

THE INTUITION OF KNOWING:
ITS BIOLOGICAL FUNCTION AND NATURAL TRIGGERING-CONDITIONS

By
Nathan Josephe Fox

Submitted in accordance with the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

The Faculty of Arts
(Philosophy)

Open University
September, 2017

Abstract

Over the last hundred years, competing and incompatible positions in relation to basic problems of knowledge and the use of the verb 'to know' have multiplied; and the prospect of a consensus solution emerging with respect to *any* of the problems has not seemed particularly good. We have a Gordian knot. Even so, I suggest that we also have a way to cut it. This will involve identifying why the cognitive mechanism that produces our intuitions of knowing evolved and was maintained (by natural selection), i.e., identifying the 'teleonomic function' of that cognitive mechanism. Also, it will involve predicting, on the basis of this teleonomic function, the triggering-conditions of these natural knowledge intuitions. In this thesis, I develop a *general theory of the origin, function and triggering-conditions of knowledge intuitions* that will allow us to cut that knot. That theory follows basic biological theory (including that which pertains to natural altruism) and also signal detection theory.

My theory identifies a number of different circumstances under which the triggering-conditions of knowledge intuitions are different. Strikingly, these different circumstances (and their associated triggering-conditions) map onto the different competing and incompatible epistemological positions to which I referred. This suggests that these positions are all correct within the boundaries of one of the circumstances that my theory identifies; and that the Gordian knot is largely the result of epistemologists claiming universal applicability of a theory that in fact only applies under particular circumstances. We cut the knot by specifying the different circumstances under which each of the different epistemological positions will hold, and the reason we should expect it to hold in just these circumstances, in light of the teleonomic function of knowledge intuitions.

For Susan, Hayley, Jonah and Rachel

Acknowledgements

I wish to thank my supervisors, Dr. Carolyn Price and Dr. Alex Barber, for the thoughtful and insightful guidance they gave me at every stage in the development of this thesis. Their objections and criticisms had no small part to play in forging the ideas presented here—although I am solely responsible for the thesis' shortcomings. Beyond theoretical and presentational guidance, I thank them for their encouragement, frankness, patience and obvious desire for this thesis to be the best thesis I could make it.

I also wish to thank all the professors and postgraduates at the Open University Philosophy Department for helping to create the warm, welcoming and collegial atmosphere that was always present.

Last, I wish to thank my wife, Susan Fox, for her love, consistent support of my work, and understanding of the commitment that I was required to make in order complete this project.

Table of Contents

Abstract	2
Acknowledgements	4
Chapter 1	8
1.1 Overview	8
1.2 Establishing a Conception of Knowledge Intuitions	10
1.2.1 Our Ordinary Picture of the World.....	10
1.2.2 The Use of 'Know'	12
1.2.3 Actions and Attitudes	13
1.3 Chapter Breakdown.....	16
Chapter 2: Two Independent Histories	18
2.1 Introduction.....	18
2.2 First History	19
2.3 Second History.....	28
2.4 Conclusion	39
Chapter 3: The Function of the Knowledge Intuition Producer	41
3.1 Idealization	41
3.2 The Function Hypothesis.....	45
3.2.1 Teleonomic vs Instrumental Functions	45
3.2.2 Cognition	47
3.2.3 The Environmental Complexity Thesis	49
3.3 A Second Function.....	54
3.4 The Hallmark of the KIP Is Flexibility	56
3.4.1 Contexts in Which Subjects Are Trusting Even with Weak Evidence	57
3.4.2 Contexts in Which Subjects Are Sceptical Even with Strong Evidence	57
3.4.3 Contexts in Which Subjects Become Abundantly Cautious	58
3.5 Second Phase of the Second History and Edward Craig.....	59
Chapter 4: Operating Principles of the Knowledge Intuition Producer	62
4.1 Introduction.....	62
4.2 What is Signal Detection Theory?	64
4.2.1 Overview.....	64
4.2.2 Factors that Determine Optimal Detection Thresholds.....	66
4.3 How Signal Detection Theory Applies to the KIP	71
4.3.1 Introduction.....	71

4.3.2	Direct Evidence	73
4.4	On Tracking Truth	96
Chapter 5:	How Do We Do It?	98
5.1	The Basic Method	98
5.2	The Adding Evidence Thesis.....	103
5.3	Basic Assumptions	107
5.4	Optimality	112
Chapter 6:	Altruistic Knowledge Intuitions.....	116
6.1	Introduction	116
6.2	Human Altruism	119
6.2.1	An Aside on Subject Sensitive Invariantism.....	122
6.2.2	An Aside on the Receiver’s Point of View.....	123
6.3	The Adding Evidence Thesis Holds for Senders	125
6.4	Communicating with Multiple Receivers.....	128
6.4.1	The Basic Method with Multiple Receivers	129
6.4.2	The Adding Evidence Thesis with Multiple Receivers.....	131
6.4.3	Alfred Hitchcock’s <i>Rear Window</i>	133
6.4.4	Broadcasts.....	135
Chapter 7:	Why Epistemology Bifurcated	142
7.1	Why the Literature Bifurcated in the Way That It Has	142
7.2	The Validity of Second History Theory.....	148
7.2.1	Reflective Access to the Process that Modulates Epistemic Standards	151
7.2.2	Interpretations of Acontextual Knowledge Attributions	152
7.2.3	An Observer Effect.....	155
7.2.4	Basic Assumptions About Knowledge.....	157
7.2.5	Three Experimental Philosophy Studies	160
7.2.6	Conclusion.....	162
Chapter 8:	Gettier-Style Cases.....	163
Chapter 9:	Conclusion.....	168
9.1	Summary.....	168
9.2	A Brief Comparison with Moderate Invariantism.....	170
9.3	Cutting the Knot.....	172
Appendix 1:	The Extent to Which the Environmental Complexity Thesis Overgeneralizes....	173
Appendix 2:	Jason Stanley’s Bank Cases	177
Appendix 3:	Biological Altruism	178

A.3.1	Helping Kin.....	178
A.3.2	Direct Reciprocation.....	179
A.3.3	Reputation-based Reciprocity.....	181
A.3.4	Commitment Theory.....	182
Appendix 4:	Cappelen & Lepore’s ‘Test (iii)’.....	185
Appendix 5:	Defined Terms.....	190
Works Cited	200

Chapter 1

1.1 Overview

The epistemological enterprise seems to be in a bad state. Over the last hundred years, theoretical positions on basic problems of knowledge, and on the use of the verb ‘to know’ and its cognates, have come apart rather than coalesced. In fact, it seems that every incompatible position has been staked out: (1) that the meaning of ‘know’ is not revealed in ordinary or everyday contexts, but it is in high-standard contexts (e.g., scientific contexts), (2) that the meaning of ‘know’ is revealed in ordinary contexts, but not in high-standard contexts, (3) that the evidential standard for knowledge is influenced by the practical interests of knowledge attributors, but not the subjects of those attributions, (4) that the standard is influenced by the practical interests of the subjects of knowledge attributions, but not the attributors’, and (5) that the standard is not influenced by the practical interests of any particular individual at all. Further, Gettier-Style Cases do not seem to fit comfortably in any theory of knowledge that has a wide base of support; and there is no agreement on how to analyse sceptical challenges (i.e., cases in which doubt is produced by the introduction of an extremely unlikely alternative hypothesis). (Title case terms, and capitalized acronyms which I use further on, receive entries in Appendix 5: Defined Terms.)

The prospect of a consensus solution or agreement emerging with respect to any of the basic problems of knowledge or the correct way to account for the use of ‘know’ has not seemed particularly good. We have a Gordian knot. Even so, I suggest that we also have a way to cut it. This will involve investigating the origin and function of knowing—or, better, the origin and function of the intuition of knowing.

Let me be clear about the extension I have in mind when I use the expression ‘intuition of knowing’. When we have such an intuition, *in that moment*, we believe without doubt, and we are certain that we would do well to act or form an attitude as if that which we take ourselves as knowing is true. Further, it seems as if these intuitions arise very often, with quick succession; and, also, very often go unnoticed. For example, this morning I took myself as knowing that my alarm clock was working properly when it went off, took myself as knowing that the clothes I planned to wear would be appropriate, took myself as knowing that each of the reports that I got from my trusted news outlet was reasonably accurate, etc. Additionally,

notice that these intuitions are, in the moment that they arise, decisive and absolute. In the moment that I received a news report from my trusted news outlet, I had no doubt that it was reasonably accurate. Unlike a degree of belief, the intuition of knowing seems to follow from a firm and final decision—even if such a decision is not reflectively accessible and even if a contradictory intuition may arise a moment later. So, one either intuits that one knows or not, either there is a doubt in the moment or not, and there is no middle ground. Call these intuitions, ‘Knowledge Intuitions’.

My investigation into the origin and function of Knowledge Intuitions is based on the assumption that, regardless of the context in which they arise, they are produced by a functionally discrete cognitive mechanism that evolved by natural selection. (This is an intentional idealization, the nature of which I discuss in §3.1.) Such a mechanism will have a Teleonomic Function, i.e., a function that explains why it evolved and was maintained, and what role it has played in our survival and reproduction. I call this cognitive device the ‘Knowledge Intuition Producer’ or KIP.

My investigation starts by hypothesizing the Teleonomic Function of the KIP. I then develop a general theory of Knowledge Intuitions based on that function. Development of the theory is consistent with basic biological theory and it draws on Signal Detection Theory (which I discuss in §4.2). My theory of Knowledge Intuitions purports to explain why, given the evidence, we have or fail to have intuitions of knowing that particular world conditions obtain, and purports to identify the triggering-conditions of these natural Knowledge Intuitions. (To be clear, the triggering-conditions indicate the conditions under which Knowledge Intuitions are predicted to be produced *under ordinary circumstances*. The process could well be derailed by extraordinary events.) Furthermore, I find (in §1.2.2) that the *use of ‘know’* is dependent on Knowledge Intuitions. So, if this last finding is correct, then my explanatory theory may serve also as a general explanatory theory of the use of ‘know’.

A key prediction of my theory is that the triggering-conditions for Knowledge Intuitions depend on certain circumstances— such as whether one is focused on one’s own practical interests or focused on helping others, and whether one is envisaging communicating to a large audience. I use the predicted effects of such circumstances to understand the state of epistemology. In particular, I find that for any one of the incompatible positions ((1) to (5) above) it is possible to describe the set of circumstances under which the position works, and other circumstances in which it does not. So, it looks like all of those basic positions are correct within the boundaries of one of my sets of circumstances. The picture emerges that epistemology seems to be in a

bad state and that we have a Gordian knot mainly because epistemologists overgeneralize, as a matter of course. They claim universal applicability of a theory that in fact applies to only a subset (though usually an important subset) of knowledge cases. We cut the knot by specifying the different circumstances under which each of the different epistemological positions will hold, and the reason we should expect it to hold in just these circumstances, in light of the Teleonomic Function of Knowledge Intuitions.

I note as well that the developed theory is consistent with all the knowledge cases that I tested it against, including Gettier-Style Cases (Chapter 8), sceptical challenges (and paradoxes) (§7.2.4), and even cases that feature rational acceptance (§4.3.2.4).

Moreover, the picture that emerges from the theory I develop suggests a framework of sorts for epistemology—a framework described by the different circumstances under which Knowledge Intuitions have different triggering-conditions. Viewing the field of epistemology through the lens of this framework allows us to see and appreciate a far greater richness in the field than we do otherwise. If one assumes that all but one of a group of competing theories is incorrect, when in fact several of those theories shed light on a range of cases, one will have failed to appreciate the insight offered by those theories that one has dismissed, and may well also dismiss valuable observations and findings on which those theories are based.

1.2 Establishing a Conception of Knowledge Intuitions

The basic conception of Knowledge Intuitions that I will present in my developed model is consistent with a large body of *prima facie* evidence. On that conception, Knowledge Intuitions are mental states that play a critical role in determining, at once (a) the picture we have of the world, (b) our use of ‘know’, and (c) our actions and attitudes.¹ In the following three sections, I present some of the *prima facie* evidence that supports this view.

1.2.1 Our Ordinary Picture of the World

¹ A Knowledge Intuition appears to be an entirely different mental state than the ‘feeling of knowing’, at least in the way many conceptualize the latter. It seems the standard view is that the latter feeling is a metamemory state. It is a feeling of having memories of a certain kind without having actually retrieved those memories.

Very many of our ordinary Knowledge Intuitions pertain to the existence of non-perceptible, though inferable, world conditions. I refer here specifically to ordinary world conditions that are relevant to our being able to navigate the world, i.e., to perform actions and form attitudes that would advance our interests. Our taking ourselves as knowing that these conditions hold greatly contributes to the ordinary picture that we have of the world. Let us look at *some* of these non-perceptible world conditions to which Knowledge Intuitions often pertain.

We may take ourselves as knowing that objects have simple, non-perceptible (though inferable) physical properties. For example, we may take ourselves as knowing that the surface and interior of the wood shelf is made from the same material, that the car parked outside is heavy (even though we have not tried to lift it), and that the glass in the doctor's office is transparent and reveals an outdoor scene (even though we have not checked to see whether it is not in fact the front of a television showing an outdoor scene). We may take ourselves as knowing that objects have complex, non-perceptible (though inferable) physical properties. For example, we may take ourselves as knowing that the watch is working properly or that the elevator has design features that make it safe to ride. We may take ourselves as knowing that objects have non-perceptible (though inferable) abstract properties. For example, we may take ourselves as knowing that the man is wearing his own shirt (as opposed to someone else's), that I am located in my country's capital city, or that at least one of several nearby gas stations is open for business at this moment. We may take ourselves as knowing that particular people have non-perceptible (though inferable) attributes. For example, we may take ourselves as knowing that the woman has my best interests at heart, that the man is looking out for himself, that the accused is guilty, that the woman knows that she made a mistake, and that the politician is truly on my side. We may take ourselves as knowing that organizations or systems have non-perceptible (though inferable) attributes. For example, we may take ourselves as knowing that the political system is rigged, or that the health care system is doing a good job. We may take ourselves as knowing that non-perceptible (though inferable) causal events have occurred or that future events will occur. For example, we may take ourselves as knowing that the steep rise in interest rates depressed the housing market, that the bank will be open this Saturday, or that the sun will rise in the East tomorrow.

As I suggested above, Knowledge Intuitions seem to follow from a firm decision about the existence of a world condition. There is no doubt *in the moment* that they arise. As a result, Knowledge Intuitions, produced in the moment by the KIP, seem to help us get an ordinary picture of the world that is relatively unambiguous—one that is filled with distinct solid and

weighty objects, filled with people who have good or bad intentions, filled with systems and organizations that may work well or badly, and filled with events that are causally connected though those connections may be non-perceptible.

