modest Evidentialism' *International Philosophical Quarterly* 46(3): 327–43.
- Alfano, Mark (2013) *Character as Moral Fiction* Cambridge University Press.
- Anderson, Elizabeth (2010) *The Imperative of Integration* Princeton University Press.
- Appiah, Kwame Anthony (1990) 'Racisms' *Anatomy of Racism* David Goldberg (ed.) University of Minnesota Press, pp. 3–17.
- Armour, Jody (1994) 'Race Ipsa Loquitur: Of Reasonable Racists, Intelligent Bayesians, and Involuntary Negrophobes' *Stanford Law Review* 46(4): 781–816.
- Arpaly, Nomy (2003) *Unprincipled Virtue: An Inquiry into Moral Agency* Oxford University Press.
- Arpaly, Nomy and Timothy Schroeder (2014) *In Praise of Desire* Oxford University Press.
- Ashton, Natalie (2015) 'Appropriate Belief without Evidence' *Teorema* 34(2): 7–28.
- Ashton, Natalie and Robin McKenna (ms) 'Situating Feminist Epistemology'.
- Basu, Rima. (ms a) 'The Wrongs of Racist Beliefs'.
- \_\_\_\_\_ (ms b) 'What We Epistemically Owe to Each Other'.
- \_\_\_\_\_ (ms c) 'The Moral Stakes of Racist Beliefs'.
- \_\_\_\_\_ (ms d) 'Motivating Ethical Demand on Belief' presentation at Athena-in-Action, Princeton University, August 2016.
- Basu, Rima and Mark Schroeder (forthcoming) 'Doxastic Wrongs' *Title TBA* Brian Kim and Matthew McGrath (eds.) Routledge.
- Begby, Endre (2013) 'The Epistemology of Prejudice' *Thought* 90–99.
- Bolinger, Renee (ms) 'The Rational Impermissibility of Accepting (Some) Racial Generalizations'.
- Buchak, Lara (2014) 'Belief, Credence, and Norms' *Philosophical Studies* 169: 285–311.
- Chignell, Andrew (2010) 'The Ethics of Belief' *Stanford Encyclopedia of Philosophy* (ed.) Edward

---

<sup>51</sup> This paper is forthcoming as Gardiner (forthcoming a).

- N. Zalta.
- Clifford, W. K. (1877) 'The Ethics of Belief' *Contemporary Review* 29: 289–309.
- Conee, Earl and Richard Feldman (2011) 'Replies' *Evidentialism and its Discontents* Trent Dougherty (ed.) Oxford University Press, pp. 283–322.
- \_\_\_\_\_ (2004) *Evidentialism* Oxford University Press.
- Crawford, Lindsay (forthcoming) 'Believing the Best: On Doxastic Partiality in Friendship' *Synthese*.
- DeRose, Keith (2004) 'Questioning Evidentialism' *Evidentialism and its Discontents* Trent Dougherty (ed.) Oxford University Press, pp. 137–46.
- Dotson, Kristie (2014) 'Conceptualizing Epistemic Oppression' *Social Epistemology* 28(2): 115–38.
- \_\_\_\_\_ (2008) 'In Search of Tanzania: Are Effective Epistemic Practices Sufficient for Just Epistemic Practices?' *The Southern Journal of Philosophy* 46: 52–64.
- Dretske, Fred (1970) 'Epistemic Operators' *Journal of Philosophy* 67(24): 1007–23.
- Eaton Daniel and Timothy Pickavance (2015) 'Evidence Against Pragmatic Encroachment' *Philosophical Studies* 172: 3135–3143.
- Enoch, David (2016) 'What's Wrong with Paternalism: Autonomy, Belief, and Action' *Proceedings of the Aristotelian Society* 116(1): 21–48.
- Fantl, Jeremy and Matthew McGrath (2002) 'Evidence, Pragmatics, and Justification' *The Philosophical Review* 111(1): 67–94.
- Feldman, Richard (2000) 'The Ethics of Belief' *Philosophy and Phenomenological Research* 60(3): 667–695.
- Feldman, Richard and Earl Conee (1985) 'Evidentialism' *Philosophical Studies* 48(1): 15–34.
- Franklin, John Hope (2005) *Mirror to America: The Autobiography of John Hope Franklin* New York, NY: Farrar, Straus and Giroux.
- Fritz, James (2017) 'From Pragmatic Encroachment to Moral Encroachment' *Pacific Philosophical Quarterly*.
- Ganson, Dorit (2008) 'Evidentialism and Pragmatic Constraints on Outright Belief' *Philosophical Studies* 139(3): 441–458.
- Gardiner, Georgi (forthcoming a) 'Evidentialism and Moral Encroachment' *Believing in Accordance with the Evidence: New Essays on Evidentialism* Kevin McCain (ed.) Springer.
- \_\_\_\_\_ (forthcoming b) 'The Burden of Proof and Statistical Evidence' *Routledge Handbook of Applied Epistemology* David Coady and James Chase (eds.) Routledge.
- \_\_\_\_\_ (2015) 'Normalcy and the Contents of Philosophical Judgements' *Inquiry* 58(7): 700–740.
- Gendler, Tamar (2011) 'On the Epistemic Costs of Implicit Bias' *Philosophical Studies* 156: 33–63.
- Gordon, Lewis (2000) *Existentialia Africana: Understanding Africana Existential Thought* Routledge.
- \_\_\_\_\_ (1995) *Bad Faith and Antiracist Racism* Humanities Press.
- Grimm, Stephen (2011) 'On Intellectualism in Epistemology' *Mind* 120(479): 705–733.
- Guerrero, Alexander (2007) 'Don't Know, Don't Kill: Moral Ignorance, Culpability, and Caution' *Philosophical Studies* 136: 59–97.
- Hájek, Alan (2007) 'The Reference Class Problem is Your Problem Too' *Synthese* 156(3): 563–585.
- Hawthorne, John (2004) *Knowledge and Lotteries* Oxford: Clarendon Press.
- Hazlett, Allan (2016) 'Intellectual Loyalty' *International Journal for the Study of Skepticism* 6: 326–350.
- \_\_\_\_\_ (2013) *A Luxury of the Understanding* Oxford University Press.



- Heil, John (1983) 'Believing What One Ought' *The Journal of Philosophy* 80(11): 752–765.
- Ichikawa, Jonathan, and Benjamin Jarvis (2009) 'Thought-Experiment Intuitions and Truth in Fiction' *Philosophical Studies* 142: 221–46.
- Ichikawa, Jonathan, Benjamin Jarvis, and Katherine Rubin (2012) 'Pragmatic Encroachment and Belief-Desire Psychology' *Analytic Philosophy* 53(4): 327–343.
- James, William (1956/1896) *The Will to Believe and Other Essays in Popular Philosophy*. New York: Dover.
- Kahneman, Daniel (2011) *Thinking Fast and Slow* Farrar, Straus and Giroux
- Kawall, Jason (2013) 'Friendship and Epistemic Norms' *Philosophical Studies* 165(2): 349–370.
- Keller, Simon (2004) 'Friendship and Belief' *Philosophical Papers* 33(3): 329–351.
- Kelly, Tom (2002) 'The Rationality of Belief and Some Other Propositional Attitudes' *Philosophical Studies* 110: 163–196.
- Kim, Brian (2017) 'Pragmatic Encroachment in Epistemology' *Philosophy Compass* 12(5): 1–14.
- Kunda, Ziva (1990) 'The Case for Motivated Reasoning' *Psychological Bulletin* 108(3): 480–98.
- Kyburg, H.E., Jr. (1961) *Probability and the Logic of Rational Belief* Wesleyan University Press.
- Leslie, Sarah-Jane (forthcoming) 'The Original Sin of Cognition: Fear, Prejudice and Generalization' *Journal of Philosophy*.
- Lewis, David (1996) 'Elusive Knowledge' *Australasian Journal of Philosophy* 74: 549–567.
- Littlejohn, Clayton (2012) *Justification and the Truth-Connection* Cambridge University Press.
- MacFarlane, John (2005) 'Knowledge Laundering: Testimony and Sensitive Invariantism' *Analysis* 65(2): 132–138.
- Marušić, Berislav (2015) *Evidence and Agency: Norms of Belief for Promising and Resolving* Oxford University Press.
- McCain, Kevin (2014) *Evidentialism and Epistemic Justification* Routledge.
- McCormick, Miriam (2015) *Believing Against the Evidence: Agency and the Ethics of Belief* New York: Routledge.
- Mills, Charles W. (1997) *The Racial Contract* Cornell University Press.
- Moss, Sarah (forthcoming) *Probabilistic Knowledge* Oxford University Press.
- Nelkin, Dana (2000) 'The Lottery Paradox, Knowledge, and Rationality' *Philosophical Review* 109(3): 373–409.
- Pace, Michael (2011) 'The Epistemic Value of Moral Considerations: Justification, Moral Encroachment, and James' 'Will To Believe' *Noûs* 45(2): 239–268.
- Piller, Christian (2016) 'Evidentialism, Transparency, and Commitments' *Philosophical Issues* 26: 332–350.
- Preston-Roedder, Ryan (2013) 'Faith in Humanity' *Philosophy and Phenomenological Research* 87(3): 664–687.
- Reichenbach, Hans (1949) *The Theory of Probability* University of California Press, Berkeley.
- Riggs, Wayne (2003) 'Balancing Our Epistemic Goals' *Noûs* 37(2): 342–352.
- Rinard, Susanna (2017) 'No Exception for Belief' *Philosophy and Phenomenological Research* 94(1): 121–143.
- \_\_\_\_\_ (2015) 'Against the New Evidentialists' *Philosophical Issues* 25(1): 208–223.
- Ross, Jacob and Mark Schroeder (2014) 'Belief, Credence, and Pragmatic Encroachment' *Philosophy and Phenomenological Research* 88(2): 259–288.
- Ryan, Sharon (2015) 'In Defense of Moral Evidentialism' *Logos and Episteme* 6(4): 405–427.
- Schaffer, Jonathan (2006) 'The Irrelevance of the Subject: Against Subject-Sensitive Invariantism' *Philosophical Studies* 127: 87–107.

- Schroeder, Mark (forthcoming) 'Rational Stability under Pragmatic Encroachment' *Episteme*.
- Shah, Nishi (2006) 'A New Argument for Evidentialism' *Philosophical Quarterly* 56: 481–498.
- Smith, Martin (2010) 'What Else Justification Could Be?' *Noûs* 44(1): 10–31.
- Stanley, Jason (2015) *How Propaganda Works* Princeton University Press.
- \_\_\_\_\_ (2005) *Knowledge and Practical Interests* Oxford University Press.
- Stroud, Sarah (2006) 'Epistemic Partiality in Friendship' *Ethics* 116(3): 498–524.
- Thomson, Judith Jarvis (1986) 'Liability and Individualized Evidence' *Law and Contemporary Problems* 49(3): 199–219.
- Venn, John (1866) *The Logic of Chance* Macmillan, London.
- Wells, Ida B. (1970) *Crusade for Justice: The Autobiography of Ida B. Wells* Alfreda M. Duster (ed.) University of Chicago Press.
- Westfall, Sandra Sobieraj (2014) 'The Obamas: How We Deal with Our Own Racist Experiences' *People Magazine*.
- Worsnip, Alex (2015) 'Two Kinds of Stakes' *Pacific Philosophical Quarterly* 96: 307–324.

## B. Teleologies and Understanding

### III. Teleologies and the Methodology of Epistemology

**Abstract.** The teleological approach to an epistemic concept investigates it by asking questions such as ‘what is the *purpose* of the concept?’, ‘What *role* has it played in the past?’, or ‘If we imagine a society without the concept, why would they feel the *need* to invent it?’ The idea behind the teleological approach is that examining the function of the concept illuminates the contours of the concept itself. This approach is a relatively new development in epistemology, and as yet there are few works examining it.

This paper aims to fill this gap and engender further understanding of the teleological method. I first contrast the teleological method with more orthodox approaches in epistemology. I then draw a three-way taxonomy of different kinds of teleological approach and provide an example of each kind. The teleological approach is often presented as antithetical to the more orthodox approaches in epistemology, and so in competition with them. I demur. I argue that the methods can be fruitfully combined in epistemological theorising, and I suggest specific ways the teleological approach can be incorporated alongside more orthodox methods in a general methodological reflective equilibrium. In the final section I respond to an objection to the teleological approach pressed by Hilary Kornblith.

#### 1. Introduction

A teleological approach to an epistemic concept investigates it by asking questions such as ‘what is the purpose of the concept?’, ‘What role has it played in the past?’, or ‘If we

imagine a society without the concept, why would they feel the need to invent it?' The idea behind the teleological approach is that examining the function of the concept illuminates the contours of the concept itself. This approach is a relatively new development in epistemology, largely introduced by Edward Craig's *Knowledge and the State of Nature*; as yet there are few works examining the method.<sup>1</sup>

I aim to fill this gap and engender greater understanding of the method. In section two I sketch some orthodox methods in epistemology and explain how the teleological approach differs from these. In section three I examine the teleological approach, draw a taxonomy of different kinds of teleologies and provide an example of each kind. Some theorists suggest the teleological approach is antithetical to orthodox approaches. I argue, by contrast, the teleological approach can be fruitfully incorporated alongside the orthodox approaches within a general methodological reflective equilibrium. I argue for this claim in two ways. Firstly, for each example of a teleology in section three, I draw out what the account suggests about the nature of knowledge, thereby illustrating how the teleological approach informs and complements standard epistemological projects. Secondly, in section four I explain specific ways that insights from the teleological method can augment the orthodox methods. In section five I respond to a recent objection to teleological approaches in epistemology, pressed by Hilary Kornblith. Whilst I focus on understanding epistemic concepts, such as knowledge, much of what I say can be fruitfully generalised to other areas of philosophy.

---

<sup>1</sup> Klemens Kappel has written a useful and insightful paper on teleologies in epistemology, see Kappel (2010). Kusch (2009) was also helpful in constructing my taxonomy. See also Craig (2007) and Fricker (1998: esp. section three) for penetrating discussion about the method. See also the articles in Henderson and Greco (2015).

## 2. Orthodox Approaches

In this section I sketch three orthodox approaches that can be usefully contrasted with the teleological approach. The descriptions are best viewed as summaries, rather than as thorough accounts. It should also be noted this list is not intended as an exhaustive survey of methods in epistemology, nor do I suggest the methods are mutually exclusive; in practice theorists often use all of the methods to varying degrees, and each generates inputs into a large reflective equilibrium. By teasing the methods apart, however, we are better placed to understand them.

### 2.1 Extension-first approach

The ‘extension-first’ method elucidates a concept using various cases to mine our judgements about the intuitive extension and to test the intuitive extension against a proposed claim or theory. A presupposition of this method is that we have intuitive access to the extension of the concept (that is, we have reliable pre-theoretic judgements about when the proposition *S knows that p* is correct). We test a claim or theory by determining whether the theory’s extension and the intuitive extension align.

Suppose our target theory is *Knowledge is true belief*. We can test this theory by devising a case: Sammy correctly guesses the result of a dice throw. The theory’s extension entails Sammy knows the result; if intuition dictates he does not know then there is a mismatch. According to the extension-first method we must either reject the intuition, and provide a debunking story to explain away the misleading intuition, or we must sanction the intuition, and amend the target theory accordingly, so that it no longer conflicts with the intuitive extension.<sup>2</sup>

---

<sup>2</sup> Following Timothy Williamson’s book *The Philosophy of Philosophy* (2007), there has been a recent surge of interest in determining precisely the content of the judgement

Some judgements about the extension might be harder to explain away than others: The Sammy vignette seems to generate a confident and unambiguous judgement about a relatively central case. Other cases might be more peripheral or obscure, and might generate less confident judgements. These latter intuitions will be easier and less costly to explain away. The idea driving the method is that thought experiments reveal our intuitions about the extension of the concept, and after a process of reflective equilibrium—in this case, a two-way reconciliation between intuitive judgements about particular cases and general principles about the nature of knowledge—we can provide an account of the concept.<sup>3</sup>

Experimental philosophy can be seen as a variation of the extension-first approach, since it employs cases to collect intuitive responses about the extension of the concept by asking whether the case exhibits the target concept. Whereas traditionally those using the extension-first method consult fewer people, and typically consult experts, such as fellow

---

we have in response to the vignettes when employing the extension-first method. Various proposals have been advanced, most of which suggest the content is some kind of modal claim. Examples include ‘necessarily, anyone who stands to p in the relation described by the vignette has a true belief that p but does not know that p’, ‘it is possible that someone stands to p as described by the vignette, and she has a true belief that p but does not know that p’, ‘normally the person described in the vignette has a true belief that p but does not know that p’, and ‘counterfactually, were someone to stand to p as described in the vignette, then she would have true belief that p but does not know that p’. See Malmgren (2011: esp. 274–281) and Gardiner (2015a) for discussion of these proposals. Each of the proposals concerns purely the extension of knowledge. Ernest Sosa (forthcoming) advances a competing theory. He suggests that our judgement on hearing a vignette is that, roughly, ‘The subject of the vignette has a true belief, but does not thereby know, that p. Since this judgement includes the ‘thereby’ relation, it is not purely extensional, and so strictly speaking falls outwith the extensional approach as outlined here.

<sup>3</sup> In using the expression ‘account of a concept’ I hope to remain ecumenical among an analysis of a concept, a Carnapian explication, a synthetic account such as Edward Craig’s, or other similar understandings of the aim of epistemological theorising. For discussion of reflective equilibrium in epistemology, see Elgin (1996).

philosophers, experimental philosophy consults large numbers of people, often the untrained ‘folk’, about the intuitive extension.

## 2.2 Intension-first Approach

The intension-first approach uses intuitive judgements about the intension of the target concept to inform an account. We believe, for example, that whether S knows that p is not typically sensitive to the day of the week.<sup>4</sup> Plausibly we access this day-of-the-week-invariance platitude directly from our intuitive understanding of the concept, without proceeding via the intuitive extension. Whilst the day-of-the-week datum is not much help in formulating a theory of knowledge, there may be other judgements about the intension that bear more significantly on our analysis. If we discover from reflecting directly on the intuitive intension whether knowledge has a conceptual connection to rational action or blameless assertion, for instance, or whether possessing knowledge is less demanding for children than for adults, then this will significantly inform our theory.

Theorists who begin with a set of platitudes and build an account of knowledge from them employ the intension-first approach. Such platitudes may include the value platitude (knowledge has value), the anti-skeptical platitude (we have some knowledge), the ability platitude (knowledge is, at least in part, to the credit of the knower), and the anti-luck platitude (knowledge enjoys some kind of modal stability). Theorists who take as inputs the normative roles that knowledge plays in our socio-cognitive environment (such as the norm of assertion, the norm of rational action, or the epistemically proper end of inquiry) can also be fruitfully seen as employing the intension-first approach.

---

<sup>4</sup> For some propositions, of course, whether S knows p does depend on the day of the week. S cannot know it is Thursday, for example, unless it is Thursday. But it is not typical for knowledge to depend on the day of the week. Thanks to Ernie Sosa for pointing this out.



### 2.3 Linguistic Approach

Another method we can usefully contrast here is the linguistic or ‘ordinary-language’ approach: When analysing epistemic concepts, we survey how people *use* epistemic terms and this guides theorising about the concepts. The linguistic approach might proceed with corpus studies, testing felicity judgements, semantic analysis, cross-linguistic studies, or simply observing, using empirical methods, how people use the term ‘knows’ and its cognates. Studying linguistic behaviour to illuminate concepts presupposes that there are deep connections between language and concepts; it is fruitful only if language provides some insight into concepts.

This method is employed in debates about epistemic contextualism, for example, as how people use ‘knows’ in ordinary language is held to support claims about which parameters affect the truth of knowledge attributions. Similarly we can examine what epistemic terms convey in conversations: ‘Apparently Mark attended the party’ conveys to competent language users that the speaker does not possess first-hand evidence, for instance.

The linguistic approach is distinct from extension- and intension-first approaches: Whereas the extension-first approach asks whether the concept intuitively applies to a case, the linguistic approach examines language use itself. Proper language use—including attributions and denials of knowledge—can diverge from judgements about the extension of knowledge. We might assert a knowledge ascription, even whilst denying the corresponding extensional claim. On hearing a piece of gossip we might exclaim ‘I knew it!’, whilst recognising we did not know. Here language use conveys something other than a literal knowledge attribution. Perhaps in this case ‘I knew it!’ conveys

something like, ‘That information fits well with my previous beliefs and suspicions’, and generates the perlocutionary force of encouraging the gossipier to reveal more gossip. We might comfort an ill friend by saying ‘I know you will survive this’ whilst recognising we do not know that she will. We might express annoyance at getting a pub quiz question wrong by asserting, ‘I knew the answer’, whilst recognising we did not. These utterances might be linguistically appropriate, even though strictly speaking false, precisely because linguistic behaviour is distinct from reporting judgements about the extension. Conversely it might be inappropriate to assert a knowledge claim, even if you recognise that knowledge obtains. Perhaps lottery cases are like this: Perhaps it is typically inappropriate to assert ‘I know I won’t win the lottery’, even though you do know, because the assertion typically implies you possess some non-statistical, inside evidence.

Another, related, difference between the linguistic approach and the extension-first approach is that the former may be particularly useful in determining various different conversational uses of the term ‘knows’, whereas the extension-first approach may be better suited to exploring just one concept of knowledge.

A fourth difference between the linguistic approach and the extension- and intension-first approaches, is that plausibly expertise plays a different role in data collection. Perhaps consulting experts is particularly helpful when generating results using the extension- and intension-first approaches; plausibly consulting experts provides more accurate and coherent responses. By contrast, plausibly when seeking information about linguistic behaviour any competent—or native—speakers would suffice, and expertise does not provide extra advantage for, and may even hinder, inquiry. If correct this suggests that insofar as theorists aim to determine the structure of a coherent concept we should ask experts, rather than the folk, and use the extension- and intension- first

methods. And insofar as we aim to tap into folk concepts we should consult the folk, and the linguistic approach may be particularly well suited for this.

Those employing the linguistic approach may either count all ordinary language uses of ‘knows’ as equally proper, and so theorise about the nature of knowledge from all uses; or they may rule that some cases are proper or central uses, whilst others are figurative, improper or mistaken. Either way, ordinary language use shapes the account of the concept being developed.

### 3. Teleological Approaches

Teleological approaches aim to illuminate a concept by asking what the *point* of that concept is; what *purpose* it fulfils, what *need* it meets, what *function* it has, or what *role* it characteristically plays. The idea is that by focusing on the needs the epistemic concept fulfils we illuminate the nature of the concept.<sup>5</sup>

Because it relies less on intuitive responses to individual cases, the success of the teleological approach does not depend on our having reliable intuitive access to the intension or extension of the target concept. It also does not require that our ordinary language behaviour track closely the contours of the concept. The approach is distinct from the others because it begins by looking at our socio-cognitive economy, and determining what concepts might be useful, rather than starting with any particular claims about a concept’s contours, such as particular instantiations or intensions.

---

<sup>5</sup> Teleological approaches are not limited to understanding concepts: they can also be used to understand, for example, artefacts or cultural phenomena. Kappel (2010) illustrates this with a practical explication of the artefact of a car. I focus here on teleologies as employed to illuminate epistemic concepts.

There has been a recent surge of interest in teleologies in epistemology, and whilst various teleologies have been proposed, the method has not been extensively explored.<sup>6</sup>

In this section I construct a taxonomy comprising three variations of the approach. For each variation I explain the approach, and provide an example for illustration. I then draw out what that example teleology, if correct, suggests about the nature of knowledge.

Many theorists who adopt the teleological approach in epistemology couple their application of it with a rejection of the project of giving an analysis of the concept of knowledge in terms of necessary and sufficient conditions, and instead judge that the project of epistemology ought to be elucidation of concepts, constructing a network analysis, or some other synthetic, non-orthodox, aim.<sup>7</sup> The guiding assumption seems to be that the teleological approach is not well suited to serving the aim of conceptual analysis.<sup>8</sup> Conversely the orthodox methods are seen as appropriate for conceptual analysis, but ill-equipped for a synthetic aim, so that those who reject the orthodox methods thereby see themselves as having reason to eschew the analytic aim in favour of

---

<sup>6</sup> For existent discussions of the methodology, see Craig (2007), Fricker (1998: esp. section three), Kappel (2010), and Kusch (2009).

<sup>7</sup> In line with Craig's terminology, by an 'analytic' aim I mean the project of rendering the concept into constituent parts, such as necessary and sufficient conditions (Craig (1990: 2)). By 'synthetic' aim I mean something more broad: any aim of better understanding knowledge, or the concept of knowledge, that does not try to analyse knowledge into constituent parts. At least, this is certainly in line with Craig's use of 'analytic'. It is slightly unclear what Craig means by 'synthesis', other than that he contrasts it with analysis. The term 'synthesis' and its cognates only appears once in Craig's monograph—twice if you count the subtitle—and not at all in Craig (1986/1987) or Craig (2007). (This is rather like, if you will permit the comparison, learning that Darwin's *Origin* only contains the term 'evolved' once.) Craig's not being wholly explicit about what he means by a synthesis may have contributed to a subsequent conflating of his method (a genealogy) with his aim (a non-analytic account of a concept).

<sup>8</sup> For explicit claims to this effect, see Fricker (1998: esp. section three; 2008: esp. 47–48), Craig (1990: esp. chapter one; 1986/1987), MacBain (2004: 193–196), and Hannon (ms). See also Williams (2002) and Kusch (2009: esp. section three). Gardiner (ms) is an example of a teleological approach in epistemology that does not aim to result in necessary and sufficient conditions or a similar analysis.

a synthetic aim.<sup>9</sup> In a similar spirit, theorists who adopt the aim of rendering knowledge into constituents, such as necessary and sufficient conditions, may doubt the usefulness of the teleological method for their project.<sup>10</sup>

I think these methodological inferences are mistaken. Against the tide of people working in teleological theorising I hold that once we have distinguished clearly the aim, such as analysing knowledge into constituents or generating a synthetic account of knowledge, from the method, namely positing teleologies or using more orthodox methods, we should see that either method can be fruitfully applied for either aim. In particular I argue the teleological method can be used alongside more orthodox methods to generate inputs for our reflective equilibrium as part of the analytic project. I advance this claim in two ways: In this section I give examples of how specific teleologies can support particular claims about the analysis of knowledge. In the next section I articulate several ways the teleological and the more orthodox approaches can be fruitfully combined.

### 3.1 Practical Explication

Practical explications are the most straightforward of the three variations. The method is synchronic (it does not probe the past) and factual (it relies on empirical facts rather speculative claims). The theorist makes the following two claims:<sup>11</sup>

---

<sup>9</sup> For claims that adopting a synthetic aim in epistemology should be paired with rejection of the orthodox methods, see Kusch (2011: 17–20) and Fricker (2008: 47–50). A notable exception is the knowledge-first movement, which employs orthodox methods whilst rejecting the aim of rendering knowledge into constituent parts.

<sup>10</sup> There are exceptions, see for example Greco (2012: esp. section four) and Pritchard (2011; 2012). Both harness Craig's teleological account as evidence to support their favoured analyses of knowledge. By using the methods in harmony to increase the inputs into a reflective equilibrium, they exemplify precisely the approach I advocate.

<sup>11</sup> I have adapted this schema from Kappel (2010: 72–73). Kappel describes the second claim as one of stipulation—we stipulate that the target concept is what fulfils the need. I think this may be better viewed as a hypothesising, however, rather than stipulating.

- i. Given a set of facts about our physical constitution, cognitive powers and environment, and a set of aims and interests, we have a particular conceptual need;<sup>12</sup>
- ii. The target concept is what actually and currently meets that need.

The practical explication method runs thus: since the target concept fulfils the identified need, determining the exact nature of the need and discerning what would satisfy that need, should illuminate the concept.

Klemens Kappel offers a practical explication of the concept of knowledge.<sup>13</sup> He notes some propositions are important for practical reasons and we are better off, from a practical point of view, if we have true beliefs about them. We thus have a practical reason to inquire. He notes inquiry costs time and resources, and has no natural stopping point—since there are always further uneliminated error possibilities, we always have a *pro tanto* reason to continue inquiring.<sup>14</sup>

The need Kappel identifies is for an inquiry stopper; a predicate that expresses, to ourselves and others, the judgement that we should take the truth of a claim for granted in practical deliberation.<sup>15</sup> In Kappel's words, 'The sound mind will be in need of a way of expressing the attitude that inquiry has now taken one far enough, and that one shouldn't worry about remaining as yet uneliminated error-possibilities'.<sup>16</sup> He

---

<sup>12</sup> As Kappel clarifies, by 'need' he means something that we are better off if we meet, rather than the stricter sense of something without which we cannot function (Kappel (2010: 79–80)).

<sup>13</sup> Chris Kelp (2011) independently posits a similar function—that of an inquiry stopper—for the concept of knowledge.

<sup>14</sup> Kappel (2010: 73–76).

<sup>15</sup> Kappel (2010: 76).

<sup>16</sup> Kappel (2010: 75).

hypothesises the concept ‘knowledge’ meets this need, and suggests S’s knowing p entails:

P, and S is in a sufficiently good epistemic position with respect to p, such that S ought to take the truth of p for granted in her practical and theoretical deliberation. Kappel (2010: 78).

We can apply the thesis that the concept ‘knowledge’ functions as an inquiry stopper to the project of analysing the nature of knowledge. If the concept of knowledge fulfils this need, as Kappel claims, this suggests pragmatic encroachment about knowledge. Pragmatic encroachment holds that practical factors affect whether the agent’s epistemic position with respect to p suffices for knowledge (or, to formulate it as a semantic thesis, practical factors affect the truth value of a knowledge attribution in a context).<sup>17</sup> Kappel’s teleology suggests pragmatic encroachment because if we care about truth for pragmatic reasons, and we care more about the truth of some propositions than others, then we should invest more time inquiring into those propositions that we have more reason to care about. So it takes less inquiry to reach the point of stopping when investigating less important propositions (if, for example, stakes are lower). Thus if knowledge marks when to stop inquiry, it is easier to know less important facts. This is the pragmatic encroachment thesis.<sup>18</sup>

---

<sup>17</sup> Hannon (2013) develops the view that Craig’s genealogy of knowledge, described below, supports pragmatic encroachment about knowledge.

<sup>18</sup> Moves are available, of course, to endorse Kappel’s teleology yet deny pragmatic encroachment. Perhaps one could hold that the concept of knowledge serves to flag when we can permissibly cease inquiry in some default, typical practical context. Sometimes, depending on the context, we must inquire more—or may inquire less—than that required for knowledge. Such views might successfully deny pragmatic encroachment whilst endorsing Kappel’s view. But Kappel’s teleology, at the very least, *prima facie* suggests, and gives evidence for, pragmatic encroachment. Thanks to Lisa Miracchi for pressing me on this point.

### 3.2 Historical Genealogical Teleology

I introduced the practical explication teleology, which is synchronic; the other two kinds are genealogical and hence diachronic. They also aim to elucidate a concept by asking teleological questions—examining the function of the concept to illuminate the concept itself—but the genealogical versions are retrospective: they look at what function the concept has played or what need it used to meet. There are two kinds of genealogical teleologies, historical and hypothetical.

The historical approach examines the function of the concept in the actual past, and from this infers facts about the present concept. It uses empirical data from historians about the role the concept played, or the way a term was used, and then analyses the conceptual need that was being fulfilled. The theorist then postulates cultural change that causes the concept to evolve over time and thereby illuminates the contemporary concept. The historical approach postulates a hereditary relationship between ancestor concept and its contemporary counterpart.

Martin Kusch provides an example of a historical genealogical teleological account.<sup>19</sup> He draws on Steven Shapin's (1994) *A Social History of Truth: Civility and Science in Seventeenth Century England*, which tracks epistemic language and concepts during the Enlightenment. Shapin notes that natural philosophers of seventeenth-century England accepted testimony was needed in order to learn about the natural world, but recognised that finding good testifiers was a difficult practical problem. Seventeenth-century literature suggests several maxims, such as 'assent to testimony which is plausible', 'assent to testimony which is multiple', and 'assent to testimony which is consistent'. As soon as a maxim was suggested, however, problems were found. Only one maxim was never

---

<sup>19</sup> Kusch (2009).



challenged: ‘assent to testimony from sources of acknowledged integrity and disinterestedness’.<sup>20</sup> But how, asks Kusch, were they supposed to identify who the disinterested reporters were? This is the conceptual need Kusch identifies: finding and predicating those sources who were disinterested and so would tell the truth.

Kusch claims this is where nobility and freedom become important in the genealogy of knowledge. He argues seventeenth century intellectuals believed ‘Gentlemen were truth-tellers because nothing could work upon them that would induce them to be otherwise.’<sup>21</sup> They were believed to be disinterested because of financial independence, and so were considered ‘knowers’, unlike labourers and females. Kusch quotes Shapin on the servants who assisted Boyle in his lab:

Whatever information the domestics produced, it became knowledge, and thus a property of the gentlemanly community of natural philosophers, only once it was vouched for by Boyle or another gentleman. (Shapin (1994: Ch.8), quoted in Kusch (2009: 87)).

Gentlemen held that employees and women did not have the financial freedom to create and pass on knowledge. Kusch quotes Shapin:

The conventions of gentlemanly experimental philosophy did not allow for anyone to openly express disbelief in a report coming from a gentleman. The situation was very different for all those who did not make the gentry grade: women, servants, ‘the poor and the mean in general’, merchants, Catholics, Continental gentry, Italians and politicians. In the cases of all of these groups, their ‘unreliable truthfulness [...] was pervasively referred to their constrained circumstances’. (Shapin (1994: 86); Kusch (2009: 87)).

---

<sup>20</sup> Shapin (1994: 212) and Kusch (2009: 86).

<sup>21</sup> Shapin (1994: 84) and Kusch (2009: 87).

According to contemporary usage these people could possess true beliefs, but not knowledge. The term ‘knower’ was reserved for someone you could trust because they would not need the financial gain made possible by lying.

What does this teleology suggest about the concept of knowledge? Kusch argues his historical genealogy is evidence that ‘to attribute knowledge is to attribute honour, freedom and social power’.<sup>22</sup> He argues that we should expect conceptual connections between the concept of knowledge and that of value, praising someone, virtue, knowledge being an honourific, and bestowing status (such as freedom from being doubted).

Kusch has independent reason, inspired by Williams (2002), for positing that knowledge ascriptions are honourific.<sup>23</sup> Following Williams, Kusch notes that in a state of nature people would free ride on the institution of testimony.<sup>24</sup> He thus identifies a need: motivating people to investigate and testify honestly, and incentivising people to not free ride on the (state of nature) institution of sharing epistemic goods within a community. He suggests that attributions of knowledge fulfil that need, and—here Kusch departs from Williams—they can do so because knowledge attributions are honourific; they bestow value and praise the subject (Kusch (2009: 79)).

### 3.3 Hypothetical Genealogical Teleology

The third kind of teleological approach is a hypothetical genealogical teleology. This approach is genealogical as explores the beginning and history of the concept, and

---

<sup>22</sup> Kusch (2009: 83–87).

<sup>23</sup> Kusch (2009: sections six to eight).

<sup>24</sup> Williams (2002: 88) and Kusch (2009: 74).

hypothetical because rather than tracking the actual history of the concept, it posits a fictional historical narrative.<sup>25</sup>

The method runs thus: First, imagine a state of nature society, and hypothesise our needs in that state of nature, then posit a proto-concept which fulfils those state-of-nature needs. We then hypothesise how the proto-concept would evolve as a response to societal change, thus illuminating our contemporary concept.

Craig's seminal *Knowledge and the State of Nature* employs this hypothetical genealogical methodology, and it paved the way for other similar approaches in epistemology, such as Kusch's and Kappel's work described above. Here I provide only a quick summary of his teleology. Craig imagines a primitive society in which people are language-using and minimally cooperative. He posits that in this state of nature humans need true beliefs about their environment. Given that inquiry is costly, and some people have better access to true beliefs than others, we have a need to tag good informants.

Craig's hypothetical genealogy claims the proto-concept of knowledge fulfils the need of tagging good informants. He then describes a process of increasing objectivity through which the concept changes over time to become less relativised to particular inquirers, so that knowledge attributions can be usefully passed among individuals. 'S knows that p' changes from expressing *S is a good informant for me about whether p* to expressing *S is a good informant for anyone about whether p*.

---

<sup>25</sup> Fricker (2008: 47) refers to this as a narrative of 'semi-fictional time'. Craig (2007: esp. 190–197) discusses to what extent his genealogy, and state of nature genealogies in general, qualify as fictional as opposed to real.

We can then ask what Craig's account suggests about knowledge. Several suggestions have been advanced; here I focus on one. Duncan Pritchard argues that the concept 'good informant' has two aspects: someone on whom we can rely, and someone who is reliable.<sup>26</sup> Pritchard argues that these correspond to modal and virtue-theoretic conditions respectively. Thus we should expect an analysis of knowledge to reveal a bipartite structure, reflecting these dual aspects of 'good informant'. Thus, he concludes, Craig's thesis that knowledge attributions tag good informants provides reason to think the modern concept 'knowledge' has the structure of anti-luck virtue epistemology (the thesis that knowledge is safe, virtuous, true belief).

#### **4. Roles for Teleologies**

I have introduced the teleological approach, explained three variations of it, and contrasted the approach with more orthodox methods in epistemology. In this section I explain specific ways teleological theorising can complement the orthodox methods. In doing so I hope to further advance the view that teleological theorising can augment, rather than compete with, the orthodox methods.

Firstly, deploying the teleological approach may reveal how precise or fine-grained we should expect a concept to be, and so delineate the limits of the other three methods. For example, the concept 'adult' serves a legal need (voting, criminal law etc.). It requires a relatively precise definition to fulfil that need—and sure enough it has one. In contrast the concept 'adolescent' fulfils needs in marketing, media, social science, and stereotyping. In order to fulfil these roles the concept does not need such precise contours; its boundaries can be more vague. Thus we have reason to think it is misguided

---

<sup>26</sup> Pritchard (2011; 2012).

to seek a precise analysis of the concept ‘adolescent’ and we ought only aim for more broad-brush analyses.

To illustrate within epistemology: those working on anti-luck epistemology employ myriad subtle variations of examples, such as variants of barn façade cases, to mine intuitions about epistemic luck’s exact nature.<sup>27</sup> These cases differ only slightly from one another; the distinctions the theorists adjudicate among are fine-grained. Perhaps exploring the function of the anti-luck aspect of knowledge will indicate that we should not expect the anti-luck condition to be quite so fine-grained, and our intuitions about the extension to be quite so discriminatory. Thus the teleological method could illuminate the limits of the other methods.