1.2.2 The Use of 'Know'

I suggest that the ordinary use of 'know', in knowledge attributions and denials, generally depends on Knowledge Intuitions; the use depends on these *mental events*. (This claim contrasts with claims that suggest that the use of 'know' reflects a particular concept, definition, or theory of knowledge.) However, I set aside entirely the use of 'know' in (a) knowledge attributions and denials that are either intended to mislead or to be ironic, and in (b) second-person and third-person knowledge attributions and denials that are made in order to explain successful behaviour, or unsuccessful behaviour respectively².

We have compelling *prima facie* evidence that the use of 'know' depends on the production of Knowledge Intuitions, i.e., it depends on not having doubt in the moment. Consider the clear oddity of the sentence, 'I know the train company will put me up in a hotel if I miss the connection, though I do have some doubts on the matter'; or the sentence, 'I do not know that the bank will be open on Saturday, though I do not have any doubt that it will'. When we have doubt, even the slightest doubt, I suggest we ordinarily either flat out deny knowing, or we hedge and avoid claiming to know. Moreover, there are many different ways in which we may imply the presence of a doubt: 'I *think* that *p*', 'I *believe* that *p*' (the emphasis on the word 'believe' suggesting only a degree of belief), 'It *may be* that *p*', 'I'm *pretty sure* that *p*', 'It is *plausible* that *p*', 'I *rationaly accept* that *p*', 'Let us *assume for the moment* that *p*', or 'I don't know that *p*, but do know that *probably p*'. Moreover, when we have, in the moment, the intuition of knowing that *p* and have no doubt whatsoever, it seems that we ordinarily either claim to know, or avoid denying that we know.

² I refer here to circumstances in which (a) a subject behaves as if *p*, (b) the behaviour proves successful, and (c) the success is explained by an observer using a second-person and third-person knowledge attribution. The observer may not intuit knowing that *p*, but may nonetheless attribute knowledge; and she does so in order to express the proposition that the success of the behaviour is a consequence of the subject being in possession of the truth. Consider, this example due to Peter Godfrey-Smith: 'Jake arrived safely because he knew the lie of the land. He thought the quicksand was to the right, and that is where it was' (1998, p.172). The observer may not have taken herself as knowing that the quicksand was to the right.

Notice too that 'know' has a linguistic characteristic that makes it a suitable linguistic sign for Knowledge Intuitions. Jason Stanley finds that the word 'know' does *not* come in 'varying degrees of strength' (2005, p.36). 'Know' is not gradable. This works hand-in-glove with Knowledge Intuitions insofar as, as I suggested earlier, the latter are decisive and absolute, and seem to follow from a firm and final decision that a world condition obtains.

The thesis that the use of 'know' depends on Knowledge Intuitions provides a basis for an explanation of data that is at the centre of an important debate in epistemology. That data is the appearance that the standard of evidence reflected in the use of 'know' moves up and down, and this modulation seems to depend on the context in which knowledge attributions and denials are uttered. What principles are these modulations based on? What does the data tell us about the meaning of 'know'? My investigation will suggest that the KIP is a cognitive device that has particular operating characteristics, which characteristics are consistent with its Teleonomic Function. So, the production of Knowledge Intuitions has these characteristics; and, if the use of 'know' depends on Knowledge Intuitions, then this use too will reflect these characteristics. I will show that these operating characteristics of the KIP are consistent with and explain the data in question.

1.2.3 Actions and Attitudes

I suggested that Knowledge Intuitions are, in the moment that they arise, decisive and absolute, and contrasted them with degrees of belief. We have no doubt at all in the moment when Knowledge Intuitions arise, whereas we have doubts that vary in intensity when we have a degree of belief. All other things being equal, the decisiveness of Knowledge Intuitions appears to allow a flexibility in our behaviour that we do not have when we have a degree of belief. In fact, the range of behaviours that may follow from knowing seems far greater. For example, if I have a degree of belief that the stock market will crash (and other conditions are met), I may sell my stocks. However, if I take myself as knowing that it will crash and have no doubt about it at all, then I might *not only* sell my stocks, but also see if I can bet against the stock market (or 'short' the market). In fact, when one takes oneself as knowing that p , and so long as one continues to take oneself as knowing that p , i.e. so long as doubts don't creep back in, it seems that there are no intrinsic limitations in relation to one's behaving as if p (beside physical ones). If I take myself as knowing and have no doubt at all that the man will attempt to kill me, my behaviour may well breach ordinary boundaries. (Further on, we find that this is a reason, ultimately, that the elimination of all doubt, under some circumstances, requires very strong

evidence.) Moreover, this allowance of behavioural possibilities without intrinsic limitation occurs for us in the quick succession with which Knowledge Intuitions arise throughout the day. For example, if I take myself as knowing that the report I received from my trusted news outlet is accurate, then that Knowledge Intuition will allow the performance of a wide range of possible actions and the formation of a wide range of possible attitudes consistent with my taking myself as knowing that the news report is accurate. So, it seems that Knowledge Intuitions have some *special* role in controlling our actions and attitudes. (The foregoing might be read as suggesting that some behaviour may be controlled by degrees of belief. Moreover, it seems that our behaviour may also be controlled by the rational acceptance of a proposition. Even so, further on, in §4.3.2.4, a close examination of these cases reveal that the behaviour in these cases depends entirely on certain ancillary Knowledge Intuitions that generally go unnoticed. However, I set this aside for the moment because, as I suggested, my intention at this stage is merely to present some *prima facie* evidence that supports the conception of Knowledge Intuitions that emerges further on.)

If my present claim that Knowledge Intuitions have a special role in controlling our actions and attitudes is correct, and my earlier claim that Knowledge Intuitions give voice to the ordinary use of 'know' is correct, *then* Knowledge Intuitions are a *common cause* of the performance of actions, the formation of attitudes and of the use of 'know'. So, all three should be consistent. The ordinary utterance 'I know that *p*' should be a highly reliable indicator of the speaker acting or having an attitude as if *p*. (I note that *attitudes* as if *p* seem to very often relate to actions. For example, it seems typical that an attitude as if *p* involves being positively disposed to performing a particular range of possible actions in the event that a different proposition or set of propositions is subsequently taken as being known. If I take myself as knowing that it is a practical possibility that my house may one day catch fire, then I may form a positive attitude about buying fire insurance. If I later learn that the cost of fire insurance is low, then I will be prepared to buy it. I look at cases of this kind in §4.3.2.4.)

The general hypothesis that there is a deep connection between the use of 'know' and the performance of actions is supported by an investigation conducted by John Hawthorne and Jason Stanley. (It is just this general hypothesis that is supported. There are a number of important differences between our accounts.) They find that 'our ordinary folk appraisals [which are reflected in ordinary discourse] of the behaviour of others suggest that the concept of knowledge [that is reflected in ordinary discourse] is intimately intertwined with the rationality of action' (2008, p.571). Further, 'it bears emphasis that (in English at least) it is considerably more natural to appraise behaviour with the verb 'know' than the phrase "justified

belief”, or even “reasonable belief” (p.573). Hawthorne & Stanley suggest, for example, that the use of ‘know’ in criticisms of action is a norm: ‘When someone acts on a belief that does not amount to knowledge, she violates the norm, and hence is subject to criticism’ (p.577). A mother berates her son who has just cancelled his health insurance on the basis that he does not know that he will not fall ill (pp.571-572).

I note also that if, in fact, Knowledge Intuitions have the special role I described in controlling our actions and attitudes, then the operating characteristics of the KIP would have a profound effect on our ability to advance our interests, and on how well we do in life. Those characteristics would determine the circumstances under which intrinsic limitations on our behaviour are lifted. This situation seems to suggest that there could be practical reasons to modulate the standard of evidence that we set for producing Knowledge Intuitions. Compare the situation of (a) an agent *A* who will perform some everyday inconsequential action if she intuits knowing some ordinary proposition, and (b) an agent *B* who will respond accordingly if she intuits knowing that the man will attempt to kill her. *B*’s circumstances would call for a higher standard of evidence. In fact, the theory I will develop (particularly in §4.2, §4.3, and §5.2)—which, again, is based on a hypothesis about the biological function of Knowledge Intuitions and on Signal Detection Theory—predicts that the KIP may modulate evidential standards substantially, and in quick succession (discussed in §3.4). Moreover, the theory predicts a particular set of circumstantial factors that together determine the standard of evidence required to produce Knowledge Intuitions. I refer to this standard as the ‘Epistemic Standard’.

However, if ‘know’ in fact gives voice to our Knowledge Intuitions, then the standard of evidence that appears to be reflected in our use of ‘know’ should be our Epistemic Standard; and our Epistemic Standard should determine the truth conditions of knowledge attributions and denials. For the sake of presentational simplicity, I will use the term ‘Epistemic Standard’ not only to refer to the standard that one’s evidence must meet in order to produce Knowledge Intuitions, but also to refer to the standard that we seem to apply when we use the word ‘know’. (As an accommodation to presentational simplicity, this is not ideal: I am, to a large extent, adopting the working hypothesis that the word ‘know’ gives voice to our Knowledge Intuitions even before developing my theory of Knowledge Intuitions. Even so, making this accommodation seems worthwhile.)

1.3 Chapter Breakdown

I suggested that theoretical positions on knowledge and the use of ‘know’ have come apart, rather than coalesced. In Chapter 2, I take a look at the epistemological literature over the last century—focusing on just this situation, and presenting the literature in a way that will facilitate later analysis and discussion based on my developed model.

Chapters 3 and 4 focus on theory that will serve as the foundation of my model. In Chapter 3, I present a formal hypothesis of the Teleonomic Function of the KIP (that explains why it evolved and was maintained). This is based on basic biological theory. In Chapter 4, I look at what the operating principles of the KIP would have to be in order for it to perform the Teleonomic Function described in Chapter 3. This work is based on Signal Detection Theory. These operating principles seem to tell us a great deal about the triggering-conditions of Knowledge Intuitions.

In Chapter 5, I focus on practical matters—particularly on how we may go about adjudicating propositions in a way that takes the KIP’s operating principles into account. I am concerned here that it is plausible that we are actually able to adjudicate a quick succession of propositions in ‘real-time’ in a way that accords with the KIP’s operating principles. Also in this chapter, I identify an additional contextual factor in the production of Knowledge Intuitions that is not predicted by Signal Detection Theory, though it follows from my hypothesis of the KIP’s function.

In Chapter 6, I add a crucial component to my model, which is based on biological findings related to human altruism. The analysis that follows leads to descriptions of different sets of circumstances under which different basic positions in epistemology hold.

I note that, over the course of Chapters 5 and 6, I test the triggering-conditions of Knowledge Intuitions that are predicted on the basis of my hypothesis about the KIP’s function. I do so by seeing whether they apply in a wide range of knowledge cases. They do; and, because it is unlikely that the factors would all be applicable if the hypothesis on which they were based was wholly false, the successful applications constitute support for my hypothesis about the KIP’s function.³

³ Broadly similar methodologies have been used by Susan Haslanger (1999), Edward Craig (1999), and Justin Fisher (2006). They develop a hypothesis about the function of the subject of their investigation, work out the implications of that hypothesis, and finally test the implications.

In Chapter 7, I return to the epistemological literature, and look at it from the perspective of my developed theory. That theory allows me to explain profound differences in the variety of theoretical views that are expressed. Also, (in §7.2) I examine a number of objections to the thesis that the Epistemic Standards we seem to set are context-dependent. I identify what appear to be critical weaknesses in those objections.

Chapter 8 is a very short analysis of Gettier-Style Cases. In it, I show how my developed theory predicts our sense that the subjects in those cases do not know, even though they have justified true belief.

In Chapter 9, I briefly look back at what I think my investigation accomplished, and say a few words about the effect that I hope the investigation would have on the epistemological enterprise.

Chapter 2: Two Independent Histories

2.1 Introduction

It seems that over the last century the epistemological literature has bifurcated—creating two separate, distinct, individually coherent and mutually incompatible histories within epistemology. I call them the First History and the Second History. My developed theory of Knowledge Intuitions explains this occurrence. The theory predicts the emergence of two separate and independent research programs within epistemology that, I suggest, correspond with these two histories.

It will be useful to separate my telling of the two histories, which I do in this chapter, from my explanation of the predicted emergence of two separate and independent research programs, which I do in Chapter 7. The former will both set up a puzzle, the apparent impasse I have just described, and situate my theory in the field. The latter explanation is made on the basis of the theory that I develop in the intervening chapters.

Even so, it will be helpful to first roughly outline the principal aims of the two researcher programs (which are defended in Chapter 7). One of these research programs has the principal aim of, roughly, helping those who broadcast messages widely identify propositions that are unlikely to prove false. The other has the principal aim of shedding light on the circumstances under which we take ourselves as knowing or not knowing that particular world conditions obtain, and also to shed light on the ordinary use of ‘know’.

In this chapter, I provide a very brief account of these two histories. To keep things focused, I emphasize the internal coherence and independence of each of the research programs, and emphasize fundamental differences in the theoretical approaches that are taken by the theorists that contribute to the different programs. This requires that I set aside a wide range of otherwise important contributions to the two histories.

2.2 First History

It seems that First History theorists generally have the sense that a certain responsibility, obligation or duty attends the claim to possess knowledge. This view and the reason for having it is brought to life in W. K. Clifford's extended metaphor about the voyage of ship filled with emigrants:

A shipowner was about to send to sea an emigrant-ship. He knew that she was old, and not overwell built at the first; that she had seen many seas and climes, and often had needed repairs. Doubts had been suggested to him that possibly she was not seaworthy. These doubts preyed upon his mind, and made him unhappy; he thought that perhaps he ought to have her thoroughly overhauled and refitted, even though this should put him at great expense. Before the ship sailed, however, he succeeded in overcoming these melancholy reflections. He said to himself that she had gone safely through so many voyages and weathered so many storms that it was idle to suppose she would not come safely home from this trip also. He would put his trust in Providence, which could hardly fail to protect all these unhappy families that were leaving their fatherland to seek for better times elsewhere. He would dismiss from his mind all ungenerous suspicions about the honesty of builders and contractors. In such ways he acquired a sincere and comfortable conviction that his vessel was thoroughly safe and seaworthy; he watched her departure with a light heart, and benevolent wishes for the success of the exiles in their strange new home that was to be; and he got his insurance-money when she went down in mid-ocean and told no tales.

What shall we say of him? Surely this, that he was verily guilty of the death of those men. It is admitted that he did sincerely believe in the soundness of his ship; but the sincerity of his conviction can in no wise help him, because he had no right to believe on such evidence as was before him. He had acquired his belief not by honestly earning it in patient investigation, but by stifling his doubts. And although in the end he may have felt so sure about it that he could not think otherwise, yet inasmuch as he had knowingly and willingly worked himself into that frame of mind, he must be held responsible for it. (1876, p.289)

Clifford draws our attention to the multiplication of harms that may follow from broadcasting false messages widely. Many may be harmed in such situations by setting doubt aside.

Even so, doubt, and particularly Cartesian doubt, about the veracity of sense perceptions, and doubt even about the existence of an outside world, undermines the very possibility of empirical knowledge; and First History theorists generally *also* have the sense that it is of critical importance that some empirical knowledge be possible, that we can come to know that many different world conditions obtain.

Early in the last century, a number of epistemologists concluded that Cartesian doubt may be set aside. Bertrand Russell, for example, explains that '[p]hilosophy cannot boast of having achieved such a degree of certainty that it can have authority to condemn the facts of experience . . . Universal scepticism, though logically irrefutable, is practically barren' (1914, p.67). Moreover, 'the immediate facts perceived by sight or touch or hearing do not need to be proved by argument, but are completely self-evident' (p.68).

Nonetheless, it seems that most of what we care to know about is subject to doubt; and this is because most of what we care to know about depends on more than merely accepting the veracity of sense perceptions. H. H. Price's alludes to the nature of the problem in this famous passage:

When I see a tomato there is much that I can doubt. I can doubt whether it is a tomato that I am seeing, and not a cleverly painted piece of wax. I can doubt whether there is any material thing there at all. Perhaps what I took to be a tomato was really a reflection . . . One thing however I cannot doubt: that there exists a red patch of a round and somewhat bulgy shape, standing out from a background of other colour-patches, and having a certain visual depth, and that this whole field of colour is directly present to my consciousness. (1932, p.3)

Seeing the red patch of a round and somewhat bulgy shape is consistent not only with a tomato being there, but also consistent with a wide variety of other states of affairs, including a painted piece of wax being there. So, empirical knowledge requires a series of correct deductive and / or inductive inferences (even if they are mostly made unconsciously). To know that there is a tomato there, Price may have to correctly infer that (a) it is possible that a real tomato could be there, (i.e., there are tomatoes where he lives, etc.), (b) it is not possible that a wax tomato could be there, or (c) it is not possible that a wax tomato could look so realistic that he could mistake it for a real tomato. Each such inductive inference introduces a source of possible error that undermines the possibility of knowing the proposition in question.

So, at this stage in First History theory, many epistemologists hold the view that beliefs that are not based on these sorts of inferences may have a firm foundation. For example, the non-inferential belief that there is a red patch of a round and somewhat bulgy shape may have a firm foundation based on sense perception. However, beliefs that are based on these sorts of inferences may not. The 'inferential belief' that there is a tomato there may not. Bertrand Russell observes that 'what does not go beyond our own personal sensible acquaintance must be for us the most certain' (1914, p.67).

C. I. Lewis (1946, p.343) identifies a general approach that we can take to brace the foundations of our inferential beliefs. He suggests that the 'congruence' or the 'coherence' of different pieces of evidence confer a much higher probability upon hypotheses than is generally assumed—high enough, in many instances, for knowledge. Lewis is particularly focused on pieces of evidence that are connected in a particular way, viz., they have all arisen as a consequence of a particular prior state or event:

A point which particularly should draw our attention, is what can happen when various consequences [i.e., potential pieces of evidence] of a single hypothesis are found to be true. If different consequences be verified, or independently confirmed, the conjoint fact of these separate confirmations may increase the probability of the hypothesis in a degree notably greater than that which any one of them alone would give . . . (p.344)

Lewis illustrates his point with an account of

relatively unreliable witnesses who independently tell the same story . . . For any one of these reports, taken singly, the extent to which it confirms what is reported may be slight. And antecedently, the probability of what is reported may also be small. But congruence of the reports establishes a high probability of what they agree upon, by principles of probability determination which are familiar: On any other hypothesis than that of truth-telling, this agreement is highly unlikely; the story any one false witness might tell being one out of so very large a number of equally possible choices . . . And the one hypothesis which itself is congruent with this agreement becomes thereby commensurably well established. (p.346)

Further, Lewis suggests that this effect is multiplied when the hypothesis is consistent with antecedent beliefs, producing a wider coherence:⁴

We could hardly be too impressed with the part which is played, in determining what we shall believe . . . by the congruence of any statement in question with our antecedent beliefs and with what has already some degree of confirmation. (p.349)

If the probability of a hypothesis is sufficiently high, then it is rational to believe it, and to take it as being known. We know it is a tomato because of a congruence or coherence of our evidence. In denying the possession of such knowledge, 'we would repudiate all thought and action and every significance of living' (Lewis, p.362).

Still, Lewis' account does not tell us *how strong* the total evidence for a proposition must be in order to properly consider it to be known. Moreover, it seems that an answer that has a comfortable place within the First History will not be easily found. Even so, A. J. Ayer proposes an answer that has a comfortable place in the Second History. (I bring it up at this juncture because the subsequent reaction from First History theorists is an important part of the First History, and helps distinguish First History from Second History theory.)

Where there are recognized criteria for deciding when one has the right to be sure, anyone who insists that their being satisfied is still not enough for knowledge may be accused . . . of misusing the verb 'to know'. (1956, p.34)

Here, Ayer suggests that the criteria that we happen to recognize as being adequate for knowledge, in fact constitutes an adequate standard. Ayer's suggestion, along with those of other contemporaries⁵, is widely viewed as reducing knowledge to justified true belief.⁶ More precisely, it is widely viewed as reducing knowledge to true belief plus whatever happens to be ordinarily viewed as sufficient justification.

This last reduction is firmly rejected by First History theorists: the coincidence of truth, belief and ordinary justification is not sufficient for knowledge. In fact, some eight years prior to Ayer's proposal, in a 1948 publication, Bertrand Russell presents what appears to be a counterexample to the reduction:

⁴ Bertrand Russell (in a paper originally published 1948) traces coherence theory to Hegel (2009, p.141).

⁵ For example, see Roderick Chisholm (1957, p.16).

⁶ For example, see Edmund Gettier (1963, p.121).

There is the man who looks at a clock which is not going, though he thinks it is, and who happens to look at it the moment when it is right; this man acquires a true belief as to the time of day, but cannot be said to have knowledge. (2009, p.140)

The man possesses true belief plus ordinary justification on the basis of seeing a clock clearly. Still, he does not have knowledge. Further, it seems that the reason that we do not intuit knowing in such cases is that the evidence fails to have an appropriate connection with that which makes the belief true. The belief is true only by accident. Look a little more closely: What makes the man's belief about the time true is that it corresponds to the actual daily progress of the earth's rotation. For example, if he believed it was noon and the belief was true, then the earth would have completed half of its daily rotation; and, at 6 p.m., it would have completed three-quarters of its daily rotation. However, the man's evidence, i.e., the time indicated by the clock, does not have the appropriate connection with the daily progress of the earth's rotation. The time indicated by the clock *is not a consequence* of the clock accurately tracking the progress of the earth's rotation.

Edmund Gettier (1963) echoes Russell's point and provides a number of additional and much discussed examples. Cases of this kind are now standardly referred to as Gettier cases. I refer to them as Gettier-Style Cases.⁷

D.M. Armstrong acknowledges that Gettier-Style Cases serve as counterexamples to the hypothesis that knowledge is true belief plus ordinary justification. However, as a First History theorist, he is puzzled that epistemologists would be excited by disproving the hypothesis. Gettier-Style Cases feature merely 'justifiably believed grounds, in the "ordinary language" sense of "justifiably believed"' (1973, p.152). He adds,

⁷ Cases in which the subject has met a very high evidential standard, such as the standard used in scientific contexts, are *not* standardly offered as examples of Gettier-Style Cases. This seems to follow from the fact that, in these cases, the probability of the questioned proposition being accidentally true is extremely low, and sometimes approaching nil. So, such cases just seem too hypothetical in the contexts in which Gettier-Style Cases come up. Consider, for example, how low the probability might be of (a) a man finding on a day like any other that four different clocks (which are all powered in different ways) agree on the time, (b) the four clocks were all stopped at the moment that the man looked at them, and (c) the four clocks all told the correct time. Or, consider how low the probability might be of the following proposition being true: the number of elementary particles per mole of a substance is 6.022×10^{23} (i.e., Avogadro's number) is accidentally true. If it were only accidentally true, then the multiple experiments that have indicated Avogadro's number are all faulty and indicated the correct number only by accident.

But because possession of such grounds could not constitute possession of *knowledge* I should have thought it obvious that they are too weak to serve as suitable grounds. It is not surprising, therefore, that Gettier is able to construct examples where a true belief is justified in an ordinary sense of 'justified', but the true belief is clearly not a case of knowledge. (pp.152-153)

Further, Armstrong investigates an intuitive remedy to the evidential problem on which Gettier-Style Cases are based. He investigates gathering additional evidence that would show that the initial evidence has the appropriate connection to that which makes the belief true. For example, the evidential remedy for the man who looks at the stopped clock would be to gather additional evidence that might confirm that the clock is working properly. However, as Armstrong point out, the remedy is attended by the same evidential problem: The possibility remains that the new evidence, like the initial evidence, may not have the appropriate connection to that which makes the belief true (e.g., the evidence that suggests that the clock is working could be Gettiered). Moreover, further additional evidence that supported the secondary evidence is also subject to the same problem. So, the situation leads to an indefinite 'regress of reasons'.

Armstrong concludes that ultimately knowledge depends on meeting a condition that knowers may not know that they have met. Focusing initially on *non-inferential* knowledge such as that 'there is something red and round over there' (1973, p.163), he writes,

what makes a true non-inferential belief a case of knowledge is some natural relation which holds between the belief-state . . . and the situation which makes the belief true. It is a matter of a certain relation holding between the believer and the world. (p.157)

Non-inferential knowledge depends on the existence of this relation; and knowers may well not know whether that relation exists. This constitutes an externalistic condition on knowledge. Let us look at this relation a little more carefully.

The relation between the belief-state and the situation that makes the belief true is satisfied only if the belief-state is *empirically reliable* (Armstrong, 1973, p.159); and to be empirically reliable, there 'must be a law-like connection' (p.166) between the belief-state and the situation that makes the belief true. Non-inferential knowledge depends on beliefs that covary with the actual situation as 'the temperature readings given by a thermometer' (p.166) covaries

with actual temperature. So, roughly, non-inferential knowledge is true belief acquired using a reliable belief formation process.

Armstrong expands his framework to include inferential knowledge (e.g., there is a tomato there). Inferential knowledge must have a law-like connection to non-inferential knowledge (which by definition has a law-like connection to that which makes the belief true). For instance, if it happened that round red objects nomically indicated tomatoes, then knowing that there is a round red object there allows me to know that there is a tomato there.⁸ So, inferential knowledge is roughly true belief acquired using a reliable belief formation process that involves both inferential and non-inferential beliefs. Barry Stroud, in a 1989 publication, observes that the

promise of the new 'externalist' strategy is that it would avoid the regress [problem by allowing that a] person who knows something does not himself have to know that what he has got in his prior evidential base amounts to knowledge in the domain in question. (2001, p.137)

Armstrong's theory confirms the possibility of knowledge. However, it suggests that we may well not know that we know in particular instances.

In the 1980s, First History theorists have sharply divergent views on whether there is an externalistic condition on knowledge, as Armstrong suggests. On one hand, many epistemologists write positively about the general approach and develop externalistic theories.⁹ On the other, there is a torrent of First History writings that goes against externalism. (I look at this more carefully in §3.5 and Chapter 7.) John Pollock, for example, suggests that externalism misses the point because 'the fundamental problem of epistemology [is helping us to decide] what to believe' (1986, p.10). Alvin Goldman points out that externalism is 'unsuitable for realizing the true and original goals of epistemology' (1999, p.271). There is new interest in 'accessibilism', a view that contradicts externalism by suggesting that knowledge depends on having reflective access to one's evidence. According to William Alston, in a 1986 paper, 'it

⁸ Armstrong suggests that there are other factors that determine whether an agent has inferential knowledge, which I have set aside here.

⁹ Prominent among these are Nozick (1981) and Dretske (1981). Dretske's externalism is reflected in his relevant alternative theory: 'the difference between a relevant and an irrelevant alternative resides, not in what we happen to regard as a real possibility (whether reasonably or not), but in the kind of possibilities that actually exist in the objective situation' (p.377). Thus, a subject knows if and only if all the alternatives that are made relevant by the objective situation have been ruled out.

looks as if there is a basic, irreducible, requirement of epistemic accessibility of ground for the belief that attaches to our concept of epistemic justification' (2001, p.103). Trent Dougherty adds that the mental states to which one 'has privileged access . . . are . . . evidential bedrock; they are our cognitive home: we start from there and must always come back to them if questioned long enough' (2011, p.4). Earl Conee and Richard Feldman develop a theory of knowledge they call 'evidentialism'; and it is based on the notion that knowledge depends on evidence that is accessible on reflection.¹⁰ Also, there is renewed interest in the deontological conception of justification, which suggests that we have a certain responsibility, obligation or duty to ensure that our beliefs are properly justified. This contradicts externalism which suggests that we will probably not be able to fulfil such an obligation. Let us look at this briefly.

Recall W. K. Clifford's extended metaphor about the ship-owner (reproduced earlier) in which an unseaworthy ship, filled with emigres, 'went down in mid-ocean and told no tales' (1876, p.289). Many First History theorists conclude that if we have a responsibility, obligation or duty in relation to what we communicate to others, then the very meaning of the term 'epistemic justification' should be rooted in this responsibility—giving rise to the 'deontological conception of justification'. In fact, Alvin Plantinga suggests that this conception is a fundamental feature of the epistemological enterprise since Descartes. He explains,

to understand the contemporary situation [in relation to present views on epistemic justification] we must take a careful look at its history, in particular at some of the ideas of Descartes, and perhaps even more importantly, Locke. And here what is of first importance is to see that for Descartes and Locke the notion of duty or obligation play a central role in the whole doxastic enterprise. (1990, p.49)

Plantinga explains that '[o]riginally and at bottom, epistemic justification is deontological justification' (1993, p.14).