Secondly, teleological accounts engender understanding by making intelligible why we have the epistemic concepts that we have. We can discern the structure of knowledge using orthodox methods, such as the extension-first approach, and then employ teleologies to explain this structure. The thesis that knowledge functions as an inquiry stopper, for instance, can help explain the role of truth in knowledge. Teleologies deepen understanding by explaining how constituents of the concept come together to perform its function. (It should be noted that in some special cases the teleological approach cannot illuminate why we have the concepts we do. Perhaps, for example, the concept’s structure came about through mistake, equivocation or misunderstanding, rather than by

---

<sup>27</sup> Examples might include Kripke’s red barn cases (Kripke (2011: chapter seven)), cases where there are forks in the road near barn façade county and only some paths would lead the protagonist to barn façade county, and cases where the façades were removed the previous day, leaving only real barns on the day S drives through. Some theorists might already worry about whether we can adjudicate whether S knows in the simple barn façade cases, and as the variations become increasingly baroque, these concerns only heighten.

being well adapted to meet a need. In such cases no functional story will explain the structure.)

As Craig points out the teleological method can explain features of knowledge, such as why the term can be found in every language, and can cast light on the appeal of various proposed analyses of knowledge. Craig (1986/1987: 212) writes,

Suppose that the problem of the analysis had been solved, so that agreed necessary and sufficient conditions for the ascription of knowledge were now on the table. Many writers make one feel that this would be a terminus, and investigation concluded. I should see this as a prolegomenon to further inquiry: why has a concept demarcated by those conditions enjoyed such widespread use?

Craig sees the traditional project as concerning the word 'know', rather than either the concept of knowledge or knowledge itself (Craig (1986/1987: 221)). Many theorists deny this claim. But the methodological point he makes is important: there is more to theorising about knowledge than simply generating necessary and sufficient conditions, or capturing the intuitive extension with an explicit intension. We also want to understand features of the phenomena, and teleological theorising can promote understanding.

Thirdly teleologies can provide independent verification of the results of other approaches. The teleological approach is less dependent on 'raw' intuition, and not in the direct way of the intension- and extension-first approaches. Comparing results from the teleological approach might indicate that our raw intuitions about cases are reliable. This can happen if, for instance, the intuitive extension comports with a proposed teleology. Again, these methods interact in reflective equilibrium, if no plausible teleology fits an

account developed using the orthodox methods, then this is a strike against the account's plausibility.

Conversely, results from orthodox methods can provide justification for, or adjudicate among, given teleologies. There is nothing bedrock about a chosen teleology or what we claim it reveals about the concept: we judge the plausibility of a teleology by determining whether the extension or intension predicted by the teleology matches the intuitive one. Someone committed to the falsity of pragmatic encroachment may well reject Kappel's teleology, for example. And they will have developed their opinion on pragmatic encroachment via the orthodox methods of intension-first, extension-first and linguistic studies. Whilst few theorists hold that intuitions about the extension and intension of a concept are infallible, or even enjoy the level of confidence that observations typically have in empirical science, they may be closer to justificatory bedrock than theses about the function of the concept, and so can justify proposed teleological accounts. In general teleological studies increase the kinds of sources that feature in our reflective equilibrium, which increases the justification of the resultant theory.

One related worry is there is not enough to constrain teleological theorising, especially regarding hypothetical and historical teleologies. Fricker writes,<sup>28</sup>

State of nature stories [...] are notorious for providing a blank canvas onto which a philosopher may paint the image of his personal theoretical predilections.

The concern expressed is that if teleological theorising is too open ended, we cannot appropriately adjudicate among proposed teleologies, and teleologies cannot verify or

---

<sup>28</sup> Fricker (1998: 164). Kappel (2010: 80–81) expresses a similar worry. This concern is frequently pressed in conversation about teleological approaches.

calibrate the results of other methods. I am sympathetic to this worry, but see, for example, Pritchard (2011) for an example of employing a given teleology to adjudicate between competing analyses, and see Kelp (2011) and the debate between Gardiner (ms) and Hannon (ms) for examples of employing extension-first and linguistic considerations to adjudicate between competing teleologies. The existence of such debates provides some sense that such adjudication is possible, and some sense of how it may proceed.

Some theorists suggest that teleological theorising illuminates epistemic value, and proceed with the following kind of reasoning: the concept of knowledge allows us to tag good informants. Tagging good informants is valuable, since this helps us acquire true beliefs and avoid false beliefs, and so this value explains the value of knowledge. This line of reasoning is mistaken; it is too quick. To see why, consider an analogous line of reasoning. The concept of sexual harassment is a valuable concept since it allows people to tag a certain kind of poor behaviour and tag people who act in those ways.<sup>29</sup> This tagging is valuable, since it helps people avoid such behaviour and perhaps encourages appropriate consequences for the behaviour. And so this explains the value of sexual harassment.

Clearly this second line of reasoning is mistaken. We must keep separate the value of the concept and the attributions from the value of the phenomenon itself. Perhaps the value of knowledge is connected to the teleological functions of the concept of knowledge and knowledge attributions but, if so, it is not the straightforward relation articulated above.

---

<sup>29</sup> The concept of sexual harassment is valuable for other reasons too, such as making sense of a certain kind of experience.



Finally, the teleological approach might play a role in responding to skeptical challenges. The teleological approach is non-skeptical in a flatfooted way. There are few candidate roles for the concept of knowledge that entail the extension of the concept is empty. There is the role of picking out the mental state, evidence base, or epistemic relation to the world that is unobtainable for people in the human epistemic condition. There is flagging an ideal or perfect epistemic position. Perhaps there is the role of describing an epistemic haven that angels and gods can enjoy, but lesser beings could never enter. These kinds of functions for the concept of knowledge would result in a skeptical conclusion: we do not possess much, if any, knowledge. But these roles for the concept of knowledge are not compelling; our behaviour and thought indicates the concept does not fulfil these roles. Instead it seems the concept of knowledge fulfils roles that entail a broad extension; we possess much of the knowledge that we typically take ourselves to have. In this way the teleological approach is a non-skeptical approach to epistemological theorising.

But the teleological approach can provide resources to engage with skepticism in less flatfooted ways too.<sup>30</sup> One fruitful way to understand the skeptical challenge is as an attempt to deny us something that we thought we possessed, and that we care about possessing. Perhaps there are some epistemic states, practices, or competences—such as Cartesian certainty about the external world, perfectly trustworthy testimony, or wholly infallible thinking, for example—that skeptical reasoning shows we do not possess. If we either do not value those things, or on reflection would recognise that we do not have them, then arguably relinquishing these things to skepticism is not so troubling or serious. By contrast there might be other states, practices, and competences that we do

---

<sup>30</sup> Thanks to Yuval Avnir, Jon Garthoff, and John Greco for helpful discussion of these issues.

value, and that we take ourselves to ordinarily possess. Examples might include the legitimacy of our practices of giving and accepting reasons for belief, typically being in a position to testify responsibly, and typically being warranted in trusting our reasoning, perceptual and memorial faculties. Relinquishing these things to skepticism would be a more serious defeat.

The teleological approach can both help explain why we care about the epistemic states, practices and competences that we value, and provide evidence that these phenomena obtain. This is because the approach can illuminate the roles these phenomena play in our social, cognitive, and practical lives, and so help explain their value whilst proving their existence. Perhaps, for example, teleological reasoning illuminates why the institution of testimony has value, and by subsequently showing that the sustainability of this institution requires the concept of knowledge to incorporate honourific force, we thereby illuminate the value of knowledge.<sup>31</sup> To put the point slightly differently: skepticism should show both that something is lacking, and that the thing missing has value. The teleological approach can both provide evidence that something obtains—if the concept of knowledge plays the role of inquiry stopper, for example, this suggests that there is some knowledge—and it can illuminate why that thing has value—because, for example, appropriate cessation of inquiry plays an indispensable or constitutive role in our cognitive and practical lives.

Williams (2002) and Nietzsche both posit genealogies to explain features of normative, value-laden phenomena. Williams explores a genealogy for epistemic value; Nietzsche explores a genealogy for morality. They both posit that the original function of the value associated with these domains is to manipulate people. Attributions of epistemic value,

---

<sup>31</sup> See Williams (2002) and Kusch (2009: 77ff.).

according to Williams, motivate people to not free ride on the practices of information sharing in the state of nature.<sup>32</sup> Attributions of moral value, according to Nietzsche, are to motivate the powerful to consider the interests of the less powerful. But whereas Nietzsche's genealogy purports to debunk or undermine the value in question, Williams's genealogy seeks to vindicate it.

### 5. Kornblith's Contention

Hilary Kornblith (2011) criticises Craig's version of the teleological approach for relying on the false assumption that humans create concepts, such as that of knowledge, to satisfy needs. Kornblith contends that Craig's project rests on this assumption, and it is false. Kornblith instead argues knowledge is a natural kind that fulfils the roles it does because of its nature; it does not have its nature because of the roles it fulfils. Kornblith thus challenges the foundation of the Craigian research programme in epistemology. In this section I defend the teleological approach against Kornblith's criticisms by suggesting that even if Kornblith is correct that knowledge is a natural kind, we can still fruitfully employ teleological theorising in epistemology.<sup>33</sup>

Kusch (2013) argues that Kornblith's understanding of natural kinds is too narrow. According to the narrow sense of natural kind, natural kinds are limited to the theoretical posits of natural science; natural kinds are classifications that play a significant explanatory and predictive role in natural science. This narrow interpretation of 'natural kind' includes chemical kinds, for example, but does not include social kinds such as Islam. Kusch argues that social kinds such as 'monarchy' and 'Islam' qualify as natural kinds, under a superior 'wide' understanding of natural kind. According to this 'wide

---

<sup>32</sup> Williams (2002: 88–93).

<sup>33</sup> This section does not appear in Gardiner (2015b).

reading' of natural kinds, non-natural kinds are limited to gerrymandered kinds, such as 'the tip of my nose, the Vienna *Hofburg* and the number 255' (Kusch 2013).<sup>34</sup>

Kusch holds that knowledge qualifies as a natural kind according to the wider reading, since it is a non-gerrymandered social kind. He argues that thinking about the genealogy of our social kind concepts is important for understanding the social phenomena themselves. He writes,

Kornblith's thesis—that concepts and intuitions are of little interest—is not very plausible, however, in the case of social kinds [...] His dismissal of folk concepts and intuitions in the case of social kinds amounts to a rejection of 'actors' categories' as important to social science. This is a highly contentious claim. Can we really make sense of, say, 'democracy', 'Islam' or 'Empiricism' as wide natural kinds in political science or history, without paying attention to how these categories were understood by the historical actors themselves? Surely only in some pretty exceptional circumstances. (Kusch 2013: 93.)

Kusch's view seems plausible for social kinds; how we think of them contributes to their causal and explanatory powers. We could not understand democracy or its history without considering the evolution of the concept and conception of democracy. But Kornblith does not agree that knowledge is a social kind. He argues that knowledge is a natural kind even in the narrow sense of 'natural kind'; on Kornblith's view, knowledge is the kind of property posited in cognitive ethology.<sup>35</sup>

In this section I consider briefly the prospects for the teleological approach in epistemology if knowledge is, as Kornblith holds, a natural kind in the narrow sense. I

---

<sup>34</sup> I will not adjudicate here how we should understand the term 'natural kind'. Although plausibly the term 'natural kind' is less useful if it is so broad that the only kinds it excludes are gerrymandered kinds such as 'the tip of my nose, the Vienna *Hofburg* and the number 255'. This consideration would favour the narrower sense of natural kind, which also excludes social and artificial kinds.

<sup>35</sup> See Kornblith (2002).

argue that even if Kornblith is correct that knowledge is a narrow natural kind, the teleological approach might still be a useful method in epistemology.

Kornblith argues that knowledge is a natural kind that fulfils the roles it does because of its nature; it does not have its nature because of the roles it fulfils.<sup>36</sup> He contends that this fact undermines Craig's approach. But note that the teleological approach can be fruitfully applied to phenomena that do not exist in order to fulfil a function, and instead serve a function because of its antecedent nature. To illustrate, suppose you learn Shakespeare used a rock as a paperweight; the rock functioned to prevent the breeze from scattering his papers. By identifying the function the rock satisfied you can infer facts about the rock. You can infer that likely it was not larger than a coffee mug, for example, nor smaller than a thimble; it was not chalk or charcoal, and would not leave a smudge. The rock was likely aesthetically pleasing and was not smelly, dirty, fragile, or sharp. We can infer all this by reflection on the function the phenomenon fulfilled even though the phenomenon was not created to satisfy the function. The rock fulfils the role it does because of its nature; it does not have the nature it does because of role it fulfils, and yet we can fruitfully apply teleological theorising to illuminate the nature of the rock.

Perhaps, then, we can similarly theorise about the roles that knowledge and knowledge attributions satisfy as a way to better illuminate the nature of knowledge. This method can be fruitful even if epistemic kinds satisfy their roles because of their nature, rather than having their nature because of their roles.

---

<sup>36</sup> Kornblith (2011: 45).

In Kornblith's view knowledge is a narrow natural kind and this contention underwrites his criticisms of the teleological approach: Craig's version of the teleological approach presupposes that,

Knowledge is not a given phenomenon, but something that we delineate by operating with a concept which we create in answer to certain needs, or in pursuit of certain ideals. The concept of water, on the other hand, is determined by the nature of water itself and our experience of it. (Craig 1990, quoted in Kornblith 2011).

Kornblith argues that since knowledge is a (narrow) natural kind the teleological approach is inappropriate. In response I suggest that the teleological approach can fruitfully illuminate phenomena that are natural kinds. This can include, for example, applying a teleological approach to natural kinds that have the nature they do because of the function they satisfy. This can be illustrated by reflection on bird nests and beehives. From knowing that beehives serve to shelter bees, we can discern that beehives are largely hollow and have entrances. We can estimate their size, location, and frequency. We can explain why, when, and where they appear. We might also predict that beehives contain chambers, since this would maximise resting space for bees, and include a larger chamber for the queen bee. We can infer all this from reflection on the function of beehives, even though they are a natural kind.

We can also use our patterns of attribution and labelling to illuminate natural kinds and our uses for them. We do not attribute all and only instances of a kind we observe; we are selective. We have purposes in attributions and language use. Water is a natural kind, for example. We tend to refer to the contents of the Raritan River as water, but not a cup of tea as water, even though the former plausibly contains less H<sub>2</sub>O than the latter. These patterns of attribution can illuminate important features of water, including its role in our lives, even though water is a natural kind. Another suggestive example is that 'relation',

‘kin’, and ‘family’ might have similar or identical extensions, yet the different functions of the attributions are illuminating.

Thus I hope to have motivated the idea that even if Kornblith is correct that knowledge is a narrow natural kind, teleological theorising can still be a useful method in epistemology.

## 6. Conclusion

My aim is to engender understanding about teleological approaches in epistemology. I outlined the teleological method and distinguished it from the extension-first, intension-first and linguistic approaches. I constructed a three-way taxonomy of the teleological approach and provided an example of each. I illustrated how the teleological approach can be fruitfully applied to the more orthodox aim of analysing the concept of knowledge by drawing out what particular teleologies suggest about knowledge. Contrary to those who believe the teleological approach is in opposition to more orthodox methods, I develop the idea that teleologies can be fruitfully employed alongside orthodox approaches when analysing knowledge. I do this, in part, by explaining some specific ways to do so. Finally, I suggest a defence of teleological theorising against a recent objection.<sup>37</sup>

## Acknowledgements

Thanks to Yuval Avnir, Peter Baumann, Matthew Chrisman, Megan Feeney, Will Fleischer, Miranda Fricker, Jon Garthoff, John Greco, David Henderson, Chris Kelp, Martin Kusch, Michael Lynch, Peter Milne, Lisa Miracchi, Duncan Pritchard, Ernie Sosa, Chris Tweedt, and Tedla Woldeyohannes for helpful comments and discussions. Thanks also to audiences at the *Episteme* Conference on Cognitive Ecology, Saint Louis University, the University of Edinburgh, the Pontifical Catholic University of Rio

---

<sup>37</sup> An earlier version of this essay was published as Gardiner (2015b); parts of this essay appear in Gardiner (2011). Please note that the section numbers in this essay differ from the section numbering system in Gardiner (2015b).

Grande do Sul, Swarthmore College, Long Island Philosophical Society, and the Pacific Division Meeting of the American Philosophical Association.

## Bibliography

- Craig, Edward (2007) 'Genealogies and the State of Nature' *Bernard Williams* Alan Thomas (ed.) Cambridge University Press, pp. 181–200.
- \_\_\_\_\_ (1990) *Knowledge and the State of Nature: An Essay in Conceptual Synthesis* Oxford University Press.
- \_\_\_\_\_ (1986/1987) 'The Practical Explication of Knowledge' *Proceedings of the Aristotelian Society* 87: 211–226.
- Elgin, Catherine (1996) *Considered Judgment* Princeton University Press.
- Fricker, Miranda (2008) 'Scepticism and the Genealogy of Knowledge: Situating Epistemology in Time' *Philosophical Papers* 37(1): 27–50.
- \_\_\_\_\_ (1998) 'Rational Authority and Social Power: Towards a Truly Social Epistemology' *Proceedings of the Aristotelian Society* 98: 159–177.
- Gardiner (ms) 'Understanding and Emulation'.
- \_\_\_\_\_ (2015a) 'Normalcy and the Contents of Philosophical Judgements' *Inquiry* 58(7): 700–740.
- \_\_\_\_\_ (2015b) 'Teleologies and the Methodology of Epistemology' *Epistemic Evaluation: Purposeful Epistemology* John Greco and David Henderson (eds.) Oxford University Press, pp. 31–45.
- \_\_\_\_\_ (2011) *A Defence of Robust Virtue Epistemology* Master's Thesis at the University of Edinburgh.
- Greco, John (2012) 'A (Different) Virtue Epistemology' *Philosophy and Phenomenological Research*, 85(1): 1–26.
- Hannon, Michael (ms) 'What's the Point of Understanding?'
- \_\_\_\_\_ (2017) 'A Solution to Knowledge's Threshold Problem' *Philosophical Studies* 174: 607–629.
- \_\_\_\_\_ (2015a) 'The Importance of Knowledge Ascriptions' *Philosophy Compass* 10(12): 856–866.
- \_\_\_\_\_ (2015b) 'The Universal Core of Knowledge' *Synthese* 192(3): 769–786.
- \_\_\_\_\_ (2013) 'The Practical Origins of Epistemic Contextualism' *Erkenntnis* 78: 899–919.
- Henderson, David and John Greco (2015) (eds.) *Epistemic Evaluation: Purposeful Epistemology* Oxford University Press.
- Kappel, Klemens (2010) 'On Saying that Someone Knows: Themes from Craig' *Social Epistemology* Adrian Haddock, Alan Millar, and Duncan Pritchard (eds.) Oxford University Press, pp. 69–88.
- Kelp, Chris (2011) 'What's the Point of "Knowledge" Anyway?' *Episteme* 8: 53–66.
- Kornblith, Hilary (2011) 'Why Should We Care about the Concept of Knowledge?' *Episteme* 8(1): 38–52.
- \_\_\_\_\_ (2002) *Knowledge and its Place in Nature* Oxford University Press.
- Kripke, Saul (2011) *Philosophical Troubles: Collected Papers, Volume One* Oxford University Press.
- Kusch, Martin (2013) 'Naturalized Epistemology and the Genealogy of Knowledge' *Contemporary Perspectives on Early Modern Philosophy: Nature and Norms in Thought* Martin Lenz and Anik Waldow (eds.) Springer, pp. 87–97.
- \_\_\_\_\_ (2011) 'Knowledge and Certainties in the Epistemic State of Nature' *Episteme* 8: 6–23.



- \_\_\_\_\_ (2009) 'Testimony and the Value of Knowledge' *Epistemic Value* Adrian Haddock, Alan Millar, and Duncan Pritchard (eds.) Oxford University Press, pp. 60–94.
- MacBain, James (2004) 'Epistemic Analysis and the Possibility of Good Informants' *Principia*, 8(2): 193–211
- Malmgren, Anna-Sara (2011) 'Rationalism and the Content of Intuitive Judgments' *Mind* 120(478): 263–327.
- Nietzsche, Friedrich (1967) *On the Genealogy of Morals* and *Ecce Homo* trans. Walter Kaufmann, New York: Vintage Books.
- Pritchard, Duncan (2012) 'Anti-Luck Virtue Epistemology' *Journal of Philosophy* 109: 247–279.
- \_\_\_\_\_ (2011) 'The Genealogy of the Concept of Knowledge and Anti-Luck Virtue Epistemology' *Der Begriff des Wissens/The Concept of Knowledge* Stefan Tolksdorf (ed.) de Gruyter, pp. 159–178.
- Schaffer, Jonathan (2006) 'The Irrelevance of the Subject: Against Subject-Sensitive Invariantism' *Philosophical Studies* 127: 87–107.
- Shapin, Steven (1994) *A Social History of Truth: Civility and Science in Seventeenth-Century England* Chicago University Press.
- Sosa, Ernest (forthcoming) 'The Metaphysical Gettier Problem and the X-Phi Critique' *Explaining Knowledge: New Essays on the Gettier Problem* Rodrigo Borges, Claudio de Almeida, and Peter Klein (eds.) Oxford University Press.
- Strawson, Peter (1992) 'Reduction or Connection? Basic Concepts' *Analysis and Metaphysics: An Introduction to Philosophy* Oxford University Press.
- Williams, Bernard (2002) *Truth and Truthfulness: An Essay in Genealogy* Princeton University Press.
- Williamson, Timothy (2007) *The Philosophy of Philosophy* Blackwell.

## IV. Understanding and Emulation

**Abstract.** In *Knowledge and the State of Nature* Edward Craig hypothesises that knowledge attributions evolved to fulfil a function; specifically, they served to tag good informants. Craig draws on this hypothesis about the function of knowledge attributions to illuminate the contours of the concept of knowledge. I develop a related idea: I propose that understanding attributions also evolved to fulfil a function, namely serving to tag those worthy of intellectual emulation. I examine some competing hypotheses both for alternative functions understanding attributions might serve, and for what else may have served this function. Finally, I explore what this hypothesis suggests about understanding and understanding attributions.

### 1. Teleological Approaches

Teleological approaches to understanding a phenomenon proceed by reflection on the purpose or function of that phenomenon, what role it fulfils, or what it characteristically does. The idea behind teleological approaches is that exploring the function of a phenomenon illuminates the phenomenon itself; we can learn about its nature, value, and place within a wider web of related phenomena. Teleological approaches can be fruitfully applied to artefacts, concepts, words, structures, systems, and other phenomena. One category of phenomena we can explore are epistemic evaluations.<sup>1</sup>

---

<sup>1</sup> Edward Craig's seminal project, which I outline below, concerns primarily epistemic evaluations, and draws conclusions about the concept of knowledge (and knowledge itself) from the function of knowledge attributions. Michael Hannon's essay 'What's the Point of Understanding' emphasises that Craig's teleological account concerns primarily

Epistemic evaluations are ubiquitous. We frequently characterise ourselves and others as irrational, rational, knowers, trustworthy, educated, insightful, understanders, ignorant, gullible, judicious, and so on. Epistemic evaluations can be explicitly expressed as epistemic attributions such as ‘Joe knows the way’ or ‘Fred understands quantum mechanics’. They are frequently implicit, such as when we advise ‘follow Joe’ or ‘ask Fred’. Epistemic evaluations can also manifest themselves in our unspoken behaviour, such as when we follow Joe, or learn from Fred.

Evaluating ourselves and others epistemically—and being the subject of epistemic evaluation by others—plays an important role in our cognitive and social lives. By exploring the roles these evaluations play, we can illuminate both the epistemic phenomena invoked by the evaluations, such as knowledge, and our broader epistemic circumstances, such as the contours of our fallibility. By examining when and why we offer and withhold certainty attributions, for example, we acquire insight into how much guarantee we want from our evidence and how much guarantee our evidence typically provides.

A teleological approach may proceed using the following schema,<sup>2</sup>

---

epistemic evaluations, rather than the epistemic phenomena itself. Note that, in addition to exploring the function of epistemic evaluations, one can also discuss the functional role of the epistemic phenomena itself.

<sup>2</sup> This schema appears in Gardiner (2015) and is adapted from Kappel (2010: 72–3). As Kappel clarifies, by ‘need’ he means something that we are better off if we meet, rather than the stricter sense of something without which we cannot function (Kappel 2010: 79–80).

- i. Given a set of facts about our physical constitution, cognitive powers and environment, and a set of aims and interests, we have a particular need;
- ii. The target phenomenon is what actually and currently meets that need.

As I describe in Gardiner (2015), this schema can be applied in several ways. Firstly, one might posit a synchronic practical explication, which identifies a *current* need and posits a current role for the target phenomenon. Klemens Kappel, for example, posits that knowledge attributions currently function to mark when to cease inquiry.<sup>3</sup> The schema can alternatively be applied as a historical genealogy, which identifies a historical need and posits that the phenomena in fact fulfilled that need in the past (and perhaps still does so). Martin Kusch, for example, explores epistemic needs amongst seventeenth-century scientists.<sup>4</sup> Thirdly, the schema can be applied as a hypothetical genealogy. A hypothetical genealogy describes a fictionalised ‘quasi-historical’ period and hypothesises about needs, and the phenomena that satisfy those needs, in this hypothetical period.<sup>5</sup> Nietzsche’s state of nature genealogy is a familiar example of a hypothetical genealogy.<sup>6</sup> In *Knowledge and the State of Nature*, Edward Craig proposes an influential hypothetical genealogy for knowledge attributions.

Craig posits that knowledge attributions fulfilled the need of tagging good informants. In the hypothetical state of nature, humans needed to attain true beliefs and avoid false beliefs, and thus required a way to keep track of which people are good sources of information about a

---

<sup>3</sup> Kappel (2010).

<sup>4</sup> See Kusch (2009) and Shapin (1994).

<sup>5</sup> Miranda Fricker employs the term ‘semi-fictional time’ when describing hypothetical genealogies.

<sup>6</sup> Nietzsche (1967).

subject matter. Craig hypothesises that knowledge attributions met that need, and so the concept of knowledge developed from the idea of a good informant. Craig writes, ‘The concept of the informant will then leave its mark, so to speak, on the concept of knowledge’.<sup>7</sup> Craig employs this hypothesis to plume the nature and value of knowledge.

Craig’s methodology has created substantial debate. Several theorists have endorsed Craig’s claim about the principal function of knowledge attributions, but harnessed this hypothesis to support competing claims about the nature and value of knowledge.<sup>8</sup> Others have endorsed Craig’s overall approach, but advanced a competing principal function for knowledge attributions.<sup>9</sup> Others have raised methodological worries about aspects of Craig’s approach.<sup>10</sup>

Whilst Craig’s research has invigorated interest in teleological approaches in epistemology, most of this interest has focused on teleological approaches to knowledge and knowledge attributions.<sup>11</sup> In this essay I explore a function for attributions of understanding. I suggest that attributions of understanding tag those worthy of intellectual emulation in that domain.

This essay advances two aims. Firstly, extending the teleological approach from knowledge attributions to other epistemic evaluations can facilitate understanding of teleological

---

<sup>7</sup> Craig (1990: 95).

<sup>8</sup> Pritchard (2011; 2012), Grimm (2015), and Greco (2010; 2015).

<sup>9</sup> Kelp (2011) and Kappel (2010).

<sup>10</sup> Kornblith (2011) and Gerken (forthcoming).

<sup>11</sup> For exceptions, see Smithies (2014), Joshua Habgood-Coote (2017), Wilkenfield et al. (2016), Hannon (ms), Emily Sullivan (ms), and Dogramaci (2012).

approaches. I thus aim to shed light on this much-misunderstood approach. Secondly, I aim to illuminate the nature and value of understanding.

In section two I describe some widely held features of understanding. In section three I suggest that these features accord well with the proposal that understanding attributions function to tag those worthy of intellectual emulation, and I outline implications this proposal suggests about understanding. (Note the distinction between considerations in favour of the claim and upshots or consequences that follow from the claim is somewhat artificial. The considerations cohere in a reflective equilibrium. But those claims that are more bedrock or independently supported appear in section two.) In section four and five I consider objections concerning whether I have identified the wrong characteristic function of understanding attributions.

## **2. Characteristics of Understanding**

Understanding is often contrasted with propositional knowledge. One distinction concerns the object of propositional knowledge and understanding. Propositional knowledge can be of simple matters such as relatively isolated atomistic facts. One can know the time, which direction is north, how many cups are on the table, and whether your neighbour's name is Bill or Tim. But objects of understanding tend to be more complex; they typically comprise an integrated collection of facts or parts that cohere or 'hang together'. We can understand a theory, language, explanation, system, machine, or structure.<sup>12</sup> Plausibly, to be an object of understanding there must be sufficient complexity or depth; there must be structural

---

<sup>12</sup> Another important species of understanding is semantic understanding, where we understand a concept, simple idea, sentence, or word.

relations to grasp, a unifying pattern to perceive, or underlying explanations or reasons to appreciate. One cannot understand a mere list of telephone numbers, for example.

A second distinction concerns acquisition. Propositional knowledge can characteristically be transmitted via testimony, and the hearer is relatively passive in acquiring knowledge this way. The hearer's roles include comprehension of the claim, possession of requisite virtues, and lacking knowledge-threatening defeaters and epistemic vices. The testifier's assertion contributes a significant part of someone's coming to know, and the hearer need not 'supply' a great deal. Understanding, by contrast, is not characteristically transmitted by testimony.<sup>13</sup> Testimony can certainly aid understanding. It can supply parts of an overall picture, convey true claims, remove blockages to understanding such as misconceptions and false beliefs, and help the hearer forge or appreciate connections amongst her ideas.<sup>14</sup> But testimony (at least typically) cannot perform a central component of generating new understanding. Recall the aphorism, close to the heart of many teachers, 'I can explain it for you, but I can't understand it for you.' To acquire understanding the person must characteristically be relatively active. The hearer must see how things hang together for herself, forge

---

<sup>13</sup> Some epistemologists argue that understanding can be transmitted by testimony. It is instructive that epistemologists take the burden of proof to reside with those who hold that understanding can be transmitted by testimony, since this burden of proof gestures at the kind of 'active' role required by the acquirer.

<sup>14</sup> See Gordon (forthcoming) for an exploration of the role of questions in acquiring understanding. Whilst exploring the role of an interlocutor—specifically a therapist—in facilitating understanding, Gordon's research underscores that the potential understander plays a relative active role in coming to understand. It is they who must see for themselves how things are.

connections in her mind and perceive coherence-making or unifying relations amongst claims.<sup>15</sup>

This distinction between knowledge and understanding is not limited to acquisition through testimony. We can observe a fact, such as that Elton John is wearing blue, and thereby come to know. The knowledge acquisition process between acquiring requisite evidence and possessing knowledge is typically relatively direct or passive. Not all propositional knowledge acquisition is so direct or passive, of course—some knowledge is challenging to acquire, even once you possess the requisite evidence—but it characteristically can be. Understanding, by contrast, typically requires more activity. The process between possessing requisite evidence and possessing corresponding understanding can be more challenging and is characteristically not immediate.

Thirdly, understanding attributions are frequently connected to appreciation, emotion, and normative domains. We talk naturally of understanding moral issues, aesthetic phenomena, and people. The notion of understanding, as contrasted with knowledge, is particularly appropriate when a person's sensibilities or appreciations are involved.<sup>16</sup> We might say 'Bob is understanding' to convey something about the emotional and moral character of Bob. (Suppose, say, that you promised Bob you would mind his cat, and now you must renege on that promise.) Knowledge attributions, by contrast, seem related to whether the claim is true and whether the person is sufficiently reliable and possesses adequate evidence. Knowledge

---

<sup>15</sup> See Gardiner (2012) for more about the kinds of coherence-making relations that are constitutive of understanding, and why they do not reduce to further propositions.

<sup>16</sup> See, for example, Grimm (forthcoming).



attributions do not seem to share the connection to appreciation, emotion, and moral domains.

Fourthly, a person might gain or lose knowledge relatively swiftly. If I tell you a fact, you characteristically promptly come to know it. A defeater, likewise, can instantly forfeit knowledge. A person can forget, and so lose knowledge. I might know the correct depth for sowing radish seeds for a few hours after reading the instructions on the packet, but no longer know merely days or weeks later, once I have confused radish seed depths with beetroots and carrots. Indeed, I might know whilst looking at the seed packet, and not know ten minutes later, once I am in my garden with my rake. Contextualism posits that knowledge attributions are unstable, and one's status as a knower can depend on all sorts of facts, such as stakes or salient alternatives. Whether or not contextualism is correct, the fact it is plausible points to the relative instability of knowledge. If a person understands, by contrast, this tends to be a relatively stable feature. Understanding characteristically takes longer to develop—perhaps years of study—and cannot typically be lost by forgetting a few facts or acquiring defeaters. Reflecting on language acuity illustrates this difference. If I understand French today, you can usually rely on my understanding French in six months. If I know the French words for household furniture today, this does not underwrite the same assurance. Similarly, if I know the historical sequence of Post-impressionist painters, I may well forget by next year. But if I understand the historical sequence of Post-impressionist painters—and thereby can see how each developed from the preceding ideas—I am much more likely to retain this understanding for years hence.

In some unusual or marginal cases we might lose understanding suddenly, by undermining some foundational facts or perhaps a paradigm shift makes us realise things are not as we had believed. Perhaps learning that your parent's seemingly happy marriage was a sham or that your partner has a secret career, for example, could undermine understanding in an instant. But this is not the usual case. Perhaps too when we initially acquire understanding, and so learn about a domain, our understanding can be unstable or fleeting. We might understand trigonometry whilst the teacher is explaining it to us, for example, but not when we attempt to work through the problem cases by ourselves. But this is not an example of *understanding* being fleeting; instead it is illustrative of the process of coming to understand.

When a person understands they can characteristically apply their insights, comprehension, and knowledge to novel cases.<sup>17</sup> They are not limited to existing knowledge. Those who understand are well-positioned to explore new areas, perceive gaps in one's picture of the world, and detect false claims and misinformation.<sup>18</sup>

Understanding is related to practical matters; it allows us to do things and affords us power. When we understand something, such as a machine, we can manipulate and control it. Riggs (2003) notes that a mechanic who understands a machine knows how much beyond its specifications it can be pushed. Someone who merely knows how to use it may not. These are the practical, useful, applied upshots of understanding. With understanding we gain insight into questions about counterfactual outcomes. A historian who understands a battle, for example, can answer questions about what would have happened if conditions had been

---

<sup>17</sup> Hills (2016).

<sup>18</sup> Kvanvig (2003).

different and can predict what would happen in similar circumstances. Mere knowledge is less apt to provide counterfactual insight and predictive capacities. These kinds of insights underwrite power. Through understanding we can plan the next battle effectively, choose our battles, and manipulate circumstances to generate a desired outcome. A person who understands a friend might realise how best to help them through a difficult time. These practical abilities might not be constitutive of, or even necessary conditions for, understanding, but they at least seem to be characteristic consequences of possessing understanding.

Finally (and more controversially) plausibly understanding has a distinctive value. It is intrinsically valuable and is an accomplishment. Consider an item of easy or trivial knowledge. We might question whether the knowledge has much value. Perhaps the person has not pursued anything of value by, for example, counting grains, or beads, or blades of grass. But for understanding this worry seems less apposite. If a person understands a topic, the topic must have a certain degree of interest—it must exhibit structure, unity, or patterns; there must be causal and explanatory connections to grasp; the pieces must fit together somehow. And the person must perceive how these pieces fit, grasp explanatory or causal relations, or appreciate some unifying or coherence-making aspects. These features are required for understanding and are valuable.<sup>19</sup> Thus the worry about trivial or non-valuable instances of the kind is less pressing for understanding than for knowledge.

---

<sup>19</sup> For an argument that understanding, unlike knowledge, is intrinsically valuable, see Brewer (2009).

These features of understanding, I suggest, accord well with the proposal that understanding attributions tag those worthy of intellectual emulation.

### 3. Understanding and Emulation

Craig notes that humans have a need to identify good sources of information. He writes,

Human beings need true beliefs about their environment, beliefs that can serve to guide their actions to a successful outcome. That being so, they need sources of information that will lead them to believe truths. They have ‘onboard’ sources, eyes and ears, powers of reasoning, which give them a primary stock of beliefs. It will be highly advantageous to them if they can also tap the primary stocks of their fellows [...] that is to say, if they act as informants for each other. On any issue, some informants will be better than others, more likely to supply a true belief [...] So any community may be presumed to have an interest in evaluating sources of information; and in connection with that interest certain concepts will be in use. The hypothesis I wish to try out is that the concept of knowledge is one of them. To put it briefly and roughly, the concept of knowledge is used to flag approved sources of information.<sup>20</sup>

Craig posits that in the hypothetical epistemic state of nature—where humans are language using and minimally cooperative but have not yet developed sophisticated society and lack basic epistemic concepts—individuals have a need to tag those who will reliably tell them correct information. Craig hypothesises that knowledge attributions satisfy that need. In this essay I leave aside whether Craig’s view of the function of knowledge attributions is plausible, correct, or helpful. Instead I develop a related idea.

Acquisition of true beliefs and avoidance of false beliefs are far from our only epistemic needs. We also have a need to improve ourselves as inquirers. We want to improve the way we think, evaluate, and explore. We would benefit from being more creative hypothesis-

---

<sup>20</sup> Craig (1990: 11).

formers, more effective problem-solvers, and better retainers of information. We want to be better at navigating complex ideas, comprehending challenging materials, and adjudicating between conflicting claims. As good inquirers we can attain true beliefs for ourselves without relying on others and we can apply information in novel ways. We will be well-positioned to detect misleading evidence, misinformation, and deceit. Improving oneself as an inquirer can lead to social power. Thus a natural expansion of Craig's approach identifies a second need: that of improving ourselves cognitively.

This raises the question: how can we improve ourselves cognitively? One effective way would be to find those with desirable and worthwhile epistemic traits and practices, and emulate them. These kinds of traits include the ability to apply existing insights to novel circumstances, the ability to perceive discrepancies, unexplored areas, unifying themes, how things fit together, and so on. A person worthy of emulation will be able to make predictions and manipulate and control phenomena. They will have a grip on the facts and see how those facts cohere. They will grasp causal and explanatory relations amongst facts and phenomena, and so expand their knowledge. They will have a good epistemic character and be in excellent cognitive contact with the world.

We can improve ourselves epistemically in non-social ways too, such as thinking by ourselves, but the Craigian methodology—by examining the function of epistemic evaluations—emphasises the social aspects of our epistemic lives. Note too that improving ourselves epistemically in non-social ways is probably relatively difficult—even more difficult than acquiring knowledge without the aid of others.

The desirable epistemic traits and practices listed above cannot be wholly acquired by testimony and—relatedly—are not simply a matter of grasping atomistic facts. Thus it is not sufficient to find a knower and ask them. Instead these features, skills, and practices are acquired over time by emulating those who already have the abilities, practices, and right kind of cognitive contact with the world. We tag those people, according to this hypothetical genealogy, by labelling them as understanders.

‘Knowledge’ is a more common term than understanding. This accords well with the proposal. We rely on all sorts of people for information, but we are more selective about those we emulate. Plausibly we frequently desire information from others, but we desire to emulate less often.

In some instances in order to emulate we must be explicitly taught; a person must explain or describe something. In other cases one can emulate without explicit teaching. We might observe a demonstration, for example.

I noted in section one three kinds of teleological approach: a practical explication, a historical genealogy, and a hypothetical genealogy. Craig’s hypothetical genealogy posits a process of objectivisation.<sup>21</sup> Craig posits that over time knowledge attributions evolve from conveying ‘S is a good informant for me about whether p’ to ‘S is a good informant for anyone about whether p’. (Note that ‘for anyone’ is understood as something like ‘for anyone who requires an informant’, that is, for inquirers who do not already possess the

---

<sup>21</sup> Craig (1990: 82) argues that the process of concept objectivisation is a general process of concept development, and not distinctive to his account.

relevant knowledge.) We can emulate this proposal by positing that understanding attributions evolve from ‘S is worthy of intellectual emulation in this domain for me’ to ‘S is worthy of intellectual emulation in this domain for anyone’.

This process of objectivisation explains why a person attributes understanding to someone that she should not emulate in a domain, such as someone whose understanding is less advanced than her own; a person who understands is worthy of emulation in a domain by an appropriate person (that is, by someone who does not yet understand, but has an intellectual interest in the domain). This is isomorphic to Craig’s contention that knowledge attributions tag good informants for an appropriate person (such as someone who does not yet know the relevant fact, but has a desire to become informed), or the claim that knowledge attributions mark when to cease inquiry for those for whom the question is relevant, such as those inquiring about whether p.<sup>22</sup>

If instead we posit a practical explication for understanding attributions, and so posit a contemporary principal (or characteristic) function, the process of objectivisation is not available. A practical explication posits the actual, current function of the phenomena. The proposal that understanding attributions tag those worthy of intellectual emulation thus faces the following challenge. People attribute understanding to those they do not deem worthy of emulation for themselves (because the attributor is more expert in the domain than the recipient of the attribution, for example).<sup>23</sup> But note that we can still deem someone worthy

---

<sup>22</sup> Kelp (2011) and Kappel (2010).

<sup>23</sup> Note that almost any function-first account that posits that epistemic attributions tag people with certain qualities will face a version of this challenge. If ‘attributions of understanding are primarily used to identify individuals who can provide us with

of emulation even if we ourselves will not emulate that person. A martial arts expert can recommend a coach to a beginner. A professor can recognise that her less advanced students should emulate her more advanced student, even whilst thinking that others, such as her colleagues, should not emulate the advanced student. And we can pick out individuals worthy of emulation, whilst recognising that others might be even more worthy of emulation.

Practical explications, since they cannot appeal to the passage of time to explain apparent divergences between the posited function and current role of the phenomenon, appear to face a methodological dilemma. Either the posited function perfectly matches the current use, or it diverges. If it diverges, the view seems to suffer from a counterexample. But if it perfectly matches the current use, the proposal is likely to resemble an analytic account—specifying necessary and sufficient conditions for the phenomena, such as necessary and sufficient conditions for apposite knowledge attributions—rather than a novel methodology. One response to this dilemma emphasises that the method is not attempting an analysis, nor a counterexample-free proposal. Instead the aim is to suggest an illuminating, central, or important function of the phenomenon. Counterexamples matter only if they suggest the posited function is not important, illuminating, or central.

The hypothesis that understanding attributions serve to identify those worthy of intellectual emulation suggests several things about the nature of understanding and the contours of the

---

explanations' as Michael Hannon posits, a worry arises about understanding attributions from those who already possess good explanations in that domain. If knowledge tags good informants, as Craig posits, the worry concerns attributions from those who already possess that information, and so on.



concept of understanding.

Firstly, we can expect that understanding attributions are an honourific or compliment, and are even more expressive of praise than saying someone knows. This is because to be worthy of intellectual emulation is a relatively high bar.

Secondly, the hypothesis suggests a deep connection between understanding and qualities of character such as epistemic and character virtues. To call someone a knower might indicate merely that the person has access to a good source of information. The weather forecaster's flatmate might reliably know the weather (and be a good informant about the weather), for example, simply by residing with the weather forecaster. His status as a (mere) knower does not indicate much about his deep and stable epistemic character; after he moves out he might well be ignorant about the weather. But to say someone understands—and so is intellectual emulation-worthy—suggests their epistemic *character* is in some way good. Relatedly, understanding attributions also suggest a degree of intellectual independence. Understanding does not typically depend on someone else for its being sustained. (Although, of course, we might well depend on other people when we pursue or expand our understanding.)

Thirdly, and more controversially, the hypothesis suggests, or at least accords with, a degree of pluralism about understanding. There is more than one way to be an understander, and the different ways might not be wholly mutually consistent. When deciding who to emulate, we have options, and a variety of choices might be reasonable. In the domain of physical health, for instance, we can legitimately emulate the approach of a dietician, a cross fit

trainer, a medical doctor, a massage therapist, a yoga instructor, and so on. Each person can teach us something about physical health and we will gain some insight from each. We might emulate more than one person, and so gain a broader understanding or we might commit to apprenticing with one specialist and so gain a deeper but narrower understanding.<sup>24</sup> If each is intellectually emulation-worthy then, according to this suggestion, we can attribute understanding to each. (Other approaches, such as that of the phrenologist, are not deemed intellectually emulation-worthy; these exhibit insufficient grasp of reality.) And note that we might emulate and attribute understanding to each even if the approaches of the yoga teacher and the cross fit trainer are not mutually consistent. This accords with the idea that there is more than one approach to understanding a domain.

The fact that two people can understand a domain, yet their views be not wholly mutually consistent might be because each person's understanding is imperfect. Perhaps if two people perfectly understand physical health, they will wholly agree. But imperfect understanding is nonetheless understanding and, owing to its ubiquity, an important epistemic kind.

We can compare this degree of relative pluralism to the relative non-pluralism about knowledge. If two people have conflicting beliefs about whether *p*, at least one of them does not know whether *p*; plausibly at least one is not a good informant about whether *p*. They might both be good informants about other things, such as ways to think about whether *p*, or how to inquire about whether *p*. But these questions are distinct from whether *p*.

---

<sup>24</sup> Given that there might be tensions amongst their approaches, we cannot simply passively receive the insights of various people. This approach risks incoherence and confusion. Instead we must actively make sense of what people can teach us. We must create coherence for ourselves. Thanks to Catherine Elgin for helpful discussions on these issues.

To put this point slightly differently, to be intellectually emulation-worthy—and to understand—one's beliefs need not be wholly true. There is still much to gain by emulating a person with false beliefs if they understand.<sup>25</sup> If a person has false beliefs about a subject matter, by contrast, they are not a good informant on that subject matter and not a knower.

In 'Teleologies and the Methodology of Epistemology' I argue that the teleological approach to understanding a concept can provide an indication of how coarse- or fine-grained we should expect the concept to be.<sup>26</sup> I illustrate with the following example. The concept 'adult' serves a legal need (in voting, criminal law, and other domains.). It requires a relatively precise definition to satisfy that need—and sure enough it has one. By contrast the concept 'adolescent' fulfills needs in marketing, media, social science, and stereotyping. In order to fulfill these roles the concept does not need such precise contours; its boundaries can be more vague. Thus we have reason to think it is misguided to seek a precise analysis of the concept 'adolescent' and we ought only aim for more broad-brush analyses.

We can now apply this insight to the proposal that the primary function of understanding attributions is to tag those who are worthy of intellectual emulation. I propose that intellectual emulation worthiness—at least compared with notions such as who is a good informant or when to cease inquiry—does not have sharp and precise boundaries. People are emulation worthy to different degrees, in different domains, and in different ways. It will

---

<sup>25</sup> See Elgin (2006; 2009).

<sup>26</sup> The teleological approach can also indicate how coarse- or fine-grained we should expect the phenomena itself to be. Perhaps, for example, there are instances in which it is indeterminate whether someone is an adolescent.

depend on who else we might emulate, and what shape other people are in general. Someone with a basic grasp of sailing might be emulation worthy, and accordingly understand sailing, in a community of people with little grasp on the mechanics of sailing, for example, but in a community of sailing experts, they will not qualify as emulation-worthy or as possessing understanding.

Thus we should expect understanding attributions to exhibit a relatively large degree of flexibility. Perhaps we should expect an analysis of understanding—or even agreement over identifying the target phenomena—to be more elusive.

Finally, the proposal illuminates or accords with the claim that understanding is a stable feature of a person. Typically we cannot or should not emulate someone whose epistemically valuable traits are unstable. If a person has a good grasp of a subject matter one day, but not on others, it seems that we should not emulate them.<sup>27</sup> But if someone knows a fact (the weather, say) on some days but not on other days they can still be a reliable informant: we can acquire the information on days our friend has spoken with his weather forecasting flatmate, for example. This fits with the idea that understanding concerns a person's

---

<sup>27</sup> There might be examples where arguably someone is worthy of emulation even though their epistemically valuable traits are unstable. Perhaps an eccentric scholar has good days and very bad days, for example, but their good days are exceptionally good. Perhaps they are usually in a haze or confusion about their field but their insights, when they occur, can advance their field. But note firstly that these cases will be marginal, and not characteristic cases. (The function of a phenomenon is related to its characteristic or normal function, not marginal or unusual cases.) And secondly, it is unclear whether these people are worthy of epistemic emulation, even if exceptionally intellectually gifted, precisely because of the instability of their epistemically valuable features, such as their grasp on reality.

character, practices, and stable cognitive contact with the world; it is not merely a matter of a relationship to a proposition or set of propositions, as knowledge might be.<sup>28</sup>

#### **4. Other Potentially Emulation-Worthy Kinds**

In this section I examine some potential criticisms of the proposal; I here focus on objections to the specific thesis about the function of understanding attributions, and bracket general objections to the overall methodological orientation of teleological accounts in epistemology. There are two central approaches to challenging the thesis. The first argues that attributions of an alternative epistemic kind pick out those we should emulate intellectually. This approach concedes that there is a need to improve one's epistemic character and, further, that this generates a need to tag those worthy of intellectual emulation. But it argues that understanding attributions do not satisfy this need; instead something else does. A second family of objections argues that understanding attributions fulfil an alternative function. This approach agrees with the overall teleological approach, but posits some other role for understanding attributions.

What other epistemic evaluations might compete for the function of tagging those worthy of intellectual emulation? Three plausible candidates are attributions of knowledge, wisdom, and knowledge how.

---

<sup>28</sup> Note that if a person is correct about the weather some days but believes they are correct about the weather every day, he is not a good informant about the domain. But this accords with the intuitive judgement that he is not a knower about the domain either. His belief-forming is not safe or sensitive.

Knowers might plausibly tag those worthy of intellectual emulation because knowledge is clearly valuable. Knowers believe truly and have a competent relationship to the truth of their belief. And knowledge brings power, such as the power to avoid dangers and locate valuable items. And perhaps when we come to believe on the basis of testimony we, in some small way, copy or mimic the person: they believe *p* and because of this we do too. But knowledge attributions do not seem appropriate for tagging those worthy of intellectual emulation. For one thing, much knowledge is extremely easy to gain. Characteristically one only need perceive an item to learn its colour, shape, size, location, and so on. Plausibly emulation would be reserved for something more challenging. Much knowledge seems trivial or unimportant. It is a stretch to call someone intellectually emulation-worthy merely because they are reliable at counting blades of grass, for example, and so the category of knower seems too broad to pick out those worthy of intellectual emulation.

Picking out those from who we can acquire knowledge seems distant from the central conception of emulation. A person might know simply by hearing another person's testimony. Recall the weather forecaster's flatmate. He knows but is not himself worthy of emulation, he just happens to live with a reliable informant. He is a good person to ask, but coming to know from testimony does not seem like a genuine instance of emulation; it is certainly not the most helpful, characteristic, or central illustration of intellectual emulation. Coming to know can be a relatively passive process, whether we learn from testimony, observation, or some other method. Thus social acquisition of knowledge is removed from the notion of intellectual emulation.

Wisdom attributions plausibly lie at another extreme. Whereas to acquire knowledge interpersonally we characteristically need less than emulation, to gain wisdom interpersonally requires more than emulation. Initially wisdom attributions seem like a good candidate for having a deep connection to intellectual emulation. Wisdom is certainly well worth pursuing, and one natural way to pursue wisdom is to identify the wise and be like them. But plausibly pursuing wisdom is so active that one cannot proceed (at least centrally, principally, or characteristically) by emulation. To gain wisdom, one must find the way oneself, and not largely follow another. One might rely on other people to ‘get you going’. And one might emulate or rely on another person for those aspects of wisdom that consist in understanding or knowledge, but one cannot follow another person to wisdom itself. So although we aim to be like the wise, we cannot accomplish this by copying or following. Wisdom must be earned through one’s own insights and experiences.

This discussion suggests a scale of how much activity and independence is associated with each epistemic kind. Instances of knowledge can be relatively passive. For large, non-atypical bodies of knowledge one must be receptive, be free from defeaters, comprehend and believe the claims, and possess minimal requisite virtues (such as not being gullible). But one need not be particularly active. This suggests a testifier-hearer model for knowledge attained interpersonally. To understand, one must be more active. A person can tell you the relevant facts, explanations, and how things hang together. But one must see for oneself those connections; the claims must cohere properly in one’s own way of thinking. A student coming to understand cannot simply be passive. You cannot, I think, understand physics simply by spending time in a physics laboratory. You would have to apply yourself, listen carefully, and think through problems. I think you can learn some items of physics

knowledge, by contrast, by simply spending time in the physics laboratory. This suggests a student-teacher model for understanding attained interpersonally. Once a person acquires understanding they thereby gain some intellectual independence; understanding empowers one to think for oneself about novel matters.

To become wise, one must be more active still: one cannot merely follow other people to the degree one can to acquire understanding. I think it is instructive that it is harder to develop a model for acquiring wisdom interpersonally; this reflects the *independent* character of the search for wisdom. Perhaps it is illuminating to think of guru-advisee, counsellor-client, mentor-mentee, or parent-adult child relationships; it is key to socially-acquired wisdom that the mentee reflects on discussions with independence of thought and lives life independently between meetings. She must reflect on her own experiences, and critically evaluate the perspective of others. She cannot acquire wisdom merely by listening receptively to the wise person and understanding their insights.

Knowledge how attributions are also a potential candidate for picking out those worthy of intellectual emulation. If a person knows how to do something, one can benefit by being like her in that domain. I have four comments in response. Firstly, one might hold that knowledge how is very similar—conceptually, linguistically, psychologically, or metaphysically—to propositional knowledge.<sup>29</sup> One upshot of this similarity is that propositional knowledge attributions and knowledge-how attributions might serve the same

---

<sup>29</sup> Stanley (2011) and Stanley and Williamson (2001).



or similar functions.<sup>30</sup> But, as I suggest above, knowledge attributions are not a good candidate for tagging those worthy of intellectual emulation. Secondly, know how seems very practical. Whilst we can learn how to do something from copying someone who knows how, intellectual emulation is plausibly reserved for something more epistemic. A person might know how to find the length of the hypotenuse by plugging numbers into his calculator in the precise way he was taught. But he may have no idea why this procedure calculates the length of the hypotenuse. It is better to emulate someone who understands trigonometry and *understands* how to calculate lengths.

Thirdly, someone who merely knows how, but lacks understanding, might be difficult to emulate. Successful emulation sometimes requires explanation, demonstration, emphasis on central or important aspects, and feedback on performance. In some cases, only with this guidance can a person emulate. Someone who merely knows how, but does not understand (how), is less well positioned to provide this kind of guidance.

To render this vivid, consider gymnastics. I learnt how to perform a cartwheel as a young child. I was about five years old when I learnt. I can barely remember not being able to cartwheel. I can perform dozens in a row and at various different speeds. I know how to do a cartwheel. I have not reflected on this skill, and would have trouble breaking it down into steps or explaining the manoeuvre. Now consider a different gymnastics skill: the handstand. I am not very good at handstands. My handstands are mediocre. They last a couple of seconds. But as an adult I have attended handstand workshops, read articles on handstands,

---

<sup>30</sup> Craig (1990) was committed to this upshot, and so argued that knowledge how attributions also tag good informants.

watched videos explaining the mechanisms of handstands, and spoken to experts. I understand handstands. Suppose cartwheels and handstands are equally challenging for a person. If you want to emulate me, and so learn a skill, you are more apt to succeed at the activity I understand, rather than the one where I have mere knowledge-how. I can help you more effectively with handstands than I can with cartwheels. And this is true even though my mere ability at the skill is higher with cartwheels.

In order to provide a vivid illustration, the chosen example is very practical. I contend the same is true for intellectual endeavours. An unreflective speaker may know how to conjugate verbs in her native language. A second speaker—perhaps a talented non-native speaker—might understand verb conjugation. If you want to acquire this insight, and emulate a speaker, you will likely succeed with the person who possesses understanding rather than mere knowledge how. And this is true even if the knowledge how is somehow better, such as if the native speaker knows how to speak more articulately in the language.

Finally, for those unconvinced by the proceeding points and who instead maintain that knowledge how attributions enjoy a tight connection to emulation-worthiness, note that the central claim I explore in this paper concerns *intellectual* emulation, rather than emulation in general.<sup>31</sup> It is plausible that intellectual emulation concerns understanding even if emulation

---

<sup>31</sup> That said, I am sympathetic to the idea that emulation in general is connected to understanding and understanding how attributions. If we want to acquire an ability, we should seek out those who understand how. Those who merely know how are, by contrast, useful as people who can get things done. Recall the contrast between knowing how to perform a cartwheel and understanding how to perform a handstand. In section two I suggested that to be an object of understanding, the thing must have a degree of complexity or depth. Propositional knowledge, by contrast, could be of simple things such as propositions. Plausibly an isomorphic distinction obtains for knowledge how and

in general concerns knowledge how. Relatedly, perhaps ‘knowledge how’ better captures emulation-worthiness for beings who do not have reflective understanding or critical reasoning. Perhaps animals and small children should, in order to develop, emulate those who know how, whilst critically reflective adult humans should emulate those who understand.

Note, of course, there are ways of acquiring knowledge, understanding, and wisdom that are independent from other people. We might gain knowledge from perception, understanding from experiments, and wisdom from adversity. But focusing on epistemic evaluations, as this methodology does, serves to direct attention to distinctly interpersonal ways to pursue these kinds.

## 5. Other Potential Functions

As noted at the beginning of section four, a second family of objections argues that understanding attributions fulfil an alternative function. One potential competing function, advanced by Michael Hannon, is that understanding attributions serve to tag those who can provide good explanations.<sup>32</sup> One weakness of this account is that—as Hannon acknowledges—it only applies to attributions of understanding why. But it would be better to posit a function that unifies various species of understanding attributions, such as attribution of understanding why, understanding topics (objectual understanding), understanding how, and attributions of being an understanding person. An account should

---

understanding how: we can know how to perform very simple, mundane activities, such as lifting a pencil. But we can only understand how to do something if it is relatively complex.

<sup>32</sup> Hannon (m/s).

at least explain how these species of understanding attribution relate to each other. Secondly, offering explanations requires linguistic or communicative faculties that plausibly are not characteristically required for understanding.

A second competing function for understanding attributions, advanced by Emily Sullivan, is that of tagging experts.<sup>33</sup> This proposal has at least two weaknesses. Firstly, expertise seems like a rarity in a community, but understanding is common. Secondly, many topics seem like characteristic objects of understanding, but are not the kind of subject matter we attribute expertise about. I might understand my friend, understand grief, or (in virtue of seeing the café patron jog the table) understand why the coffee spilt, but it seems odd to attribute expertise in these matters.

Finally, even if understanding attributions can, in some cases, function to tag experts or those who can provide good explanations, there is plausibly nonetheless something distinctive and privileged about the function of tagging those worthy of intellectual emulation. Offering good explanations or being an expert in a domain are ways of being emulation-worthy—perhaps ways which correspond to some subsets of understanding attributions—but tagging emulation-worthiness is privileged as a unifying and general characteristic of understanding attributions.

### **Acknowledgements**

Many thanks to David Black, Rodrigo Borges, Catherine Elgin, Will Fleisher, Michael Hannon, Jonathan Schaffer, Ernest Sosa, and Beth Valentine for comments on an earlier draft of this essay. Thanks also to audiences at the University of Houston, the Pontifical

---

<sup>33</sup> Sullivan (m/s).

Catholic University of Rio Grande do Sul, Underwood International College at Yonsei University, and Rutgers University.

## Bibliography

- Brewer, Talbot (2009) *The Retrieval of Ethics* Oxford University Press.
- Craig, Edward (2007) 'Genealogies and the State of Nature' *Bernard Williams* Alan Thomas (ed.) Cambridge University Press, pp. 181–200.
- \_\_\_\_\_ (1990) *Knowledge and the State of Nature: An Essay in Conceptual Synthesis* Oxford University Press.
- \_\_\_\_\_ (1986/1987) 'The Practical Explication of Knowledge' *Proceedings of the Aristotelian Society* 87: 211–226.
- Elgin, Catherine (2009) 'Is Understanding Factive?' *Epistemic Value* Adrian Haddock, Alan Millar and Duncan Pritchard (eds.) Oxford University Press, pp. 322–30.
- \_\_\_\_\_ (2006) 'From Knowledge to Understanding' *Epistemology Futures* Stephen Hetherington (ed.) Oxford: Clarendon Press, pp. 199–215.
- Fricker, Miranda (2008) 'Scepticism and the Genealogy of Knowledge: Situating Epistemology in Time' *Philosophical Papers* 37(1): 27–50.
- \_\_\_\_\_ (1998) 'Rational Authority and Social Power: Towards a Truly Social Epistemology' *Proceedings of the Aristotelian Society* 98: 159–177.
- Gardiner, Georgi (2015) 'Teleologies and the Methodology of Epistemology' *Epistemic Evaluation: Purposeful Epistemology* John Greco and David Henderson (eds.) Oxford University Press, pp. 31–45.
- \_\_\_\_\_ (2012) 'Understanding, Integration, and Epistemic Value' *Acta Analytica* 27(2): 163–181.
- Gerken, Mikkel (forthcoming) *On Folk Epistemology* Oxford University Press.
- Gordon, Emma (forthcoming) 'Social Epistemology and the Acquisition of Understanding' *Explaining Understanding* Routledge.
- Greco, John (2015) 'Testimonial Knowledge and the Flow of Information' *Epistemic Evaluation: Purposeful Epistemology* John Greco and David Henderson (eds.) Oxford University Press, pp. 274–290.
- Grimm, Stephen (forthcoming) 'Understanding as an Intellectual Virtue' *Routledge Handbook of Virtue Epistemology* Heather Battaly (ed.) Routledge.
- \_\_\_\_\_ (2015) 'Knowledge Practical Interests, and Rising Tides' *Epistemic Evaluation: Purposeful Epistemology* John Greco and David Henderson (eds.) Oxford University Press, pp. 117–137.
- Habgood-Coote, Joshua (2017) *Knowledge-How: Linguistic and Philosophical Considerations* Doctoral dissertation at St. Andrews University.
- Hannon, Michael (ms) 'What's the Point of Understanding?'
- \_\_\_\_\_ (2013) 'The Practical Origins of Epistemic Contextualism' *Erkenntnis* 78: 899–919.
- Hills, Alison (2016) 'Understanding Why' *Noûs* 50 (4): 661–688.
- James, William (1987) 'The Will to Believe' *The Will to Believe and Other Essays* NY: Longmans.
- Kappel, Klemens (2010) 'On Saying that Someone Knows: Themes from Craig' *Social Epistemology* Adrian Haddock, Alan Millar, and Duncan Pritchard (eds.) Oxford University Press, pp. 69–88.
- Kelp, Chris (2011) 'What's the Point of "Knowledge" Anyway?' *Episteme* 8: 53–66.

- Kornblith, Hilary (2011) 'Why Should We Care about the Concept of Knowledge?' *Episteme* 8(1): 38–52
- Kusch, Martin (2011) 'Knowledge and Certainties in the Epistemic State of Nature' *Episteme* 8 (1): 6–23.
- \_\_\_\_\_ (2009) 'Testimony and the Value of Knowledge' *Epistemic Value* Adrian Haddock, Alan Millar, and Duncan Pritchard (eds.) Oxford University Press, pp. 60–94.
- Kvanvig, Jonathan (2003) *The Value of Knowledge and the Pursuit of Understanding* Cambridge University Press.
- Nietzsche, Friedrich (1967) *On the Genealogy of Morals and Ecce Homo*, trans. Walter Kaufmann, New York: Vintage Books.
- Pritchard, Duncan (2012) 'Anti-Luck Virtue Epistemology' *Journal of Philosophy* 109: 247–279.
- \_\_\_\_\_ (2011) 'The Genealogy of the Concept of Knowledge and Anti-Luck Virtue Epistemology' *Der Begriff des Wissens/The Concept of Knowledge* Stefan Tolksdorf (ed.) de Gruyter, pp. 159–178.
- Riggs, Wayne (2003) 'Understanding "Virtue" and the Virtue of Understanding' *Intellectual Virtue: Perspectives from Ethics and Epistemology* Michael DePaul and Linda Zagzebski (eds.) Oxford University Press, pp. 203–226.
- Shapin, Steven (1994) *A Social History of Truth: Civility and Science in Seventeenth-Century England* Chicago University Press.
- Smithies, Declan (2015) 'Why Justification Matters' *Epistemic Evaluation: Purposeful Epistemology* John Greco and David Henderson (eds.) Oxford University Press.
- Stanley, Jason (2011) *Know How* Oxford University Press.
- Stanley, Jason and Timothy Williamson (2001) 'Knowing How' *Journal of Philosophy* 98(8): 411–444.
- Sullivan, Emily (ms) 'Valuable Understanding without Achievement'
- Williams, Bernard (2002) *Truth and Truthfulness: An Essay in Genealogy* Princeton University Press.

### C. The Role of Normalcy in Thought

## V. Normalcy and the Contents of Philosophical Judgements

**Abstract.** Thought experiments as counterexamples are a familiar tool in philosophy. Frequently understanding a vignette seems to generate a challenge to a target theory. In this paper I explore the content of the judgements that we have in response to these vignettes. I first introduce several competing proposals for the content of our judgements, and explain why they are inadequate. I then advance an alternative view. I argue that when we hear vignettes we consider the normal instances of the vignette. If the normal instance of the vignette is a counter-instance, the vignette constitutes a challenge to the target theory. I argue that this proposal shows how responses to vignettes are a kind of ordinary, everyday judgement, and I explain how the proposal avoids the problems generated by competing views. Finally, I argue this ‘normalcy proposal’ most naturally accords with our understanding of the method.

### 1. Introduction

Thought experiments as counterexamples are a familiar method in philosophy. Amongst the most paradigmatic include Edmund Gettier’s use of vignettes to argue that knowledge is not justified true belief, Frank Jackson’s Mary vignette against physicalism about conscious experience, and the ‘murderer at the door’ vignette against the view that lying is always all-things-considered impermissible.<sup>1</sup> In typical instances of the ‘vignette as counterexample’ method, we take a theory or claim, such as:

---

<sup>1</sup> These examples are from Gettier (1963), Jackson (1982), and Constant (1797) respectively.



(1) Knowledge is justified true belief.

or

(2) Necessarily, knowledge is justified true belief.

We then present a vignette, short story or anecdote challenges the target theory or claim. An example of a vignette that challenges claims (1) and (2) is,<sup>2</sup>

ORLANDO. A fraudulent bookseller fakes evidence which appears to show that a particular book once belonged to Virginia Woolf. Orlando, deceived by the forgery, purchases the book for a large sum. He thereby comes to believe that he owns a book that once belonged to Virginia Woolf. As it happens, Orlando does own a book that once belonged to Virginia Woolf, since another book he owns—one he does not associate with her—once belonged to Woolf.

If the method is successful, then considering the vignette provides reason and motivation to abandon, revise, reduce confidence in, or reconsider the target theory or claim. The ubiquity and apparent efficacy of thought experiments as counterexamples in philosophy raises several questions, such as (how) does this method work, what kind of judgement do we have in response to these vignettes, and what is the epistemic status of these judgements?<sup>3</sup> In this paper I focus on the content of the judgement we have in response to these cases. Since this judgement seems to both motivate and legitimate certain inferences—such as casting doubt on claims (1) and (2)—understanding the content of this judgement is important. And

---

<sup>2</sup> This vignette is found in Williamson (2007: 183).

<sup>3</sup> Note that judgements in response to vignettes are not, by any means, the only kind of philosophical judgement, and counterexamples are not the only role vignettes play in philosophical theorising. See Gardiner (2015b) for discussion of various kinds of philosophical judgement.

resolving the ‘content question’ illuminates other questions about this central method in philosophy.<sup>4</sup>

## 2. Two Initial Proposals

To see why this question is tricky, consider some proposals. Consider first the view that the content of the judgement we have in response to the vignette is simply the negation of the target claim. Call this the ‘direct proposal’ since it posits that upon hearing the vignette we directly grasp, unmediated by a bridge proposition, the falsity of the target claim.<sup>5</sup> In the above example the target claim is:

- (2) Necessarily, knowledge is justified true belief.

We can understand this claim as entailing claim (3):

- (3) Necessarily, for any subject *x* and proposition *p*, *x* knows *p* iff *x* has a justified true belief in *p*.

To pin down a precise formulation of the target claim, we can formalise (3) as:

---

<sup>4</sup> Williamson (2007), Malmgren (2011), and Ichikawa and Jarvis (2009; 2013) advance views about how the content question bears on other questions. Williamson (2007) argues that his view allows these judgements to be continuous with other everyday judgements; Malmgren (2011) and Ichikawa and Jarvis (2009; 2013) argue that a virtue of their account is that it renders these judgements apriori. Whilst in this paper I bracket the question of whether the judgements are apriori or aposteriori, I argue that answering the content question helps vindicate these judgements, and so legitimates some uses of ‘thought experiments as counterexample’ in philosophy.

<sup>5</sup> Sosa (2006) appears sympathetic to the direct proposal, since he suggests that considering a vignette can lead us directly to endorsing claims such as claim (5).

$$(4) \quad \Box \forall x \forall p (K(x, p) \equiv JTB(x, p))$$

Call claim (4) the JTB view of knowledge (strictly speaking, it is an entailment of the JTB view<sup>6</sup>). The vignette challenges claim (4) by seeming to present us with a counter-instance—a case where someone possesses justified true belief that *p* without knowledge that *p*. Accordingly the direct proposal holds that the judgement has the following content:

$$(5) \quad \text{Someone could have a justified true belief that } p, \text{ yet not know that } p.$$

We can formalise this possibility as:

$$(6) \quad \Diamond \exists x \exists p (JTB(x, p) \& \neg K(x, p))$$

The first proposal, then, is that when we hear a vignette, such as the Orlando vignette, we form a judgement with the content articulated by (6). This suggestion has several virtues. Firstly, (6) directly contradicts (4), and so this proposal explains why the judgement in response to these cases constitutes a problem for the target theory. Secondly, it captures the immediacy that the judgement often has: hearing a successful vignette can quickly produce a grasp of the falsity of the target claim. (Or, less boldly, we immediately grasp the tension between the vignette and the target claim.) The challenge is immediately apparent. By

---

<sup>6</sup> Focusing on claim (3), formalised as (4), as the target claim is useful because most accounts that aim to understand knowledge as justified true belief will entail (3). It is important to note, however, that claim (3) concerns necessary and sufficient conditions, whereas an account of knowledge is typically a claim about constitution, identity, metaphysical dependence, reduction, conceptual equivalence, or linguistic synonymy. See also Williamson (2007: 183).

contradicting the target claim, the proposed content (6) vindicates this.

There are problems, however, with the direct proposal. Firstly, the proposed judgement is not *about* the vignette presented. But it seems that our response to a vignette is about the vignette and includes content specific to the particular vignette. The possibility judgement (6) seems to arrive via, and be mediated by, some judgement that concerns details of the case: It seems as though our bridge judgement has something to do with Orlando, the justification of testimonial beliefs, whether an alternative book renders Orlando's belief true etc., and from this bridge judgement we infer (6).

Relatedly, the proposed content does not posit difference in content between the judgements we have in response to the Orlando case and an alternative Gettier case, such as:<sup>7</sup>

SHEEP. Driving past a field, Roddy looks at a sheep-shaped, white object in normal daylight across a field. As a naturalist with keen eyesight, Roddy is reliable at recognising sheep. His visual inputs accord with his background belief that there are sometimes sheep grazing in this area. Roddy accordingly forms the belief 'there is a sheep in the field'. As it happens, however, what Roddy saw is not a sheep but a sheep-shaped rock. Unbeknownst to him, behind the rock a sheep is grazing.

This vignette also challenges the target claim. As with the Orlando vignette, the case seems to exhibit the possibility of justified true belief without knowledge. But according to the direct proposal, the judgement in response to the sheep vignette is identical to the judgement in response to the Orlando vignette. This is implausible.<sup>8</sup>

---

<sup>7</sup> Pritchard (2012: 251), adapted from Chisholm (1977: 105).

<sup>8</sup> Malmgren (2011: 284) independently makes a similar point.

Thirdly, the direct proposal does not provide much explanatory power. It does not generate resources to explain how the vignettes precipitate or vindicate judgement (6); it does not explain, in other words, why rejecting or revising the target theory can be the correct response to a vignette. The proposed content also does not explain why some instances of forming the judgement upon hearing a vignette are mistaken or illegitimate. Suppose we hear a vignette that is incoherent, has nothing to do with belief formation, or does not describe an instance of justified belief. In these cases, if a hearer were to judge in response ‘someone could have a justified true belief that *p*, yet not know that *p*’, i.e., claim (6), her response does not seem appropriate. It is only an appropriate conclusion when the vignette warrants the judgement. But this proposed content does not seem to offer resources to explain when the judgement is, and is not, appropriate.

One response here argues that the response judgement (6) is appropriate only when it is well-grounded, based on sufficient relevant evidence, reliably formed etc. According to this response, judgement (6) sometimes is inappropriate, even though it is always true; whether it is apt depends on connections between the possibility judgement and the vignette. The response maintains, however, that the judgement’s content is simply (6). The success of this proposal demands discussion of how the possibility judgement might relate to the vignette. This essay can be seen as doing this work: My positive proposal, advanced in section seven, can explain how the possibility judgement (6) can be legitimately arrived at via considering the vignette. And so although I do not think (6) is the correct proposal, the insights in this essay can be harnessed by advocates of this proposal, as a way to explain how and when (6) is a legitimate and reasonable response to considering a vignette.

An alternative proposal posits that the judgement concerns the vignette itself more directly.

The proposal holds that the content is something like,

- (7) Orlando has a justified true belief but does not know that he owns a book that once belonged to Virginia Woolf.

This proposal, unlike the ‘direct proposal’ above, captures that the operative judgement concerns the vignette. The judgement is about Orlando. And this proposal provides resources for explaining the difference between those judgements that are justified, proper, and appropriate responses, and those that are not. Roughly, the judgement is appropriate whenever the vignette describes the protagonist’s having a justified true belief that is not knowledge.

But there are weaknesses. Firstly, it is not clear how the judgement appropriately attributes properties to Orlando; it is common knowledge that Orlando does not exist. Theorists who discuss thought experiments typically realise the vignettes are fictitious, and so cannot judge (7) as a straightforward matter of fact. (Some vignettes deployed as counterexamples are, or purport to be, veridical. In these cases, content (7) seems more plausible. But these cases constitute the small minority of vignettes in philosophy, and the method is successful regardless of whether the vignettes purport to be true.) So if the judgement resembles (7), more must be said about how the content features a non-existent person, and whether or how the acknowledged fictionality of the vignette plays a role in the judgement.

In light of this worry, an alternative proposal holds that the content is something like, ‘within the vignette, Orlando has a justified true belief that  $p$ , but does not know that  $p$ ’. One concern about this proposal is that vignettes can be inconsistent and perhaps even conceptually incoherent. Advocates of this proposal must explain how Orlando’s epistemic states within the vignette bear on which epistemic states are possible. The proposed judgement, in other words, does not itself threaten (4), since it is *prima facie* possible that impossible things can be embedded in a ‘within the vignette’ operator.<sup>9</sup>

### 3. The Necessity Proposal

Given the problems faced by the first two proposals, it seems worthwhile to examine other possibilities. To better articulate other proposals, I first introduce some terminology.<sup>10</sup>

$$(8) \quad V(x, p)$$

This terminology expresses that every sentence in the vignette is true of a person and proposition. It can be understood informally as:

---

<sup>9</sup> Thanks to Daniel Rubio for helpful discussion of these issues.

<sup>10</sup> This terminology resembles Williamson’s ‘GC ( $x, p$ )’, where ‘GC’ stands for ‘Gettier case’ (Williamson (2007: 184)). I prefer to not use Williamson’s terminology for several reasons. Firstly, it is ambiguous between the mere *sentences* of the vignette being true of a person, and the (more full-bodied, richer) fiction being true of a person. See Ichikawa (2009) and Ichikawa and Jarvis (2009; 2013) for more on this distinction. Secondly, ‘Gettier case’ seems like a success term: something is a Gettier case only if it successfully exhibits justified true belief without knowledge. But the mechanisms underlying the methodology are present whether or not the vignette is auspicious. A third reason is that my account of the method generalises to all thought experiments as counterexamples, and thought experiments in other roles, and so there is reason for the terminology introduced to apply smoothly outwith Gettier cases.

(9a) The vignette is true of person  $x$  and proposition  $p$ .

or

(9b)  $x$  stands to  $p$  as described by the sentences in the vignette.

or

(9c)  $x$  and  $p$  satisfy the vignette.

I take claims (9a-c) to provide roughly equivalent glosses of  $V(x, p)$ .<sup>11</sup> I understand the vignette to be the set of claims used to describe a case. An instance of a vignette is any situation in which someone satisfies the vignette.