Additionally, it seems that many First History theorists conclude that if we have a responsibility, obligation or duty in relation to what we communicate to others, then our beliefs and denials should reflect that responsibility. William Alston explains that epistemic 'principles will *forbid* beliefs formed in such a way as to be likely to be false and either *permit* or *require* beliefs formed in such a way as to be likely to be true' (original italicization, 1988, pp.258-259).¹¹

¹⁰ A collection of their essays on the subject may be found in Conee and Feldman (2004).

¹¹ See for example Kornblith (1983), Bonjour (1985), and Feldman (1988).

So, for many First History theorists, externalism fails to offer an acceptable remedy to the problem first recognized decades earlier: the inferences that need to be made to secure inferential knowledge are fallible, and so the foundation for inferential knowledge is not firm. For many, there seems to be only one way forward; and it involves abandoning the thesis that knowledge depends on making infallible inferences, i.e., abandoning foundationalism. They could pick up where C. I. Lewis and others left off, further investigating what Lewis called 'congruence' or 'coherence'. Perhaps knowledge depends on the coherence of an entire system of belief. Laurence Bonjour suggests that coherentism is 'the only remaining possibility for a non-sceptical account of empirical knowledge' (1985, p.87). He continues,

Epistemic reflection, according to such a theory, *begins* from a (perhaps tacit) representation of myself as having (approximately) such and such a specific system of beliefs: only relative to such a representation can questions of justification be meaningfully raised and answered. (p.104)

However, the challenge—in fact, the principal challenge—for coherentist theory is to establish proper criteria for coherence. Such criteria would ensure that beliefs that are part of a system that is deemed to be coherent have a very low likelihood of proving false. Addressing that challenge, Bonjour suggests that it is critical that there be 'various sorts of inferential, evidential and explanatory relations which obtain among the various members of a system of beliefs' (1985, p.93). Further,

coherence is not to be equated with mere consistency; . . . coherence . . . has to do with the mutual inferability of the beliefs in the system [and] relations of explanation are one central ingredient in coherence . . . (p.95)

Coherentism seems to provide a promising approach to identifying beliefs that will not prove to be false. Unfortunately, producing a detailed account of coherence has proven to be a tricky undertaking. The challenge largely relates to identifying universal criteria for a coherent set of beliefs; and, given that a set of beliefs may be more coherent relative to some criteria than others, identifying the relative importance of these criteria, and the threshold for coherence.

2.3 Second History

Let us now look at the Second History—which, I suggest, follows from a research programs that has as its principal aim shedding light on the circumstances under which we take ourselves as knowing or not knowing that particular world conditions obtain, and also to shed light on the ordinary use of ‘know’.

I pointed out that early in the last century many epistemologists hold the view that non-inferential beliefs (e.g., that there is a red patch of a round and somewhat bulgy shape) generally have firm foundations, whereas inferential beliefs (e.g., there is a tomato there) generally do not. Further, I gestured at C. I. Lewis’ (1946) response to this basic difficulty, viz., to try to find different pieces of evidence that are mutually congruent or coherent, and that are congruent or coherent with established beliefs. I start my account of the Second History with G. E. Moore’s radically different response to the basic difficulty, which he presents in a 1940 paper. He suggests that the foundations of many ordinary inferential beliefs are firm—particularly, those inferential beliefs that involve identifications of perceptible and unambiguous objects, states or events (e.g., there is a tomato there). He proposes that we set aside the suggestions made by the likes of Bertrand Russell and H. H. Price that those latter foundations are not firm—on the basis that their views are overly rationalistic and stringent. Moore explains,

Russell’s view that I do not know for certain that this is a pencil . . . rests, if I am right, on no less than four distinct assumptions: (1) That I don’t know these things immediately [on the basis only of sensory perception]; (2) That they don’t follow logically from any thing or things that I do know immediately; (3) That, *if* (1) and (2) are true, my belief in or knowledge of them must be ‘based on an analogical or inductive argument’; and (4) That what is so based cannot be *certain knowledge*. . . I do not think it is *rational* to be as certain of any one of these four propositions, as of the proposition that I do know that this is a pencil. (1959, p.222)

Here, Moore seems to adopt an attitude that is similar to that which Russell, Price and others took when they set aside the earlier Cartesian suggestion that we may doubt the veracity of sense perceptions and even the existence of an outside world. Where Russell suggested that ‘[p]hilosophy cannot boast of having achieved such a degree of certainty that it can have authority to condemn the facts of experience’ (1914, p.67), Moore seems to suggest that

philosophy cannot boast of having achieved such a degree of certainty that it can have authority to condemn identifications of unambiguous objects.

Following Moore, J. L. Austin, in a 1946 investigation, begins with the assumption that an ordinary person may well know a 'particular, current, empirical fact' (1979, p.77)—such as that the object in their view is a goldfinch. He then sets out to answer the question 'what sort of thing does actually happen when ordinary people are asked "How do you know?"'(p.77) Over the course of his rather informal and brief investigation, Austin manages to (1) outline a rough theory of relevant alternatives (RA Theory) which holds that ruling out the relevant alternatives to a questioned hypothesis is a sufficient condition for knowledge; (2) suggest a general approach to responding to sceptical challenges, such as whether the tomato is not actually a cleverly painted piece of wax; and (3) defend a fallibilist conception of knowledge which holds that it is possible to possess knowledge without possessing entailing evidence.

Austin finds that knowing that p involves ruling out just those alternatives to p that are relevant, and not those that are irrelevant. For example, in relation to the proposition that the bird is a goldfinch, 'enough evidence' means,

enough to show that (within reason, and for present intents and purposes) it 'can't' be anything else, there is no room for an alternative, competing description of it.
[However, it] does not mean, for example, enough to show it isn't a stuffed [or fake] goldfinch. (1979, p.84)

So, this early RA Theory postulates how it is that we ordinarily set evidential standards for using 'know': the standard is set just high enough to rule out the alternatives that are relevant 'within reason, and for present intents and purposes'. With this, Austin gestures toward a knowledge theory under which evidential standards depend on context. Moreover, his RA Theory suggests a response to classic sceptical challenges—such as that one may not in fact know that the object is a goldfinch because of the possibility that it is a fake. Such a possibility is not a cause for concern because ordinarily it is not relevant—within reason, and for present intents and purposes.

Still, on Austin's account, it is possible, though unlikely, that an irrelevant alternative to p is true, and so it is possible that p is false. However, we would not have been wrong to have asserted that p , even if it is a practical possibility that it is false:

If we have made sure [that the bird is] a goldfinch, and a real goldfinch [by ruling out the relevant alternatives], and then in the future it does something outrageous (explodes, quotes Mrs. Woolf, or what not), we don't say we were wrong to say it was a goldfinch, *we don't know what to say*. Words literally fail us . . . [However, if we have ruled out the relevant alternatives, we cannot] be proved wrong, whatever happens . . . [On the other hand, what] the future can always do, is to make us *revise our* ideas about goldfinches or real goldfinches . . . (original italicization, 1979, pp.88-89)

Austin's suggestion that, when we have ruled out all of the relevant alternatives, we cannot be proved wrong, seems to constitute a step toward the fallibilist hypothesis that a knowledge attribution that follows from ruling out the relevant alternatives may be deemed to be true—even if it is *possible* that it will later prove to be false.

RA Theory is not limited to the Second History. It is also advanced in First History philosophy. However, there is a profound general difference between First History and Second History versions. Many alternatives that are deemed to be relevant in First History versions (and therefore needing to be ruled out) are *not* deemed to be relevant in Second History versions (and not needing to be ruled out). In fact, under First History RA Theory, alternatives are generally *deemed to be relevant by default*, i.e., requiring good reasons to be deemed to be irrelevant; whereas, under Second History RA Theory, alternatives are generally *deemed to be irrelevant by default*, i.e., requiring good reasons to be deemed relevant. Let us look at this more carefully.

The standard First History view that alternatives are relevant unless there are good reasons to deem them irrelevant is suggested by Keith Lehrer: 'no hypothesis should be rejected as unjustified without argument against it' (1971, p.293). On Dretske's account, alternatives are irrelevant only if they are not 'possibilities that . . . exist in the objective situation' (1981, p.377). He presents a case in which a birdwatcher spots what looks exactly like a gadwall. However, the bird also looks similar to a grebe. On Dretske's account, the possibility that the bird is a grebe would only be irrelevant if in fact (a) grebes do not habitat in her location and (b) cannot migrate.

In contrast, Second History RA theorists take alternatives to be irrelevant by default, i.e., good reasons are required to make them relevant. This view is reflected in Austin's explanation of why a sceptical alternative, such as that the goldfinch is a fake, is not relevant:

The doubt or question 'But is it a *real* one?' has always (*must* have) a special basis, there must be some 'reason for suggesting' that it isn't real, in the sense of some specific way, or limited number of specific ways, in which it is suggested that this experience or item may be phony. (1979, p.87)

In a more developed Second History RA Theory, Stewart Cohen suggests that an alternative is relevant only if one or more of three particular conditions are met. Cohen's *first condition* is that the objective probability of the alternative is high (1988, p.102). The *second condition* is that the subject's *internal statistical evidence* indicates that the probability of an alternative is high (p.102). So, an alternative possibility is deemed relevant if the subject assesses the probability of it as being high. Cohen's *third condition* for an alternative being relevant is that it is salient in the context of attribution (p.106). Here, Cohen echoes Alvin Goldman: 'If the speaker is in a class where Descartes's evil demon has just been discussed, or Russell's five-minute-old-world hypothesis, he may think of alternatives he would not otherwise think of and will perhaps treat them seriously' (1976, p.776).

Notice that Cohen's third condition suggests that our standard for how probable an alternative must be to make it relevant depends on context. This constitutes an early contextualist condition on knowledge.¹² He explains,

In normal everyday contexts, the standards [for how probable an alternative must be to make it relevant] are such that skeptical alternatives are not relevant (unless, of course, the circumstances are such that skeptical alternatives are highly probable . . .). This explains our confidence in the truth of our everyday attributions of knowledge. However, when we are confronted with skeptical arguments, we may come to consider skeptical alternatives as relevant, thereby lowering our standards [for how probable an alternative must be to be relevant]. (p.96)

Like Austin's earlier theory, Cohen's theory of knowledge incorporates a fallibilist thesis: 'a fallibilist theory allows that S can know q on the basis of r where r only makes q probable' (Cohen, 1988, p.91). We may know that q without entailing evidence.

Before continuing in our look at the Second History, there is a small conceptual matter that it would be helpful to address. Many Second History theorists defend the thesis that whether or

¹² G. C. Stine (1976) and Peter Unger (1986) offer two earlier contextualist accounts of knowledge.

not an individual knows that p can depend on certain contextual factors that are not truth-conducive, particularly factors that relate to the practical circumstances. However, these theorists have different views on what the nature of the standard is that ultimately determines whether we know. For example, Cohen suggests that the standard is *how probable an alternative must be* to make it relevant. This probability standard may move up or down when the practical circumstances change. Other epistemologists suggest that the standard is *how strong the evidence must be* in order to be able to properly use the word 'know', which I refer to as the Epistemic Standard. Again, the Epistemic Standard may seem to move up or down when the practical circumstances change. In contrast, Keith DeRose (1992)¹³ suggests that the standard relates to the knower's 'epistemic position'. He explains that this is determined by a combination of the strength of the knower's evidence and the range of relevant alternatives in the context. (I look at this more closely in a moment.) Finally, notice that Cohen's standard is negatively correlated with both the Epistemic Standard and DeRose's standard that relates to the knower's epistemic position. For example, a *low* probability standard for the relevance of an alternative tends to produce more relevant alternatives that need to be ruled out; and so this calls for stronger evidence to rule out a greater number of alternatives, i.e., it calls for a high Epistemic Standard.

Let us return to our look at the Second History. Keith DeRose recognizes that Cohen's theory is unable to account for the use of 'know' in an important range of ordinary cases. He demonstrates this with his famous low-stakes and high-stakes bank cases.

Bank Case A. My wife and I are driving home on a Friday afternoon. We plan to stop at the bank on the way home to deposit our paycheques. But as we drive past the bank, we notice that the lines inside are very long, as they often are on Friday afternoons. Although we generally like to deposit our paycheques as soon as possible, it is not especially important in this case that they be deposited right away, so I suggest that we drive straight home and deposit our paycheques on Saturday morning. My wife says, "Maybe the bank won't be open tomorrow. Lots of banks are closed on Saturdays." I reply, "No, I know it'll be open. I was just there two weeks ago on Saturday. It's open until noon."

¹³ I focus on this older paper in order to better represent the historical progression of theory in the Second History.

Bank Case B. My wife and I drive past the bank on a Friday afternoon, as in Case A, and notice the long lines. I again suggest that we deposit our paycheques on Saturday morning, explaining that I was at the bank on Saturday morning only two weeks ago and discovered that it was open until noon. But in this case, we have just written a very large and very important check. If our paycheques are not deposited into our checking account before Monday morning, the important check we wrote will bounce, leaving us in a very bad situation. And, of course, the bank is not open on Sunday. My wife reminds me of these facts. She then says, "Banks do change their hours. Do you know the bank will be open tomorrow?" Remaining as confident as I was before that the bank will be open then, still, I reply, "Well, no. I'd better go in and make sure". (1992, p.913)

In DeRose's analysis, knowledge attributors set different standards for the knower's 'epistemic position' in different contexts. These standards, in turn, track the stakes for the attributor, or *the importance to the attributor of being right*.¹⁴ DeRose attributes knowledge to himself in Bank Case A (taking himself as knowing that the bank will be open), but not in Bank Case B—because the importance of being right to the attributor (to himself) is greater in the second case.

In a number of writings that follow DeRose's contextualist theory, it is suggested that it is not *the attributor's* practical interests that determine standards, but, rather, *the subject's* (which writings I discuss in a moment). However, DeRose's bank cases feature first-person knowledge attributions, in which attributor and subject are one and the same. As a result, DeRose's bank cases seem to support both DeRose's thesis (which looks to the attributor's practical interests), and the competing thesis (which looks to the subject's practical interests). So, in order to draw a sharp contrast between DeRose's thesis and the competing thesis it is helpful to consider cases that feature second-person or third-person knowledge attributions, in which the attributor is not the subject. Accordingly, let us look now at cases of this kind *in support of* DeRose's contextualist thesis. The following original cases are based on a knowledge attribution that Jason Stanley uses as an example: 'Herman knows at 1:30 p.m. on September 24, 2004 that Hillary Clinton is a Democrat' (2005, p.86).