We can employ this apparatus to articulate a third proposal for the content of our operative judgements in response to vignettes. The third proposal posits that the judgement is ‘for any  $x$  and  $p$  that satisfy the vignette,  $x$  has a justified true belief that  $p$ , but does not know that  $p$ ’. Put slightly differently, ‘anyone who stands to a proposition as described by the vignette has a justified true belief in that proposition, but does not know it’. This content can be expressed:

(10) Necessarily, if the vignette applies to a person and proposition, then that person possesses a justified true belief in that proposition, but the belief is not knowledge.

---

<sup>11</sup> Note that this terminology is different from Ichikawa and Jarvis’s ‘truth in fiction’ apparatus. See Ichikawa and Jarvis (2009: esp. 227; 2013, esp. 204–5). This is because some facts might be true in the fiction, yet not articulated in the vignette itself and so irrelevant to whether  $V(x, p)$ . According to Ichikawa and Jarvis, it is true in the fiction that Orlando has two eyes, for example (Ichikawa and Jarvis (2009: 227)). But the question of eye-count does not arise for determining whether the vignette is true of somebody: the vignette does not mention eye-count. Thus if someone,  $S$ , has one eye, then it is possible for  $V(x, p)$  to hold for  $S$ , but not possible for Ichikawa and Jarvis’s ‘true in the fiction’ relation to hold for  $S$ .



Claim (10) can be formalised as:

$$(11) \quad \Box \forall x \forall p [V(x, p) \rightarrow (JTB(x, p) \& \neg K(x, p))]$$

By itself the truth of claim (11) does not threaten the target theory. The two claims are compatible provided that the vignette could not be satisfied; if the vignette is impossible, then  $V(x, p)$  could never obtain, and (11) is trivially true. But we can combine (11) with the further premise that the vignette could be true of some  $x$  and  $p$ :

$$(12) \quad \Diamond \exists x \exists p V(x, p)$$

This proposal, which I call the necessity proposal, has many virtues. The conjunction of claims (11) and (12) conflicts with the target claim (6), and seems to explain several features of the methodology. Upon hearing a vignette the judgement that the vignette challenges the target theory is often quick. As soon as we comprehend the vignette we perceive the problem. We need not consider a particular variation or specification of the case, such as the possible variation where Orlando purchases the other Woolf-owned book over the internet or has possessed it for decades. This speed and effortlessness might be explained by the necessity proposal: if every case that satisfies the vignette is a counterexample to the target theory, the person need not think of a particular instance—any instance will suffice. We need not consider a specific instance of the vignette, and this proposal can explain that phenomenology.<sup>12</sup>

---

<sup>12</sup> I think this first putative advantage is spurious. This is because the thing it putatively explains—the speed of the judgement—is not well explained by the proposed content.

Relatedly the proposal might capture the generality that we understand the method to have. When we comprehend that the vignette constitutes a problem for the target theory, we also comprehend that various specific details do not affect its dialectical force. We recognise that it does not matter whether the bookseller has brown or blonde hair, for example, or whether Orlando is married. In this sense, the methodology has a kind of recognised generality.<sup>13</sup> Arguably this proposal captures the generality by recording in the content of the judgement that any instance satisfying the vignette is a problem for the view. Note that it is not simply that there is generality—i.e. that Orlando’s marital status is irrelevant—we also can recognise this generality. And the necessity proposal arguably captures this; the proposal holds that this generality is located in the *content* of the judgement.

The central problems with this proposal are that the proposed content, judgement (11), is false, and obviously so. A vignette cannot hope to provide a comprehensive accounting of a case—vignettes are typically between 3 and 12 lines long—and inevitably leaves a great deal unspecified. Many unarticulated facts are consistent with the vignette;  $V(x, p)$  obtains in instances where the protagonist has brown hair, or blonde hair, or is bald. And it is not only incidental facts about Orlando’s hair colour and marital status:  $V(x, p)$  also holds—the vignette is satisfied—by instances where the given vignette is supplemented with epistemically relevant facts. Consider for example the following two vignettes:

---

The proposed content, if formed responsibly, requires that we make a judgement covering *every* possible version of the case. This process, even if possible, would not be immediate, contra the conjectured advantage.

<sup>13</sup> See Malmgren (2011: especially 281 ff) for discussion of the generality of the method.

GULLIBLE ORLANDO. A fraudulent bookseller fakes evidence which appears to show that a particular book once belonged to Virginia Woolf. Orlando, deceived by the forgery, purchases the book for a large sum. He thereby comes to believe that he owns a book that once belonged to Virginia Woolf. Orlando should have known better, since just the previous evening he watched the local news describe the bookseller as a fraud, liar, and manipulator, who was selling worthless books as historical artefacts. Orlando recognised the man from the exposé, but as he entered the shop Orlando decided to disregard the exposé. As it happens, Orlando does own a book that once belonged to Virginia Woolf, since another book he owns—one he does not associate with her—once belonged to Woolf.

MARRIED ORLANDO. A fraudulent bookseller fakes evidence which appears to show that a particular book once belonged to Virginia Woolf. Orlando, deceived by the forgery, purchases the book for a large sum. He thereby comes to believe that he owns a book that once belonged to Virginia Woolf. As it happens, Orlando does own a book that once belonged to Virginia Woolf, since another book he owns—one he does not associate with her—once belonged to Woolf. He phones his wife to describe the purchase, and she—a competent and rigorous bookseller herself—tells him about the other Woolf-owned book.

These variants are consistent with the original Orlando vignette. They merely augment the vignette with further details. Some instances such that  $V(x, p)$ —i.e. some instances in which the original vignette is satisfied—are also instances that satisfy the gullible Orlando or married Orlando vignettes. But the former intersection—situations that satisfy both the Orlando vignette and the gullible Orlando vignette—includes many instances in which Orlando is not justified in his belief. He ignored relevant evidence. The second intersection includes many instances where Orlando possesses knowledge—Orlando's reliable wife informed him that  $p$ . Thus there are instances where  $V(x, p)$ , but where  $x$  lacks JTB or  $x$  possesses knowledge. Judgement (11) is false.

There are many possible ways to deviantly satisfy the Orlando vignette, such as instances in which Orlando is obsessed with Woolf, and so is disposed to believe that all kinds of his purchases once belonged to Woolf, or in which he resides in a country where it is common

knowledge that booksellers are terrible frauds, or in which Orlando lives one thousand years after Woolf, and so any books once owned by her would be prohibitively expensive. Any of these ‘deviant variations’ render his belief unjustified. There are also many deviant variations that would render his belief knowledge. Call the falsity of claim (11) owing to deviant variations of the vignette the ‘deviant variation’ problem.<sup>14</sup>

The falsity of judgement (11) is a problem for the necessity proposal because if employment of counterexamples to challenge target theories rests on a false judgement, then our resultant beliefs are unjustified. We should not, for example, have rejected the JTB view of knowledge based on Gettier vignettes. The necessity proposal undermines, rather than vindicates, the method.

But the *existence* of deviant variations is not the only problem with the method. Perhaps, after all, our inference from hearing a Gettier vignette to rethinking the JTB view of knowledge is a flawed inference, and maybe a clear-eyed understanding of the method would reveal this. If the method is flawed in this way, the falsity of (11) is not a strike against the proposal. This would be a surprising conclusion, but not one we should dismiss outright. The second problem is that (11) is obviously false, yet many careful thinkers are convinced by Gettier cases and other vignettes. It is implausible that so many thinkers commit such an elementary error. Even if we are wrong, we are not foolish.<sup>15</sup>

---

<sup>14</sup> The problem of deviant variations for the necessity proposal is discussed in Williamson (2005: 6–7; 2007: 185ff), Malmgren (2011: 275 ff), Grundmann and Horvath (2014), and Ichikawa and Jarvis (2009: 223–5; 2013: 200–1).

<sup>15</sup> See also Malmgren (2011: 276).

A further problem with the necessity proposal is that when we hear a compelling vignette, we continue to think it poses a challenge to the target theory even after we also realise that there are possible deviant ways to augment the case.<sup>16</sup> We simply acknowledge that these deviant variations, whilst consistent with the text, are not the normal understanding of the text. But according to the necessity proposal, acknowledging these possible deviant variations should cause thinkers to judge the vignette benign; they realise claim (11) is false.

One could attempt to defend the necessity proposal by ruling out deviant variations. A flatfooted strategy posits a (tacit) ‘that’s all folks’ clause at the end of every vignette used in theorising. This clause conveys that no other facts should be imported; the story is complete as told. This defence of the necessity proposal holds that whenever theorists consider a vignette as a counterexample, they implicitly acknowledge that they cannot import any further facts. Thus the deviant cases described above, since they rely on importing ‘extra’ facts, are ruled out.

But this flatfooted proposal will not work. We cannot consider a vignette without importing *some* facts not explicitly appearing in the vignette. These include that Orlando is alive, for example, and that the book is not on fire. If we are limited to explicit claims in the vignette, the case might be literally unimaginable.

---

<sup>16</sup> Malmgren (2011: 276) also makes this point, although she offers a different diagnosis—appealing to the vignette designer’s intentions—to explain why considering deviant variations does not undermine the perceived potency of the vignette.

A slightly more refined defence argues that we hold the epistemically relevant facts fixed (or, if the target theory concerns ethics or aesthetics, the ethical or aesthetic facts respectively). This proposal allows that we are free to round out the story as needed to comprehend the case, but proscribes importing facts that might bear on whether the protagonist possesses justified true belief or knowledge. But this strategy also lacks promise. Inevitably there are epistemically relevant facts not mentioned in the vignette itself. These include that Orlando is not a fool, that the bookseller is not a well-known fraud, and that Orlando did not later receive testimony about his other Woolf-owned book. Other facts, such as that Orlando does not own every book in the world, and that his recent purchase was not published within the decade, are not explicitly epistemic but are epistemically relevant in this case because they bear on the normative status of Orlando's doxastic attitudes. We must import these extra facts into our understanding of the story: They are not explicitly stated in the vignette, but are expected. They fill out how we normally understand the case and help constitute typical instances of the vignette. I develop the role of expectations, normal understandings, and typicality judgements in section seven.

#### **4. The Possibility Proposal**

The obvious falseness of proposed judgement (11) suggests the necessity proposal does not capture the operative judgement we have in response to vignettes. But perhaps this is unsurprising: judgement (11) is far stronger than required. Rather than requiring every instance that satisfies the vignette to exhibit justified true belief without knowledge, (plausibly) all that is required is just one possible instance where a person possesses justified

true belief without knowledge. This is sufficient to challenge the target theory.<sup>17</sup> Realising this suggests a natural alternative proposal. This fourth proposal holds that the judgement the vignette elicits is that at least one possible instance of the vignette exhibits justified true belief without knowledge. This proposal is advanced by Anna-Sara Malmgren, and can be expressed:<sup>18</sup>

- (13) It is possible that some  $x$  and  $p$  satisfy the vignette, and  $x$  has a justified belief that  $p$  but does not know that  $p$ .

This can be formulated as:

$$(14) \Diamond \exists x \exists p [V(x, p) \& (JTB(x, p) \& \neg K(x, p))]$$

This possibility claim entails the falsity of the target claim. Note this proposed content entails the possibility claim (12)  $\Diamond \exists x \exists p V(x, p)$ , which states that the vignette is possible, and so claim (12) need not be independently endorsed.

The possibility proposal has virtues. Firstly, it is psychologically plausible; whilst it does not seem that we consider or judge every possible instance that satisfies the vignette, it is plausible that we ‘latch onto’ one instance. Secondly, supposing the right conditions, the judgement is true. The judgement is false if either i.) the target claim is true, and so anyone

---

<sup>17</sup> One possible counter-instance is sufficient to challenge a target theory that is, or entails, a necessity claim. One possible counter-instance is unlikely to challenge a target claim concerning what is typical, normal, or characteristic, such as ‘knowledge is normally justified true belief’ or knowledge is characteristically justified true belief’. I return to this issue in section 7.2.

<sup>18</sup> Malmgren (2011).

who has justified true belief possesses knowledge, or if ii.) the vignette is abysmal, such that no augmenting of the vignette can lead to an instance that exhibits knowledge without justified true belief. An example of an abysmal vignette might be:

NO ONE. All the humans and animals have died. Earth was the only planet with life, and the explosion destroyed all living things. Gradually the buildings decay and collapse. A mug filled with rainwater falls through a rotted picnic table and smashes on the ground.

Perhaps no instance of this vignette, no matter how creatively augmented, can exhibit justified true belief without knowledge. This is because any candidates for Gettierisation are stipulated dead. But with either of these disjuncts—the truth of the target theory or an abysmal vignette—we ought *not* judge that the vignette generates a counter-instance to the theory; in both cases it cannot. Thus the possibility proposal, unlike the necessity proposal, seems to generate the right results about when the proposed judgement is true.<sup>19</sup>

The possibility proposal has been criticised for not vindicating the generality of the judgements we have in response to vignettes.<sup>20</sup> There are several ways to unpack this criticism. One way is to note that it does not matter for the methodology whether the protagonist has two children, blonde hair etc. But the proposal says we fix on some particular instance, and so focus on an instance in which hair colour, etc. is determined. We should be able to recognise that the judgement is legitimate about a range of variations; it does not matter what hair colour Orlando has, for example. The implicit generality, it is

---

<sup>19</sup> I argue below that this virtue is spurious. The proposed content is true about vignettes that are not genuine challenges to the target theory, and so is not a plausible candidate for the content of our judgement.

<sup>20</sup> Malmgren (2011: 281–5).



argued, is important to our understanding of the method and so should be reflected in the judgement. But the judgement posited by the possibility proposal does not contain or allow any generality. Another way to unpack this criticism is to note that if a theorist thinks that the Orlando vignette is a challenge to the target theory when Orlando is brunette, but not when he is blonde, she has misunderstood the method. As Malmgren describes the criticism,<sup>21</sup>

[If I hold these two beliefs, i.e. that Orlando has justified true belief without knowledge when he is brunette, but not when he is blonde] I am being confused or inconsistent: there is a problematic discord between my earlier and later judgement [...] Now, the natural way to account for this—the natural diagnosis of my apparent rationality failure—is to say that I am here *contradicting* my original intuitive judgement (or that I am contradicting something that obviously follows from it).

But suppose that the content of our judgement is, as Malmgren suggests, the possibility proposal. Then the person who judges that blonde Orlando knows, but brunette Orlando does not, has not judged anything inconsistent. The contents of the judgements concern different particular possible instances of the vignette. Since there is no generality in the judgement, the criticism contends, the proposal cannot explain the fault with this pair of judgements.

Malmgren responds to the generality objection. She distinguishes between generality in the content of a judgement and generality in the grounds for a judgement. She argues that the objection requires that the generality—and associated tension—be located in the content of the judgement. Malmgren concedes that if tension between the judgements—that brunette Orlando has justified true belief without knowledge and blonde Orlando does not—is in the

---

<sup>21</sup> Malmgren (2011: 285, emphasis in original).

content of the judgement, she cannot explain this. Her proposal cannot explain why these two judgements are *contradictions*.<sup>22</sup> But, Malmgren argues, the tension is actually in the *grounds* for holding the judgement: when someone judges that brunette Orlando has justified true belief without knowledge, this rationally commits her to the same judgement when something irrelevant, such as hair colour, is altered.<sup>23</sup> I will not adjudicate whether Malmgren's response is persuasive. We might note the dialectical force of the putative objection as a conditional: *if* the person who judges that 'brunette Orlando, but not blonde Orlando, has justified true belief without knowledge' has thereby contradicted herself, this is a serious problem for the possibility view. If she has simply violated rational demands, but not contradicted herself, this is not a problem for the possibility proposal.

The possibility proposal suffers from other problems. Consider someone, Ann, who believes there can be justified true belief that is not knowledge.<sup>24</sup> Ann antecedently denies the target claim. Perhaps she is convinced by thinking about lottery cases, kinds of testimony, or barn façade cases, or perhaps other reasoning has led her to think that knowledge differs from justified true belief. Now suppose Ann hears the Orlando vignette, but does not have the 'Gettier intuition'; she does not judge that the vignette as told describes a case of justified true belief without knowledge. Given Ann antecedently believes that justified true belief and knowledge can diverge, she should agree that *some* instance that satisfies the vignette exhibits justified true belief without knowledge. All Ann must do is provide additional details that transform the vignette into a lottery case, for example, and judge that the variation exhibits

---

<sup>22</sup> Malmgren (2011: 291).

<sup>23</sup> Malmgren (2011: 290–4).

<sup>24</sup> This objection is also found in Ichikawa and Jarvis (2013: 203).

justified true belief without knowledge. Whilst she does not think the case presented is compelling, Ann does agree that:

(13) It is possible that some  $x$  and  $p$  satisfy the vignette, and  $x$  has a justified belief that  $p$  but does not know that  $p$ .

(14)  $\Diamond \exists x \exists p [V(x, p) \& (JTB(x, p) \& \neg K(x, p))]$

Ann will, in other words, upon hearing the Orlando vignette, agree to the content proposed by the possibility proposal. This suggests the possibility proposal cannot capture disagreement about putative counterexamples. Consider someone else, Bob, who judges that the Orlando vignette is a successful Gettier case. Bob believes that the natural understanding of the vignette describes someone who has justified true belief but not knowledge. Surely Ann and Bob disagree in their judgement about the case. But they agree about claim (14). If claim (14) represents the judgement we have about the case when we take it to challenge the target claim, this renders Ann and Bob's agreement about (14) confusing. An account of the judgement in response to vignettes should represent Ann and Bob as having contrary responses to the vignette: one straightforwardly judges the vignette to exhibit justified true belief without knowledge, whilst the other does not. But, according to the possibility proposal, Ann assents to Bob's judgement. This is a problem for the possibility proposal.

One potential response holds that whilst Ann and Bob agree on claim (14), their grounds for assenting differ. Bob's ground for (14) is reflection on the case as presented. Ann's grounds are an antecedent commitment to the possibility of justified true belief without knowledge combined with the judgement that nothing in the presented vignette precludes augmentation

into such a case. (Contrast the ‘no one’ vignette, above, which cannot be transformed into a case of justified true belief without knowledge.) But this response is unlikely to work: Ann and Bob disagree about the vignette itself and whether the presented case is a problem for the target theory. Very plausibly they have different judgements about the case, and this should be reflected in the proposed content. The possibility proposal cannot do this.

One diagnosis of the possibility proposal’s disagreement problem suggests that the operative judgement makes reference to whether knowledge obtains in specific versions of the vignette, not simply ‘some possible variation of the vignette’: Ann and Bob do disagree about more specific versions of the vignette.

There is a further problem for the possibility proposal. Consider this vignette,

ANEMIC VIGNETTE. Ted looks at a tree; he then forms a belief.

Note that the Anemic Vignette, like the Orlando vignette, can be filled out with supplementary details. We can augment the case with details about what Ted believes and how. Supposing that knowledge diverges from justified true belief, a hearer can provide supplementary details so that Ted possesses a justified true belief that is not knowledge. Since such augmentation is possible, and the vignette is possible, claim (14) is true of the anemic vignette. A hearer can endorse judgement (14), in other words, about the Ted vignette. But Ted provides no evidential weight as a counterexample to the target theory of knowledge. Before we are moved by the anemic vignette, we must learn more about how Ted’s belief is justified, whether it is true, and what might impede Ted’s knowledge. The

existing vignette is not persuasive. The mere fact that we can (assuming knowledge is not justified true belief) endorse claim (14) about the vignette is neither here nor there in epistemological theorising.

The same problem applies to botched vignettes. A botched vignette is when the speaker, who has a compelling vignette in mind, bungles important details in the delivery, so that the normal interpretation of the vignette is not one which exhibits justified true belief without knowledge. An example of a botched vignette is, ‘Roddy was looking at a field, and he saw lots of sheep. One the things he thought was a sheep was actually a rock. Roddy formed the belief “there are sheep in the field”.’ Since botched vignettes can often be augmented into successful counterexamples, the possibility judgment (14) is usually true of botched vignettes, but this does not seem like the right result. Botched vignettes do not present a challenge to the target theory.

This might seem like a strange criticism of the possibility proposal—after all, anyone who thinks the anemic or botched vignette as it stands constitutes a problem for the target theory is mistaken. But I think it illuminates a problem with the possibility proposal: the proposal cannot vindicate the evidential significance of vignettes as counterexamples. Vignettes generate potent, compelling evidential weight. One explanation of their effectiveness is that when we hear a vignette we thereby *confront* an instance of, for example, justified true belief that fails to be knowledge. The vignette does not merely make us realise that such as instance is (merely) possible. The compelling vignette itself presents us with an instance that is counter to the target theory. If the Orlando vignette merely describes something that *could* be a case of justified true belief without knowledge, as the possibility proposal suggests, this

provides significantly less evidential force against the target theory. But the possibility proposal does not paint this picture. The possibility proposal only says that we judge that some version or other of the vignette is justified true belief without knowledge. The proposal does not quite capture the way that we are *faced* with a counter-instance.

Note too that vignettes play a social role: theorists agree, disagree, and persuade each other about vignettes. The possibility proposal might render it mysterious how vignettes fulfil this social role since, according to the possibility proposal, the counter-instance itself is not communicated by the vignette. Only the judgement that ‘a counter instance based on or consistent with this vignette is possible’ is evinced by the vignette. But it seems that when we recount vignettes we convey more than that a possible variation of the vignette is a case of justified true belief without knowledge. Instead we seem to convey something more like particular illustrations of justified true belief without knowledge.

To illuminate this social role, suppose a theorist told an anemic or botched vignette, and a second theorist, on that basis, judged the target theory false. This would be a concern. It does not seem like good philosophical communication or a legitimate basis on which to reject the target theory. If we learn that the intermediary judgement was merely (13) ‘it is possible that some  $x$  and  $p$  satisfy the vignette, and  $x$  has a justified belief that  $p$  but does not know that  $p$ ’, this will not allay our concerns. To judge that the theorists were working well, we want to know about the specific variant of the vignette each had in mind, to see whether it was a lottery case, a Gettier case, or some other case. But the possibility proposal does not generate resources to explain why this interaction is ineffective. According to the possibility proposal, this ‘possibility’ content is all we convey, enable or evince when we employ

vignettes as counterexamples.

To summarise this strand of objection: Insofar as (14) captures the content of our judgement, then to that extent the vignette is not dialectically compelling. A vignette that generates a stronger judgement would be more compelling. But vignettes in philosophy are remarkably compelling; given how compelling vignettes are, (14) cannot adequately capture our reaction to them. (Indeed perhaps less compelling counterexamples do generate judgement (14). Perhaps sometimes we hear a case and judge ‘there’s a counterexample lurking in there’, or ‘some version of this case is a counterexample’. But such counterexamples do not typically legitimate jettisoning established theories. Instead they are an invitation to construct a better vignette.) Given the method’s effectiveness, the content of our response to the vignettes must be stronger than the possibility proposal.

A general difficulty with the possibility proposal is that judgement (14) is too weak. Given that vignettes can be augmented, and knowledge and justified true belief diverge, the judgement is often true of a vignette. It is true of *all* possible vignettes other than those that stipulate either i.) that the person knows p, ii.) that the person does not have a justified true belief that p, or iii.) there are no potential candidates for Gettierisation, such as in the ‘no one’ vignette.<sup>25</sup> Unless one of these conditions holds, the vignette can be transformed into a case of justified true belief without knowledge, such as a Gettier case, barn façade case or

---

<sup>25</sup> Perhaps more precisely the first two conditions should read i.) that every person knows everything and ii.) that no person has a justified true belief. These conditions entirely preclude judgement (14) being true of an augmented variation of the vignette.

lottery case. Accordingly claim (14) is true of almost every vignette. This suggests the judgement we have in response to cases is stronger.<sup>26</sup>

## 5. Williamson's Counterfactual Proposal

The fact that we are *struck* by vignettes as counterexamples suggests the judgement elicited by the vignettes concerns a more specific instance of the vignette. We are confronted by a problem instance, and this should be reflected in the content of our judgement. A fifth proposal for the content of our operative judgement is advanced by Timothy Williamson. He holds that the content of the judgement is the subjunctive:<sup>27</sup>

- (15) Were the vignette to be satisfied by an *x* and *p*, then she would have justified true belief, but not knowledge, that *p*.

Williamson uses a Lewisian possible worlds understanding of counterfactuals to interpret claim (15) as something like ‘in the nearest possible worlds where the vignette is satisfied, *x* has a justified true belief that *p*, but does not know *p*’.<sup>28</sup> Since this proposed counterfactual content, unlike the possibility content explored above, does not entail the possibility of the vignettes being satisfied, this proposal requires that we also judge:

---

<sup>26</sup> As I explain in section 7.2, an account of the content of our judgements in response to vignettes should explain other roles that vignettes play in philosophy, such as motivating claims, conveying ideas, illustrating distinctions, and homing in a topic or debate. I argue below that the possibility proposal is ill-equipped to do this.

<sup>27</sup> Formalising the precise content of Williamson's counterfactual proposal is complicated by the problem of donkey anaphora for counterfactuals. For discussion see Williamson (2007: 195–9 and appendix two), Ichikawa and Jarvis (2009: 225, n.4; 2013: 201, n.6), and Malmgren (2011: 277; 306).

<sup>28</sup> Lewis (1973).



$$(12) \Diamond \exists x \exists p V(x, p)$$

which expresses that the vignette could be satisfied by some  $x$  and  $p$ . The judgement posited by Williamson's counterfactual proposal is less demanding than the judgement expressed by (11), the necessity proposal. This is because it does not require that every instance satisfying the vignette manifests justified true belief without knowledge. It thus avoids the 'deviant variations' problem faced by the necessity proposal. The proposed judgement is more demanding, on the other hand, than that expressed by (14), the possibility proposal. This is because the truth of the counterfactual (15) requires more than that one possible instance of justified true belief satisfies the vignette and is not knowledge; it requires that those possible instances would obtain, were the vignette true. This allows the proposal to avoid the 'anemic vignette' and 'botched vignette' problems facing the possibility proposal and better explain the disagreement between Ann and Ben. Were the 'anemic vignette' to occur, it almost certainly would not be a Gettier case—since most beliefs are not Gettiered, it would be a remarkable coincidence if Ted's belief were a Gettier case. And Ann would not endorse judgement (15). So Williamson's approach posits a claim about which Ann and Ben disagree. Perhaps then, this fifth proposal about the content of the judgement is 'just right', laying between judgements that are too weak and too strong.

Judgement (15) seems true about paradigmatic successful counterexamples: were the Orlando vignette to occur, it would be an instance of justified true belief without knowledge. Intuitively, given the way the world is, the variations on the Orlando case that do not exhibit justified true belief without knowledge are likely more modally distant possibilities than variations in which the Orlando case is a Gettier case. Williamson's counterfactual proposal

avoids the ‘deviant variations’ objection; the mere possibility of deviant variations are irrelevant to the counterfactual judgement, and so do not threaten the proposal.

This proposal also plausibly accords with the phenomenology of forming a judgement upon hearing a vignette. Plausibly we consider, ‘were that to occur, would it be justified true belief without knowledge?’ And the proposal seems well-equipped to explain the speed of judgements about successful cases: we need only think about a relatively straightforward counterfactual question.

According to Williamson’s proposal, the world provides the unasserted details.<sup>29</sup> Details such as whether Orlando’s wife informs him of the other book or whether the bookseller was featured in the previous day’s news are furnished by whether, were the vignette to occur, these features obtain. These counterfactual facts depend on facts about the actual world. This feature seems like a virtue of the view: it is plausible that (in at least some ways) the world determines unarticulated details of cases. Suppose that in the actual world booksellers were often fraudulent, and this was common knowledge. Suppose bookselling were a devious profession with lax regulations, ubiquitous corruption, and inadequate methods of verifying the authenticity of books. In this scenario, bookselling resembles the practices of fortune tellers, relic touters, and email spammers. If this were the circumstance in the actual world, it seems that were the Orlando story true, Orlando would not have a justified belief. And were this the actual world, if the original Orlando vignette were told hearers would likely not think it causes a problem for the target view, because Orlando’s belief would be

---

<sup>29</sup> Williamson (2007: 186).

unjustified. Thus Williamson's counterfactual approach captures that the judgements depend on contingent facts about the world.<sup>30</sup>

This proposal might also explain our ambivalence about cases such as TrueTemp and the Reliable Clairvoyant.<sup>31</sup> These cases are employed to test theories in epistemology, but many people simply cannot adjudicate whether they qualify as cases of knowledge (or whether there is a fact of the matter), and so whether they constitute a challenge to the target theory. Perhaps this uncertainty stems from difficulties forming the relevant counterfactual judgement about the vignette: it is too hard to determine what would be known in a world so epistemologically alien that reliable clairvoyance, or brain scans that determine belief contents, existed. Similarly if a vignette involves many intricacies or implausible psychology it might be harder to determine the relevant counterfactual, and so the vignettes are not perceived as challenges to the target theory. It would be a substantial theoretical virtue if Williamson's proposal not only predicts which cases are problem cases for a theory, but also explains why with some vignettes it is hard to determine whether they challenge the target theory.

Although the proposal has many virtues, there are also problems. The first worry is whether the world really can supply the relevant background details. It seems implausible that the world fills in facts about Orlando and the bookseller; they are fictional characters (or fictional characters *at best*, since perhaps they do not enjoy enough facets to qualify as

---

<sup>30</sup> As I explain in section 7.1, Williamson's counterfactual view overstates the ways that the truth of the judgement, and the success of a counterexample vignette, depends on contingent facts about the world.

<sup>31</sup> For the TrueTemp case, see Lehrer (1990: 163–4). For the Reliable Clairvoyant case, see Bonjour (1980).

genuine characters). Given that Orlando is merely fictional, it might be indeterminate whether the nearer worlds are ones where Orlando watches an exposé about the charlatan, is married to another bookseller, or is dreadfully gullible.

We might suppose that most people are not, and so Orlando is likely not, married to a bookseller or terribly gullible. But Orlando is not an actual person, so it is unclear whether this induction is legitimate; it is unclear how the actual world can provide counterfactual facts for non-existent, fictional entities. If I tell an untrue story about Stephen Fry, by contrast, but omit explicit facts about his psychology, the world might fill in that he is not foolishly gullible in my story, because he is not so in real life. Were my story true, it would be true that Fry would not be gullible. But it is not clear how a similar mechanism can apply to fictional characters.

It seems more plausible that the actual world helps determine what the hearer takes for granted and what defaults can be supposed, and so helps determine what the vignette-teller should mention and what she can leave unarticulated. The actual world affects, in other words, how we understand the case. In the actual world, booksellers are not typically fiendish frauds, and so hearers can assume Orlando has default licence to trust the seller without the vignette-teller needing to make this explicit. And most people are not hopelessly gullible, so the fact Orlando is not hopelessly gullible need not be mentioned. But most people do not fortuitously possess books that once belonged to Virginia Woolf, so the teller must explicitly mention this feature. Hearers would not otherwise assume it.

Facts about the actual world inform, in other words, our storytelling norms: if Orlando is gullible and disregarded a televised exposé, this is relevant to the story and is an atypical and abnormal feature, and so the storyteller would mention it. It would not be an expected part of a backstory, and so must be articulated. If something would be abnormal, but is not articulated in the vignette, this licenses hearers to assume the more normal version. Since the vignette-teller did not say Orlando was unusually gullible, for example, the listener can suppose that he is not. But this mechanism is not the actual world directly determining which Orlando-backstory is a closer possibility, as Williamson's counterfactual view holds. Instead narrative norms—which are informed by the nature of the actual world—help determine what hearers should infer and how they understand the case.

A second worry about the proposal is that the counterfactual judgement might be too demanding.<sup>32</sup> Suppose that, unbeknownst to the vignette-teller, in the actual world the Orlando vignette is in fact satisfied, but it is satisfied in a deviant way, such that it is not justified true belief or is knowledge. In this eventuality, the actual world is one where something like 'gullible Orlando' or 'married Orlando' obtains. This renders the judgement expressed by (15) false. (The same would happen if the Orlando vignette is not satisfied in this world, but in the nearest worlds where it were satisfied, a deviant realisation obtains. This would also render judgement (15) false.)

But, the worry continues, this does not seem like the right result. It seems that even if the actual world is, surprisingly, a 'married Orlando' world, and so the putatively fictional

---

<sup>32</sup> See Williamson (2007: 200), Malmgren (2011: 279), Grundmann and Horvath (2014), Ichikawa and Jarvis (2009; 2013), and Ichikawa (2009) for discussion of the deviant realisations problem.

Orlando case is actually deviantly satisfied, the Orlando vignette remains a successful challenge to the target theory.

Williamson argues that this objection is mistaken. He argues that if the nearest instance that satisfies the vignette is not a case of justified true belief without knowledge—if, in other words, judgement (15) is false—then the vignette is *not* a successful counterexample, even if no one ever realises this counterfactual. In these cases, he says, our judgement that the vignette challenges the target theory is mistaken. Williamson notes that thinkers are often reluctant to admit error, and claims that if we insist on the potency of a putative counterexample after learning that the nearest satisfaction of the vignette is deviant then we manifest this stubbornness.<sup>33</sup>

Williamson's response does not sound correct to me. The methodology of using vignettes to illustrate problems with theories requires that those vignettes be sensitive to acknowledged or systematic features of the world. The vignettes are less compelling to the extent that they depart from our ordinary understanding of how the world is, or violate scientific (including psychological) theories. But the method is not hostage to unexpected quirks of the actual world. If, unbeknownst to the vignette-teller, it happens that the Orlando vignette is satisfied by someone who knows that *p*, this should not impugn their using the vignette as evidence against the target theory. A bizarre anomaly in the actual world does not affect whether the Orlando vignette succeeds in challenging the target theory. Whilst the world fills in the gaps—by informing storytelling norms and shaping our expectations about systematic acknowledged features, such as whether clairvoyance is reliable and whether booksellers are

---

<sup>33</sup> Williamson (2007: 200–1).

often fraudulent—the specific vignettes should not be held hostage to unusual quirks and isolated contingent features of the actual world.<sup>34</sup>

Note that the kind of anomalies Williamson has in mind—the ones that can, according to Williamson, annul a vignette—are truly unusual. They are not systematic fixtures of the world such as the truth of clairvoyance or whether booksellers are typically charlatans. And they are not widely known contingent features, such as whether Virginia Woolf owned books. Instead they are one-off anomalies, including ones that are unknown and unknowable.<sup>35</sup>

Suppose in philosophy class a professor tells the sheep vignette to challenge the view that justified true belief entails knowledge. A student raises his hand and comments that this occurrence happened to him once, but explains that ‘in my case there were some other sheep in the field I saw too; it’s just that the first one I thought I saw was actually a rock with a sheep behind it’. In this case the class plausibly learns that the sheep vignette has a deviant actual realisation.<sup>36</sup> It seems that many students in the class will, despite the interjection, see

---

<sup>34</sup> As Ichikawa and Jarvis note, if we reject the JTB account of knowledge based on a vignette, and judgement (11), but the vignette in fact has a deviant, non-Gettier realisation, the rejection is itself a Gettier case (Ichikawa and Jarvis (2009: 226; 2013: 203)).

<sup>35</sup> Unpacking the distinction between isolated contingent anomalies and systematic features will take some work. But I think we have a good enough understanding of the distinction to cautiously proceed.

<sup>36</sup> There are difficulties: it is not clear whether a vignette that is intended to feature wholly fictional characters can be satisfied by real people. As Kripke argues, perhaps fictional protagonists are essentially fictional (Kripke 1980; 2000). As Williamson notes, and Ichikawa and Jarvis develop, we might understand the names in these vignettes as simply placeholders for variables. See Williamson (2007: 184) and Ichikawa and Jarvis (2009: 229, n. 13; 2013: 207).

how the vignette challenges the target theory. They perceive that justified true belief and knowledge can diverge.

According to Williamson's counterfactual view, these students would be mistaken. But it seems like a perfectly fine piece of reasoning. In fact, if a classmate dwells on the deviant case, and so does not see how the initial vignette challenges the target theory, it seems she manifests less supple and sophisticated thought. The teacher might draw attention to the natural, normal understanding of the vignette and explain how, given the content of the vignette, the most natural interpretation is one where no other sheep were seen.<sup>37</sup> If the classmate still cannot see any problem for the target theory, because of the deviant realisation described—but would have perceived the problem had the student never spoken up—it seems that she does not understand the method. The vignette can and should have dialectical force, despite acknowledging a deviant realisation; its success is not hostage to abnormal isolated occurrences in the world.<sup>38</sup>

Note too that if another student subsequently contributes, 'this happened to me too, but I did not see any other sheep. I just continued to believe the rock was a sheep. I discovered

---

<sup>37</sup> Suppose the teacher adds 'no other sheep were seen' after the student pipes up. Plausibly according to Williamson's view the teacher has now presented a different case—the vignette picks out different possible worlds. But a more natural understanding of this is that the teacher has simply further described or clarified the original case. It is the same case, but with more description.

<sup>38</sup> Ichikawa notes that if we learn about a deviant instance of a vignette, we might change, expound on, or further describe the vignette, so that it precludes this deviance. Whilst this might appear to support Williamson's counterfactual proposal, because deviant actual realisations cause us to adjust the vignette, Ichikawa notes that we also do this with *merely mentioned* possible deviant instances. This suggests we amend or further describe the cases to be clearer or more convincing in our presentation, rather than to generate the right truth conditions for the counterfactual judgement. See Ichikawa (2009: 439).



the truth months later' and thereby describes a non-deviant instance of the vignette, plausibly judgement (15) is nonetheless false. This illustrates a further problem with the proposed content: the semantics of counterfactuals plausibly render (15) false if there are some actual deviant realisations and some actual non-deviant realisations. But this seems like a bad result for the counterfactual proposal: If the class hears a successful-sounding vignette, *and* hears compelling evidence that some actual realisations of the vignette are justified true belief without knowledge, this presents a hefty challenge to the target theory. But in this case judgement (15) is false. Since judgement (15) is false whenever actual or modally nearby instances are divided between deviant and non-deviant realisations, proponents of the counterfactual view must provide an explanation of why these vignettes constitute a challenge to the target theory.

The fact that when reasoning with vignettes we can ignore deviant actual instances, as the students in the above class should, provides evidence that discussing vignettes does not commit theorists to transient, quirky, unusual, isolated or unexpected features of the world. Instead it seems we can focus on a normal understanding or typical instance of the vignette. The teacher might even say of the student's own (deviant) sheep-rock experience 'that isn't the *typical* case', even though it is an actual instance. The actual might not be typical, and the typical might not be actual.

Jonathan Ichikawa and Benjamin Jarvis have argued that a consequence of Williamson's counterfactual proposal is that we can never know whether a vignette is a successful

challenge to the target theory.<sup>39</sup> This is because, they suggest, the relevant counterfactual is too difficult to know. Ichikawa writes,<sup>40</sup>

The worry is that in too many cases, it is not plausible that we know the relevant counterfactuals, because it is not plausible that we know whether we are situated deviantly in modal space. The world is a big place, and we should not be at all surprised if [a deviant instance obtains]. And I see few prospects for learning whether [the counterfactual judgement is true].

I am not convinced: we do in general know many true counterfactuals, even when we cannot rule out their being false (just as in general we know many truths, even when we cannot rule out their being false). And so perhaps we do know that, for example, were a person deceived by an unscrupulous bookseller in the way described by the vignette, but happened to own a book of that description, they would have a justified true belief that is not knowledge. The facts that would make the protagonist's belief knowledge or unjustified are unlikely, and so we can know the counterfactual. We are not in a position to rule out that a deviant instance would obtain, but it is sufficiently unlikely, so we can nonetheless possess the relevant knowledge. Thus I am not convinced that Williamson's counterfactual proposal posits content that cannot be known.

But I have a related worry: The proposal makes the wrong kind of factors epistemically relevant. Whether we are correct that the vignette constitutes a challenge to the target theory should depend on whether the vignette exhibits poor understanding of natural laws or faulty science, is conceptually confused, or relies on implausible claims. According to Williamson's proposal, by contrast, whether a successful-seeming vignette challenges a target theory

---

<sup>39</sup> Ichikawa and Jarvis (2009: 226; 2013: 202) and Ichikawa (2009: 440).

<sup>40</sup> Ichikawa (2009: 440).

depends on isolated quirks of the actual world. This does not seem correct: if a sophisticated thinker has devised a vignette, such as the sheep or Orlando vignette, and they have not made conceptual errors, or violated our understanding of the world, then they are in a position to know the falsity of the target theory, regardless of strange and unlikely happenstance. Even if the nearest world happens to be deviant, we still know that the vignette constitutes a challenge to the target theory. This is not to say that we can never be wrong about when a vignette constitutes a successful challenge; but Williamson's view generates the wrong sources of error.

Another weakness of Williamson's counterfactual proposal stems from the 'anemic' and 'botched' vignettes. Recall the anemic vignette,

ANEMIC VIGNETTE. Ted looks at a tree; he then forms a belief.

Claim (15) might well be true of the anemic vignette. It might be that given facts about this world, were Ted to believe something it would be justified true belief without knowledge. And so according to Williamson's counterfactual proposal it is possible—if the world were thus situated in modal space—that the anemic vignette presents a challenge to the target theory. But this cannot be right. The anemic vignette is not, and could not be, a challenge to the target theory. The vignette should not be a successful vignette, in other words, but according to the counterfactual view it might be. As I shall argue, the anemic vignette is not a challenge to the target theory because the most natural understanding of the vignette does not provide an instance of justified true belief without knowledge.

Recall that a ‘botched’ vignette fails to describe a counter-instance. Suppose that a botched vignette happens to describe what, in the nearest world where it obtains, is in fact a Gettier case. Williamson’s proposal says the botched vignette constitutes a challenge to the target theory. But that does not sound right.<sup>41</sup> Whilst the deviant realisations problem suggests that the counterfactual proposal is too demanding, the anemic vignette and botched vignette problems suggest it is too permissive: Judgement (11) is true in cases where the vignette should not be understood as a challenge to the target theory.

## 6. The Emendation Approach

The problems with the counterfactual proposal have lead theorists to suggest a different approach to understanding the vignette method. Grundmann and Horvath argue that the necessity proposal can be defended by amending the vignettes. Grundmann and Horvath focus specifically on Gettier cases. They argue that since deviant instances of Gettier-style vignettes can be sorted into two kinds—those where the protagonist fails to have justification, and those where the protagonist knows via some other route—there is a simple and systematic way to amend the vignettes.<sup>42</sup> They suggest theorists stipulate that the belief is justified. They argue that since theorists already usually stipulate that the proposition is believed and true, this amendment is benign. Then we stipulate that if the belief is known it is known only in virtue of the ways described by the vignette. This precludes independent sources of knowledge, such as that described by the ‘married Orlando’ augmentation.

---

<sup>41</sup> Proponents of the counterfactual view do have some resources available. They might argue that whilst the anemic or botched vignette is in fact a successful counterexample to the target theory, theorists lack grounds for thinking that it is. In the absence of good grounds, a theorist who (correctly) takes the anemic or botched vignette to challenge the target theory is making an error.

<sup>42</sup> Grundmann and Horvath (2014: 428).

Thomas Grundmann and Joachim Horvath illustrate this proposal with an example of a typical vignette in epistemology.<sup>43</sup>

FORD. Smith believes that Jones owns a Ford, on the basis of seeing Jones drive a Ford to work and remembering that Jones always drove a Ford in the past. From this, Smith infers that someone in his office owns a Ford. In fact, someone in Smith's office does own a Ford—but it is not Jones, it is Brown. (Jones sold his car and now drives a rented Ford.)

Grundmann and Horvath provide a version amended according to their proposal. Following their convention, I have underlined their modifications.<sup>44</sup>

AUGMENTED FORD. Smith justifiedly believes that Jones owns a Ford, on the basis of seeing Jones drive a Ford to work and remembering that Jones always drove a Ford in the past. From this belief alone, Smith logically infers, at time t, to the justified belief that someone in his office owns a Ford, which provides his only justification for that belief at t. In fact, someone in Smith's office does own a Ford, so that Smith's latter belief is true—but it is not Jones, it is Brown, and so Smith's initial belief was false. (Jones sold his car and now drives a rented Ford.) Also, if Smith knows at t someone in his office owns a Ford, then he knows this at t only in virtue of the facts described.

But the 'emendation proposal' has problems. Firstly, it is not how vignettes are actually presented, and so the proposal is perhaps best understood as guide to revising philosophical practice, rather than a proposal for understanding the existing method. Some of Grundmann and Horvath's comments suggest their intention is prescriptive rather than descriptive. They say, for example,<sup>45</sup>

---

<sup>43</sup> Grundmann and Horvath (2014: 525).

<sup>44</sup> Grundmann and Horvath (2014: 529–30). For ease of presentation, I collate the underlining from pages 529 and 530.

<sup>45</sup> Grundmann and Horvath (2014: 532).

Improved case descriptions might nevertheless help prevent philosophical laypersons from misinterpreting the relevant descriptions. In this way, an improved case description like [augmented Ford] might help to avoid potential confounds in experimental studies that test the Gettier judgments of philosophical laypersons.

If intervention, improvement, or prescription is the aim of the emendation proposal, it responds to a different question from the other proposals: Grundmann and Horvath are not asking how the method works, instead they ask how it could be improved. Accordingly the proposal is not a genuine alternative to the others.<sup>46</sup> But Grundmann and Horvath suggest that ‘augmented Ford’ is how experts implicitly understand the vignettes, and so their proposal is descriptive rather than revisionary. They write,<sup>47</sup>

something along the lines of [augmented Ford] seems like a good reconstruction of how the relevant experts implicitly interpret the original case description (the Ford vignette) when they perform the thought experiment. Arguably, no professional epistemologist would be seriously tempted to interpret [the Ford vignette] in any of the deviant ways that are left open by this description. Moreover, the way in which we amended the initial description (the Ford vignette) should be readily available and transparent to every professional epistemologist. It therefore seems quite plausible that professional epistemologists tacitly interpret [the Ford vignette] in a way that roughly corresponds to our improved description (augmented Ford).

Experts do not interpret vignettes as deviant variations. But is this because they respond to an emended vignette, as Grundmann and Horvath suggest, or because they simply interpret the given vignette in a more narrow and less literal way? There is reason to think the latter: Firstly, Gettier-style vignettes are almost never described in the emended way. Given how experts in philosophy are keen to make the implicit explicit and articulate each assumption carefully, this suggests that the amendments to the vignette are not required to successfully

---

<sup>46</sup> Ichikawa and Jarvis (2009: 224) makes a related point.

<sup>47</sup> Grundmann and Horvath (2014: 531–2). I have adjusted the way the Ford vignette is referred to, to better match my nomenclature.

elicit the Gettier judgement, and that emending the vignette as described by Grundmann and Horvath is not part of the method. Instead it seems that theorists respond to the vignette *as presented* (not the emended version), but do so in a more sophisticated way than simply quantifying over the instances where  $V(x, p)$ .

Secondly, laypeople judge in response to vignettes. These responses seem akin to expert responses. But it is unlikely that laypeople respond to the given vignette by ‘translating’ it into the amended vignette, even implicitly. Laypeople rarely use the expression ‘justified belief’, for example, and ‘in virtue of’ is a relatively technical notion.<sup>48</sup> The amount and content of emendation required varies by vignette in relatively sophisticated ways, furthermore. The above example uses time-indexed beliefs, ‘logically infers’, and other notions that different vignettes do not require. Explicit emendation is challenging.

This suggests that the emendation proposal does not capture what theorists and lay people do. It seems that rather than translating the vignette into the amended vignette, and responding to the amended vignette, people perform something different. In the next section I argue that they respond to the vignette itself, but they do so by interpreting it with good judgement by considering the normal version of the vignette.

Perhaps one of the most powerful and compelling features of the vignette method is that a simple, mundane story elicits a powerful response. The vignette contains few technical or explicitly epistemological details. Yet the response often is a distinct and confident epistemological judgement, one compelling enough to unseat established theories. One

---

<sup>48</sup> Caton (m/s).

weakness of Grundmann and Horvath's proposal is that it generates a more confusing and complicated vignette to respond to, and so arguably erodes the potency of our response. Grundmann and Horvath argue that since theorists often already stipulate that the proposition is true and believed, it is not damaging to also stipulate that the belief is justified. I demur. One reason the vignettes are so successful is that those very same (largely not explicitly epistemic) features that hearers recognise as making the belief justified nonetheless fail to make the belief known.<sup>49</sup> That hearers actively realise that the vignette describes a justified belief—rather than being passively told within the vignette—might well contribute to the efficacy of the method.

Finally, as Grundmann and Horvath acknowledge, the method does not readily generalise. Since the augmentation proposal is specific to Gettier-style cases, it does not generalise to vignettes as counterexamples to other necessity claims. It also does not illuminate the other roles that vignettes play in philosophical reasoning. As I argue in section 7.2, my proposal has the advantage of illuminating and explaining various roles of vignettes in philosophy.

## 7. The Normalcy Proposal

Rather than positing a simple quantification over the situations in which the vignette is true, or considering what is described in the most literal sense by a vignette, I propose that upon hearing a vignette, we instead think of what constitutes a normal, typical, generic, or natural instance of the vignette.

---

<sup>49</sup> Many thanks to Lisa Miracchi for emphasising this point.



When we reason, we often use the notions of typicality and normality. It is very natural to think things like ‘lying is usually wrong’, ‘we normally meet on Wednesdays’, and ‘Scandinavians typically buy their first home in their early thirties’. It is, I think, less typical to think in terms of strict necessity or mere possibility. Thinking in terms of normalcy and typicality are important and ubiquitous parts of both ordinary and scientific thought, and such thoughts are not easily translated or reduced to claims of probability, prevalence, counterfactuals, or other forms of quantification.

My proposal is that when we hear a vignette we consider the normal understanding of that vignette. If the normal versions of the vignette exhibit justified true belief without knowledge, then the vignette constitutes a challenge to the target theory. We judge something like,

- (16) Normally, whenever a person is related to a proposition such that the vignette obtains, then that person will have JTB without knowledge.

This is formalised below, where NORM denotes the operator ‘normally’:

$$(17) \text{ NORM } \forall x \forall p [V(x, p) \rightarrow (JTB(x, p) \ \& \ \neg K(x, p))]^{50}$$

---

<sup>50</sup> To emphasise the departure from the usual approaches to vignette satisfaction, we might also proceed using the apparatus ‘NORM (x, p)’ which conveys ‘x and p satisfy a/the normal understanding of the vignette’. We can then posit the judgement, ‘necessarily, anyone who stands to a proposition p as in the normal vignette has a justified true belief that p but does not know that p’. We can formalise this as ‘ $\Box \forall x \forall p [NORM(x, p) \rightarrow (JTB(x, p) \ \& \ \neg K(x, p))]$ ’.

Note that the judgement that the vignette is possible must also be adjusted from ‘the vignette could be true of someone’, ( $\Diamond \exists x \exists p V(x, p)$ ), to ‘normally, the vignette could be true of someone’ or ‘a normal version of the vignette could be true of someone’.

What we judge to be a normal instance of the vignette is influenced our background understanding about science, laws, politics etc.<sup>51</sup> It is also influenced by storytelling norms. As described in section five, if something surprising or abnormal happens, it would be mentioned by the vignette-teller. The omission from the vignette thus allows the hearer to infer that the more normal and expected version occurred. Given that it is a vignette about belief-formation, if the protagonist has another route to knowledge, for example, it would be mentioned by the vignette-teller. But this inference occurs without the hearer explicitly thinking about things like knowledge or storytelling norms; it is simply how we in fact understand anecdotes, explanations, and stories.

The normal understanding of a vignette does not simply interpret the vignette as though it were a fiction, as suggested by Ichikawa and Jarvis. Ichikawa and Jarvis suggest that when we consider a successful vignette our response is:<sup>52</sup>

$$(18) \quad \Box (p \rightarrow \exists x (x \text{ has justified true belief without knowledge}))$$

---

<sup>51</sup> Given the importance of background beliefs in judgements concerning the normal or typical understanding of a vignette, there are parallels with modelling non-monotonic logics, default logics, and other attempts to model natural thought and *ceteris paribus* clauses.

<sup>52</sup> Ichikawa and Jarvis (2013: 207; 2009: 229).

where  $p$  stands for the proposition ‘the members of *story* are true’. *Story* is the set of sentences that are true in the fiction. This is an infinite set, and is based on the given vignette and our interpretation of it. The interpretation provides something richer than the mere vignette, and is constructed using our ability to recognise truth in fiction. Ichikawa and Jarvis provide an example: it is not mentioned in the vignette that Orlando has two eyes, for example, but it is true in the fiction that he does.<sup>53</sup> So if someone,  $S$ , has one eye, then it is possible for  $V(x, p)$  to hold for  $S$ , but not possible for Ichikawa and Jarvis’s ‘true in the fiction’ relation to hold for  $S$ .

Whilst our proposals share some features, I have several concerns about the truth-in-fiction proposal: Firstly, it is hard to see how this view captures agreement and disagreement about vignettes. This is because every hearer will round out the vignette slightly differently. Since they create different stories from the vignette, perhaps their judgements (18) will not have the same content. When hearers seem to agree and disagree about vignettes, in other words, they might have different propositions in mind, which would render agreement or disagreement illusory.<sup>54</sup> The normal version of the vignette, by contrast, supplies fewer details and is therefore interpersonally available: people can agree on the normal version of a vignette, even if they do not agree on the fictional version.

---

<sup>53</sup> Ichikawa and Jarvis (2009: 227).

<sup>54</sup> See also Williamson (2009: 468). The same concern arises for Ichikawa and Jarvis’s ‘demonstrative’ suggestion; the concern is that theorists might be picking out different things by ‘that’, rendering agreement illusory (Ichikawa and Jarvis (2009: 232)).

Secondly, it is not clear how we produce the set *story*.<sup>55</sup> We do not seem to generate a genuine full-blooded fiction from hearing vignettes: when we hear a trolley case, for example, we do not imagine details concerning what kind of train it is, where it is going, or how the driver is feeling (unless we deem these things relevant to either the normal version of the vignette or the philosophical question at stake).<sup>56</sup> The normalcy approach does not require that we fill out these kinds of details; according to the normalcy approach we supply some extra details in addition to the vignette, but the case remains largely schematic. The view seems to posit a relatively complicated operative judgement, furthermore, since it appears to invoke the notion of sets and use a meta-language.

Thirdly, the proposal seems too narrow: some vignettes describe acknowledged non-fiction, such as anecdotes or historical events. It seems strange that we would interpret these using our capacity to understand fictions, as the truth-in-fiction proposal suggests. The normalcy approach, by contrast, captures how we interpret both fictional and non-fictional vignettes.

Fourthly, the proposal seems too broad: fictions, and our interpretation of fictions, can be inconsistent or conceptually incoherent. Ichikawa and Jarvis suggest that philosophical vignettes are a special genre of fiction, one that ensures no inconsistent or incoherent fillings-out obtain.<sup>57</sup> They thereby narrow down the content of the judgement. But by adding these kinds of restrictions, the ‘truth-in-fiction’ operator begins to resemble simply *understanding* the vignette, rather than a specific kind of fictional interpretation as such.

---

<sup>55</sup> See also Williamson (2009: 467–8).

<sup>56</sup> Williamson argues that the truth-in-fiction proposal is psychologically implausible because it suggests that we imagine the vignette, via daydreaming and imagery (Williamson (2009: 467)).

<sup>57</sup> Ichikawa and Jarvis (2013: 211, n. 25).

Thus whilst my view shares features with Ichikawa and Jarvis's proposal it is less wedded to the notion of truth-in-fiction, and accordingly avoids their proposal's weaknesses. The supplementary details generated by the normal understanding of the vignette 'fill out' the vignette—and provide more detail than simply  $V(x,p)$ —but the normal understanding does not demand as much filling out as the truth-in-fiction proposal.<sup>58</sup>

Judgements about normalcy are related to counterfactual judgements. Indeed, our understanding of what is normal, typical, or generic affects our judgements about counterfactuals, and is plausibly explanatorily prior to them. We can judge counterfactuals in virtue of judgements about normalcy. The normal understanding of a vignette will often approximate what would obtain were the vignette true. The judgements differ only when either the unexpected or abnormal is a modally nearby possibility, such as when—surprisingly—the actual world turns out to be a gullible Orlando world, rather than a world where were the Orlando vignette satisfied, the normal version would obtain. Or where our judgement of what is typical and ordinary is in error, such as when our scientific theories are mistaken.<sup>59</sup>

### 7.1 Virtues of the Normalcy Proposal

Since whether something constitutes the normal instance of a vignette is not subject to unexpected contingent anomalies in the actual world, the normalcy approach avoids the

---

<sup>58</sup> For further criticisms of the truth-in-fiction proposal, see Williamson (2009) and Malmgren (2011).

<sup>59</sup> In section 7.4 I return to the question of to what extent the truth conditions of the 'normally' operator are constrained by reality, and to what extent they are influenced by our beliefs, including false beliefs.

‘deviant realisations’ problem faced by Williamson’s counterfactual approach. Even after we have learnt of a deviant actual or modally nearby realisation, we might well disregard it as an abnormal version of the vignette. The abnormal can be actual. This is illustrated by the class discussing the sheep vignette and its variants. What is normal—and our judgements concerning what is normal—has stability. It does not fluctuate as a result of quirky one-off features of the world.<sup>60</sup> And so the normalcy proposal captures how the potency of a given vignette is not hostage to random quirks of the world.

What qualifies as a normal version of a vignette is somewhat stable. It is not contingent on fluky isolated facts about the world, as suggested by Williamson’s counterfactual approach. But it will be influenced by, for example, changes in systematic and familiar features of the world. What is normal can change over time. Suppose in the future there is rampant bookseller fraud, for example, perfect replication technology allows fakers to create undetectable forgeries, there is inadequate regulation in the market, and much demand for old books. Were the Orlando vignette told in this future, the normal version of the Orlando vignette would not be one where his belief is justified. In this unfortunate future, the Orlando vignette as stated does not provide compelling evidence against the target theory. (The vignette-teller would at least have to augment the vignette to motivate the thought that Orlando was justified in believing the bookseller. Perhaps, for example, by adding that this bookseller was widely regarded as honest.) In this respect the account captures both the modal stability and the contingency of the vignette method.

---

<sup>60</sup> This is not to say that we can never be wrong about which vignettes challenge the target theory. (Williamson (2009) emphasises that we must countenance the fact that we can be wrong about this kind of philosophical reasoning, and accuses Ichikawa and Jarvis’s view of having the result that we are too often correct.) We might think that a normal version of a vignette exhibits justified true belief without knowledge, but be mistaken.

Ichikawa and Jarvis argue for, and I think overstate, the non-contingency of the success of a given vignette. They write,<sup>61</sup>

It is not right to think of Williamson's project as a defense of a traditional understanding of philosophical methodology, for the judgments with which he identifies 'intuitions' are scarcely recognizable as the things traditional philosophers had in mind. In particular, on Williamson's view, the propositional contents of the intuitions in question are contingent. This is problematic, from a traditional standpoint [...] tradition has it that intuitions like the Gettier intuition have necessarily true contents; Williamson's counterfactuals, however, will vary in truth value between possible worlds.

But, surely the judgement—and corresponding success—for a particular thought experiment, should depend on contingent facts; consider the Orlando vignette told in a world with widespread book forgery. Williamson (2007; 2009) by contrast argues for, and I think misattributes, the contingency. He thinks quirky, isolated, unknown features can render a vignette impotent. I think the correct view is that facts about the actual world should only affect the success of a vignette insofar as they influence the truth of claims about normalcy, such as claim (16).

Unlike the necessity proposal, the normalcy proposal does not quantify over all instances that satisfy the vignette. It narrows down to only those instances that are normal versions of the vignette. And, as I explain in section 7.4, variations of the normalcy proposal narrow down the quantification even further, to just some normal variations, a normal variation, or the most normal variations. Thus the proposal avoids the 'deviant variations' problem that

---

<sup>61</sup> Ichikawa and Jarvis (2009: esp. 223). See also Ichikawa and Jarvis (2013: chapter nine).

plagued the necessity proposal. The deviant variations all involve abnormal understandings of the vignette. Whenever the vignette constitutes a compelling problem for the target theory, the normal version (or most normal versions) of the vignette exhibit a counter-instance to the target theory.

The normalcy approach also avoids the ‘anemic vignette’ and ‘botched vignette’ problems. These problems afflict the possibility and counterfactual proposals because, according to these views, patently inadequate vignettes could qualify as challenges to the target theory. The normalcy proposal avoids this problem because if the vignette presented fails to be one where the normal version of that vignette is justified true belief without knowledge, then judgement (16) is false about that vignette.<sup>62</sup> The proposal thus explains why these vignettes could not constitute a challenge to the target theory.

Unlike the possibility proposal, the normalcy proposal does not require only that some variation of the vignette is a counter-instance. It must be a normal instance. Since we are typically first struck by the normal understanding of the vignette, this explains why we are struck by the problem instance, and so why the method is so compelling. Unlike the possibility proposal, which faces trouble explaining the forcefulness of the method, the normalcy proposal can explain why the vignettes are potent: the normal understanding of them presents us with a counter-instance.

---

<sup>62</sup> Or ‘a’ or ‘any’ normal version—I return to the question of quantification in section 7.4.



Williamson notes that one virtue of his counterfactual view is the continuity between thought experiment judgements and normal everyday judgements.<sup>63</sup> His view renders thought experiment judgements non-mysterious, because it understands them as a kind of ordinary everyday widely-recognised judgement. The normalcy proposal also enjoys this advantage. Judgements about normalcy are ubiquitous and mundane. They are more familiar to us than judgements about what is a possible or necessary interpretation of a vignette. Normalcy judgements might be even more basic than counterfactual judgements, since plausibly we require the former in order to make the latter.

The normalcy proposal has some similarities with the emendation proposal. The emendation proposal aims to reduce the kinds of variations of the vignette that are relevant to the judgement, and does so by amending the vignette with additional content. The normalcy proposal also restricts which variations of the vignette are relevant, by saying that only the normal variations are relevant. The normalcy proposal, unlike the emendation proposal, is not revisionary about the method. Instead, I posit, it simply describes what people actually do when they hear a vignette. The content is simpler: There was a worry that the experts do not consider the vignettes as described by Grundmann and Horvath. If they did, they would have made the amendments explicit. And there was a worry that the folk do not either: they do not possess the specific conceptual resources to easily think of the vignettes as proposed by Grundmann and Horvath. But laypeople use the vignette method with ease. The content posited by the normalcy proposal is simple. Indeed, considering the normal version of something is a hallmark of normal everyday thinking.

---

<sup>63</sup> Williamson (2007: 192).

Another virtue of the normalcy proposal is it provides resources to explain some cross-cultural variation in response to vignettes. Weinberg, Nichols and Stich (2011) claim that Asian respondents judge that the protagonist possesses knowledge in the ‘American car’ Gettier vignette, for example, whereas American respondents judge the protagonist lacks knowledge. Perhaps Asian respondents fill in as a normal background condition that someone who once drove an American car will continue to do so, and so can be known to drive an American car. This background assumption affects their judgement of the normalcy claim. My proposal accords with this kind of explanation of cross-cultural variation.<sup>64</sup>

## 7.2 A Unified Account

Unlike Grundmann and Horvath’s proposal, the normalcy proposal readily generalises to other vignettes as counterexamples, and other roles that vignettes play in philosophical theorising. Firstly, the emendation suggestion was particular to Gettier cases, and so does not readily generalise to vignettes as counterexamples against other necessity claims. The normalcy proposal readily generalises to other vignettes as counterexamples. But additionally, the normalcy proposal offers a unified account of other roles that vignettes play. Vignettes can motivate claims and principles, illustrate distinctions, communicate ideas, help pin down the phenomena of interest, and frame research questions. And they provide evidential weight against claims that are not, or do not entail, necessity claims. An account of the judgement is better if it unifies these roles. In what follows I illustrate ways the normalcy account can better capture the range of roles that vignettes play.

---

<sup>64</sup> See also Sosa (2009).

It is hard to see how the possibility proposal could explain how vignettes help to build a theory, motivate a principle, or illustrate a distinction. If a theorist says that ‘F’s are G’s’, for example, and provides a vignette, it does not much motivate the theorist’s claim if hearers thereby judge only that some possible version of the vignette exhibits an F that is G. Providing a vignette and judging that in *normal* instances of the vignette F’s are G’s is more compelling. It is a good starting place for developing a theory.

Similarly it is hard to build a theory or begin to motivate claims if vignettes, to be potent, must be such that every variant is an instance of an F that is a G. The necessity proposal, by requiring this, renders vignette construction too demanding to play theory-building roles.

Whilst the discussion thus far focuses on vignettes as a challenge to necessity claims, vignettes are also employed to challenge claims about what is typical, normal, characteristic, generic, or usual. A theorist might claim, for example, ‘typically F’s are G’s’, and we might put pressure on that claim by providing an everyday, mundane vignette in which F’s are not G’s. The more everyday and normal the vignette, the better. Suppose a theorist claims that ‘typically testimonial beliefs from strangers are unjustified’. An interlocutor replies with a vignette about a man, Fred, visiting Manhattan who politely asks for directions to the Flatiron Building. The woman he asks, Jenny, is a local and confidently provides directions. This challenge to the target claim is successful, in large part because it is an ordinary, mundane circumstance, and the ordinary, typical instances of the vignette exhibit justified belief. A vignette about an unusual occurrence would be less successful against the target claim. If the vignette stipulated that the woman was Manhattan’s chief cartographer and lives in the Flatiron building, the vignette would be less successful against the target claim. And if

the initial vignette elicits the judgement only that some possible variation of the vignette — such as the cartographer augmentation—exhibits justified belief, this is also less potent against the target typicality claim.

I submit that this kind of dialectic—where a vignette challenges a claim about what is generic, characteristic, normal, or typical—is common in philosophy. But the existential proposal has trouble vindicating it: Some possible variant of a vignette’s being a counter-instance would not harm the target claim. In order to challenge a claim about what is typical, generic, normal, or characteristic, the vignette-based judgement must concern a normal instance of a normal vignette.

There is some debate in philosophy about whether philosophical inquiry concerns conceptual facts, metaphysical facts, some other kind of fact, or a combination. This meta-philosophical debate concerns, for example, whether we are investigating what knowledge is or the concept of knowledge. Williamson notes that his favoured proposal, the counterfactual proposal, is ill-suited where philosophy claims concern conceptual, rather than metaphysical, possibility.<sup>65</sup> He subsequently argues that the metaphysical question is of interest to philosophers, and the conceptual question is not. One virtue of the normalcy approach is its applicability to both families of questions. If the target claim (3) concerns the metaphysical nature of knowledge, we can investigate it by thinking of the typical, normal versions of vignettes. If the target claim concerns the concept *knowledge*, we can also investigate it by thinking of normal versions of the vignette, since what we take to be normal

---

<sup>65</sup> Williamson (2007: 205–7).

is influenced by our conceptual scheme; our judgements about normalcy provide a window into our concepts.

### 7.3 Resistance to the Proposal

There has been resistance to this kind of proposal. Williamson discusses the problem of deviant variations that plagues the necessity proposal. He says in a footnote,<sup>66</sup>

Merely adding the stipulation that *x* and *p* constitute a normal instance of the Gettier case is unlikely to solve the problem, for the relevant notion of normality is an epistemological one that violates the supposed neutrality of the initial description of the case.

By ‘neutrality’ he means that in telling the vignette we do not bias the judgement with regard to the target claim. In the case of Gettier vignettes, for example, it is not made explicit whether the protagonist knows or is justified. Williamson writes,<sup>67</sup>

Suppose that we fix on a particular Gettier-style story (the one about Orlando would do), henceforth ‘the Gettier case’, told in neutral terms, without prejudice to the target analysis. For instance, it is not explicitly part of the story that Orlando does not know that he owns a book which once belonged to Virginia Woolf.

Malmgren uses the term ‘neutrality’ in a similar way. She writes,<sup>68</sup>

[A vignette] exhibits a certain neutrality—it does not specify that the subject does, or does not, know the relevant proposition; nor does it specify that she does, or does not, have a justified true belief in it. (But note that it *does* specify that she truly believes it.) More generally, a canonical description of a problem case is neutral with respect to the distribution of the ‘test properties’—the properties whose modal connection is at issue in the thought experiment: it does not explicitly stipulate that the test properties are or that

---

<sup>66</sup> Williamson (2007: 204, n. 22).

<sup>67</sup> Williamson (2007: 184). This is all Williamson says about neutrality in Williamson (2007).

<sup>68</sup> Malmgren (2011: 274–5, emphasis in original).

they are not instantiated in the given case (nor does it stipulate anything that transparently entails that they are/are not instantiated).

It is not clear why asserting that it is a normal case of the vignette violates neutrality understood in this way. Calling it a normal instance of a vignette does not in general suggest whether the person has a justified belief or knowledge. It is still left to the hearer to determine these facts.<sup>69</sup>

Malmgren also considers and dismisses a view similar to the normalcy proposal. Like Williamson, she reflects on the problem of deviant variations for the necessity proposal. She considers weakening the content of the necessity proposal, whilst retaining the necessity operator, by specifying that the relevant subset of cases is narrower than the set of cases where  $V(x, p)$ . Malmgren writes,<sup>70</sup>

There are different ways to elaborate on this suggestion, but the basic thought is just to weaken [the judgement proposed by the necessity proposal, (11)] by restricting the scope of the necessity operator to some specific subset of possible worlds that satisfy the given case description [...] Let us use ‘the *intended* Gettier case’ as a placeholder name for that subset of worlds. We can then express the suggestion by saying that the content of the Gettier judgement is the claim that,

- (18) Necessarily, anyone who stands to a proposition  $p$  as in the intended Gettier case has a justified true belief that  $p$  but does not know that  $p$

Malmgren then criticises this idea,<sup>71</sup>

---

<sup>69</sup> Note that Grundmann and Horvath’s proposal does violate neutrality with regard to whether the belief is justified. Sosa (2006: 642) suggests neutrality may not be an important feature of the vignette method.

<sup>70</sup> Malmgren (2011: 285, emphasis in original).

<sup>71</sup> Malmgren (2011: 286–7).

First, (18) needs to be spelled out in more detail—in particular, the placeholder (‘the intended case’) must be eliminated—and this turns out to be extremely difficult [...] As stated, (18) is way too schematic; we cannot properly evaluate it [...] We need to be provided with a more informative characterization of the set of worlds that makes up the intended Gettier case [...] a characterization that captures all and only non-deviant realizations of the Gettier case, but that does not itself contain terms like ‘deviance’ and ‘intended’ [...] But it turns out to be very difficult to provide a characterization that meets these constraints.

Malmgren argues that formulation (18) must be explained more. She argues that advocates of (18) must translate ‘intended’ into new terms, and may not use terms such as ‘intended’ or ‘deviant’. Presumably ‘normalcy’, ‘typicality’, ‘generically’ and similar terms are also prohibited. Malmgren argues that this constitutes a too difficult task, and resembles the problem of *ceteris paribus* clauses.<sup>72</sup>

I agree the formulation demands elaboration. An advocate of the view would need to say more about the notion ‘intended case’, such as what kinds of details are included in the intended case that are absent from the articulated vignette itself, and whether the intended case can include inconsistencies or conceptual incoherence. But I disagree that this elaboration requires an elimination of the ‘intended case’ placeholder where the illumination cannot include terms such as ‘deviant’, ‘normal’, ‘intended’ etc.

Even if we were committed to thinking those terms cannot be used in descriptions of natural phenomena, such as in chemical and physical theories, the content question concerns how we *think*.<sup>73</sup> Thinking often uses ideas such as deviance, intended, normally, typically, *ceteris*

---

<sup>72</sup> Malmgren (2011: 289).

<sup>73</sup> I expect that it is also mistaken to try to understand many other phenomena without using these notions. Arguably we cannot understand causation without these kinds of notions. Perhaps too we cannot understand many natural phenomena, such as biological

paribus, generically, and characteristically. We do not need to reduce or eliminate these notions in order to think well. These ideas are successfully employed in thought; understanding things as intended, normal or deviant instances of a kind is an important and apparently irreducible aspect of thought. It would be misguided to attempt to capture the content of thought in general—and the responses to vignettes in particular—without employing these notions.

I do not think ‘intended’ is the right modifier, and so do not endorse (18) as the correct content. This is because I do not think speaker intentions are dispositive about how we should understand the vignettes. (There are many reasons to think speaker intentions should not play such a large role. One is that if the speaker botches recounting a vignette, so that the vignette does not describe justified true belief without knowledge, the recounted vignette should not qualify as successful just because the vignette that the speaker intended to tell would have exhibited justified true belief without knowledge.) But I think that Malmgren’s reasons for rejecting the view are unconvincing.

Similarly Ichikawa and Jarvis express reservations about this line of thought. They write,<sup>74</sup>

Richard Heck has proposed to us in conversation another move on behalf of Williamson. Instead of a counterfactual, one could render [the judgement] as a generic like ‘Fire engines are red,’ which is true not because all fire engines are red (we’ve seen some green ones), but because fire engines typically are red. [The judgement] then becomes: in cases that satisfy the Gettier text, people (typically) have justified true belief

---

and psychological kinds, without these notions. Thus we should not be averse to them within science, let alone in an account of human thought; if we should allow these kinds of notions within scientific theories, they should certainly be permitted in theories about the content of responses to vignettes.

<sup>74</sup> Ichikawa and Jarvis (2013: 204, n. 12).



without knowledge. The conclusion we arrive at, then, is that typical cases that satisfy the Gettier text are ones in which people have justified true belief without knowledge. We find this approach plausible, although we're hesitant to rest too much on such generics, as their interpretation is a vexed issue [...] It may be that this approach is ultimately similar to the one we will go on to develop.

I am skeptical that the fact that generics are a vexed issue provides much reason to avoid using them in an account of the content of our judgements. Firstly, we are able to reason well using generics, *ceteris paribus* clauses etc., despite lacking a full account of them. And the project is to provide an account of *reasoning*. We do not, in general, need an account of something for it to play a successful role in thought. Secondly, people are not skeptical about the existence of something akin to generics. I can understand avoiding positing something where there is some contention about its existence. But in this case theorists typically assume that generics exist but debate their semantics. Thirdly, often in philosophy we make use of ideas in our theories without fully having an account of them. We commonly use notions such as luck, belief, causation, and truth, despite their natures being controversial. Provided we understand them well enough for our purposes, and are not accidentally using an incoherent or misleading notion, this is unproblematic.

Thus whilst there has been some resistance to this kind of suggestion, I do not think this resistance is well-grounded.

#### 7.4 Decision Points

The normalcy proposal as stated generates some choice points—places where differing theories can be developed from the basic idea. Perhaps the most obvious concerns the relevant quantification. The normalcy proposal demotes the importance of which situations are literally satisfied by the vignette, and instead focuses on how we *understand* vignettes. We

understand them via their normal and typical instances. When I introduced the normalcy proposal I used universal quantification, and so offered an account that—in that regard—resembles the necessity proposal. I characterised the normalcy proposal as:

(16) Normally, whenever a person is related to a proposition such that the vignette obtains, then that person will have JTB without knowledge.

(17)  $\text{NORM } \forall x \forall p [V(x, p) \rightarrow (JTB(x, p) \ \& \ \neg K(x, p))]$

But universal quantification is not essential to the view. The view could be understood using something akin to the existential quantification of the possibility approach, by positing that the judgement is ‘some normal variations of the vignette exhibit justified true belief without knowledge’. According to this proposal, ‘normally’ is understood as described above, but the judgement is true even if some normal variations of the vignette are not counter-instances to the target claim.

A third alternative uses the formulation ‘most normal variations of the vignette exhibit justified true belief without knowledge’, where ‘most’ is understood quantitatively, i.e. meaning ‘more than half’ or some high proportion. This is akin to formulations that employ the possible worlds framework and say ‘most nearby possible worlds’. We might instead posit that the judgement is ‘The most normal variations of the vignette exhibit justified true belief without knowledge’, where ‘the most normal’ is understood qualitatively. According to this variation of the view, versions of the vignette can be ordered, with some being more normal than others, and the ones relevant to the judgement are the most normal ones. Or we might posit that the judgement is simply ‘normal variations of the vignette exhibit

justified true belief without knowledge', where the advocate proposes that we have a good grasp of what this means independently from a detailed spelling out.

Perhaps too advocates of the counterfactual proposal can adopt a subjunctive version of the normalcy approach. The proposed content would be something like, 'were the vignette satisfied normally, x would have a justified true belief, but not know, that p'. Employing Lewis's possible worlds framework for interpreting counterfactuals yields, 'in the nearest possible worlds in which the vignette is satisfied in a normal way, x has a justified true belief that p, but does not know p'. This proposed content uses the structure of the counterfactual proposal, but content from the normalcy proposal.