¹⁴ Robert Hambourger (1987) defended essentially this thesis some five years earlier in a carefully detailed paper. However, the thesis is presented as a part of an overall theory of knowledge that is otherwise considerably less developed than is DeRose's.

Case 2-1: Low Stakes Contextualist Case

Frank wants to get a sense of the American political scene just so that he can better understand and enjoy political newscasts. At 1:30 p.m. on September 24, 2004, Herman tells Frank that he has learned from a newspaper article that Clinton is the Democrat Senator from New York. Frank thinks 'Herman knows at 1:30 p.m. on September 24, 2004 that Hillary Clinton is a Democrat'.

Case 2-2: High Stakes Contextualist Case

Frank is a Democratic Party strategist and is worried that Hillary Clinton quietly switched party affiliations, from the Democrats to the Republicans. If true, this would be very bad news. At 1:30 p.m. on September 24, 2004, Herman tells Frank that he has learned from a newspaper article that Clinton is the Democrat Senator from New York. Frank thinks 'Herman does not know at 1:30 p.m. on September 24, 2004 that Hillary Clinton is a Democrat'.

Frank finds that Herman knows that Clinton is a Democrat in the low stakes case, but not in the high stakes case. This state of affairs seems to be explained by the difference in the importance of being right for Frank, the attributor in the cases.

DeRose explains that the importance of being right is a factor that is internal to the attributor, perhaps psychological in nature (1992, p.919). He suggests that it depends on matters such as the attributor's interests and values and calls it an 'attributor factor'. Further, he suggests that there are additional 'external' factors that may determine whether the attributor knows. He calls these 'subject factors'. Echoing Goldman (1976, pp.772-773), DeRose explains that the context determines the range of relevant alternatives. Referring to his bank cases (above), he writes:

If very many nearby banks have discontinued their Saturday hours in the last two weeks, then it seems that my original claim to know may well have been false, and if I admit that I did not know after this surprising fact about local banks is called to my attention, I will be taking back and contradicting my earlier claim to have known. (1992, p.921)

If very many nearby banks have closed, then (a) the probability that DeRose's bank will be closed increases, (b) the possibility that the bank will be closed becomes a relevant alternative that needs to be ruled out, and (c) DeRose's *epistemic position* is weakened. Thus, DeRose

suggests that our epistemic position depends both on the strength of our evidence and on the range of relevant alternatives. (I say a little more about this in a moment.)

According to DeRose, (internal) attributor factors (particularly the importance of being right) determine the epistemic position we must be in, in order to know; and so, these determine the *meaning* of 'know'; whereas (external) subject factors determine the range of relevant alternatives, and so, along with the strength of our evidence, determines our epistemic position. DeRose explains,

[a]ttributor factors set a certain standard the putative subject of knowledge must live up to in order to make the knowledge attribution true: They affect how good an epistemic position the putative knower must be in to count as knowing. They thereby affect the truth conditions and the content or meaning of the attribution. Subject factors, on the other hand, determine whether or not the putative subject lives up to the standards that have been set, and thereby can affect the truth value of the attribution without affecting its content: They affect how good an epistemic position the putative knower actually is in. (1992, pp.921-922)

Additionally, if attributor factors in a case are such that the attributor sets a particularly high standard for her epistemic position (say because the importance of being right to the attributor is high), then the standard for the relevance of an alternative is lowered—so that, potentially, more alternatives may be deemed to be relevant, and need to be ruled out. For example, it may be that in a scientific context, a high standard in relation to the attributor's epistemic position would lead to a low standard for the relevance of an alternative; and, as a result, even relatively unlikely alternatives may need to be ruled out. However, if subject factors are such that there are simply no probable alternatives suggested in the context, then the agent may be in a strong epistemic position even with weak evidence, and may be able to meet her high standard in relation to epistemic position with weak evidence.

At this point, it will be useful to define the term 'Contextualism' in a way that facilitates later discussion; and to do so in a way that accommodates DeRose's theory. Let us say that Contextualism is the view that the truth conditions for knowledge attributions vary across contexts in accordance with the attributor's practical interests.

While DeRose's account is consistent with an important range of ordinary cases, there are nonetheless many other important ordinary cases that his account is not consistent with. I refer

particularly to cases in which (a) the knowledge attributer is not the subject of the attribution, (because the cases feature second-person or third-person attributions), and (b) the Epistemic Standard that seems to be reflected in the knowledge attribution seems to depend, not on the importance of being right for the attributor, but on the importance of being right for the subject.

Jeremy Fantl and Mathew McGrath develop a theory of knowledge that seems to be able to account for these cases. It is based on a view that is often referred to as Subject Sensitive Invariantism (SSI). They suggest that the attributor is sensitive to the action that the subject *S* would perform on the basis of what she intuitively knows. Further, *S* 'is justified in believing that *p* only if *S* is rational [i.e., seeking to advance her interests] to act as if *p*' (2002, p.78).

Look again at Jason Stanley's example: 'Herman knows at 1:30 p.m. on September 24, 2004 that Hillary Clinton is a Democrat' (2005, p.86). Stanley, who defends a version of SSI, says that this knowledge attribution reflects an Epistemic Standard that depends on Herman's practical situation; and it does not depend on the practical circumstances of the attributor. So, it 'expresses the same proposition relative to every context of use [i.e., from any attributor's point of view]' (p.86). However, Stanley does not provide a context for his example—which in my view is necessary in order to assess the example¹⁵. So, I offer both low stakes and high stakes cases that incorporate this knowledge attribution, in support of SSI.

Case 2-3: Low Stakes SSI Case

Frank knows that his brother Herman wants to get a sense of the American political scene just so that he can better understand and enjoy political newscasts. At 1:30 p.m. on September 24, 2004, Herman tells Frank that he has learned from a newspaper article that Clinton is the Democrat Senator from New York. Frank thinks, 'Herman knows at 1:30 p.m. on September 24, 2004 that Hillary Clinton is a Democrat'.

Case 2-4: High Stakes SSI Case

Frank's brother Herman is a Democratic Party strategist. Frank knows that Herman is worried that Hillary Clinton quietly switched party affiliations, from the Democratic Party to the Republican Party. Frank knows that this would be very bad news for Herman. At 1:30 p.m. on September 24, 2004, Herman tells Frank that he has learned from a newspaper article that Clinton is the Democrat Senator from New York. Frank

¹⁵ I defend this suggestion about the necessity of context in §7.2.2.

thinks, 'Herman does not know at 1:30 p.m. on September 24, 2004 that Hillary Clinton is a Democrat'.

It seems reasonable that Frank would take Herman as knowing that Clinton is a Democrat in the low stakes case, but not in the high stakes case; and this difference seems to be explained by the stakes being higher for the subject of the knowledge attribution, Herman, in the high stakes case.

Like Contextualist theories, SSI theories are generally consistent with cases in which the attributor is the subject of knowledge attributions, i.e., first-person cases. In such cases, the attributor's sensitivity to her own practical circumstances constitutes a sensitivity to the subject's practical circumstances. Even so, SSI theories are not consistent with the cases I offered in support of Contextualism, Case 2-1 and Case 2-2, which feature a third-person attribution and a third-person denial respectively. Oppositely, Contextualist theories are not consistent with the cases I offered in support of SSI, Case 2-3 and Case 2-4, which also feature a third-person attribution and a third-person denial respectively.

Subject Sensitive Invariantism suggests that the Epistemic Standard that seems to have been set does not depend on the attributor's practical circumstances. Even so, that Epistemic Standard still varies. It varies in accordance with the subject's practical circumstances. So, both SSI and Contextualism support the *general* thesis that knowledge depends on practical circumstances (the attributor's or the subject's). I refer to this thesis as 'Epistemic Standard Variabilism'. This thesis is standardly rejected by First History theorists. They generally hold the view that the evidential standard for knowledge is very high and not dependent on practical circumstances at all. I refer to this view as 'Classic Invariantism'. Moreover, Classic Invariantists have attacked Epistemic Standard Variabilism in a variety of ways (and I look at this further on, in §7.2). However, a recent Second History approach to knowledge theory is not at all subject to these attacks. The approach suggests that the evidential standard for knowledge is not dependent on practical circumstances at all, and suggests that that standard is low enough that we know much of what we think we know. The approach is referred to as Moderate Invariantism.¹⁶

A basic component of Moderate Invariantist theory is an explanation of the elevated evidential standard that seems to be reflected in high stakes cases—such as those that are used to support SSI and Contextualism. The theory posits that these standards are suggested, not by the

¹⁶ See Kent Bach (2005), Patrick Rysiew (2005), and Jessica Brown (2006).

literal meaning of knowledge attributions or denials, but by propositions that the latter pragmatically convey. For example, as I read her, Jessica Brown suggests that very often the importance of being right and the evidential standard that this calls for are relevant and obvious in the context; and, as a result, this importance and evidential standard is reflected in the proposition that is pragmatically conveyed by knowledge attributions and denials. For example, in DeRose's high stakes bank case (which I reproduced earlier), when,

DeRose's wife points out that it's practically very important that they deposit the cheques before Monday (or there'll be big trouble with the bank) . . . she makes it clear that what's relevant to the conversation is a very strong epistemic position; in other words what's relevant to the conversation is whether DeRose's belief matches the facts not only at the actual and nearby worlds, but also at some further away worlds which are ordinarily too far away to undermine knowledge. Specifically, given that DeRose's wife mentions the possibility that the bank has changed its hours since DeRose's last visit, what's relevant is whether DeRose's belief matches the facts out to the possible world in which the bank has changed its hours since DeRose's last visit. As a result, . . . given the context, DeRose's assertion 'Ok, I don't know that the bank is open on Saturday' pragmatically conveys, though it does not literally state, that DeRose is not in a very strong epistemic position with respect to the relevant proposition, that his belief does not match the facts out to the world in which the bank has changed its hours. (2006, p.426)¹⁷

Thus, Moderate Invariantism can account for the use of 'know' in DeRose's cases, and, in general, the cases that are used to support Contextualism (e.g., Case 2-1 and Case 2-2). Moreover, the approach seems to account for our use of 'know' in cases that support SSI (e.g., Case 2-3 and Case 2-4). Moderate Invariantism does so in these latter cases in virtue of its assertion that the proposition pragmatically conveyed by a knowledge attribution or denial is responsive to the *practical circumstances that are relevant* in the context, and it is *the subject's* practical circumstances that are relevant in these cases. For example, in Case 2-4, it is the subject's, Herman's, practical circumstances that are relevant. So, the proposition that is conveyed by Frank's denial, 'Herman does not know at 1:30 p.m. on September 24, 2004 that Hillary Clinton is a Democrat' is responsive to Herman's practical circumstances; and that proposition being pragmatically conveyed suggests an elevated evidential standard. Thus, not

¹⁷ The account that is presented in this (2006) article is intended to fill out Patrick Rysiew's (2001) account.

only is Moderate Invariantism not subject to the First History attacks on Epistemic Standard Variabilism to which I alluded, it accounts for the cases that are used to support Contextualism and the cases that are used to support SSI.

I view the theory of Knowledge Intuitions that I develop in Chapters 3 through 8 as a Second History theory, and one that competes directly with Moderate Invariantism. In my final chapter, in §9.2, I briefly compare my theory with the latter and find that my theory seems to have greater explanatory and predictive adequacy than does the latter—which, if correct, would weigh in favour of my theory.

*

I suggest that this brief telling of the Second History generally reveals an effort to shed light on the circumstances under which we take ourselves as knowing or not knowing that particular world conditions obtain, and also to shed light on the ordinary use of ‘know’. I started the history with Moore’s suggestion that the epistemology of his day was overly rationalistic and stringent; and with Austin’s answer to the question ‘what sort of thing does actually happen when ordinary people are asked “How do you know?”’ (1979, p.77). However, the trajectory of the Second History has been such that a three-way impasse developed—one in which Contextualists assert that Epistemic Standards depend on the importance of being right to the attributor, Subject Sensitive Invariantists assert that these standards depend on the importance of being right to the subject of knowledge attributions, and Moderate Invariantists assert that the standard is not dependent on practical circumstances at all. Further on (in §3.5), I will suggest that all this constitutes a first phase in the Second History, and briefly describe a possible second phase.

2.4 Conclusion

The three-way impasse that I described in the Second History expands into a four-way impasse when Classic Invariantism is brought into the picture. Like Moderate Invariantism, Classic Invariantism suggests that the truth of knowledge attributions does not depend on the importance of being right—for either the attributor or the subject. However, Classic Invariantists suggest that the Epistemic Standard is significantly higher than do Moderate Invariantists. Thus, the picture emerges that not only has the epistemological literature

bifurcated over the last century, but multiple impasses have developed; and, it seems, almost every possible incompatible position has been staked out (as suggested in §1.1 where I outline five such positions). Even so, the theory of Knowledge Intuitions that I will develop in Chapters 3, 4, 5 and 6 suggests an explanation of why the literature would have bifurcated in the way that it has, why multiple theoretical impasses may have arisen, and it suggests a means by which this Gordian knot may be cut.

Chapter 3: The Function of the Knowledge Intuition Producer

The last chapter revealed multiple theoretical impasses and the existence of two mutually incompatible histories in the epistemological literature. I begin now in earnest to address this situation. To start, recall my working hypothesis that Knowledge Intuitions, regardless of the context in which they arise, are produced by a functionally discrete cognitive mechanism that evolved by natural selection. I called the device, the Knowledge Intuition Producer, or KIP. In this chapter, I develop a hypothesis about the Teleonomic Function of the KIP, and consider some of the implications of that function. In subsequent chapters, I predict the triggering-conditions of Knowledge Intuitions on the basis of that function. This will lead to the prediction that the triggering-conditions for Knowledge Intuitions depend on certain circumstances; and that, for many of the basic positions that epistemologists have staked out, it is possible to describe the set of circumstances under which the position works, and other circumstances in which it does not. We may cut the Gordian knot in epistemology by specifying the different circumstances under which each of the different epistemological positions will hold, and the reason we should expect it to hold in just these circumstances.

Also, at the end of this chapter, I note that my hypothesized function of the KIP may serve as the basis for *explaining* the appearance of variation in Epistemic Standards; and then briefly discuss the place that explanatory accounts have in the Second History (outlined in §2.3).