The appropriate quantification depends partly on how we understand 'normal'. If 'normal version of a vignette' denotes a narrow range of instances, then 'every normal variation' will be a more appropriate quantification. Conversely if 'normal version of a vignette' allows wider variation, then something like the quantitative 'most normal variations' or the qualitative 'the most normal variations' might be more appropriate.

One strategy to adjudicate amongst these quantification options is to consider the unsuccessful vignettes. If a vignette does not provide a challenge to the target theory simply because some normal variations of the vignette are not counter-instances, this suggests that the universal quantification 'all normal variations' plays a role in our judgement. If the vignette only fails when no normal understanding of the vignette provides a counter-instance, this suggests that the relevant judgement is 'some normal variations'. And if the vignette fails when some normal variations, but not the most normal variations, constitute a

counter-instance, this suggests the judgement is something like ‘the most normal variations’, which picks out the qualitatively most normal.

A second decision point concerns what to say about atypical vignettes. If a vignette is too weird, plausibly there is no normal version of the vignette.<sup>75</sup> This raises the question of how judgment (16) applies to ‘far out’ vignettes. One route maintains that the most normal versions are relevant, even if they are still somewhat abnormal. This is akin to when modally distant possibilities are raised in the antecedent of subjunctives, and so we must consider very distant possibilities to evaluate the subjunctive.<sup>76</sup> We might evaluate what would normally happen if a person were a reliable but unreflective clairvoyant, for example, even whilst recognising that this is a deeply abnormal situation. An alternative route maintains that if there is no normal understanding of the vignette, the vignette cannot present a challenge to the target theory. Advocates of this approach might argue that weird vignettes cannot be well understood, and so lack dialectical force. Alternatively one could advance some combination of strategies: the judgement concerns the most normal version of the vignette, even if somewhat abnormal, until a threshold after which, since there is no sufficiently normal version of the case, judgement (16) cannot be true of the vignette and so the vignette does not pose a challenge to the target theory.

This decision will be influenced, in part, by whether employing weird cases should be vindicated by an account of the method. If we independently think weird vignettes are methodologically suspect, then the normalcy proposal provides resources for explaining

---

<sup>75</sup> Thanks to Bob Beddor, David Black, Will Fleisher, and Ernest Sosa for discussion of these issues.

<sup>76</sup> Lewis (1973).

why. (Needless to say, a full treatment of this topic requires articulating different ways a vignette might be ‘weird’ or ‘abnormal’, and which kinds of abnormality are potentially problematic for philosophical theorising.)

Another decision point concerns to what extent, and in what ways, ‘normalcy’ is influenced by our epistemic position, and how much it is simply a function of how the world in fact is. To illustrate this distinction, consider this case,

Sarah is teaching an epistemology class. She intends to explain how knowledge and justified true belief can diverge and tells a vignette. She says,

‘Joey walked into a room. He saw something lush and verdant in a pot and so formed a belief that there was a plant on the table. In fact what he looked at was an abstract sculpture that resembles a plant. Unbeknownst to him, however, there was a mushroom growing in a pot behind the pile of books on the table. And so by chance, Joey’s belief was true.’

Sarah’s class are impressed. They all judge that knowledge and justified true belief can indeed diverge, and proceed to have an intelligent and thoughtful conversation about different ways to modify an account of knowledge to overcome the problem posed by the vignette.

Here is the rub: Sarah—the vignette-teller—is mistaken. Mushrooms are not plants.<sup>77</sup> But no one in the class knew; few people know mushrooms are not plants. (We can vary the vignette: Sarah could have made a mistake that *nobody*, scientists included, realises is mistaken.) So, is Sarah’s vignette a successful Gettier case?

If the answer is no—Sarah’s case does not provide a successful challenge to the target theory of knowledge—this suggests that what is normal is an objective, mind-independent feature

---

<sup>77</sup> Phylogenetically, mushrooms are closer to humans than to plants.

of the world. There is no normal variation of the vignette where Joey has a justified true belief without knowledge, because there is no normal version of the vignette where his belief is true. ('How can it be normal if it is biologically impossible?', the thought goes.)

If the answer is yes, this suggests that a normal version of a vignette can be in part determined by what we *take* to be normal. (Perhaps 'normal' is understood as something epistemic, such as 'normal understanding'.) The thought might continue that it is a persuasive vignette, the students are not being irrational, and if everyone thinks that mushrooms are plants that is sufficient for legitimating the dialectic.<sup>78</sup>

Perhaps both views are partially correct: perhaps Sarah's vignette is problematic in some ways, and would be improved if it reflects biological fact. But the vignette is not valueless; it has some dialectical force. Perhaps a correct understanding of 'normalcy', reflecting both objective and subjective factors, can capture this.<sup>79</sup>

Thus there is some flexibility in the view advanced. I hope to have motivated the view that the normalcy proposal merits further investigation and development. I have explained how the view avoids the problems faced by competing proposals, and have articulated its virtues.

---

<sup>78</sup> If what is normal is partly a function of expectations, this raises the question of how socially determined the expectations and normalcy operator are, versus how private and individual. It should not be wholly private, I think. A person narrating an abysmal or botched vignette to himself should not qualify as a challenge to a target theory just because he takes normal versions of it to exhibit a counter-instance. Even if 'normalcy' is partly a function of expectations and beliefs, there should be some objective or social constraints on what qualifies as a normal instance of a vignette.

<sup>79</sup> A related decision point is to what extent, and in what ways, the normal version of a vignette is influenced by speaker intention. It is an open question in the philosophy of fiction whether, and how, the intentions of the author affect what is true in a given fiction.

The normalcy proposal provides a unified account of the role that vignettes play in philosophical theorising, and captures how we normally think. I suggest the normalcy view best captures the content of the operative judgements we have in response to vignettes.<sup>80</sup>

### Acknowledgements

Many thanks to Bob Beddor, David Black, Will Fleisher, Jon Garthoff, Jonathan Schaffer, Susanna Schellenberg, and Ernest Sosa for comments on an earlier draft of this paper. Thanks also to Frances Egan, Thony Gillies, John Greco, Lisa Miracchi, and Daniel Rubio for helpful discussion, and to audiences at Rutgers University and at the Central European University.

### Bibliography

- BonJour, Laurence (1980) 'Externalist Theories of Empirical Knowledge' *Midwest Studies in Philosophy* 5: 53–73.
- Caton, Jacob (ms) 'Is "Justification" an Ordinary Term?'
- Chisholm, Roderick (1977) *Theory of Knowledge*, Englewood Cliffs, NJ: Prentice-Hall
- Constant, Benjamin (1797) 'On Political Reactions' *France* 5 (1).
- Elgin, Catherine (2014) 'Fiction as Thought Experiment' *Perspectives on Science* 22: 55–75.
- \_\_\_\_\_ (1996) *Considered Judgement* Oxford University Press.
- Gardiner, Georgi (2015a) 'Normalcy and the Contents of Philosophical Judgements' *Inquiry* 58(7): 700–740.
- \_\_\_\_\_ (2015b) 'Teleologies and the Methodology of Epistemology' *Epistemic Evaluation: Purposeful Epistemology* John Greco and David Henderson (eds.) Oxford University Press, pp. 31–45.
- Geddes, Alexander (forthcoming) 'Judgments about Thought Experiments' *Mind*.
- Gettier, Edmund (1963) 'Is Justified True Belief Knowledge?' *Analysis* 23: 121–3.
- Grundmann, Thomas and Joachim Horvath (2014) 'Thought Experiments and the Problem of Deviant Realizations' *Philosophical Studies* 170: 525–33.
- Ichikawa, Jonathan (2009) 'Knowing the Intuition and Knowing the Counterfactual' *Philosophical Studies* 145: 435–43.
- Ichikawa, Jonathan and Benjamin Jarvis (2013) *The Rules of Thought* Oxford University Press.
- \_\_\_\_\_ (2009) 'Thought-experiment Intuitions and Truth in Fiction' *Philosophical Studies* 142: 221–46.
- Jackson, Frank (1982). 'Epiphenomenal Qualia' *Philosophical Quarterly* 32: 127–36.
- Kitcher, Philip (2000) 'A Priori Knowledge Revisited' *New Essays on the A Priori* Paul Boghossian and Christopher Peacocke (eds.) Clarendon Press, pp. 65–91.
- Kripke, Saul (1980) *Naming and Necessity* Harvard University Press.

---

<sup>80</sup> This essay is published as Gardiner (2015a). Alexander Geddes (forthcoming) independently advances a related idea.

- Lehrer, Keith (1990) *Theory of Knowledge* Westview Press.
- Lewis, David (1973) *Counterfactuals* Harvard University Press.
- Malmgren, Anna-Sara (2011) 'Rationalism and the Content of Intuitive Judgements' *Mind* 120: 263–327.
- Pritchard, Duncan (2012) 'Anti-Luck Virtue Epistemology' *Journal of Philosophy* 109: 247–79.
- Sosa, David (2006) 'Scepticism About Intuitions' *Philosophy* 81: 633–48.
- Sosa, Ernest (2009) 'A Defense of the Use of Intuitions in Philosophy' *Stich and His Critics*, Dominic Murphy and Michael Bishop (eds.) Wiley-Blackwell, pp. 101–112.
- Weinberg, Jonathan, Shaun Nichols and Stephen Stich (2001) 'Normativity and Epistemic Intuitions' *Philosophical Topics* 29: 429–60.
- Williamson, Timothy (2009) 'Replies to Ichikawa, Martin and Weinberg' *Philosophical Studies* 145: 465–76.
- \_\_\_\_\_ (2007) *The Philosophy of Philosophy* Blackwell.
- \_\_\_\_\_ (2005) 'Armchair Philosophy, Metaphysical Modality and Counterfactual Thinking' *Proceedings of The Aristotelian Society* 105: 1–23.



## VI. The Burden of Proof and Statistical Evidence

**Abstract.** In order to perform certain actions—such as incarcerating a person or revoking parental rights—the state must establish certain facts to a particular standard of proof. These standards—such as preponderance of evidence and beyond reasonable doubt—are often interpreted as likelihoods or epistemic confidences. Many theorists construe them numerically; beyond reasonable doubt, for example, is often construed as 90 to 95% confidence in the guilt of the defendant.

A family of influential cases suggests standards of proof should not be interpreted numerically. These ‘proof paradoxes’ illustrate that purely statistical evidence can warrant high credence in a disputed fact without satisfying the relevant legal standard. In this essay I evaluate three influential attempts to explain why merely statistical evidence cannot satisfy legal standards.

### 1. The Golden Thread

The presumption of innocence is a keystone of many criminal justice systems. In criminal trials, the defendant is formally considered ‘innocent until proven guilty’. The burden of proof is on the state to prove the defendant committed all elements of the crime; the defendant, by contrast, need not prove his innocence. This asymmetry was famously dubbed ‘the golden thread’ of criminal law.<sup>1</sup> Correspondingly in many legal systems, such as Anglo-

---

<sup>1</sup> Lord Sankey in *Woolmington v DPP* (1935) AC 462.

American common law, the defendant in civil cases enjoys a presumption of non-culpability; the burden falls to the plaintiff to establish the facts of the case.<sup>2</sup>

A legal burden of proof generates a formal standard of proof. This is the standard to which the facts must be established to satisfy the burden. Formal standards include the ‘reasonable suspicion’ standard in the US, which is required for a brief stop and search.<sup>3</sup> A slightly higher standard, ‘probable cause’, is required for detentions, arrests and indictments, and for more substantial searches of persons and property. When quantified, the thresholds for satisfying the standards are typically glossed as around 10–50% and 30–60% confidence, respectively, that the person is participating in criminal activity.<sup>4</sup>

The ‘preponderance of evidence’ standard, usually understood as requiring that a claim is more likely true than not, operates in civil and family courts. The standard is typically quantified as exceeding 50% likelihood that the litigated facts obtain.<sup>5</sup> The ‘clear and convincing evidence’ standard, frequently quantified as around 65–75% confidence, is more demanding. This standard is employed in equity cases such as right-to-die hearings, wills, libel, child custody, paternity disputes, and commitment to mental institutions.<sup>6</sup>

---

<sup>2</sup> The burden of proof for some claims—such as affirmative defences or immunities from prosecution—falls to the defendant.

<sup>3</sup> *Terry v. Ohio*, 392 U.S. 1 (1968).

<sup>4</sup> McCauliff (1982), Lerner (2003), Bacigal (2004), and Goldberg (2013). Note there can be substantial difference between an officer’s confidence that crime occurred and her confidence that a specific individual committed that crime.

<sup>5</sup> This standard is also known as the ‘balance of probabilities’ standard (*In re Winship*, 397 U.S. at 371–72; House of Lords in *Re B (A Child)* (2008) UKHL 35). For information about quantifying the standard, see McCauliff (1982) and Simon (1969).

<sup>6</sup> See McCauliff (1982); Vars (2010), Simon and Mahan (1971), *United States v. Fatico*, 458 F Supp. 388 (E.D.N.Y. 1978), and Sand and Rose (2003).

The proof ‘beyond reasonable doubt’ standard, which is required for criminal conviction in many jurisdictions, is more demanding still. It is one of the most familiar aspects of criminal law and is commonly understood to be around 90–95% confidence in the guilt of the defendant.<sup>7</sup>

Burdens of proof and their associated standards raise many epistemological questions. Questions concern whether the presumption of innocence is a propositional attitude, the nature of this attitude or stance, and whether it is appropriate. There are questions about how to interpret, and what legitimates, the various standards of proof.<sup>8</sup> In this essay I examine whether these standards can be properly understood in quantitative terms and how this affects the kinds of evidence that can satisfy the standard. In section two I introduce examples that suggest legal standards of proof cannot be properly interpreted in quantitative terms. In sections three, four, and five I evaluate three competing proposals for a non-quantitative epistemic condition on satisfying legal burdens of proof.

## 2. Statistical Evidence

Consider the following cases:<sup>9</sup>

---

<sup>7</sup> See, for instance, Tribe (1971), Laudan (2011), Walen (2015), Mulrine (1997), and Lippke (2010).

<sup>8</sup> See, for example, Laudan (2011; 2012) and rebuttals by Risinger (2010) and Gardiner (2017a). See also Lerner (2003), Walen (2015), Mulrine (1997), Nesson (1979), Pardo and Allen (2008), Sand and Rose (2003), and the 2006 exchanges in *Law, Probability, and Risk* (vol. 5), including Newman (2006), Franklin (2006), Tillers and Gottfried (2006), and Weinstein and Dewsbury (2006).

<sup>9</sup> I include five kinds of example partly because readers unconvinced by one might find a different example compelling, and partly to provide an example for each standard of proof outlined above. For related examples and discussion, see Tribe (1971), Redmayne (2008), Koehler (2001), Schneps and Colmez (2013), Thomson (1986), Buchak (2014), and *People v. Collins* 68 Cal.2d 319, 66 Cal.Rptr. 497 (1968).

Underage Alcohol. Reliable studies establish that at most undergraduate house parties, underage drinking is ubiquitous. The local police force decides this base rate evidence satisfies the reasonable suspicion standard. On this basis they search underage individuals arriving at undergraduate parties.<sup>10</sup>

Gatecrasher. A music venue sells seats at its event but does not issue, or record who purchases, tickets. One day the gate is left open and, taking advantage of the lack of ticketing system, many people gatecrash. The managers realise only 10 seats were sold, but 80 people attended. They call the police, who arrest some attendees for gatecrashing. The police reason that there is probable cause to arrest any attendee, since it is overwhelmingly likely they committed a crime.<sup>11</sup>

Red Taxi. A vehicle hit Jeanette late one night. She could determine it was a taxi, but could not discern the colour. The Red Taxi company operates 75% of taxis in town. The remaining 25% are operated by the Green Taxi company. Jeanette sues the Red Taxi company. Using only the evidence described here, Jeanette reasons, she will win the case, because it is more likely than not that Red Taxi is liable.<sup>12</sup>

Uncertain Paternity. Magda does not know who fathered her child. She knows it is one of two lovers, and asks them to undergo DNA testing. Owing to disorganisation at the laboratory, both samples become adulterated. Statistical analysis cannot generate the high certainty characteristic of flawless DNA testing but reports a 90% likelihood that Tom is the father. Since paternity tests are expensive, Magda reasons this evidence ought suffice for her paternity suit, which only requires ‘clear and convincing evidence’.

Prison Yard. One hundred prisoners exercise in the prison yard. Ninety-nine prisoners together initiate a premeditated attack on a guard. Security footage reveals one prisoner standing against the wall refusing to participate. There is no evidence indicating who refused to participate. The prison officials decide that since for each prisoner it is 99% likely they are guilty, they have adequate evidence to successfully prosecute individual prisoners for assault.<sup>13</sup>

---

<sup>10</sup> This case is inspired by Kerr (2012). To satisfy the reasonable suspicion standard the officer must have a specific, articulable, and individualised suspicion that crime is afoot before they conduct a brief search (*Terry v. Ohio*, 392 U.S. 1 (1968)). Police might reason the statistic-based suspicion is individualised: those particular underage individuals enter an undergraduate party.

<sup>11</sup> Cohen (1977). Note the example is usually employed to illuminate the preponderance of evidence standard rather than probable cause.

<sup>12</sup> This case was first discussed in print in Tribe (1971). Tribe refers to the case as a ‘famous chestnut’ (p. 1341). See also Thomson (1986), Nesson (1985), *Smith v. Rapid Transit* 317 Mass. 469, 58 N.E.2d 754 (1945), and *Kaminsky v. Hertz Corp.* 288 N.W. 2d 426 (Mich. Ct. App., 1980).

<sup>13</sup> Nesson (1979).

In each case brute statistical evidence plausibly licenses a relatively high credence in the target claim. If the burden of proof can be satisfied by statistical evidence alone, these vignettes exemplify satisfying the relevant burdens. But there is something dubious about these cases. Law courts would not adjudicate in favour of the claimants and the police would likely violate their code of conduct were they to act on this basis.<sup>14</sup>

For each vignette, furthermore, we can construct a comparison case in which non-statistical, ‘individualised’ evidence is employed instead. The overall likelihood of the disputed fact given the evidence might be lower, yet the burden seems satisfied. Consider the following case:

Red Taxi Testimony. A vehicle hit Jeanette late one night. The Red Taxi company operates 50% of taxis in town. The Green Taxi company operates the remaining 50%. Jeanette claims she saw the taxi was red. Eyewitness reports are notoriously unreliable, and so tests are performed to determine Jeanette’s reliability under the relevant conditions. She discerns the correct colour 70% of the time. Her eyewitness testimony thus suggests a red taxi caused the accident. With this evidence Jeanette sues the Red Taxi company.<sup>15</sup>

In some sense, this eyewitness evidence is less determinative: plausibly a red taxi is more likely at fault given the mere market share evidence in the original example. But finding liability in the second case seems more legitimate and courts are significantly more likely to find in the plaintiff’s favour. Similar counterparts can be constructed for each of the five vignettes above.

---

<sup>14</sup> Koehler (2001). See also Gardiner (forthcoming a), which casts doubt on the adequacy of merely statistical evidence in non-legal contexts.

<sup>15</sup> Thomson (1986). For related discussion see Tversky and Kahneman (1982).

These cases suggest the standard of proof is not a quantitative measure, such as a credence or statistical likelihood. If ‘preponderance of evidence’ could be captured numerically, for example, it would be lower than 75%. And this would mean, implausibly, the fact finder should find liability in the Red Taxi case. Similarly if ‘beyond reasonable doubt’ can be understood quantitatively, it must be lower than 99% confidence. This would entail, implausibly, statistical evidence in the prison yard case suffices for conviction. Questions about whether merely statistical evidence can satisfy a burden of proof are increasingly pressing, as cold hit DNA or fingerprint matches and other statistical approaches to evidence are becoming more prevalent.<sup>16</sup> (‘Cold hits’ are when a database search connects a hitherto unconnected perpetrator to a crime.) This raises the question: if the standard of proof is not simply a numerical likelihood, and if merely statistical evidence does not satisfy that standard, what other conditions are required to satisfy the burden of proof?<sup>17</sup>

### 3. Causal Relations and Guarantees

Legal theorists note that vignettes like the first five above lack ‘individualised’ evidence against the defendant. But this notion is itself opaque. Judith Jarvis Thomson aims to illuminate the nature and value of individualised evidence.<sup>18</sup> She argues individualised evidence is that which is appropriately causally connected to the target claim. This causal relation might be that the evidence caused the crime. Gambling debt, for example, might be individualised evidence of guilt if debt motivated the theft. Thomson calls this ‘forward-

---

<sup>16</sup> For recent advances in statistical approaches to legal evidence, see Pardo (2013), Schneps and Colmez (2013), Koehler (2001), Kaye (2009), Roth (2010), Slatkin, Song and Murphy (2009); Tribe (1971), and Goldberg (2013).

<sup>17</sup> Di Bello (2013) discusses the relationship between the claims that standards of proof are quantitative and that purely statistical evidence satisfies the standard.

<sup>18</sup> Thomson (1986).

looking' individualised evidence. Alternatively crime can cause evidence. The crime's occurring, for example, causes incriminating fingerprints. Thomson calls this 'backward-looking' individualised evidence. Finally, evidence and crime might have a common cause. Suppose the defendant verbally abused an officer one evening. This might be evidence the defendant later attacked a pedestrian because it indicates a common cause—the defendant was enraged that evening—for both incidents.

Evidence with an appropriate causal relation, Thomson argues, supplies epistemic import distinct from merely raising the likelihood of the disputed claim. Individualised evidence, in her view, 'guarantees' the claim is true.<sup>19</sup> In the Red Taxi Testimony case, the taxi's being red causally explains the eyewitness evidence. Eyewitness evidence can therefore guarantee, argues Thomson, the litigated taxi-colour claim. Thomson writes, 'to require individualized evidence of guilt just is to be requiring a guarantee'.<sup>20</sup>

Thomson notes the similarity to lottery cases in epistemology.<sup>21</sup> Statistical evidence about lotteries can highly probabilify the claim that a specific ticket lost, but does not guarantee it. Lottery results printed in the newspaper, by contrast, can guarantee the ticket lost. As with individualised evidence of a litigated fact, newspaper evidence can, in Thomson's view, guarantee the ticket lost even if the likelihood of false belief from newspaper evidence is higher than the likelihood of false belief via merely statistical evidence. This can happen if,

---

<sup>19</sup> Thomson (1986: 206, 208–9, 214).

<sup>20</sup> Thomson (1986: 214).

<sup>21</sup> Thomson (1986: 207–208). Discussion of lottery cases in epistemology stems from Kyburg (1961). See also Nelkin (2000) and Hawthorne (2004).

for example, the newspaper-reported lottery is smaller and there is some chance the newspaper misprinted the numbers.

Thomson's suggestion has virtues. Statistical evidence, at least in the above kinds of cases, characteristically lacks a causal relation to the disputed claim. And statistical evidence plausibly fails to provide a guarantee: even excellent statistical evidence is straightforwardly consistent with the target claim's falsity. Individualised evidence, by contrast, typically has a causal relation and can characteristically—in some sense—generate a guarantee. If a person is convicted even though no evidence is causally related to the putative crime, furthermore, this seems inconsistent with justice; the defendant's crime has no bearing on whether he is convicted.<sup>22</sup>

But there are weaknesses in Thomson's account. Firstly, absent a more precise conception of causation, the account remains schematic. Causation is a notoriously 'murky', elastic notion, and Thomson does not specify the causal relation her account employs.<sup>23</sup>

Some base rates are plausibly causally related to some litigated claims, furthermore, which suggests that lack of causal relation fails to capture the inadequacy of merely statistical evidence. If local gang membership rates are sufficiently high, for example, this base rate can causally contribute to an individual participating in an illegal gang initiation rite. Since it is causal, the base rate is a candidate for forward-looking evidence, but nonetheless is

---

<sup>22</sup> I sometimes, for concision, use terminology specific to one kind of legal burden, such as criminal trials, even though the ideas apply equally to other burdens of proof.

<sup>23</sup> Enoch et al. (2012), Enoch and Fisher (2013), Redmayne (2008), and Blome-Tillmann (2015) raise similar concerns. The term 'murky' appears in Redmayne (2008: 300).



statistical. The statistical evidence in the uncertain paternity case above is plausibly caused by paternity; if so, this is backward-looking causal evidence that is nonetheless statistical.<sup>24</sup> Base rates can be causally related to a common cause. Perhaps in some cases religious tenets causally contribute both to rates of domestic violence within that religion and a particular instance of domestic violence. Despite these causal relations, base rate evidence seems illegitimate for satisfying the burden of proof and is (arguably) unindividualised.

Realistic social base rate evidence does not typically generate high statistical likelihood (and tends to be complex). In the following fictional examples base rate statistics are high and so are a potential candidate for sufficient evidence.<sup>25</sup>

**First Gatecrasher.** A music venue sells seats at its event, but does not issue, or record who purchases, tickets. Fern notices the gate is open, and decides to gatecrash. Other people see her gatecrashing and decide to follow. The managers realise only 10 seats were sold but 80 people attended, and summon the police. Fern happens to be one of the people apprehended. The statistical evidence against Fern is deemed sufficient to satisfy the burden of proof, since it is overwhelmingly likely she gatecrashed.

**Opportunistic Gatecrasher.** A music venue sells seats at its event, but does not issue, or record who purchases tickets. One day the gate is left open and, taking advantage of the lack of ticketing system, many people gatecrash. Oppy walks past the venue and realises many people are gatecrashing. Oppy sees this as opportunity to gatecrash, and so joins in. The managers realise only 10 seats were sold, but 80 people attended. Oppy is one of those arrested. The statistical evidence against Oppy is deemed sufficient to satisfy the burden of proof, since it is overwhelmingly likely he gatecrashed.

---

<sup>24</sup> For more on DNA evidence and the distinction between statistical and individualised evidence, see Di Bello (2013), Pritchard (forthcoming), Stein (2005), and Kaye and Sensabaugh (2000).

<sup>25</sup> The first gatecrasher and opportunistic gatecrasher cases are from Blome-Tillmann (2015). Blome-Tillmann mistakenly claims Thomson's account is not vulnerable to the opportunistic gatecrasher case (Blome-Tillmann 2015: 110). This is because Blome-Tillmann overlooks forward-looking evidence in Thomson's account. Blome-Tillmann's account of Thomson's view cannot be correct, since it omits forward-looking individualised evidence such as evidence of motive.

In these cases, arguably the statistical evidence is causally related to the person's guilt, yet is not individualised, does not generate a guarantee, and should not satisfy a burden of proof. But according to the *letter* of Thomson's account, the statistical evidence, by being causally related, qualifies as individualised. Since the statistics render guilt likely, Thomson's account cannot explain why they do not suffice for conviction. This objection is pressed in Blome-Tillmann (2015).

Responses are available. By refining the account of causation, one might deny Fern caused the statistic. There were, after all, many intervening causes. It is implausible, however, that an account of causation is available where gambling debt causes the theft but the high proportion of gatecrashers does not cause Oppy's opportunistic gatecrashing. A second response posits the way the statistic causes Oppy's crime, by providing motivation or cover, is importantly disconnected from how it epistemically supports guilt, namely by establishing a base rate likelihood. Perhaps, then, the case does not undermine the *spirit* of Thomson's proposal. Thirdly, advocates of a causal account of the epistemic difference between statistical and individualised evidence might argue the statistical evidence against Oppy is not itself causal, perhaps because it is abstract. For the evidence to be causal one must incorporate that Oppy notices the high proportion of gatecrashers, which motivates him. But then, the response continues, the evidence qualifies as causal precisely because one includes an individualised, not purely statistical, aspect.

An alternative line of response argues that the existence of cases illustrating that *sometimes* brute statistical evidence is causally related to the target fact does not undermine Thomson's account; Thomson identifies why brute statistical evidence is *characteristically* inadequate for

conviction. Thomson's account explains that merely statistical evidence characteristically cannot guarantee the disputed fact, but individualised evidence—evidence with appropriate causal relations to the fact—characteristically can. Thus a refined version of Thomson's causal account may avoid these criticisms.

This response raises a second set of concerns about Thomson's account. It is unclear how causal relations generate a guarantee, or indeed what Thomson means by 'guarantee'. A witness might testify that although she thinks the perpetrator wore sandals, she cannot recall the incident well. A causal relationship obtains—the witness believes this because the perpetrator wore sandals—but the evidence cannot provide a guarantee. We might speculate that if Jones stole, his debt was a cause. But we are unsure whether he stole and, if so, whether debt was a motive. These pieces of evidence offer no guarantee. It is unclear from Thomson's account why combining evidence of this kind is epistemically or legally superior to statistical evidence. Perhaps lottery-style, statistical evidence provides a greater guarantee than combining many unreliable, but individualised, considerations.

To further cleave guarantees from causally connected evidence, note there can be non-causal facts that generate guarantees. These might include facts of legal definition, such as if the accused performed certain actions he thereby committed an offence, or medical diagnosis, such as if the witness has specific traits then he has a particular disorder. A lawyer might articulate entailment relations such as if three people conspired as alleged, then at least three people knew the plan in advance. Evidence can include mathematical testimony from expert witnesses. Cold hit DNA evidence is typically characterised as statistical and, if the DNA is

unmixed and non-degraded, produces a virtual certainty.<sup>26</sup> These examples of non-causal evidence do not, by themselves, satisfy a burden of proof. The prosecutor must determine, for example, that the DNA would not have been at the crime scene were the defendant innocent. But although DNA evidence has a causal element, the statistical reasoning guarantees some facts of the case, such as whether the accused's DNA was found at the scene. These examples illustrate that non-causal evidence can provide guarantees during a trial, which raises the question of whether and how causal relations generate a distinctive guarantee.

Perhaps the guarantee arises because causal relations can only obtain if the crime occurred. Jones's committing the crime did not cause eyewitness beliefs, for example, unless he committed the crime. But this is a trivial sense of guarantee and cannot be what Thomson intended. It is divorced from our epistemic state; fact finders do not know whether the causal relation between the evidence and the litigated fact holds unless they know whether the litigated fact obtains. So the logic of causation does not provide reasoners with a guarantee. The relation holds objectively, and is not one to which we have independent epistemic access.

To develop Thomson's account, more must be said about the relation between individualised evidence and guarantees. In section five I articulate one way we might understand this connection.

---

<sup>26</sup> Devik (2006); Schneps and Colmez (2013: 69). If the test is performed correctly and if—as geneticists contend—the locations of the 13 peaks in genetic code are mutually independent, the probability that two people share a DNA match is 1/400 trillion. Enoch and Fisher (2013) also note some probative evidence is non-causal.

As noted above, a causal relation does not obtain unless the relata exist. This generates a further problem for Thomson's account.<sup>27</sup> If individualised evidence requires a causal relation, how is individualised evidence possible when the putative fact does not obtain or where the evidence misleads? If the defendant is innocent, his (putative) guilt cannot stand in any causal relationship to evidence. But then, according to Thomson's account, there is no individualised evidence. This is problematic for at least two reasons. Firstly, evidence can be individualised, as opposed to merely statistical, regardless of whether it misleads and whether the accused is innocent. Secondly, it is a desideratum of an account of legal evidence that—if evidence is compelling but misleading—the burden can be satisfied even if the judgement is false. A police officer can have legitimate reasonable suspicion, for example, when no crime is afoot and an innocent defendant can appear guilty beyond reasonable doubt. The causal account owes an explanation of how misleading evidence can be individualised.

To conclude this section I articulate two further worries concerning Thomson's account of the connections amongst individualised, causally-related, and guarantee-providing evidence. Thomson holds evidence is individualised only when uniquely identifying. She writes,<sup>28</sup>

Mrs. Smith believes she saw a one-legged, left-handed, entirely bald, and extremely tall man kill Bloggs. That is individualized evidence that a man with those four features killed Bloggs, for the (putative) fact that a man with those four features killed Bloggs would causally explain Mrs. Smith's believing she saw a man with those four features kill Bloggs [...] There are [99 other men in the world in addition to Mullins] who have all four features, so *getting individualized evidence against Mullins requires getting some fact in respect of which an appropriate causal role is played by a feature which distinguishes Mullins from [all other men in the world]*.

---

<sup>27</sup> Blome-Tillmann (2015) raises a version of this worry.

<sup>28</sup> Thomson (1986: 217, emphasis mine).

In the revised version of the case, Mrs. Smith believes she saw a one-legged, left-handed, entirely bald, extremely tall, and one-eyed man kill Bloggs [...] [O]ur further evidence suggests that only Mullins has all five features [...] To the extent to which we believe that further evidence, then, we shall take ourselves to have individualized evidence, not merely that a man with those five features killed Bloggs, but that Mullins did—he being the only available candidate with the five features.

The italicised condition would help underwrite a guarantee of guilt. But as a condition on individualised evidence, it is too strong. Evidence can be individualised and against the defendant without being uniquely identifying or conclusive. This point echoes Donnellan's and Kripke's objection to descriptivism as a semantics for proper names: we can have a *de re*, individualising thought about something without being able to uniquely discriminate it.<sup>29</sup> This underscores that individualised evidence is far from a guarantee of guilt, at least in the strong, ordinary sense of 'guarantee'.

Secondly, the notion of 'guarantee' arguably sits uncomfortably with lower standards. Thomson extends her account from criminal to civil burdens. According to Thomson, various strengths of standard of proof correspond to how sure fact finders are of having a guarantee of culpability. She writes,<sup>30</sup>

Our law requires the jury in a criminal case to be sure beyond a reasonable doubt that the defendant is guilty before imposing liability on him; the friend of individualized evidence may be taken to say that the jury must be sure beyond a reasonable doubt that the defendant is guilty because of being sure beyond a reasonable doubt that there are facts available to it which guarantee that the defendant is guilty. Our law requires the jury in a case in tort to believe no more than that it is more probable than not that the defendant is guilty; the friend of individualized evidence may be taken to say that the jury must believe it is more probable than not that the defendant is guilty because of believing it more probable than not that there are facts available to it which guarantee that the defendant is guilty.

---

<sup>29</sup> Donnellan (1972) and Kripke (1972). Thanks to Jon Garthoff for helpful discussion of these issues.

<sup>30</sup> Thomson (1986: 215, 219).

We might further extend this treatment: to satisfy a ‘clear and compelling evidence’ standard fact finders must believe they have clear and compelling evidence that that there are facts available to them which guarantee the defendant is liable; to satisfy a ‘reasonable suspicion’ standard the officer must believe she has reasonable suspicion that that there are facts available to her which guarantee that crime is afoot.

But our ordinary notion of ‘guarantee’ seems inappropriate for understanding standards lower than beyond reasonable doubt. Plausibly finding fault in civil cases does not require any evidence or chance of a *guarantee* of liability, for example, instead simply requiring a certain level of epistemic support. Perhaps a court can appropriately find in favour of the plaintiff, in other words, even if no evidence purports to guarantee liability, but the preponderance of available evidence indicates liability.

Thomson’s account of the inadequacy of statistical evidence appears unsuccessful as it stands. To render the view more compelling, we need a better understanding of the relevant notion of guarantee, the causal relation, and how causal relations underwrite guarantees.

#### **4. Sensitivity and Incentives**

An alternative view holds that merely statistical evidence cannot satisfy a legal burden because legal standards of proof require one’s judgement to be sensitive to the truth of the

litigated facts.<sup>31</sup> Enoch et al. (2012) and Enoch and Fisher (2013) argue the deficiency of merely statistical evidence is that it obtains regardless of the specific crime, and so a belief based on statistical evidence would be held even were the accused innocent. Individualised evidence, by contrast, characteristically engenders sensitive beliefs; if the defendant did not commit the crime, the individualised evidence would not obtain.<sup>32</sup>

Sensitivity of Beliefs. S's belief that p is sensitive iff had it not been the case that p, S (likely) would not have believed p.

Extending the notion of sensitivity of beliefs suggests a corresponding 'sensitivity of evidence': Evidence is sensitive to a claim iff (roughly speaking) were the claim false, the evidence would likely not obtain. (The evidence must be understood with some imprecision. It is not that 70 people gatecrashed, for example, it is that many people gatecrashed.)<sup>33</sup>

Enoch et al.'s suggestion enjoys a great deal of plausibility. It is plausibly a principle of justice and good reasoning that a person should not be found culpable unless evidence used to convict is counterfactually dependent on the crime. The conviction should be sensitive to the transgression. The court should be able to assert that if the accused did not the commit the crime, he would not have been convicted. Thomson's suggestion employs the causation relation to unpack this idea; the sensitivity approach employs modal conditions.

---

<sup>31</sup> For information on sensitivity see Nozick (1981) and DeRose (1996). For a sensitivity-based approach to lottery cases in epistemology see Dretske (1971) and DeRose (forthcoming: chapter 5).

<sup>32</sup> At least, those beliefs are sensitive in the good case; see discussion below.

<sup>33</sup> See also Blome-Tillmann (2015).



Enoch et al. argue, however, that sensitivity as an epistemic value does not have legal value and courts should not care about the sensitivity of their judgements.<sup>34</sup> Instead they advance a non-epistemic vindication of the legal value of sensitivity, and thus a practical vindication of the distinction between individualised and statistical evidence. They argue that relying on statistical evidence undermines incentives to obey the law. If evidence sufficient for finding culpable obtains regardless of whether the person transgresses, Enoch et al. reason, the person might as well transgress because ‘whatever he decides will have negligible influence on the likelihood of his being punished’.<sup>35</sup> If purely statistical evidence suffices to satisfy the burden of proof, then a person deciding whether to purchase a ticket or gatecrash, for example, has no incentive (stemming from evidence law) to purchase the ticket. Similarly, taxi drivers lose their incentive to drive cautiously.

I think this incentive-based explanation of the need for individualised evidence is mistaken. Firstly, consider the gatecrasher case, which is the example they select to illustrate their claims about incentive structures. In this case there remains an incentive to obey the law by not entering the venue; by not entering, the person avoids arrest. A more significant objection is that cases with the appropriate incentive structure are rare. Enoch et al.’s explanation only applies when the crime occurs regardless of whether the person participates, the person is already in the class of suspects to be investigated and will remain so regardless of whether he participates, the person must decide whether to participate in the law-breaking, *and* an investigation or other legal action will occur regardless of whether the person participates. Cases with this structure might include activities such as small-fry tax

---

<sup>34</sup> Enoch et al. (2013: 211–215) and Enoch and Fisher (2013: 577–581).

<sup>35</sup> Enoch and Fisher (2013: 583).

fraud, littering, and television licence evasion. If defendants for these infractions were identified *and* prosecuted with purely statistical evidence, the incentive from legal evidence would arguably be diminished.<sup>36</sup> But typically decisions about whether to transgress concern cases where the crime or harm to be investigated would not otherwise occur, and so no investigation would occur but for the person's transgressing. This includes almost all thefts, assaults, harassment, negligent accidents, embezzlement etc. So even with merely statistical evidence, incentive is provided: if the person had not transgressed, there would be no investigation or risk of punishment.

Enoch et al. acknowledge this concern. They write,<sup>37</sup>

The incentive story thus has different implications across the two categories of cases: those in which the act would likely be performed by others regardless of whether the would-be perpetrator decides not to engage in it, and personal-context cases [such as spousal abuse], in which the act will not likely be carried out by anyone else. In the latter type of case, the statistical evidence against the defendant ought to be admissible at trial.

What Enoch et al. fail to appreciate, I think, is that 'the latter type of case' includes most crimes and liabilities. If Jones does not coldcock Williams on his way home, probably no one else will. Their incentive-based account, rather than vindicating the legal distinction between individualised and merely statistical evidence, instead vindicates the adequacy of statistical evidence in almost every instance.

---

<sup>36</sup> Although I do not have space to develop this claim here, I believe Enoch et al. elide the distinction between statistical evidence being used to identify and convict the accused. If statistical evidence were used to identify suspects from the general population to prosecute, this *might* reduce incentives (from some aspects of evidence law) to not transgress. (Although even here, for reasons articulated below, I think the incentive-based account cannot capture the nature of the wrong.) But the use of statistics proposed is *convicting* the suspect once identified. Given this, the incentive-based account is less successful.

<sup>37</sup> Enoch and Fisher (2013: 608); see also Enoch et al. (2012: 219).

Allowing purely statistical evidence to satisfy the burden, furthermore, might well disincentivise crime: It lifts a restriction on admissible evidence, and thereby renders conviction more likely.

Arguably the strongest case for Enoch et al.'s incentive-based account concerns drivers employed by the smaller taxi service. Drivers of green taxis will not be found liable in any case adjudicated solely by market share evidence. So perhaps they lose incentive (from legal evidence) if purely statistical evidence can satisfy a burden. But even in this instance the incentive-based account is unconvincing. To see why, note the difference between statistical evidence sufficing for a finding and statistical evidence being the only admissible evidence. If it were the only admissible evidence, green taxi drivers would have no incentive (from legal evidence) to drive cautiously: they would never be found liable. But since—regardless of whether statistical evidence suffices—other kinds of evidence are admissible, drivers retain incentives from legal evidence: for any particular accident, there is likely to be particularised incriminating evidence. Only in unusual cases would the court rely wholly on market share evidence.

A further weakness of the incentive-based account is that even if merely statistical evidence is employed persons retain incentives from other sources, such as the motivation to not steal, feel guilty, be injured, or injure others. Enoch et al. do not mention these other incentives, but they dilute the importance of incentives stemming from legal evidence.

Finally, Enoch et al.'s proposal, which concerns influencing one aspect of citizens' incentive structures in some kinds of cases, does not provide adequate warrant for our response to the prospect of satisfying the burden of proof with purely statistical evidence. When we consider the cases described in section two, there seems something *unjust* about convicting, finding liable, arresting, searching, or detaining on purely statistical evidence. Plausibly those people are *wronged*. But Enoch et al.'s incentive-based account, which is rooted in practical policy considerations, is not the right kind of explanation to vindicate those moral reactions.

We can set aside Enoch et al.'s incentive-based explication of the distinction between individualised and statistical evidence. The question remains whether a sensitivity-based account explains the inadequacy of statistical evidence. As noted above, it seems remiss for a court to convict unless they would have acquitted were the defendant innocent. Plausibly this sensitivity-based epistemic condition itself has legal value. The idea merits investigation.

The sensitivity account raises some similar concerns as the causal account. Some statistical evidence—or beliefs based on statistical evidence—is sensitive to the crime's occurring, for example. Recall Fern the first gatecrasher and Oppy the opportunistic gatecrasher. In these cases, statistical evidence generates the relevant sensitivity counterfactuals. If Fern did not gatecrash, the statistical evidence would likely not have obtained. If Oppy had not gatecrashed, the courts would not have found him liable.<sup>38</sup>

---

<sup>38</sup> Blome-Tillmann (2015: 106–107). Blome-Tillman holds that the latter counterfactual is true since had Oppy not gatecrashed he would not have entered the venue. Many thanks to Martin Smith for helpful comments on these issues.

Some of the same responses are available: Like advocates of the causal account, advocates of the sensitivity account might argue the account identifies what is characteristically inadequate about employing statistical evidence to satisfy a burden, and the account is thereby consistent with some marginal instances in which statistical evidence generates sensitive beliefs. In response note that, assuming Tom is the father, statistical evidence in the uncertainty paternity case generates sensitive beliefs, and underwrites a sufficiently high likelihood, yet is inadequate for finding liability. This suggests that mundane, central kinds of statistical evidence engender sensitive beliefs.<sup>39</sup>

A further problem for the proposal is that much individualised evidence is insensitive: regardless of whether the accused is culpable, there is incriminating individualised evidence against him. This can occur with misleading or inconclusive evidence. Perhaps, for example, the accused's fingerprints are at the scene because he is an acquaintance of the victim. The fingerprints can feature in the prosecution's narrative as individualised evidence, but may not be sensitive to the crime.

Note too that sensitivity is factive; a belief cannot be sensitive unless true. This creates a challenge for a sensitivity account of the legal value of individualised evidence. To see why, consider a case in which the accused is innocent but there is compelling, misleading, incriminating, individualised evidence. To determine whether the fact finder's belief is sensitive we must ask, 'Were the accused innocent, would the fact finder believe he is innocent?' Since the defendant is actually innocent, the answer is no. No false conviction is

---

<sup>39</sup> We must assume Tom is the father because sensitivity is factive: no false belief is sensitive. Thanks to an anonymous reviewer and Laura Callahan for helpful comments on these issues.

sensitive. If the burden of proof requires that the judgement is sensitive, as the sensitivity account suggests, no false conviction satisfies the burden. But, as noted above, it is a desideratum of legal epistemology that some false convictions satisfy the burden.<sup>40</sup>

Enoch et al. respond to these kinds of objections by arguing sensitivity is a hallmark of *good* individualised evidence, not just any individualised evidence; statistical evidence, by contrast, is insensitive even when it is good statistical evidence.<sup>41</sup> But it bears noting that as an account of the characteristic virtue of individualised evidence over statistical evidence, it is a weakness of the sensitivity account that much individualised evidence is insensitive. In section five I explore whether a relation known as ‘normic support’, or a kindred relation, is a better hallmark of individualised evidence and can better explain its legal and epistemic value.

## 5. Normic Support

Many epistemologists hold that epistemic support concerns probabilification.<sup>42</sup> They hold that a body of evidence *E* epistemically supports proposition *p* iff *p* is likely, given *E*, to be true; a body of evidence *E*<sub>1</sub> epistemically supports proposition *p* more than *E*<sub>2</sub> does iff *p* is more likely given *E*<sub>1</sub> than given *E*<sub>2</sub>.

Martin Smith (2010) challenges this orthodoxy. He notes that in the lottery case, purely statistical evidence can establish that it is extremely likely a particular ticket lost, but cannot justify outright belief that the ticket lost. Reading lottery results in a newspaper, on the other

---

<sup>40</sup> This point is also made in Blome-Tillmann (2015) and Smith (2010).

<sup>41</sup> Enoch et al. (2012: 209).

<sup>42</sup> See, for example, citations in section one of Smith (2010).

hand, can justify outright belief. Individualised evidence can justify belief even if the claim is more likely to be false than if it were based on purely statistical evidence.

Smith argues that individualised evidence, but not statistical evidence, ‘normically supports’ the target claim. Evidence normically supports a conclusion when, roughly speaking, given that evidence, *p* would normally be true. This notion of normalcy does not reduce to statistical frequency.<sup>43</sup> If the evidence obtains yet *p* is false, some abnormality or malfunction has occurred. The error demands explanation. When evidence statistically indicates *p* but *p* is false, by contrast, it is not really an error; it is simply a case of ‘you win some you lose some’.<sup>44</sup>

Smith employs a possible worlds framework for understanding normic support. He writes,<sup>45</sup>

A body of evidence *E* normically supports a proposition *p* just in case *p* is true in all the most normal worlds in which *E* is true. Alternately, *E* normically supports *p* just in case the most normal worlds in which *E* is true and *p* is true are more normal than any world in which *E* is true and *p* is false [...] A body of evidence *E* normically supports a proposition *p* *more strongly* than it normically supports a proposition *q* just in case the most normal worlds in which *E* is true and *q* is false are more normal than any world in which *E* is true and *p* is false.

There are concerns with this account of normic support and, in particular, the possible world framework proposed. One concern is that more normal claims receive more normically-supportive evidence *just in virtue of* being normal. If *p* is adequately normal an intuitively irrelevant fact *F* will qualify as normically supporting *p*. Suppose *p* obtains in

---

<sup>43</sup> Smith (2010: 16).

<sup>44</sup> The expression ‘you win some you lose some’ appears in Enoch et al. (2012: 208).

<sup>45</sup> Smith (ms: 20–21). See also Smith (2010:16–17). The actual world may not be maximally normal; misleading individualised evidence is possible.

almost all normal worlds ( $p$  is ‘grass is green’, for example). Now consider an intuitively irrelevant fact  $F$ , such as ‘polar bear hind paws are elongated’. The normal  $F$ -worlds are  $p$ -worlds; the worlds where  $F$  is true and  $p$  is true are more normal than any world where  $F$  is true and  $p$  is false. This is because not- $p$  worlds are almost all abnormal. And so the seemingly irrelevant fact  $F$  normically supports  $p$ . A related worry—one common to possible world analyses—concerns logical truths. Logical truth  $T$ , which obtains in all worlds and so obtains in all normal worlds, is normically supported by an irrelevant fact  $F$ , because normal  $F$ -worlds are  $T$ -worlds.

Even if Smith’s understanding of the normic support relation is not correct, perhaps an epistemic support relation similar to Smith’s normic support can vindicate the distinction between individualised and merely statistical evidence.<sup>46</sup> In what follows I use ‘normic support’ to denote epistemic support relations that share key features with Smith’s proposal but may vary in the details. (In particular one might reject analysing normic support in terms of possible worlds.<sup>47</sup>) Consider the first five vignettes in section two. Reflecting on what normally follows from the evidence plausibly illuminates various judgements about the cases. Insofar as it might seem permissible to search partying underage undergraduates, for example, this is because given the evidence they would *normally* have alcohol and their not having alcohol would be an abnormal, surprising fact. This might obtain, for instance, if

---

<sup>46</sup> Smith (ms: 41ff.) applies normic support to legal epistemology.

<sup>47</sup> For alternative discussions of normalcy, see Pettit (1999), Millikan (1984), Nickel (2008), Schurz (2001), Haslanger (2014), Morreau (1997), Spohn (2014), and Hauska (2008). Gardiner (2015) discusses the relationship between possible worlds and normalcy. Normalcy is not usually discussed by philosophers as a subject matter, but instead the notion of normalcy is used to illuminate other areas, such as character in ethical theory, populations in philosophy of biology, illness in philosophy of disability, conditionals in metaphysics, ceteris paribus laws in philosophy of science, generics in philosophy of language, and in the semantics for nonmonotonic logic. I am very grateful to Bernhard Nickel, Liz Camp, and Martin Smith for helpful insights on these topics.



underage partygoers exhibit hallmarks of drinking, such as vomiting in bushes. If a police search is conducted purely on statistical grounds, by contrast, it seems illegitimate. The uncertain paternity vignette stipulates the samples were mixed. Given this laboratory error, it is normal and does not demand further explanation if the results are misleading. And so we cannot rely on those results in court. If the tests were conducted with non-mixed DNA, by contrast, the match would normically support the paternity claim since error would demand explanation. Perhaps a normic support relation, then, can vindicate the epistemic and legal distinction between individualised and statistical evidence.

If Smith's normic support, or a kindred relation, is a condition on satisfying burdens of proof, this may illuminate some ethical concerns about 'stop and frisk' tactics. Suppose many members of a demographic commit a particular crime. For any individual, given only demographic evidence, it is statistically relatively likely they commit the crime. And now suppose the police frisk an individual based solely on demographic evidence. If burdens of proof require only statistical likelihoods, the officer's frisk would convey, 'I judge (with reasonable suspicion, given only demographic evidence) it is statistically relatively likely the person committed the crime.' Given the stipulation, the judgement could be true and well-grounded; arguably the judgement does not capture the wrong of the frisk. If instead the burden of proof concerns normalcy, the officer's frisk conveys, 'I judge (with reasonable suspicion, given the evidence) it is normal that the person committed the crime; if the person has not committed the crime it is a departure from normalcy'. Plausibly this judgement, unlike the first, is a condemnation of the individual or their group. It is wrong to judge with only demographic evidence that someone would normally transgress. Thus part

of the moral seriousness of ‘stop and frisk’ may stem from the nature of the burden of proof.

Appealing to the idea of normalcy may illuminate promising aspects of the causal and sensitivity accounts. Causal relations often generate normic support relations. If an accident causes eyewitness evidence indicating the taxi was red, for example, normally the taxi is red. If the liable taxi is not red, this is abnormal and demands explanation. Evidence that normically supports a conclusion typically underwrites the belief’s sensitivity. Enoch et al. note that sensitivity and normic support often correlate, but remain agnostic about which is explanatorily more basic. They suggest normic support is less basic because the notion of ‘which errors call for explanation’ seems unilluminating and opaque.<sup>48</sup> Enoch et al. are correct that ‘what calls for explanation’ is a poor candidate for an explanatory foundation. But this is not the heart of normic support; it is simply a characteristic feature. Normic support is a relation of epistemic support; plausibly it generates facts about which beliefs are sensitive and about which errors demand explanation. If a belief is insensitive, in other words, plausibly this is because the evidence fails to normically support the belief. Normic support seems the more basic notion.<sup>49</sup>

Thomson argues that an epistemically valuable feature of individualised evidence is that it—unlike statistical evidence—can provide a guarantee. I noted above this notion was underdescribed in Thomson’s account. Perhaps a normic support relation can illuminate this ‘guarantee’. If evidence normically supports *p*, and yet not *p*, something is amiss and is not

---

<sup>48</sup> Enoch et al. (2012: 210).

<sup>49</sup> Gardiner (2017b) develops the claim that epistemic modal conditions supervene on other epistemic facts, and are explanatorily less basic than those facts.

as it seems. The error demands explanation because something abnormal has occurred. Perhaps this is the kind of guarantee demanded by the burden of proof.

Enoch et al. demur. They write,<sup>50</sup>

Why should the law especially care about avoiding mistakes that call for explanation? Mistakes that do not call for explanation seem—absent some storytelling otherwise, at least—just as harmful to the relevant party, just as detrimental to the relevant social interests, and so on, as mistakes that do call for explanation.

But this point seems mistaken. If a court convicts a defendant erroneously, and the only response available is ‘you win some you lose some’ because the court relied on statistical evidence to satisfy the burden, the court has wronged the defendant. If a person is convicted, found liable, searched, or arrested without participating in the alleged activity, this error ought arise from some abnormal feature. Plausibly justice—also public trust in the legal system—demands this.

Smith’s account of the relationship between error and normic support, if correct, can further illuminate Thomson’s guarantee. He writes,<sup>51</sup>

If one believes that a proposition *p* is true, based upon evidence that normically supports it then, while one’s belief is not assured to be true, this much is assured: If one’s belief turns out to be false, then the error has to be *explicable* in terms of disobliging environmental conditions, deceit, cognitive or perceptual malfunction or some other interfering factor. In short, the error must be attributable to *mitigating circumstances* and thus *excusable*, after a fashion. Errors that do not fall into this category are naturally regarded as errors for which one must bear *full responsibility*—errors for which there is no excuse.

According to Smith, if one’s evidence normically supports *p*, and one believes *p* based on

---

<sup>50</sup> Enoch et al. (2013: 214).

<sup>51</sup> Smith (ms: 18, emphasis in original).

that evidence, the person is not responsible for the error if  $p$  is false. I am doubtful of Smith's claim concerning epistemic responsibility, since it seems evidence could normically support  $p$ ,  $p$  not obtain, yet the error accrue to the person's reasoning and be something the reasoner is responsible for. Suppose, for example, Oliver believes  $p$ , where  $p$  is 'Jones murdered Jill'. He believes this because Jones's clothes are bloodstained. Oliver's belief is thus based upon evidence that normically supports  $p$ . But Jones believes  $p$  because he is confident that only butchers wear bloody clothes and butchers are angry and murderous. In this case, the evidence normically supports  $p$ , and Oliver believes  $p$  because of that evidence, but the error nonetheless accrues to Oliver's poor reasoning. If Oliver had reasoned better, he may well have realised Jones's innocence.

Even if Smith's claim about the relationship between normic support and epistemic responsibility is incorrect, a normalcy-based epistemic support relation might plausibly illuminate the sense of 'guarantee' generated by individualised evidence: When a conclusion normally follows from the evidence, perhaps the evidence does not merely probabilify the conclusion, it also entails that if the evidence is true, and yet the conclusion false, something is amiss and abnormal. Perhaps this approaches the kind of guarantee demanded by legal burdens of proof.<sup>52</sup>

### **Acknowledgements**

Many thanks to Boris Babic, Austin Baker, Laura Callahan, Liz Camp, Marcello Di Bello, David Black, Branden Fitelson, Will Fleisher, Doug Husak, Jon Garthoff, Bernhard Nickel, Ernest Sosa, Alec Walen, and an anonymous reviewer for very helpful comments on this paper. I am particularly indebted to Martin Smith for his generous and insightful comments. Thanks also to James Chase and David Coady for their excellent editorial work.

---

<sup>52</sup> This essay is published as Gardiner (forthcoming b).

## Bibliography

- Armour, Jody (1994) 'Race Ipsa Loquitur: Of Reasonable Racists, Intelligent Bayesians, and Involuntary Negrophobes' *Stanford Law Review* 46(4): 781–816.
- Bacigal, Ronald (2004) 'Making the Right Gamble: The Odds on Probable Cause' *Mississippi Law Journal* 74: 279–340.
- Blome-Tillmann, Michael (2015) 'Sensitivity, Causality, and Statistical Evidence in Courts of Law' *Thought* 4(2): 102–112.
- Buchak, Lara (2014) 'Belief, Credence, and Norms' *Philosophical Studies* 169(2): 285–311.
- Cohen, Jonathan (1977) *The Probable and the Provable* Oxford University Press.
- Colyvan, Mark, Helen Regan and Scott Ferson (2001) 'Is it a Crime to Belong to a Reference Class?' *Journal of Political Philosophy* 9(2): 168–181.
- DeRose, Keith (forthcoming) *The Appearance of Ignorance: Knowledge, Skepticism, and Context, Volume Two* Oxford University Press.
- \_\_\_\_\_ (1996) 'Knowledge, Assertion, and Lotteries' *Australasian Journal of Philosophy* 74: 568–580.
- Devik, Keith (2006) 'Statisticians Not Wanted' blog post  
[https://www.maa.org/external\\_archive/devlin/devlin\\_09\\_06.html](https://www.maa.org/external_archive/devlin/devlin_09_06.html) (accessed March 2016).
- Di Bello, Marcello (2013) *Statistics And Probability in Criminal Trials: The Good, the Bad and the Ugly* Doctoral Dissertation, Stanford University.
- Donnellan, Keith (1972) 'Proper Names and Identifying Descriptions' *Semantics of Natural Language* Donald Davidson and Gilbert Harman (eds.) Dordrecht: D. Reidel, pp. 356–379.
- Dretske, Fred (1971) 'Conclusive Reasons' *Australasian Journal of Philosophy* 49: 1–22.
- Enoch, David and Talia Fisher (2013) 'Sense And "Sensitivity": Epistemic And Instrumental Approaches To Statistical Evidence' *Stanford Law Review* 67: 557–611.
- Enoch, David, Levi Spectre, and Talia Fisher (2012) 'Statistical Evidence, Sensitivity, and the Legal Value of Knowledge' *Philosophy and Public Affairs* 40(3): 197–224.
- Franklin, James (2006) 'Case Comment—United States v. Copeland, 369 F. Supp. 2d 275 (E.D.N.Y. 2005): Quantification of the 'Proof Beyond Reasonable Doubt' Standard' *Law, Probability, and Risk* 5: 159–165.
- Gardiner, Georgi (forthcoming a) 'Evidentialism and Moral Encroachment' *Believing in Accordance with the Evidence: New Essays on Evidentialism* Kevin McCain (ed.) Springer.
- \_\_\_\_\_ (forthcoming b) 'The Burden of Proof and Statistical Evidence' *Routledge Handbook of Applied Epistemology* David Coady and James Chase (eds.) Routledge.
- \_\_\_\_\_ (2017a) 'In Defence of Reasonable Doubt' *Journal of Applied Philosophy* 34(2): 221–241.
- \_\_\_\_\_ (2017b) 'Safety's Swamp: Against the Value of Modal Stability' *American Philosophical Quarterly* 54(2): 119–129.
- \_\_\_\_\_ (2015) 'Normalcy and the Contents of Philosophical Judgements' *Inquiry* 58(7): 700–740.
- Goldberg, Erica (2013) 'Getting Beyond Intuition in the Probable Cause Inquiry' *Lewis & Clark Law Review* 17: 789–850.
- Haslanger, Sally (2014) 'The Normal, the Natural and the Good: Generics and Ideology' *Politica & Società* 3: 365–392.
- Hauska, Jan (2008) 'Dispositions and Normal Conditions' *Philosophical Studies* 139: 219–232.

- Hawthorne, John (2004) *Knowledge and Lotteries* Oxford: Clarendon Press.
- Kaye, David (2009) 'Rounding up the Usual Suspects: A Legal and Logical Analysis of DNA Trawling Cases' *North Carolina Law Review* 87: 425–503.
- Kaye, David and George Sensabaugh (2011) 'Reference Guide on DNA Identification Evidence' in *Reference Manual on Scientific Evidence, 3<sup>rd</sup> Edition* Federal Judicial Center, pp. 129–210.
- Kerr, Orin (2012) 'Why Courts Should Not Quantify Probable Cause' *The Political Heart of Criminal Procedure: Essays on Themes of William J. Stuntz* Michael Klarman, David Skeel, and Carol Steiker (eds.) Cambridge University Press, pp. 131–143.
- Koehler, Jonathan (2001) 'When Do Courts Think Base Rate Statistics are Relevant?' *Jurimetrics Journal* 42: 373–402.
- Kripke, Saul (1972) *Naming and Necessity* Harvard University Press.
- Kyburg, H.E., Jr. (1961) *Probability and the Logic of Rational Belief* Wesleyan University Press.
- Laudan, Larry (2012) 'Is it Finally Time to Put "Proof Beyond a Reasonable Doubt" Out to Pasture?' *Routledge Companion To Philosophy Of Law* Andrei Marmor (ed.) New York: Routledge, pp. 317–332
- \_\_\_\_\_ (2011) 'The Rules of Trial, Political Morality and the Costs of Error: Or, Is Proof Beyond a Reasonable Doubt Doing More Harm Than Good?' *Oxford Studies in Philosophy of Law* Leslie Green and Brian Leiter (eds.) Oxford University Press, pp. 195–227.
- Lerner, Craig (2003) 'The Reasonableness of Probable Cause' *Texas Law Review* 81: 951–1029.
- Lippke, Richard (2010) 'Punishing the Guilty, Not Punishing the Innocent' *Journal of Moral Philosophy* 7(4): 462–88.
- Littlejohn, Clayton (ms) 'Truth, Knowledge, and the Standard of Proof in Criminal Law'.
- McCauliff, C.M.A., (1982) 'Burdens of Proof: Degrees of Belief, Quanta of Evidence, or Constitutional Guarantees?' *Vanderbilt Law Review* 35: 1293–1335.
- Millikan, Ruth (1984) *Language, Thought, and Other Biological Categories: New Foundations for Realism* MIT Press.
- Morreau, Michael (1997) 'Fainthearted Conditionals' *Journal of Philosophy* 94(4): 187–211.
- Mulrine, Thomas (1997) 'Reasonable Doubt: How in the World Is It Defined?' *American University Journal of International Law and Policy* 12(1): 195–225.
- Nelkin, Dana (2000). 'The Lottery Paradox, Knowledge, and Rationality' *Philosophical Review* 109(3): 373–409.
- Nesson, Charles (1985) 'The Evidence or the Event? On Judicial Proof and the Acceptability of Verdicts' *Harvard Law Review* 98: 1357–1392.
- \_\_\_\_\_ (1979) 'Reasonable Doubt and Permissive Inferences: The Value of Complexity' *Harvard Law Review* 92(6): 1187–1225.
- Newman, Jon (2006) 'Quantifying the Standard of Proof Beyond a Reasonable Doubt: A Comment on Three Comments' *Law, Probability, and Risk* 5: 267–269
- Nickel, Bernhard (2008) 'Generics and the Ways of Normality' *Linguistics and Philosophy* 31(6): 629–648.
- Nozick, Robert (1981) *Philosophical Explanations* Oxford University Press.
- Pardo, M.S. (2013) 'The Nature and Purpose of Evidence Theory' *Vanderbilt Law Review* 66: 547–613.
- Pardo, M.S. and R.J. Allen (2008) 'Juridical Proof and the Best Explanation' *Law and Philosophy* 27: 223–268.

- Pettit, Philip (1999) 'A Theory of Normal and Ideal Conditions' *Philosophical Studies* 96: 21–44.
- Pritchard, Duncan (forthcoming) 'Legal Risk, Legal Evidence and the Arithmetic of Criminal Justice' *Jurisprudence*.
- Redmayne, Mike (2008) 'Exploring the Proof Paradoxes' *Legal Theory* 14: 281–309.
- Risinger, D. Michael (2010) 'Tragic Consequences of Deadly Dilemmas: A Response to Allen and Laudan' *Seton Hall Law Review* 40(3): 991–1021.
- Roth, Andrea (2010) 'Safety in Numbers? Deciding When DNA Alone is Enough to Convict' *New York University Law Review* 85: 1130–1185.
- Sand, Leonard and Danielle Rose (2003) 'Proof Beyond All Possible Doubt: Is There a Need for a Higher Burden of Proof when the Sentence may be Death?' *Chicago-Kent Law Review* 78: 1359–75.
- Schmalbeck, Richard (1986) 'The Trouble with Statistical Evidence' *Law and Contemporary Problems* 49(3): 221–36.
- Schneeps, Leila and Coralie Colmez (2013) *Math on Trial* Basic Books.
- Schurz, Gerhard (2001) 'What is "Normal"? An Evolution-Theoretic Foundation for Normic Laws and Their Relation to Statistical Normality' *Philosophy of Science* 68(4): 476–97.
- Simon, Rita James (1969) 'Judges' Translations of Burdens of Proof into Statements of Probability' *The Trial Lawyer's Guide* John Kennelly and James Chapman (eds.) Mundelein: Callaghan, pp. 103–114.
- Simon, Rita James and Linda Mahan (1971) 'Quantifying Burdens of Proof: A View from the Bench, the Jury, and the Classroom' *Law and Society Review* 5(3): 319–330.
- Slatkin, Montgomery, Yun S. Song and Erin Murphy (2009) 'Average Probability that a "Cold Hit" in a DNA Database Search Results in an Erroneous Attribution' *Journal of Forensic Sciences* 54(1): 22–27.
- Smith, Martin (ms) 'Justification, Normalcy and Evidential Probability' <http://philpapers.org/rec/SMIJNA-2> (accessed March 2016).
- \_\_\_\_\_ (2016) *Between Probability and Certainty: What Justifies Belief* Oxford University Press.
- \_\_\_\_\_ (2010) 'What Else Justification Could Be?' *Noûs* 44(1): 10–31.
- Spohn, Wolfgang (2014) 'The Epistemic Account of Ceteris Paribus Conditions' *European Journal for the Philosophy of Science* 4(3): 385–408.
- Staffel, Julia (forthcoming) 'Three Puzzles about Lotteries' *Lotteries, Knowledge, and Rational Belief: Essays on the Lottery Paradox* Igor Douven (ed.) Cambridge University Press.
- Stein, Alex (2005) *Foundations of Evidence Law* Oxford University Press.
- Thomson, Judith Jarvis (1986) 'Liability and Individualized Evidence' *Law and Contemporary Problems* 49(3): 199–219.
- Tillers, Peter and Jonathan Gottfried (2006) 'Case comment—United States v. Copeland, 369 F. Supp. 2d275 (E.D.N.Y. 2005): A Collateral Attack on the Legal Maxim That Proof Beyond A Reasonable Doubt Is Unquantifiable?' *Law, Probability, and Risk* 5: 135–157.
- Tribe, Laurence (1971) 'Trial By Mathematics: Precision and Ritual in the Legal Process' *Harvard Law Review* 84(6): 1329–1393.
- Tversky, Amos and Daniel Kahneman (1982) 'Evidential Impact of Base Rates' *Judgment Under Uncertainty: Heuristics and Biases* Daniel Kahneman, Paul Slovic and Amos Tversky (eds.) Cambridge University Press, pp. 153–162.

- Vars, Frederick (2010) 'Toward a General Theory of Standards of Proof' *Catholic University Law Review* 60(1): 1–46.
- Volokh, Alexander (1997) 'n Guilty Men' *University of Pennsylvania Law Review* 146: 173–216
- Walen, Alec (2015) 'Proof Beyond a Reasonable Doubt: A Balanced Retributive Account' *Louisiana Law Review* 76(2): 355–446.
- Weinstein, Jack and Ian Dewsbury (2006) 'Comment on the Meaning of "Proof Beyond a Reasonable Doubt"' *Law, Probability and Risk* 5: 167–173.



## Conclusion

In this conclusion I tentatively draw some further connections amongst the themes in these essays. These are not connections I judge to be established by the preceding papers; instead they are speculative connections or ideas for future research that further connects these themes. (I also articulate some connections amongst the papers in the dissertation abstract, dissertation introduction, and within the papers themselves.)

One connection concerns the relationship between pursuing teleological approaches in philosophy and emphasising understanding as a central aim of philosophical inquiry.<sup>1</sup> Craig (1986/1987: 212) writes,

Suppose that the problem of the analysis had been solved, so that agreed necessary and sufficient conditions for the ascription of knowledge were now on the table. Many writers make one feel that this would be a terminus, and investigation concluded. I should see this as a prolegomenon to further inquiry: why has a concept demarcated by those conditions enjoyed such widespread use?

One way to interpret Craig's claim is that identifying necessary and sufficient conditions is not an adequate end to philosophical theorising because we should pursue *understanding* in philosophy. We want to understand the phenomena we investigate. 'Understanding why' is central to understanding. Understanding involves understanding why something is the case, why it has the contours it does, and why it exists at all. It involves appreciating, amongst other things, causal and functional relationships. Explanations are crucial to understanding.

---

<sup>1</sup> I gesture at this connection in 'Teleologies and the Methodology of Epistemology'.

Teleological approaches in epistemology—with their emphasis on explanation, cause, and genealogy—can be an important part of the philosophical project of pursuing understanding.

A second tentative connection amongst the essay's themes concerns the relationship between understanding and intellectualist positions such as evidentialism.<sup>2</sup> As Wayne Riggs and Michael Pace emphasise, we have dual aims in our epistemic practice. We aim to both attain truth and avoid error. The level of evidence or other epistemic support that suffices to epistemically justify a belief might well depend on trade-offs between these dual aims. If believing falsehoods is epistemically egregious and believing truths is not very valuable, for example, then plausibly the amount of evidence required to justify belief will be high. If false belief is relatively benign and true belief has a lot of value, by contrast, the threshold will be higher.

Advocates of pragmatic encroachment (and, potentially, moral encroachment) argue that intellectualist positions do not have much to say about how to weight these dual aims, and thus how to interpret the threshold for justified belief. They argue that encroachment offers a way to establish the threshold.<sup>3</sup> Perhaps instead understanding can determine the threshold. Perhaps, that is, the relative values of attaining truth and avoiding error can be weighted by how the beliefs contribute to, or impede, understanding.<sup>4</sup> This proposal is

---

<sup>2</sup> I briefly sketch this connection in 'Evidentialism and Moral Encroachment'.

<sup>3</sup> It is worth noting that merely positing that another factor—stakes or moral considerations—can influence this threshold does not itself make progress on the question of interpreting the threshold. An expected utility account of when a belief is justified can help illuminate the threshold, but mainstream pragmatic encroachment plausibly simply adds another factor without illuminating what the thresholds are. The problem remains.

<sup>4</sup> See Riggs (2003).

embryonic, and a lot of theorising is required to establish the contours of the idea. But it offers a suggestive connection between intellectualism about justification and understanding.

The threshold problem—that is, the problem of determining how much evidence and other epistemic support is required to justify a belief—is discussed in relation to Craig’s teleological approach.<sup>5</sup> Theorists debate whether, for example, proposals about the function of knowledge attributions entail contextualism or subject sensitive invariantism. Thus another connection amongst these themes that I shall explore in future research is whether competing teleological accounts of knowledge and understanding lend support to intellectualism, pragmatic encroachment, or moral encroachment.

In ‘Normalcy and the Contents of Philosophical Judgements’ I argue that we cannot understand the role of vignettes in philosophical theorising without appeal to the notion of normalcy. We flesh out vignettes by considering normal background conditions. In the essay I focus mainly on vignettes in their role as counterexamples to philosophical claims. As I note in the essay, vignettes can also motivate claims and principles, illustrate distinctions, communicate ideas, help pin down the phenomena of interest, and frame research questions. They provide evidential weight against claims that are not, or do not entail, necessity claims. These claims are not subject to straightforward objection by counterexample. The inability to describe a compelling example can sometimes indicate something is amiss in the theory.<sup>6</sup> An account of the judgement is better if it unifies these roles. One future project examines

---

<sup>5</sup> See Grimm (2015), Greco (2015), Hannon (2017), and Schaffer (2006).

<sup>6</sup> In Gardiner (2016) I cast doubt on Poston’s posit of ‘empty symmetrical evidence’. I argue that the difficulty of describing a plausible example of empty symmetrical evidence suggests the posit is not possible. (See Poston 2014; 2016.)

further how the normalcy proposal illuminates other roles vignettes play in pursuing philosophical understanding, and explores other roles normalcy plays in philosophical theorising.

Thought experiments also play roles in ordinary, legal, and scientific thought. They can aid understanding in these domains and can prove or disprove disputed claims. One related question concerns how interpreting fictions in general requires judgements about normalcy.<sup>7</sup> Another asks how a person's sense of what is normal influences their understanding of legal evidence, such as witness testimony.<sup>8</sup>

In 'Normalcy and the Contents of Philosophical Judgements' I discuss to what extent normalcy itself (rather than judgements of normalcy) depends on our beliefs and expectations, rather than on mind-independent features of the world. If normalcy is not a mind-independent feature of the world, what is normal depends on our understanding. This forges a further relationship between normalcy and understanding.

Another set of connections concerns the relationship between normalcy and teleology. Teleological approaches examine the function of an epistemic kind. I think there are several links between functions and normal conditions. Firstly, the function of an epistemic attribution (or other epistemic kind) is plausibly revealed best by its normal uses. We can better understand the characteristic functions of knowledge attributions, for example, by

---

<sup>7</sup> See Elgin (2014) and Ichikawa and Jarvis (2009) for discussion about the relationship between fictions and thought experiments.

<sup>8</sup> Given that judgements of normalcy vary amongst people, the role of normalcy in interpreting witness statements can lead to injustice—epistemic or otherwise—in the courtroom.

examining their normal contexts of use. (In general something's function might be best understood by reference to its operation in its normal conditions of use.)

We might also understand normalcy in part through the idea of functions. Perhaps a central kind of abnormal circumstance is a *malfunction*. Illuminating functions, then, might illuminate normalcy, and vice versa.

Theorists posit various functional roles for knowledge. These include knowledge's role in the norms of assertion and action, its role in providing assurance, countering doubts, ending inquiry, underwriting blame, praise, and other reactive attitudes, legitimating punishment, giving credit for true belief, and tagging good informants. In adjudicating amongst these roles theorists commonly draw attention to circumstances when the amount of epistemic support required for these various functional roles diverge from each other and from the amount required for knowledge. Other theorists focus on the extent to which the norms *converge*. They argue that the amount of epistemic support required for knowledge *normally* coincides with the amount that is necessary and sufficient for these other roles. Thus they invoke the idea of 'normal coincidence': the various functions of knowledge normally coincide and only come apart in abnormal circumstances.<sup>9</sup> Theorising about normalcy can thus illuminate ideas about the functional roles of knowledge and knowledge attributions.

Many theorists advocate a knowledge norm for belief. As Moss (forthcoming: 165) articulates,

---

<sup>9</sup> For discussion of related ideas, see Hannon (2015), Gerken (2015), and Douven (2016).

The knowledge norm of belief is an especially simple account of [the relationship between what you know and what you should believe]. Belief aims at knowledge. Knowledge sets the standard for belief. These proverbial claims are elucidated by the thesis that knowledge is both necessary and sufficient for permissible belief.

Similarly Williamson (2005: 108) claims,

if one knows that *p*, then one can hardly be wrong to believe that *p*; conversely, given that one does not know that *p*, it arguably is wrong to believe that *p*.

The knowledge norm for belief might illuminate an epistemic flaw in some of the beliefs discussed in ‘Evidentialism and Moral Encroachment’.<sup>10</sup> Most of the cases advocates of moral encroachment use to motivate their view involve a person forming a belief about someone based on base rate evidence. A person believes Laura is an administrative assistant on the grounds that the large majority of women in the office are administrative assistants, for example.

I argue that outright belief based on purely statistical evidence is an epistemic error, regardless of whether the belief concerns a person.<sup>11</sup> The knowledge norm of belief might illuminate this flaw. Lottery cases suggest that beliefs based on purely statistical evidence cannot amount to knowledge, even when the beliefs are true. If the belief cannot amount to knowledge, the knowledge norm of belief can explain why beliefs based on purely statistical evidence are epistemically flawed: they violate the knowledge norm. Beliefs based on testimony, for example, do not share this flaw, since beliefs based on testimonial evidence

---

<sup>10</sup> I sketch this idea in a footnote in ‘Evidentialism and Moral Encroachment’.

<sup>11</sup> I note that extremely-probabilifying merely statistical evidence might justify outright belief.

can qualify as knowledge. Similar considerations might well illuminate proof paradoxes in legal contexts.