3.1 Idealization

It seems that the only practical way to investigate the function of the KIP is to first ‘idealize’ it, to imagine a simplified model of it. Roman Frigg and Stephan Hartmann explain that in the sciences,

idealization is a deliberate simplification of something complicated with the objective of making it more tractable. Frictionless planes, point masses, infinite velocities, isolated systems, omniscient agents, and markets in perfect equilibrium are but some well-known examples. (2012, section 1.1)

Catherine Elgin describes the utility of the strategy:

idealization [of studied phenomena] affords epistemic access to those features [that are believed to be worthwhile investigating], and enables us to explore them and their consequences by prescinding from complications that overshadow the features in real cases. It is valuable because it equips us to recognize these features, appreciate their significance, and tease out subtle consequences that might be obscured in the welter of confounding factors that obtain in fact . . . (2007, p.40)

I idealize the KIP in two significant ways. First, I assume that it is a functionally discrete mechanism, and that it has the job of setting an Epistemic Standard in the circumstances and determining whether the evidence at hand is strong enough to meet that standard. If so, Knowledge Intuitions are produced. Further, it carries out its job in a way that allows it to fulfil its underlying biological function, which we will consider further on.

A noteworthy aspect of this idealization is that the KIP is viewed as being functionally separated from the Cognitive mechanism that generates the propositions that the KIP evaluates in the first place. For example, the KIP is the mechanism that adjudicates the proposition that the bank will be open this Saturday; and is separate from the mechanism that generated that proposition. So, very roughly, whereas the KIP is concerned with evidence and analysis, the proposition generating mechanism is concerned with generating or identifying potentially true propositions.¹⁸

Nonetheless, it would appear that these two mechanisms are fully interdependent. *On the one hand*, the KIP's adjudications of propositions *depend* on the generation of alternative propositions. The process of adjudicating a proposition p often involves considering a small number of alternative propositions that are consistent with the evidence, but not consistent with p . For example, in Stanley's High Stakes Case (which is reproduced, along with Stanley's Low Stakes Case, in Appendix 2), the adjudication of the proposition that the bank will be open on Saturday involves generating and considering the alternative proposition that the bank changed its hours within the last two weeks. The notion that the KIP's adjudications of propositions *depend* on the generation of alternative propositions is consistent with RA Theory

¹⁸ It may be that a more accurate conception of what I am taking to be two separate mechanisms would have it that a single mechanism has two separate functions. I will set aside such a possibility.

generally, if it is not explicitly suggested by most versions of it. *On the other hand*, the reverse situation seems also to hold: The operation of the proposition generating mechanism depends on the KIP; or, more precisely, a critical aspect of the operation of the proposition generating mechanism is controlled by the KIP. Let us look at this.

Consider that the generation of alternatives seems sometimes limited to only those that are very likely; and, at other times (though far less frequently) hardly limited at all, e.g., when extremely unlikely possibilities are considered. I suggest that the KIP controls how likely the propositions that are generated (and considered) should be. It seems that when the KIP sets an Epistemic Standard, it effectively controls the generation of alternatives to p . As the Epistemic Standard is raised, the threshold likelihood of truth for generated alternatives is lowered, so as to potentially admit more, and less likely, alternatives (which according to RA Theory need to be ruled out if we are to know). For example, in Jason Stanley's Low Stakes Case, it seems that neither Hannah nor Sarah consider any alternative possibilities; while in Stanley's High Stakes Case, they consider the possibility that the bank changed its hours. Certainly the generation of the alternative in the High Stakes Case is not accidental: Generating that possibility in the low stakes case would have seemed odd and unnecessary given the practical circumstances, whereas generating the possibility in the high stakes case seems appropriate. Our Epistemic Standard seems to have been set higher in the high stakes case, and this seems to have called for the generation of less likely alternatives than were called for in the low stakes case. So, it seems that the KIP controls the production of alternative possibilities. Moreover, this control seems critical to the operation of the proposition generating mechanism. Without it, we would be inundated with alternative possibilities that are not worth considering, e.g., I should abandon the idea of going to the bank altogether because while I am at the bank it will be subject to an armed robbery during which I will be shot.

So, the KIP and the proposition generating mechanism are critical to the efficient and efficacious operation of each other. However, I suggest that even though such a relationship holds between them, I may conduct my investigation into the KIP without being too concerned about the nature and operation of the proposition generating mechanism.

My second idealization is to assume that the KIP adjudicates certain kinds of propositions but not others. My intention is to set aside a class of propositions in order to facilitate my investigation (and I do *not* suggest that the categorization of propositions I will propose mirrors different natural kinds of propositions). Even so, the distinction I will make is similar to certain distinctions that come up in the literature. *Roughly*, I *set aside* propositions that identify

perceptible objects and that are made under favourable conditions for observation (e.g., this is a tomato)—which propositions are viewed by Second History theorists as being knowable solely on the basis of sensory evidence; whereas I *include* in my investigation propositions that feature non-perceptible elements and that may well be subject to ordinary conscious assessments of whether the strength of the evidence meets an Epistemic Standard. From a different perspective, I am assuming that the KIP does not deliver adjudications of ordinary perceptions that are accepted by default.

Let us look more closely at the kind of proposition that I assume the KIP *does not* adjudicate. These are propositions that (a) identify unambiguous and perceptible objects, states, events or conjunctions of states or events, (b) locate those objects or events in space and time, or that (c) attribute unambiguous and perceptible properties to objects, states or events. (I do not suggest that this is a comprehensive characterization.) I am stipulating, for example, that the following statements (made under good observational conditions) are *not* primarily adjudicated by the KIP: (i) the object is a bird, (ii) the cat is on the mat, (iii) the tomato is red, and (iv) the red billiard ball hit the black one. I call these ‘Elementary Propositions’.

On the other hand, the KIP *does* adjudicate propositions that (a) identify *ambiguous or non-perceptible* objects, states or events, (b) locate those objects, states or events in space or time, (c) attribute *ambiguous or non-perceptible* properties to objects, states or events, (d) attribute *causes* to states and events, or (e) make *predictions* about future states and events. (This too is not intended to be a comprehensive characterization.) For example, I assume that the KIP adjudicates the following statements: (1) the sun will rise in the East tomorrow; (2) my bank will be open this Saturday; (3) she is a loyal employee; (4) my car (which I cannot see at the moment) is parked outside on the street; (5) that car over there is heavy (even though I have not tried to lift it); (6) he has the public’s best interests at heart; (7) the bridge is safe; (8) the bird is a gadwall; (9) anthropogenic global warming is real; (10) the politician’s policies directly led to the problems to which you referred; (11) the accused is guilty; (12) the watch works fine; (13) that is Jonah’s sweater; and (14) Paris is the capital of France. I call these ‘Advanced Propositions’.

So, the KIP flags Advanced Propositions as being safe to treat as true (even though they may sometimes be false). Let us appreciate how important such a task is to us. Compare, for example, the KIP’s importance to one of the several Cognitive devices that contribute to the

acuity of depth perception¹⁹, such as the Cognitive device that determines the extent to which the eyes converge when they focus on objects. A tendency to err in the production of that assessment may diminish depth perception acuity to some extent in some cases. Even so, such a tendency to err may go unnoticed—though it may rule out a career as a professional baseball player. In contrast, a tendency to err in our adjudications of Advanced Propositions would produce incorrect attributions of non-perceptible properties, incorrect predictions of all kinds, and incorrect everyday causal hypotheses. That would be profoundly debilitating.

3.2 The Function Hypothesis

In this section, I develop a hypothesis about the Teleonomic Function of the KIP. However, that work will require that we clearly distinguish such a function from an entirely different kind of function called an instrumental function. So, before starting that work in earnest, I discuss these function types (in §3.2.1). Further, my development of a Teleonomic Function hypothesis will be based on a hypothesis about the Teleonomic Function of cognition generally. However, the prospects of this approach working depends entirely on whether it is plausible that the KIP plays a key role in cognition generally. So, I will examine this matter also (in §3.2.2) before delving directly into the KIP's Teleonomic Function.

3.2.1 Teleonomic vs Instrumental Functions

My suggestion that the KIP evolved by natural selection is implicit in several writings. For example, Alvin Goldman proposed that a

fundamental facet of animate life, both human and infra-human, is telling things apart, distinguishing predator from prey, for example, or a protective habitat from a threatening one. The concept of knowledge has its roots in this kind of cognitive activity. (1976, p.791)

And Edward Craig suggests that '[h]uman beings need true beliefs about their environment, beliefs that can serve to guide their actions to a successful outcome' (1999, p.11).

¹⁹ Simon Grondin summarizes major research findings on depth perception cues (2016, pp.103-110).

If the KIP evolved by natural selection, it should have a Teleonomic Function²⁰, i.e., there should be something it does that explains, in biological terms, why we have it. Peter Godfrey-Smith suggests that such an explanation may refer either to (a) the ancient *origin* of a trait, or (b) a more recent history in which the trait was maintained or persisted (1994, p.344). For example, it is suggested that the original selection pressure for feathers was produced in virtue of their insulation value. Feathers originally evolved to help regulate the temperature of *Archaeopteryx* animals. Even so, feathers were maintained in recent history by selection pressure for flight (Ostrom, 1979, cited in Gould and Vrba, 1982, p.7). Similarly, it is argued that original selection pressure for bones in fish was produced by their ability to store phosphates. Even so, bones are *maintained* in more recent history by selection pressure for the structural support it gives land-dwelling vertebrates (Halstead, 1969, cited in Godfrey-Smith, 1994, p.358). Whether the KIP originally had a different function than it has had in recent history, or not, I will focus only on recent history.

Let us be careful not to confuse the KIP's Teleonomic Function with non-teleonomic functions that are generally referred to as 'instrumental functions'. These are functions that relate strictly to how the trait is utilized by the organism. For example, we utilize depth perception when we take in a beautiful landscape. That trait helps create aesthetically pleasing views. So, depth perception may be said to have this instrumental function, creating beautiful scenes in our minds. This is contrasted with what is plausibly the primary Teleonomic Function of depth perception—to get some sense of the spatial relations of visible objects, which sense allowed our ancestors to make a variety of fitness-enhancing inferences, such as whether objects were moving toward or away from them. Further, some traits may have an instrumental function, but not a Teleonomic Function. The Cognitive mechanisms that mediate such traits are mere by-products of mechanisms that mediate traits that have a Teleonomic Function, and that evolved by natural selection. Literacy and performing calculus are two examples that are often suggested in this context. It does not seem as if these traits could have helped our distant ancestors survive and reproduce. Additionally, John Dewey suggests that Cognitive traits that involve only the production and appreciation of aesthetic objects only have an instrumental function.

As we exercise caution in not confusing the KIP's Teleonomic Function with any of its instrumental functions, let us we bear in mind how easy it may be to make that error—as a

²⁰ See C. S. Pittendrigh (1958) and G. C. Williams (1966).

result of our generally knowing very little about the changing conditions under which traits evolved.²¹ For example, perhaps a plausible adaptationist account could have been developed for *literacy*, positing a Teleonomic Function for that trait. However, we have avoided that error in virtue of happening to know that our brain has not evolved significantly since a time before writing was invented.

With this mind, my proposed hypothesis of the Teleonomic Function of the KIP will be highly tentative. However, I do not simply stand on the hypothesis I develop about the function of the KIP. As I suggested earlier (in §1.3), my hypothesis is ultimately tested via the predictions it makes and whether those predictions are consistent with a wide range of knowledge cases.

3.2.2 Cognition

I suggest that we would do well to develop a hypothesis of the Teleonomic Function of the KIP on the basis of a hypothesis of the Teleonomic Function of Cognition generally. In principle, such a move may be made (a) if the function of the lesser device or process (in our case, the KIP) *plays a key role* in the operation of the greater device or process (in our case, Cognition), and (b) if it does not seem as if the lesser device or process has any other significant function. For example, the function of depth perception (a lesser process) follows from the function of the trait of being able to respond effectively to our physical environment (the greater process). Depth perception gives us a sense of the spatial relations of visible objects, and this *plays a key role* in our being able to respond effectively to our physical environment. We can duck if an object is flying toward us. Moreover, depth perception does not seem to have any other significant Teleonomic Function. So, let us look at whether it is plausible that the KIP plays a key role in the operation of Cognition.

There is one small matter to clear up first. I use the definition of Cognition that is used in biology, which seems to be somewhat narrower than that which is used in everyday discourse. Specifically, Cognition is constituted by ‘those higher mental processes in humans and animals, such as the formation of associations, concept formation and insight, whose existence can only be inferred and not directly observed’ (Henderson’s Dictionary of Biology, 14th edn, 2008, p.131).

²¹ S. J. Gould and R. C. Lewontin (1979) describe how these hypotheses may fail.

So, what role does the KIP play in cognition? I suggested earlier that the KIP seems to limit the generation of Advanced Propositions to those that are reasonably likely given the evidence (otherwise we would be inundated with highly unlikely alternative possibilities); and suggested that this would seem to be critical to the efficient and efficacious operation of the proposition generating mechanism. So, a creature that did not possess a KIP would seem to have no need for a proposition generating mechanism; and if a creature did not possess either a KIP or a proposition generating mechanism, it would not have intuitions of knowing or of not knowing any Advanced Propositions. It would not have intuitions of knowing or of not knowing any causes of present events, and not have intuitions of knowing or of not knowing that particular future events will arise.

It is hard to see how a creature that had neither a KIP nor a proposition generating mechanism would be able to respond to the environmental conditions it encounters *in a way that is flexible and still efficacious*. It would have no basis upon which to produce calculated deliberate responses. So, not only does the proposition generation mechanism depend on the KIP, but also *those devices that produce and control calculated deliberate responses* depend on the KIP.

Of course, organisms that have neither a KIP nor a proposition generating mechanism may very well still adjudicate Elementary Propositions. They may be able to identify unambiguous, perceptible objects, states and events; and such an organism may well be able to produce a range of responses to the environmental conditions it encounters that allow it to survive and thrive. However, these responses would seem to be limited to those that are reflexive, autonomic, and stimulus-responsive (such as startle responses, defensive physical responses, and the like). These are rule-based behaviours—stimulus A triggers response B.

Notice that the mechanisms that mediate rule-based behaviour may be *relatively* simple. Plants are capable of it, e.g., the leaves of *Mimosa pudica* defensively fold inward and droop when touched. This is not to suggest that such behaviour is not sometimes astonishing. Several bird, insect and fish species are remarkably accurate navigators over relatively great distances—variously using the earth's magnetic field, the location of the sun, smells, and chemical concentrations in water (Gallistel, 1990, pp.35-56).