Theorising about knowledge often focuses on certainty, conclusive reasons, eliminating error possibilities, and even responding to skeptical challenges. Accounts of understanding, by contrast, are more permissive with regard to certainty and even truth. One might understand a domain even if many of one's beliefs are false or one's evidence is inconclusive. Focusing on understanding might be a helpful tool to illuminate how we should think and adjudicate in domains where knowledge is limited, ill-understood, or inadequate. Perhaps in the domains where certainty or conclusive evidence is unavailable, understanding is the appropriate epistemic aim. Aiming for knowledge in these domains—which might well include philosophy—might be inappropriate because it is unattainable.

Finally, some of these considerations might point to an 'understanding first' epistemology. In 'Evidentialism and Moral Encroachment' I argue that the moral significance of a belief depends on the understanding in which it is embedded. Plausibly the epistemic evaluation of a belief often depends on the broader understanding too. As noted above, understanding might explain how much epistemic support suffices for epistemic justification. An understanding first epistemology would maintain that understanding is not merely a distinct epistemic value—as I argue in 'Understanding, Integration, and Epistemic Value'—but it would also argue that the epistemic value of truth and knowledge can be explained by reference to the epistemic value of understanding. In 'Understanding and Emulation' I argue that a central epistemic pursuit is to improve ourselves as inquirers, and that this pursuit is intimately connected to understanding. Less controversially, an understanding first

epistemology would emphasise that understanding is the principal epistemic aim in many important domains. I do not explicitly pursue an understanding first epistemology here, but perhaps I have helped lay some foundations for a contemporary resurgence of understanding first epistemology.

### Bibliography

- Douven, Igor (2006) 'Assertion, Knowledge, and Rational Credibility' *Philosophical Review* 115(4): 449–85.
- Elgin, Catherine (2014) 'Fiction as Thought Experiment' *Perspectives on Science* 22: 55–75.
- Gardiner, Georgi (2016) 'Coherence Without Conservation' *Syndicate Philosophy*.
- Gerken, Mikkel (2015) 'The Roles of Knowledge Ascriptions in Epistemic Assessment' *European Journal of Philosophy* 23(1): 141–61.
- Greco, John (2015) 'Testimonial Knowledge and the Flow of Information' *Epistemic Evaluation: Purposeful Epistemology* John Greco and David Henderson (eds.) Oxford University Press, pp. 274–290.
- Grimm, Stephen (2015) 'Knowledge, Practical Interests, and Rising Tides' *Epistemic Evaluation: Purposeful Epistemology* John Greco and David Henderson (eds.) Oxford University Press, pp. 274–290.
- Hannon, Michael (2017) 'A Solution to Knowledge's Threshold Problem' *Philosophical Studies* 174: 607–629.
- \_\_\_\_\_ (2015) 'The Importance of Knowledge Ascriptions' *Philosophy Compass* 10(12): 856–66.
- Ichikawa, Jonathan (2009) 'Knowing the Intuition and Knowing the Counterfactual' *Philosophical Studies* 145: 435–43.
- Moss, Sarah (forthcoming) *Probabilistic Knowledge* Oxford University Press.
- Poston, Ted (2014) *Reason and Explanation* Palgrave Macmillan.
- Riggs, Wayne (2003) 'Balancing Our Epistemic Goals' *Noûs* 37(2): 342–352.
- Schaffer, Jonathan (2006) 'The Irrelevance of the Subject: Against Subject-Sensitive Invariantism' *Philosophical Studies* 127: 87–107.
- Williamson, Timothy (2005) 'Contextualism, Subject-Sensitive Invariantism, and Knowledge of Knowledge' *Philosophical Quarterly* 55(219): 213–35.



## Bibliography

- Adler, Jonathan (2002) *Belief's Own Ethics* Cambridge, MA: Bradford/MIT.
- Aikin, Scott (2008) 'Evidentialism and James's Argument from Friendship' *Southwest Philosophy Review* 24(1): 173–80.
- \_\_\_\_\_ (2006) 'Modest Evidentialism' *International Philosophical Quarterly* 46(3): 327–43.
- Alfano, Mark (2013) *Character as Moral Fiction* Cambridge University Press.
- Alston, William (2005) *Beyond Justification* Cornell University Press.
- Anderson, Elizabeth (2010) *The Imperative of Integration* Princeton University Press.
- Appiah, Kwame Anthony (1990) 'Racisms' *Anatomy of Racism* David Goldberg (ed.) University of Minnesota Press, pp. 3–17.
- Armour, Jody (1994) 'Race Ipsa Loquitur: Of Reasonable Racists, Intelligent Bayesians, and Involuntary Negrophobes' *Stanford Law Review* 46(4): 781–816.
- Armstrong, David (1978) *Nominalism and Realism, Volume One* Cambridge University Press.
- Arpaly, Nomy (2003) *Unprincipled Virtue: An Inquiry into Moral Agency* Oxford University Press.
- Arpaly, Nomy and Timothy Schroeder (2014) *In Praise of Desire* Oxford University Press.
- Ashton, Natalie (2015) 'Appropriate Belief without Evidence' *Teorema* 34(2): 7–28.
- Ashton, Natalie and Robin McKenna (ms) 'Situating Feminist Epistemology'.
- Bacigal, Ronald (2004) 'Making the Right Gamble: The Odds on Probable Cause' *Mississippi Law Journal* 74: 279–340.
- Basu, Rima. (ms a) 'The Wrongs of Racist Beliefs'.
- \_\_\_\_\_ (ms b) 'What We Epistemically Owe to Each Other'.
- \_\_\_\_\_ (ms c) 'The Moral Stakes of Racist Beliefs'.
- \_\_\_\_\_ (ms d) 'Motivating Ethical Demand on Belief' presentation at Athena-in-Action, Princeton University, August 2016.
- Basu, Rima and Mark Schroeder (forthcoming) 'Doxastic Wrongs' *Title TBA* Brian Kim and Matthew McGrath (eds.) Routledge.
- Begby, Endre (2013) 'The Epistemology of Prejudice' *Thought* 90–99.
- Blome-Tillmann, Michael (2015) 'Sensitivity, Causality, and Statistical Evidence in Courts of Law' *Thought* 4(2): 102–112.
- Bolinger, Renee (ms) 'The Rational Impermissibility of Accepting (Some) Racial Generalizations'.
- BonJour, Laurence (1980) 'Externalist Theories of Empirical Knowledge' *Midwest Studies in Philosophy* 5: 53–73.
- Brewer, Talbot (2009) *The Retrieval of Ethics* Oxford University Press.
- Buchak, Lara (2014) 'Belief, Credence, and Norms' *Philosophical Studies* 169(2): 285–311.
- Carroll, Lewis (1895) 'What the Tortoise said to Achilles' *Mind* 4(14): 278–80
- Caton, Jacob (ms) 'Is "Justification" an Ordinary Term?'
- Chignell, Andrew (2010) 'The Ethics of Belief' *Stanford Encyclopedia of Philosophy* (ed.) Edward N. Zalta.
- Chisholm, Roderick (1977) *Theory of Knowledge*, Englewood Cliffs, NJ: Prentice-Hall
- Clifford, W. K. (1877) 'The Ethics of Belief' *Contemporary Review* 29: 289–309.
- Cohen, Jonathan (1977) *The Probable and the Provable* Oxford University Press.
- Colyvan, Mark, Helen Regan and Scott Ferson (2001) 'Is it a Crime to Belong to a Reference Class?' *Journal of Political Philosophy* 9(2): 168–181.
- Conee, Earl and Richard Feldman (2011) 'Replies' *Evidentialism and its Discontents* Trent Dougherty (ed.) Oxford University Press, pp. 283–322.

- \_\_\_\_\_ (2004) *Evidentialism* Oxford University Press.
- Constant, Benjamin (1797) 'On Political Reactions' *France* 5 (1).
- Crawford, Lindsay (forthcoming) 'Believing the Best: On Doxastic Partiality in Friendship' *Synthese*.
- Craig, Edward (2007) 'Genealogies and the State of Nature' *Bernard Williams* Alan Thomas (ed.) Cambridge University Press, pp. 181–200.
- \_\_\_\_\_ (1990) *Knowledge and the State of Nature. An Essay in Conceptual Synthesis* Oxford University Press.
- \_\_\_\_\_ (1986/1987) 'The Practical Explication of Knowledge' *Proceedings of the Aristotelian Society* 87: 211–226.
- DeRose, Keith (forthcoming) *The Appearance of Ignorance: Knowledge, Skepticism, and Context, Volume Two* Oxford University Press.
- \_\_\_\_\_ (2004) 'Questioning Evidentialism' *Evidentialism and its Discontents* Trent Dougherty (ed.) Oxford University Press, pp. 137–46.
- \_\_\_\_\_ (1996) 'Knowledge, Assertion, and Lotteries' *Australasian Journal of Philosophy* 74: 568–580.
- Devik, Keith (2006) 'Statisticians Not Wanted' blog post  
[https://www.maa.org/external\\_archive/devlin/devlin\\_09\\_06.html](https://www.maa.org/external_archive/devlin/devlin_09_06.html) (accessed March 2016).
- Di Bello, Marcello (2013) *Statistics And Probability in Criminal Trials: The Good, the Bad and the Ugly* Doctoral Dissertation, Stanford University.
- Donnellan, Keith (1972) 'Proper Names and Identifying Descriptions' *Semantics of Natural Language* Donald Davidson and Gilbert Harman (eds.) Dordrecht: D. Reidel, pp. 356–379.
- Dotson, Kristie (2014) 'Conceptualizing Epistemic Oppression' *Social Epistemology* 28(2): 115–38.
- \_\_\_\_\_ (2008) 'In Search of Tanzania: Are Effective Epistemic Practices Sufficient for Just Epistemic Practices?' *The Southern Journal of Philosophy* 46: 52–64.
- Douven, Igor (2006) 'Assertion, Knowledge, and Rational Credibility' *Philosophical Review* 115(4): 449–85.
- Dretske, Fred (1971) 'Conclusive Reasons' *Australasian Journal of Philosophy* 49: 1–22.
- \_\_\_\_\_ (1970) 'Epistemic Operators' *Journal of Philosophy* 67(24): 1007–23.
- Eaton Daniel and Timothy Pickavance (2015) 'Evidence Against Pragmatic Encroachment' *Philosophical Studies* 172: 3135–3143.
- Elgin, Catherine (2014) 'Fiction as Thought Experiment' *Perspectives on Science* 22: 55–75.
- \_\_\_\_\_ (2009) 'Is Understanding Factive?' *Epistemic Value* Adrian Haddock, Alan Millar, and Duncan Pritchard (eds.) Oxford University Press, pp. 322–30.
- \_\_\_\_\_ (2006) 'From Knowledge to Understanding' *Epistemology Futures* Stephen Hetherington (ed.) Oxford: Clarendon Press, pp. 199–215.
- \_\_\_\_\_ (1996) *Considered Judgment* Princeton University Press.
- Enoch, David (2016) 'What's Wrong with Paternalism: Autonomy, Belief, and Action' *Proceedings of the Aristotelian Society* 116(1): 21–48.
- Enoch, David and Talia Fisher (2013) 'Sense And "Sensitivity": Epistemic And Instrumental Approaches To Statistical Evidence' *Stanford Law Review* 67: 557–611.
- Enoch, David, Levi Spectre, and Talia Fisher (2012) 'Statistical Evidence, Sensitivity, and the Legal Value of Knowledge' *Philosophy and Public Affairs* 40(3): 197–224.
- Fantl, Jeremy and Matthew McGrath (2002) 'Evidence, Pragmatics, and Justification' *The Philosophical Review* 111(1): 67–94.

- Feldman, Richard (2000) 'The Ethics of Belief' *Philosophy and Phenomenological Research* 60(3): 667–695.
- Feldman, Richard and Earl Conee (1985) 'Evidentialism' *Philosophical Studies* 48(1): 15–34.
- Franklin, James (2006) 'Case Comment—United States v. Copeland, 369 F. Supp. 2d 275 (E.D.N.Y. 2005): Quantification of the 'Proof Beyond Reasonable Doubt' Standard' *Law, Probability, and Risk* 5: 159–165.
- Franklin, John Hope (2005) *Mirror to America: The Autobiography of John Hope Franklin* New York, NY: Farrar, Straus and Giroux.
- Fricker, Miranda (2008) 'Scepticism and the Genealogy of Knowledge: Situating Epistemology in Time' *Philosophical Papers* 37(1): 27–50.
- (1998) 'Rational Authority and Social Power: Towards a Truly Social Epistemology' *Proceedings of the Aristotelian Society* 98: 159–177.
- Fritz, James (2017) 'From Pragmatic Encroachment to Moral Encroachment' *Pacific Philosophical Quarterly*.
- Ganson, Dorit (2008) 'Evidentialism and Pragmatic Constraints on Outright Belief' *Philosophical Studies* 139(3): 441–458.
- Gardiner, Georgi (ms) 'Understanding and Emulation'.
- (forthcoming a) 'Evidentialism and Moral Encroachment' *Believing in Accordance with the Evidence: New Essays on Evidentialism* Kevin McCain (ed.) Springer.
- (forthcoming b) 'The Burden of Proof and Statistical Evidence' *Routledge Handbook of Applied Epistemology* David Coady and James Chase (eds.) Routledge.
- (2017a) 'In Defence of Reasonable Doubt' *Journal of Applied Philosophy* 34(2): 221–241.
- (2017b) 'Safety's Swamp: Against the Value of Modal Stability' *American Philosophical Quarterly* 54(2): 119–129.
- (2016) 'Coherence Without Conservation' *Syndicate Philosophy*.
- (2015a) 'Normalcy and the Contents of Philosophical Judgements' *Inquiry* 58(7): 700–740.
- (2015b) 'Teleologies and the Methodology of Epistemology' *Epistemic Evaluation: Purposeful Epistemology* John Greco and David Henderson (eds.) Oxford University Press, pp. 31–45.
- (2012) 'Understanding, Integration, and Epistemic Value' *Acta Analytica* 27(2): 163–81.
- (2011) *A Defence of Robust Virtue Epistemology* Master's Thesis at the University of Edinburgh.
- Geddes, Alexander (forthcoming) 'Judgments about Thought Experiments' *Mind*.
- Gendler, Tamar (2011) 'On the Epistemic Costs of Implicit Bias' *Philosophical Studies* 156: 33–63.
- Gerken, Mikkel (forthcoming) *On Folk Epistemology* Oxford University Press.
- (2015) 'The Roles of Knowledge Ascriptions in Epistemic Assessment' *European Journal of Philosophy* 23(1): 141–61.
- Gettier, Edmund (1963) 'Is Justified True Belief Knowledge?' *Analysis* 23: 121–3.
- Goldberg, Erica (2013) 'Getting Beyond Intuition in the Probable Cause Inquiry' *Lewis & Clark Law Review* 17: 789–850.
- Goldman, Alvin (1993) 'Epistemic Folkways and Scientific Epistemology' *Philosophical Issues* 3: 271–85.
- Gordon, Emma (forthcoming) 'Social Epistemology and the Acquisition of Understanding' *Explaining Understanding* Routledge.

- Gordon, Lewis (2000) *Existentialia Africana: Understanding Africana Existential Thought* Routledge.
- \_\_\_\_\_ (1995) *Bad Faith and Antiracist Racism* Humanities Press.
- Greco, John (2015) 'Testimonial Knowledge and the Flow of Information' *Epistemic Evaluation: Purposeful Epistemology* John Greco and David Henderson (eds.) Oxford University Press, pp. 274–290.
- \_\_\_\_\_ (2012) 'A (Different) Virtue Epistemology' *Philosophy and Phenomenological Research*, 85(1): 1–26.
- \_\_\_\_\_ (2010) *Achieving Knowledge* Oxford University Press.
- \_\_\_\_\_ (2009) 'The Value Problem' *Epistemic Value* Adrian Haddock, Alan Millar and Duncan Pritchard (eds.) Oxford University Press, pp. 313–21
- \_\_\_\_\_ (1999) 'Agent Reliabilism' *Philosophical Perspectives* 13: 273–96.
- Grimm, Stephen (forthcoming) 'Understanding as an Intellectual Virtue' *Routledge Handbook of Virtue Epistemology* Heather Battaly (ed.) Routledge.
- \_\_\_\_\_ (2015) 'Knowledge Practical Interests, and Rising Tides' *Epistemic Evaluation: Purposeful Epistemology* John Greco and David Henderson (eds.) Oxford University Press, pp. 117–137.
- \_\_\_\_\_ (2012) 'The Value of Understanding' *Philosophy Compass* 7(2): 103–17.
- \_\_\_\_\_ (2011) 'On Intellectualism in Epistemology' *Mind* 120(479): 705–733.
- Grundmann, Thomas and Joachim Horvath (2014) 'Thought Experiments and the Problem of Deviant Realizations' *Philosophical Studies* 170: 525–33.
- Guerrero, Alexander (2007) 'Don't Know, Don't Kill: Moral Ignorance, Culpability, and Caution' *Philosophical Studies* 136: 59–97.
- Habgood-Coote, Joshua (2017) *Knowledge-How: Linguistic and Philosophical Considerations* Doctoral dissertation at St. Andrews University.
- Hájek, Alan (2007) 'The Reference Class Problem is Your Problem Too' *Synthese* 156(3): 563–585.
- Hannon, Michael (ms) 'What's the Point of Understanding?'
- \_\_\_\_\_ (2017) 'A Solution to Knowledge's Threshold Problem' *Philosophical Studies* 174: 607–629.
- \_\_\_\_\_ (2015a) 'The Importance of Knowledge Ascriptions' *Philosophy Compass* 10(12): 856–866.
- \_\_\_\_\_ (2015b) 'The Universal Core of Knowledge' *Synthese* 192(3): 769–786.
- \_\_\_\_\_ (2013) 'The Practical Origins of Epistemic Contextualism' *Erkenntnis* 78: 899–919.
- Haslanger, Sally (2014) 'The Normal, the Natural and the Good: Generics and Ideology' *Politica & Società* 3: 365–392.
- Hauska, Jan (2008) 'Dispositions and Normal Conditions' *Philosophical Studies* 139: 219–232.
- Hawthorne, John (2004) *Knowledge and Lotteries* Oxford: Clarendon Press.
- Hazlett, Allan (2016) 'Intellectual Loyalty' *International Journal for the Study of Skepticism* 6: 326–350.
- \_\_\_\_\_ (2013) *A Luxury of the Understanding* Oxford University Press.
- Heil, John (1983) 'Believing What One Ought' *The Journal of Philosophy* 80(11): 752–765.
- Henderson, David and John Greco (2015) (eds.) *Epistemic Evaluation: Purposeful Epistemology* Oxford University Press.
- Hills, Alison (2016) 'Understanding Why' *Noûs* 50(4): 661–88.
- Ichikawa, Jonathan (2009) 'Knowing the Intuition and Knowing the Counterfactual' *Philosophical Studies* 145: 435–43.
- Ichikawa, Jonathan and Benjamin Jarvis (2013) *The Rules of Thought* Oxford University Press.

- \_\_\_\_\_ (2009b) 'Thought-Experiment Intuitions and Truth in Fiction' *Philosophical Studies* 142: 221–46.
- Ichikawa, Jonathan, Benjamin Jarvis, and Katherine Rubin (2012) 'Pragmatic Encroachment and Belief-Desire Psychology' *Analytic Philosophy* 53(4): 327–343.
- Jackson, Frank (1982). 'Epiphenomenal Qualia' *Philosophical Quarterly* 32: 127–36.
- James, William (1956/1896) *The Will to Believe and Other Essays in Popular Philosophy*. New York: Dover.
- Kahneman, Daniel (2011) *Thinking Fast and Slow* Farrar, Straus and Giroux
- Kappel, Klemens (2010) 'On Saying that Someone Knows: Themes from Craig' *Social Epistemology* Adrian Haddock, Alan Millar, and Duncan Pritchard (eds.) Oxford University Press, pp. 69–88.
- Kawall, Jason (2013) 'Friendship and Epistemic Norms' *Philosophical Studies* 165(2): 349–370.
- Kaye, David (2009) 'Rounding up the Usual Suspects: A Legal and Logical Analysis of DNA Trawling Cases' *North Carolina Law Review* 87: 425–503.
- Kaye, David and George Sensabaugh (2011) 'Reference Guide on DNA Identification Evidence' in *Reference Manual on Scientific Evidence, 3<sup>rd</sup> Edition* Federal Judicial Center, pp. 129–210.
- Keller, Simon (2004) 'Friendship and Belief' *Philosophical Papers* 33(3): 329–351.
- Kelly, Tom (2002) 'The Rationality of Belief and Some Other Propositional Attitudes' *Philosophical Studies* 110: 163–196.
- Kelp, Chris (2011) 'What's the Point of "Knowledge" Anyway?' *Episteme* 8: 53–66.
- Kerr, Orin (2012) 'Why Courts Should Not Quantify Probable Cause' *The Political Heart of Criminal Procedure: Essays on Themes of William J. Stuntz* Michael Klarman, David Skeel, and Carol Steiker (eds.) Cambridge University Press, pp. 131–143.
- Kim, Brian (2017) 'Pragmatic Encroachment in Epistemology' *Philosophy Compass* 12(5): 1–14.
- Kitcher, Philip (2000) 'A Priori Knowledge Revisited' *New Essays on the A Priori* Paul Boghossian and Christopher Peacocke (eds.) Clarendon Press, pp. 65–91.
- Koehler, Jonathan (2001) 'When Do Courts Think Base Rate Statistics are Relevant?' *Jurimetrics Journal* 42: 373–402.
- Kornblith, Hilary (2011) 'Why Should We Care about the Concept of Knowledge?' *Episteme* 8(1): 38–52.
- \_\_\_\_\_ (2002) *Knowledge and its Place in Nature* Oxford University Press.
- Korsgaard, Christine (1983) 'Two Distinctions in Goodness' *Philosophical Review* 92(2): 169–95.
- Kripke, Saul (2011) *Philosophical Troubles: Collected Papers, Volume One* Oxford University Press.
- \_\_\_\_\_ (1972) *Naming and Necessity* Harvard University Press.
- Kunda, Ziva (1990) 'The Case for Motivated Reasoning' *Psychological Bulletin* 108(3): 480–98.
- Kusch, Martin (2013) 'Naturalized Epistemology and the Genealogy of Knowledge' *Contemporary Perspectives on Early Modern Philosophy: Nature and Norms in Thought* Martin Lenz and Anik Waldow (eds.) Springer, pp. 87–97.
- \_\_\_\_\_ (2011) 'Knowledge and Certainties in the Epistemic State of Nature' *Episteme* 8: 6–23.
- \_\_\_\_\_ (2009) 'Testimony and the Value of Knowledge' *Epistemic Value* Adrian Haddock, Alan Millar, and Duncan Pritchard (eds.) Oxford University Press, pp. 60–94.
- Kvanvig, Jonathan (2009) 'The Value of Understanding' *Epistemic Value* Adrian Haddock, Alan Millar, and Duncan Pritchard (eds.) Oxford University Press, pp. 95–112.
- \_\_\_\_\_ (2008) 'Pointless Truth' *Midwest Studies In Philosophy* 32(1): 199–212.

- \_\_\_\_\_ (2003) *The Value of Knowledge and the Pursuit of Understanding* Cambridge University Press.
- Kyburg, H.E., Jr. (1961) *Probability and the Logic of Rational Belief* Wesleyan University Press.
- Laudan, Larry (2012) 'Is it Finally Time to Put "Proof Beyond a Reasonable Doubt" Out to Pasture?' *Routledge Companion To Philosophy Of Law* Andrei Marmor (ed.) New York: Routledge, pp. 317–332
- \_\_\_\_\_ (2011) 'The Rules of Trial, Political Morality and the Costs of Error: Or, Is Proof Beyond a Reasonable Doubt Doing More Harm Than Good?' *Oxford Studies in Philosophy of Law* Leslie Green and Brian Leiter (eds.) Oxford University Press, pp. 195–227.
- Lehrer, Keith (1990) *Theory of Knowledge* Westview Press.
- Lerner, Craig (2003) 'The Reasonableness of Probable Cause' *Texas Law Review* 81: 951–1029.
- Leslie, Sarah-Jane (forthcoming) 'The Original Sin of Cognition: Fear, Prejudice and Generalization' *Journal of Philosophy*.
- Lewis, David (1996) 'Elusive Knowledge' *Australasian Journal of Philosophy* 74: 549–567.
- \_\_\_\_\_ (1973) *Counterfactuals* Harvard University Press.
- Lippke, Richard (2010) 'Punishing the Guilty, Not Punishing the Innocent' *Journal of Moral Philosophy* 7(4): 462–88.
- Littlejohn, Clayton (ms) 'Truth, Knowledge, and the Standard of Proof in Criminal Law'.
- \_\_\_\_\_ (2012) *Justification and the Truth-Connection* Cambridge University Press.
- MacBain, James (2004) 'Epistemic Analysis and the Possibility of Good Informants' *Principia*, 8(2): 193–211
- MacFarlane, John (2005) 'Knowledge Laundering: Testimony and Sensitive Invariantism' *Analysis* 65(2): 132–138.
- Malmgren, Anna-Sara (2011) 'Rationalism and the Content of Intuitive Judgments' *Mind* 120(478): 263–327.
- Marušić, Berislav (2015) *Evidence and Agency: Norms of Belief for Promising and Resolving* Oxford University Press.
- McCain, Kevin (2014) *Evidentialism and Epistemic Justification* Routledge.
- McCauliff, C.M.A., (1982) 'Burdens of Proof: Degrees of Belief, Quanta of Evidence, or Constitutional Guarantees?' *Vanderbilt Law Review* 35: 1293–1335.
- McCormick, Miriam (2015) *Believing Against the Evidence: Agency and the Ethics of Belief* New York: Routledge.
- Millikan, Ruth (1984) *Language, Thought, and Other Biological Categories: New Foundations for Realism* MIT Press.
- Mills, Charles W. (1997) *The Racial Contract* Cornell University Press.
- Morreau, Michael (1997) 'Fainthearted Conditionals' *Journal of Philosophy* 94(4): 187–211.
- Moss, Sarah (forthcoming) *Probabilistic Knowledge* Oxford University Press.
- Mulrine, Thomas (1997) 'Reasonable Doubt: How in the World Is It Defined?' *American University Journal of International Law and Policy* 12(1): 195–225.
- Nagel, Jennifer (2007) 'Epistemic Intuitions' *Philosophy Compass* 2(6): 792–819.
- Nelkin, Dana (2000). 'The Lottery Paradox, Knowledge, and Rationality' *Philosophical Review* 109(3): 373–409.
- Nesson, Charles (1985) 'The Evidence or the Event? On Judicial Proof and the Acceptability of Verdicts' *Harvard Law Review* 98: 1357–1392.
- \_\_\_\_\_ (1979) 'Reasonable Doubt and Permissive Inferences: The Value of

- Complexity' *Harvard Law Review* 92(6): 1187–1225.
- Newman, Jon (2006) 'Quantifying the Standard of Proof Beyond a Reasonable Doubt: A Comment on Three Comments' *Law, Probability, and Risk* 5: 267–269
- Nickel, Bernhard (2008) 'Generics and the Ways of Normality' *Linguistics and Philosophy* 31(6): 629–648.
- Nietzsche, Friedrich (1967) *On the Genealogy of Morals* and *Ecce Homo* trans. Walter Kaufmann, New York: Vintage Books.
- Nozick, Robert (1981) *Philosophical Explanations* Oxford University Press.
- Pace, Michael (2011) 'The Epistemic Value of Moral Considerations: Justification, Moral Encroachment, and James' 'Will To Believe' *Noûs* 45(2): 239–268.
- Pardo, M.S. (2013) 'The Nature and Purpose of Evidence Theory' *Vanderbilt Law Review* 66: 547–613.
- Pardo, M.S. and R.J. Allen (2008) 'Juridical Proof and the Best Explanation' *Law and Philosophy* 27: 223–268.
- Pettit, Philip (1999) 'A Theory of Normal and Ideal Conditions' *Philosophical Studies* 96: 21–44.
- Piller, Christian (2016) 'Evidentialism, Transparency, and Commitments' *Philosophical Issues* 26: 332–350.
- Poston, Ted (2014) *Reason and Explanation* Palgrave Macmillan.
- Preston-Roedder, Ryan (2013) 'Faith in Humanity' *Philosophy and Phenomenological Research* 87(3): 664–687.
- Pritchard, Duncan (forthcoming) 'Legal Risk, Legal Evidence and the Arithmetic of Criminal Justice' *Jurisprudence*.
- \_\_\_\_\_ (2012) 'Anti-Luck Virtue Epistemology' *Journal of Philosophy* 109: 247–279.
- \_\_\_\_\_ (2011a) 'The Genealogy of the Concept of Knowledge and Anti-Luck Virtue Epistemology' *Der Begriff des Wissens/The Concept of Knowledge* Stefan Tolkendorf (ed.) de Gruyter, pp. 159–178.
- \_\_\_\_\_ (2011b) 'What is the Swamping Problem?' *Reasons for Belief* Andrew Reisner and Asbjørn Steglich-Petersen (eds.) Cambridge University Press, pp. 244–59.
- \_\_\_\_\_ (2009) 'Knowledge' *Central Issues of Philosophy* John Shand (ed.) Wiley-Blackwell.
- \_\_\_\_\_ (2007) 'Recent Work on Epistemic Value' *American Philosophical Quarterly* 44: 85–110.
- Redmayne, Mike (2008) 'Exploring the Proof Paradoxes' *Legal Theory* 14: 281–309.
- Reichenbach, Hans (1949) *The Theory of Probability* University of California Press, Berkeley.
- Riggs, Wayne (2009) 'Knowledge, Understanding, and the Meno Requirement' *Epistemic Value* Adrian Haddock, Alan Millar, and Duncan Pritchard (eds.) Oxford University Press, pp. 331–8.
- \_\_\_\_\_ (2004) Review of 'The Value of Knowledge and the Pursuit of Understanding' by Jonathan Kvanvig, *Notre Dame Philosophical Reviews*.
- \_\_\_\_\_ (2003a) 'Balancing Our Epistemic Goals' *Noûs* 37(2): 342–352.
- \_\_\_\_\_ (2003b) 'Understanding "Virtue" and the Virtue of Understanding' *Intellectual Virtue: Perspectives from Ethics and Epistemology* Michael DePaul and Linda Zagzebski (eds.) Oxford University Press, pp. 203–226.
- \_\_\_\_\_ (2002) 'Reliability and the Value of Knowledge' *Philosophy and Phenomenological Research* 64: 79–96.
- Rinard, Susanna (2017) 'No Exception for Belief' *Philosophy and Phenomenological Research* 94(1): 121–143.
- \_\_\_\_\_ (2015) 'Against the New Evidentialists' *Philosophical Issues* 25(1): 208–223.

- Risinger, D. Michael (2010) 'Tragic Consequences of Deadly Dilemmas: A Response to Allen and Laudan' *Seton Hall Law Review* 40(3): 991–1021.
- Ross, Jacob and Mark Schroeder (2014) 'Belief, Credence, and Pragmatic Encroachment' *Philosophy and Phenomenological Research* 88(2): 259–288.
- Roth, Andrea (2010) 'Safety in Numbers? Deciding When DNA Alone is Enough to Convict' *New York University Law Review* 85: 1130–1185.
- Ryan, Sharon (2015) 'In Defense of Moral Evidentialism' *Logos and Episteme* 6(4): 405–427.
- Sand, Leonard and Danielle Rose (2003) 'Proof Beyond All Possible Doubt: Is There a Need for a Higher Burden of Proof when the Sentence may be Death?' *Chicago-Kent Law Review* 78: 1359–75.
- Schaffer, Jonathan (2006) 'The Irrelevance of the Subject: Against Subject-Sensitive Invariantism' *Philosophical Studies* 127: 87–107.
- Schmalbeck, Richard (1986) 'The Trouble with Statistical Evidence' *Law and Contemporary Problems* 49(3): 221–36.
- Schneeps, Leila and Coralie Colmez (2013) *Math on Trial* Basic Books.
- Schroeder, Mark (forthcoming) 'Rational Stability under Pragmatic Encroachment' *Episteme*.
- Schurz, Gerhard (2001) 'What is "Normal"? An Evolution-Theoretic Foundation for Normic Laws and Their Relation to Statistical Normality' *Philosophy of Science* 68(4): 476–97.
- Shah, Nishi (2006) 'A New Argument for Evidentialism' *Philosophical Quarterly* 56: 481–498.
- Shapin, Steven (1994) *A Social History of Truth: Civility and Science in Seventeenth-Century England* Chicago University Press.
- Simon, Rita James (1969) 'Judges' Translations of Burdens of Proof into Statements of Probability' *The Trial Lawyer's Guide* John Kennelly and James Chapman (eds.) Mundelein: Callaghan, pp. 103–114.
- Simon, Rita James and Linda Mahan (1971) 'Quantifying Burdens of Proof: A View from the Bench, the Jury, and the Classroom' *Law and Society Review* 5(3): 319–330.
- Slatkin, Montgomery, Yun S. Song and Erin Murphy (2009) 'Average Probability that a "Cold Hit" in a DNA Database Search Results in an Erroneous Attribution' *Journal of Forensic Sciences* 54(1): 22–27.
- Smith, Martin (ms) 'Justification, Normalcy and Evidential Probability' <http://philpapers.org/rec/SMIJNA-2> (accessed March 2016).
- \_\_\_\_\_ (2016) *Between Probability and Certainty: What Justifies Belief* Oxford University Press.
- \_\_\_\_\_ (2010) 'What Else Justification Could Be?' *Noûs* 44(1): 10–31.
- Smithies, Declan (2015) 'Why Justification Matters' *Epistemic Evaluation: Purposeful Epistemology* John Greco and David Henderson (eds.) Oxford University Press.
- Sosa, David (2006) 'Scepticism About Intuitions' *Philosophy* 81: 633–48.
- Sosa, Ernest (forthcoming) 'The Metaphysical Gettier Problem and the X-Phi Critique' *Explaining Knowledge: New Essays on the Gettier Problem* Rodrigo Borges, Claudio de Almeida, and Peter Klein (eds.) Oxford University Press.
- \_\_\_\_\_ (2009) 'A Defense of the Use of Intuitions in Philosophy' *Stitch and His Critics*, Dominic Murphy and Michael Bishop (eds.) Wiley-Blackwell, pp. 101–112.
- Spohn, Wolfgang (2014) 'The Epistemic Account of Ceteris Paribus Conditions' *European Journal for the Philosophy of Science* 4(3): 385–408.
- Staffel, Julia (forthcoming) 'Three Puzzles about Lotteries' *Lotteries, Knowledge, and Rational Belief: Essays on the Lottery Paradox* Igor Douven (ed.) Cambridge University Press.
- Stanley, Jason (2015) *How Propaganda Works* Princeton University Press.



- \_\_\_\_\_ (2011) *Know How* Oxford University Press.
- \_\_\_\_\_ (2005) *Knowledge and Practical Interests* Oxford University Press.
- Stanley, Jason and Timothy Williamson (2001) 'Knowing How' *Journal of Philosophy* 98(8): 411–444.
- Stein, Alex (2005) *Foundations of Evidence Law* Oxford University Press.
- Strawson, Peter (1992) 'Reduction or Connection? Basic Concepts' *Analysis and Metaphysics: An Introduction to Philosophy* Oxford University Press.
- Stroud, Barry (1979) 'Inference Belief and Understanding' *Mind* 86: 179–196.
- Stroud, Sarah (2006) 'Epistemic Partiality in Friendship' *Ethics* 116(3): 498–524.
- Sullivan, Emily (ms) 'Valuable Understanding without Achievement'
- Thomson, Judith Jarvis (1986) 'Liability and Individualized Evidence' *Law and Contemporary Problems* 49(3): 199–219.
- Tillers, Peter and Jonathan Gottfried (2006) 'Case comment—United States v. Copeland, 369 F. Supp. 2d275 (E.D.N.Y. 2005): A Collateral Attack on the Legal Maxim That Proof Beyond A Reasonable Doubt Is Unquantifiable?' *Law, Probability, and Risk* 5: 135–157.
- Tribe, Laurence (1971) 'Trial By Mathematics: Precision and Ritual in the Legal Process' *Harvard Law Review* 84(6): 1329–1393.
- Tversky, Amos and Daniel Kahneman (1982) 'Evidential Impact of Base Rates' *Judgment Under Uncertainty: Heuristics and Biases* Daniel Kahneman, Paul Slovic and Amos Tversky (eds.) Cambridge University Press, pp. 153–162.
- Vars, Frederick (2010) 'Toward a General Theory of Standards of Proof' *Catholic University Law Review* 60(1): 1–46.
- Venn, John (1866) *The Logic of Chance* Macmillan, London.
- Volokh, Alexander (1997) 'n Guilty Men' *University of Pennsylvania Law Review* 146: 173–216
- Walen, Alec (2015) 'Proof Beyond a Reasonable Doubt: A Balanced Retributive Account' *Louisiana Law Review* 76(2): 355–446.
- Weinberg, Jonathan, Shaun Nichols and Stephen Stich (2001) 'Normativity and Epistemic Intuitions' *Philosophical Topics* 29: 429–60.
- Weinstein, Jack and Ian Dewsbury (2006) 'Comment on the Meaning of "Proof Beyond a Reasonable Doubt"' *Law, Probability and Risk* 5: 167–173.
- Wells, Ida B. (1970) *Crusade for Justice: The Autobiography of Ida B. Wells* Alfreda M. Duster (ed.) University of Chicago Press.
- Westfall, Sandra Sobieraj (2014) 'The Obamas: How We Deal with Our Own Racist Experiences' *People Magazine*.
- Williams, Bernard (2002) *Truth and Truthfulness: An Essay in Genealogy* Princeton University Press.
- Williamson, Timothy (2009) 'Replies to Ichikawa, Martin and Weinberg' *Philosophical Studies* 145: 465–76.
- \_\_\_\_\_ (2007) *The Philosophy of Philosophy* Blackwell.
- \_\_\_\_\_ (2005) 'Armchair Philosophy, Metaphysical Modality and Counterfactual Thinking' *Proceedings of The Aristotelian Society* 105: 1–23.
- \_\_\_\_\_ (2005) 'Contextualism, Subject-Sensitive Invariantism, and Knowledge of Knowledge' *Philosophical Quarterly* 55(219): 213–35.
- \_\_\_\_\_ (2000) *Knowledge and its Limits* Oxford University Press.
- Worsnip, Alex (2015) 'Two Kinds of Stakes' *Pacific Philosophical Quarterly* 96: 307–324.
- Zagzebski, Linda (1996) *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge* Cambridge University Press.