Nevertheless, rule-based behaviour is, critically, not able to produce the behavioural flexibility that is facilitated by the adjudication of Advanced Propositions. Daniel Dennett's observations about the intelligence of lower species are apposite,

Such a smart doggie, but can he figure out how to unwind his leash when he runs around a tree or a lamppost? . . . And dolphins, for all their intelligence, are strangely unable to figure out that they could easily leap over the surrounding tuna net to safety'. (1996, pp.115-116)

Given that the mechanisms that mediate rule-based behaviour may be *relatively* simple and that the term, Cognition, refers to higher mental processes; it seems that the combination of (a) the proposition generation mechanism, (b) the KIP, and (c) all *those devices that together produce and control calculated deliberate* responses take up a very large part of Cognition. William James' view of Cognition seems broadly consistent with such a conclusion:

Cognition, in short, is incomplete until discharged in act. And although it is true that the later mental development, which attains its maximum through the hypertrophied cerebrum of man, gives birth to a vast amount of theoretic activity over and above that which is immediately ministerial to practice, yet the earlier claim is only postponed, not effaced, and the active nature asserts its rights to the end. (1966, p.18)

Moreover, I suggested earlier that (a) the proposition generation mechanism, and (b) all *those devices that together produce and control calculated deliberate* responses depend on the KIP.

Thus, this initial and cursory look at the situation suggests that *the KIP may very well play a key role in Cognition generally*. So, (just as the function of depth perception follows from the function of the trait of being able to respond effectively to our physical environment) I take it that it would be worthwhile to explore the possibility that the KIP's primary Teleonomic Function follows from Cognition's Teleonomic Function. So, I will now attempt to identify a plausible hypothesis about Cognition's Teleonomic Function; and then see if that suggests a Teleonomic Function for the KIP.

3.2.3 The Environmental Complexity Thesis

Peter Godfrey-Smith (1998) identifies 'a general family of ideas' that pertain to the function of Cognition. These ideas all suggest that (a) the environment naturally selects Cognitive characteristics, and (b) that Environmental Complexity produces selection pressure for complex minds. Let us look more closely at (b).

Environmental Complexity is an environmental parameter that is relative to a particular organism. It has to do with whether an organism's Inclusive Fitness is affected differently when it performs particular behaviours or adopts particular Phenotypes under either different environmental conditions that it encounters (i.e., under a type of environmental complexity called 'environmental diversity'), or at different times (i.e., under a type of environmental complexity called 'environmental variability'). If an organism's behaviour or adopted Phenotype positively affects its Inclusive Fitness under some of environmental conditions it encounters but not others, or at some times but not others, its environment is viewed as being complex. For example, the environment for the plant *Hieracium umbellatum* is viewed as complex insofar as the plant benefits significantly by (i) adopting its bushy, broad leaf Phenotype when it grows on rocky, sea-side cliffs, and (ii) adopting its narrow leaf Phenotype when it grows on sand dunes.²² As I say, there is a general family of ideas that share the view that such Environmental Complexity produces selection pressure for complex minds.

I do not intend to suggest that the causal arrow between Environmental Complexity and the features of Cognition only points in one direction. Certain Cognitive adaptations may have the effect of increasing Environmental Complexity. For example, these adaptations may facilitate new ways to exploit an environment and in the process make new features of the environment relevant. So, these adaptations *increase* Environmental Complexity.²³ Still, this state of affairs does not undermine the thesis that complex environments produce selection pressure for complex minds.²⁴

²² More basic examples of environmental diversity are reflected when different foods (available at a given time) have different nutritional properties and when different predators (that are active at a given time) call for different defences. Environmental variability is reflected when (a) different foraging practices are called for in different seasons, (b) different behaviours are called for in the event of significant changes in the populations of predators or prey due to disease, and (c) opportunistic invasions into neighbouring territories are called for under certain world conditions.

²³ For more see Kim Sterelny (1997, p.555).

²⁴ The thesis that complex environments produce selection pressure for complex minds is externalistic: It looks to environmental explanations for an organism's internal mechanisms. However, Kim Sterelny (1997), Elliot Sober (1997) and James Maclaurin (1998) all suggest that internalistic explanations should not be unduly discounted. Maclaurin explains:

A lineage that evolves colour vision puts itself in the way of new selection pressures. Individuals will be likely to evolve traits that make use of the more complex information provided by their new-found discriminatory abilities. Explanations for some of those adaptations will focus on a feature of the lineage (its ability to perceive in colour) at least as much as they will point to features of the environment. (p.493)

Nonetheless, when organisms evolve traits that make use of traits they already have, those newly evolved traits may still constitute a response to external selection pressures. Moreover, Godfrey-Smith maintains that externalistic explanations tend to produce greater explanatory pay-offs than do internalistic explanations (1998, p.53).

Godfrey-Smith defends a particular thesis that belongs to this family of ideas that I have referenced pertaining to the function of Cognition. He calls it the ‘Environmental Complexity Thesis’ (ECT)—and neatly expresses it as follows:

ECT: ‘[t]he function of cognition is to enable the agent to deal with environmental complexity’ (1998, p.3).

How does Cognition enable the agent to deal with Environmental Complexity? On Godfrey-Smith’s account, it does so *by enabling behavioural flexibility*. In complex environments, individuals benefit by deploying a *multiplicity of specialized* behaviours (e.g., responding differently to the different behaviours of different predators and prey) rather than by deploying fewer *general-purpose* behaviours. Those specialized behaviours may be more efficacious or more efficient under the different world conditions that individuals encounter than are general-purpose behaviours. (In the human case, these behaviours may include any sort of action, attitude or choice.) Cognition facilitates the performance of these specialized behaviours. This suggests a corollary of ECT that is more informative and particularly applicable in the human case:

ECT’: The function of Cognition is to create or select specialized behaviours that tend to be more efficacious and efficient in the different world conditions that the individual encounters than would general-purpose behaviours.

One of the main worries about the ECT is that it significantly overgeneralizes, that there may be important Cognitive traits that it does not describe.²⁵ Even so, a close look seems to suggest

However, S. J. Gould and R. C. Lewontin take a stronger position in favour of internalism:

organisms must be analysed as integrated wholes, with *Bauplane* so constrained by phyletic heritage, pathways of development and general architecture that the constraints themselves become more interesting and more important in delimiting pathways of change than the selective force that may mediate change when it occurs. (1979, p.581)

Even so, Daniel Dennett points out that Gould & Lewontin fail to recognize that some selective explanations may be very complex, and that the details of such complex accounts may be either unknown or incorrect. He asks, “Is it objectionable when meteorologists say, begging the question about supernatural forces, that there must be a purely physical explanation for the birth of hurricanes, even if many of the details so far elude them?” (1995, p.250). Even so, most biologists hold the conservative view that adaptationist explanations are, at least, a good basis for generating hypotheses for subsequent testing and investigation.

²⁵ See for example Kim Sterelny (1997), Elliott Sober (1997) and D. M. Walsh (1997).

that it is plausible that the ECT covers enough of Cognition (and perhaps even all of Cognition) that it can indeed help us to identify the function of the KIP (which I present in Appendix 1).

ECT' does suggest a Teleonomic Function for the KIP, viz., *flagging the world conditions under which our specialized behaviours are to be deployed*—or, more generally, flagging when we may behave as if a proposition about a world condition is true. Here is a first pass hypothesis of the function of the KIP (which I express in terms of the function of Knowledge Intuitions):

First Pass: In a situation *S*, the function of intuiting that one knows that a proposition *p* is true is to indicate to oneself that, on the total evidence, behaving as if *p* in *S* will tend to enhance one's Inclusive Fitness.

However, it would be hard to make the case that our responses to the situations we encounter are all oriented toward the goal of enhancing Inclusive Fitness. Further on (in §4.3.2.3), I find that a closer look at the situation suggests that it is more plausible that an individual's calculated, deliberate actions are oriented toward the goal of producing consequences that the individual *values*; and the situation suggests that it is not implausible that our innate and acquired values, in turn, very broadly tend to facilitate the goal of enhancing Inclusive Fitness. (If they did not do so, then it would be hard to see how our species would have done as well as it has.) This suggests a modification to *First Pass*:

Second Pass: In a situation *S*, the function of intuiting that one knows that a proposition *p* is true is to indicate to oneself that, on the total evidence, behaving as if *p* in *S* will tend to produce consequences that one values.

Second Pass suggests that we assess whether behaving as if *p* will produce valued consequences in a two-step process: (1) anticipate the consequences of behaving as if *p*, and then (2) see whether we value those consequences. Consider very briefly a particular feature of step (2), valuing consequences.

It seems that the consequences of behaviour that we value are those that are an improvement over what would have followed from some alternative behaviour. For example, we value the anticipated consequences of an action if we believe that the action (a) will bring about *an extremely good situation* in the place of what would otherwise have been only *a good situation*, or (b) will bring about a *moderately bad situation* in the place of what would otherwise have been an *extremely bad one*. We do not value the anticipated consequences of an action, if it will

not make things better; and we value negatively the anticipated consequences of an action if it will make things worse. So, in general terms, we compare the consequences of behaving as if p with the consequences of not behaving as if p ; and value the former if we view it to be better than the latter. This situation suggests a further modification to Second Pass, and the final form of my hypothesis of the function of Knowledge Intuitions (and, by extension, the KIP):

F1: In a situation S , the function of intuiting that one knows that a proposition p is true is to indicate to oneself that, on the total evidence, behaving as if p in S will tend to produce better consequences than that of not behaving as if p in S .

For clarity: The phrase, 'behaving as if p ', refers to a sense that the agent has, in the moment, of a certain rough range of behaviours that she may perform if she took herself as knowing that p . The phrase, 'not behaving as if p ' refers to a sense that the agent has, in the moment, of a certain rough range of behaviours that she may perform if she did not take herself as knowing that p . Furthermore, the rough range of behaviour referred to by 'behaving as if p ' is unrestricted (aside from it being consistent with p). This feature is consistent with my earlier observation (in §1.2.3) that there are no intrinsic limitations (beyond physical ones) imposed on the behaviour that follows from Knowledge Intuitions. If I take myself as knowing and have no doubt at all that the man will attempt to kill me, my behaviour may well breach ordinary boundaries.

Notice that F1 receives a measure of support from John Hawthorne's and Jason Stanley's (2008) investigation which found that knowledge is *a norm* for deliberate action (which I discussed in §1.2.3). They point out in that same paper, that knowledge appears also to be a norm for further belief, making assertions and making choices.²⁶ They conclude their investigation,

[j]udging by our folk appraisals, then, knowledge and action are intimately related . . . Moreover, discussions of knowledge are in many cases silent about this connection. This is a shame, since if there is such a connection it would seem to constitute one of the most fundamental roles for knowledge. (p.574)

²⁶ They suggest, respectively, (1) 'Treat the proposition that p as a reason for believing q only if one knows that p ' (p.577), (2) 'one should assert only if one knows' (p.573), and (3) '[w]here one's choice is p-dependent, it is appropriate to treat the proposition that p as a reason for acting iff you know that p ' (p.578).

(Even so, there are a number of important differences between our accounts—which I will set aside here.)

However, even if F1 follows from the ECT, even if it is consistent with my initial observations (in Chapter 1) about the role that Knowledge Intuitions seems to play in our lives, and even if it is supported by Hawthorne & Stanley’s investigation, F1 is counterintuitive. In particular, it may sound more like a formulation for rational acceptance than for Knowledge Intuitions.

Nonetheless, as I have said, F1 is tested in Chapters 5 and 6 by seeing whether it is consistent with a wide range of knowledge cases, and it will pass that test; and further on (in §4.3.2.4), I look at rational acceptance cases closely and find that the mental states that feature in rational acceptance cases are distinguishable from Knowledge Intuitions, and also depend upon prior Knowledge Intuitions. Additionally, we will find that F1 helps us to explain why we have two independent histories in epistemology (Chapter 7) and explain our responses to Gettier-Style Cases (in Chapter 8).

3.3 A Second Function

Biological research favours the hypothesis that we inherit certain altruistic tendencies. (This research is reviewed in Appendix 3.) It seems that under a range of circumstances helping others tends to enhance one’s Inclusive Fitness. This state of affairs has profound implications for the knowledge theory that I develop on the basis of F1. I will say a little about it here, and examine it more closely in Chapter 6.

However, before looking at it, let us be clear that the term ‘altruism’ has a different meaning in biology than it does in ordinary use. The main difference is that, in biology, motivations are irrelevant. The sole criterion for whether a behaviour is altruistic is whether its *immediate (non-accidental) effect is to decrease the donor’s fitness and increase the recipient’s*. So, for example, incurring a cost in the process of helping others would be described as ‘altruistic’ in biology—even if that action follows from the egoistic motivation of creating an indebtedness from the recipient to the donor, burnishing the donor’s reputation, or winning God’s grace. I use this biological meaning of ‘altruism’ in my investigation in order to be consistent with the biological usage.

The biological research on altruism supports the thesis that the valuations of the consequences of behaviours that F1 refers to may in many cases be oriented toward helping others. Those valuations may be made in relation to the practical interests of those who we intend to help or envisage helping. So, whereas our valuations of consequences of behaviour may often be entirely self-oriented and self-interested, under some circumstances our valuations may be other-oriented and other-interested. Accordingly, we may distinguish the Knowledge Intuitions that follow from the former, from those that follow from the latter. I refer to the former as ‘Egoistic Knowledge Intuitions’, and the latter as ‘Altruistic Knowledge Intuitions’.

Often, Altruistic Knowledge Intuitions arise when we intend or envisage *communicating* questioned propositions to others for the purpose of helping them. Our valuations of consequences are made in relation to the practical interests of the intended recipients of those messages. In fact, Altruistic communication is a common trait among a wide range of species. For example, honey bees perform a wagging dance to communicate the location of nectar or pollen; receiver bees collect the nectar and pollen; and that nectar and pollen helps feed the sender’s hive. I reference *Altruistic* senders and *their intended* receivers using defined terms ‘Senders’ and ‘Receivers’ respectively.

The cases that I will examine in which Altruistic Knowledge Intuitions arise are limited to those that involve communications between Senders and Receivers. The function of those Knowledge Intuitions are described by F1. However, my discussion of communication cases will be facilitated by having a separate formal hypothesis that is specific to the Altruistic Knowledge Intuitions produced in these communication cases, and that is expressed in terms of Senders and Receivers. It will be a special case of F1:

F2: The function of a Sender’s Altruistic Knowledge Intuition in relation to p is to indicate to herself that communicating that p will tend to help her Receiver(s) more than would not communicating that p .

So, for example, say a Sender S has been asked by a Receiver R for directions to the nearest open tube station and S adjudicates a proposition p about how to get to that tube station on the basis of certain evidence. F2 suggests that, for S , the function of taking herself as knowing that p is to indicate to herself that R is better off acting or having an attitude as if p than not doing so. *The Sender’s Knowledge Intuitions thus facilitate her being able to help her Receivers.*

My account of the Second History (§2.3) ended with an impasse in which Contextualists asserted that Epistemic Standards depend on the importance of being right to the attributor, and Subject Sensitive Invariantists asserted that these standards depend on the importance of being right to the subject of knowledge attributions. Also, we found that in cases involving second-person and third-person knowledge attributions, e.g., ‘She knows that p ’, Contextualism was unable to explain changes in the Epistemic Standards that we seem to set in SSI cases, whereas Subject Sensitive Invariantism was unable to explain changes in the Epistemic Standards that we seem to set in Contextualism cases. I suggest that the impasse follows from a failure to distinguish Egoistic Knowledge Intuitions from Altruistic Knowledge Intuitions. The cases on which Contextualists focus exclusively feature Egoistic Knowledge Intuitions (see, for example, Case 2-1 and Case 2-2); while the cases on which Subject Sensitive Invariantists focus include those that feature Altruistic Knowledge Intuitions (see, for example, Case 2-3 and 2-4). This claim that I make will form one small part of the theoretical picture that I develop over the next several chapters.

3.4 The Hallmark of the KIP Is Flexibility

F1 and F2 suggest that a hallmark of the KIP is flexibility in determining whether the evidence is good enough that Knowledge Intuitions should be produced. If true, this would *not* constitute some particularly extraordinary trait. Consider that a certain pragmatic flexibility is evident in perceptual processes. While our perceptions may very often provide a faithful picture of our situation, they are also subject to systematic distortions and biases that help to protect us. Dangerous heights are overestimated by 60% when seen from above, and only 29% when seen from below (Stefanucci and Proffitt, 2009, p.427). This bias reduces the incidence of falling injuries (Jackson and Cormack, 2007, p.353). Also, the time-to-collision of approaching objects is underestimated for menacing objects (e.g., snakes and spiders) compared to non-menacing objects (e.g., butterflies and rabbits) (Vagnoni et al., 2012, p.826). Avoidance behaviour for menacing objects that is earlier than it needs to be, tends to serve us better than avoidance behaviour that is too late.

F1 and F2 suggest that there are some contexts in which we tend to do well to be trusting even with weak evidence, other contexts in which we tend to do well to be sceptical even with strong

evidence, and still other contexts in which we tend to do well to be abundantly cautious about taking ourselves as knowing propositions to be true. I describe these contexts in turn below.

3.4.1 Contexts in Which Subjects Are Trusting Even with Weak Evidence

F1 suggests that if we determine that we would do better to (a) intuit knowing that p (and to behave as if p) even if not- p , than to (b) not intuit knowing that p even if p , we may (c) intuit knowing that p even if our evidence for p is weak. Often, in these contexts, we value positively the consequences that are anticipated to follow from acting or forming an attitude as if p , when p is true; and we do not disvalue, or at least not strongly disvalue, the consequences that are anticipated to follow from acting or forming an attitude as if p , when p is false. Under these conditions, intuiting that one knows that p may be a kind of ‘best bet’. For example, F1 predicts that we may take ourselves as knowing to be true the proposition that the email solicitation is a scam even with weak evidence.

James Sage alludes to contexts of this type when he points out that when dangers lurk ‘cautious belief-forming processes (say, those based on weak inductive generalizations) can systematically generate false beliefs, but still be [adaptive]’ (2004, pp.102-103). It can pay to ‘over detect’ dangerous predators and there are advantages to ‘believing that all spotted mushrooms are poisonous (because consumption of spotted mushrooms was once followed by illness)’ (p.103).

3.4.2 Contexts in Which Subjects Are Sceptical Even with Strong Evidence

F1 suggests that if we determine that we would do better to (a) not intuit knowing that p (and not behaving as if p) even if p , than to (b) intuit knowing that p even if not- p , we may (c) not intuit knowing that p even if our evidence for p is very strong. Often, in these contexts, we value positively the consequences that are anticipated to follow from not intuiting that we know that p , when p is false; and do not disvalue, or at least not strongly disvalue, the consequences that are anticipated to follow from not intuiting that we know that p , when p is true. Under these conditions, not intuiting that we know that p is a kind of ‘best bet’.

The sociologist Ted Goertzel alludes to contexts of this type. In those contexts, individuals derive important psychological benefits from their knowledge denials. These denials give ‘meaning to their lives and [enable] them to shape the world’ (1992, p.36). He refers to historical revisionists, and particularly to deniers of genocides (Armenian, Jewish, and Cambodian). These revisionists find that the evidence in favour of certain events consistently fails to meet their standards. ‘With this kind of abuse of sceptical principles, doubt can be cast on almost anything, at least in the minds of those who are eager to deny an unpleasant reality’ (p.318). He finds the strategy of these deniers comparable to the strategy of those who ‘deny the scientific evidence of human evolution’ (p.318).

3.4.3 Contexts in Which Subjects Become Abundantly Cautious

My hypothesis, F2, suggests certain conditions under which Senders would require very strong evidence in order to take themselves as knowing—even though they otherwise aim to avoid being sceptical. In these contexts, Altruistic Senders intend or envisage passing messages to large numbers of Receivers—even thousands or millions via the written word or other technological means—for the immediate purpose of helping them. I call such Senders ‘Broadcasters’. (I look at this in detail in §6.4.4.)

As the number of Receivers increases, the potential *cumulative benefit* produced by a true positive, and the potential *cumulative harm* produced by a false positive, may be multiplied very many times over. In §6.4.4, I show how Broadcasters could improve the consequences of their communications, from their Receivers’ point of view, by sending messages that are not likely to prove false for any of their Receivers. They may do this by requiring very strong evidence to support the propositions that they communicate. Particularly, we find that F2 predicts that in the process of trying to produce benefits for their many Receivers, Broadcasters tend to find the following:

- (1) It is worthwhile gathering additional evidence that only modestly increases the predictive validity of their total evidence.

- (2) It is worthwhile incurring a high cost to gather additional evidence.

Prediction (1) is supported, for example, by the fact that, in most every evidence-based disciplines, multiple researchers undertake independent studies to confirm a given thesis. For

example, at least nine different kinds of experiments have been performed to confirm the value of Avogadro's Constant. Prediction (2) is supported whenever researchers—including say historians, archaeologists or botanists—expend a great deal of time and energy, travel great distances, or risk life and limb gathering additional evidence. It is further supported by the existence of 'big science', such as the research that involves the Large Hadron Collider or one of the giant telescopes.

Moreover, in §6.4.4, I show that these tendencies would tend to be amplified as Broadcasters' improve their capacities, technologies, resources and strategies for increasing the predictive validity of their total evidence.

3.5 Second Phase of the Second History and Edward Craig

Notice that F1 may serve as the basis for explanations of the appearance of variation in Epistemic Standards. Such an explanation, in a given situation *S*, would suggest that an individual *X* has a Knowledge Intuition (or, in an alternative scenario, failed to have a Knowledge Intuition), which was mediated by an unconscious setting of *X*'s Epistemic Standard in the moment to a level that *X*'s evidence would meet (or, in the alternative scenario, would not meet); the Epistemic Standard was set to this level because *X* anticipated (unconsciously) that having the Knowledge Intuition (or, in the alternative scenario, not having the Knowledge Intuition) would tend to produce consequences that *X* valued more highly than the consequences *X* anticipated would tend to follow from not having the Knowledge Intuition (or, in the alternative scenario, having the Knowledge Intuition). In this brief section, I say a few words about the place that explanatory accounts have in the Second History (which I outlined in §2.3).

Much Second History theory *analyses* and *describes* the appearance of variation in Epistemic Standards. The literature suggests factors that seem to predict the appearance of this variation. This theoretical work may be part of only the first phase of the Second History. In a second phase, epistemologists could aim to *explain* the findings produced in the first phase. This would involve explaining the ordinary use of 'know'. (Even so, such explanatory accounts may have analytical and descriptive value, and could supplant some first phase theory.) Edward Craig hoped to see just such phasing fully realized in epistemology:

Let us suppose, however optimistically, that the problem of the analysis of the everyday meaning of 'know' had both been shown to exist and subsequently solved, so that agreed necessary and sufficient conditions for the ascription of knowledge were now on the table . . . I hope that philosophers would not regard it as a terminus . . . I should like it to be seen as a prolegomenon to a further inquiry: why has a concept demarcated by those conditions enjoyed such widespread use? (Craig, 1999, p.2)

Let us look briefly at Edward Craig's second phase theory.

Craig's explanatory account, like my own, is based on identifying the function of knowing. According to Craig, the 'function of "know" is to flag good sources of true belief' (1999, p.17). (So, 'know' is viewed as having a significantly more modest role than what is suggested by F1, i.e., roughly to flag behaviour that would tend to produce consequences that we value.) Craig pictures himself as an inquirer who is narrowly focused on his present practical circumstances and whose knowledge ascriptions flag good informants. Further, he takes a 'totally subjectivist stance, the pure here and now for me as I am here and now' (p.83). This may describe the position of individuals who in the moment have Egoistic Knowledge Intuitions.

However, whereas this 'subjectivist stance' is natural on Craig's account, somehow it may not actually arise, or, if it arises, it will be short-lived. He explains that 'if [the subjectivist stance] exists, at any time, or in any individual, it will develop in the direction of objectivisation' (1999, p.84). According to Craig, objectivisation follows from practical pressures that inevitably cause us to move away from consideration of 'what any particular person wants at any particular time or place' (p.84) and toward consideration of the inquirer's future needs (p.84) and the needs of others (p.88). *Under objectivisation*, the function of knowledge is to flag

someone who is a good informant as to whether *p* whatever the particular circumstances of the inquirer, whatever rewards and penalties hang over him and whatever his attitude to them. That means someone with a very high degree of reliability. (p.91)

Thus, Craig ultimately presents a theory that defends a very high and invariant Epistemic Standard. As a result, on my reading, it does not at all solve what he called 'the problem of the analysis of the everyday meaning of 'know'' (p.2). Even so, my thesis follows in Craig's footsteps in aiming to contribute to the second phase of the Second History, i.e., aiming to *explain* the ordinary use of 'know'. Moreover, I assume as Craig does that the production of Knowledge

Intuitions is a human trait that has an important role to play in our lives. I even employ Craig's basic methodology: We both develop a hypothesis of the function of knowing, work out the implications of that hypothesis, and then test the implications. However, the trajectories of our investigations are entirely different; and this is largely due to our different initial hypotheses about the function of knowing. While Craig emphasizes inquiry, I emphasize a more general matter, behaviour, which I do on the basis of my investigation of the biological literature.

Chapter 4: Operating Principles of the Knowledge Intuition Producer


A new engineering student might be asked to suggest possible basic operating principles of a jet engine and to do so without looking inside such an engine. She would need to infer how, in principle, it manages to do what it does. The student might, for example, hypothesize (correctly) that an air and fuel mixture is ignited, causing burning gases to expand, shoot out the back of the engine and produce forward thrust. So too, one may try to work out the operating principles of the KIP—under the assumption that it performs the functions described by F1 and F2. I try to do so in this chapter. This is critical because these operating principles can suggest a set of contextual factors that determine the Epistemic Standards that we seem to set—which, in turn, would form a major part of our picture of the triggering-conditions of Knowledge Intuitions. In this chapter I propose operating principles of the KIP and propose contextual factors that determine our Epistemic Standards.

4.1 Introduction

Much of the work that would need to be done to get a sense of the operating principles of the KIP has already been done by others—and especially by Godfrey-Smith (1991; 1998). However, his focus was not on the KIP *per se*, but rather on Cognition; and the task he set for himself was to infer operating principles of Cognition under the assumption that the ECT is true. Even so, much of that work is applicable to the KIP.

Godfrey-Smith's first step was to identify the *principal* operating challenge for Cognition, under the assumption that the ECT is true. He suggests that that challenge is to respond to world conditions based on evidence about the nature of those conditions that is very often equivocal or deceptive, and to respond to them in a way that tends to advance the organism's Inclusive Fitness. I represent the challenge diagrammatically in Figure 4-1.


Figure 4-1: Cognition's Principal Challenge



The principal input of Cognition is equivocal or deceptive evidence. Its principal output is the governance of behavioural responses to world conditions. If ECT is true, then these last responses produce a tendency to advance the organism's Inclusive Fitness.

However, it seems that, in the human case, much of Cognition's principal challenge falls to the KIP. More precisely, the challenge of producing calculated deliberate responses to world conditions that tend to advance our interests falls to the KIP. Accordingly, we may reassign Cognition's principal challenge to the KIP (setting aside stimulus-responsive, or rule-based behaviour). Call it 'Principal Challenge of the KIP'. It is as follows: (1) in response to equivocal or deceptive evidence, (2) the KIP produces Knowledge Intuitions, (3) that guide our deliberate and calculated responses to world conditions in a way that (4) has a tendency to produce consequences that we value. (Further on, in §4.3.2.3, I explain that Inclusive Fitness may only be *indirectly* relevant to whether Knowledge Intuitions are produced, and that what is *directly* relevant is, instead, how we value the anticipated consequences of our behaviour.) I show the challenge diagrammatically in Figure 4-2.

Figure 4-2: Principal Challenge of the KIP



After having identified Cognition's principal challenge, Godfrey-Smith identifies a theoretical approach to addressing that challenge—which in turn suggests Cognition's operating principles. In particular, he points out that Signal Detection Theory (SDT) indicates how evidence might be assessed in a way that best advances the individual's interests. He suggests that it may be that the operating principles of Cognition follow SDT principles. Now, we reassigned Cognition's

principal challenge to the KIP. So, we have the suggestion, via Godfrey-Smith, that *the operating principles of the KIP follow SDT principles*. In this chapter, I investigate whether this is plausible.

4.2 What is Signal Detection Theory?

4.2.1 Overview

The basic principles of Signal Detection Theory were developed in the 1940's in a push to design effective radar systems. One of the main challenges for designers of those systems was how to interpret radar data that is very often equivocal.²⁷ That data consists of an electrical stream produced as a radar dish spins on its axis. Different patterns in those electrical streams suggest different probabilities of aircraft passing through the system's airspace. The question was how compelling the data should be, or what the standard of data strength should be, for flagging that data as a sign of an aircraft. An overly low standard would produce too many false detections, i.e., false positives; while an overly high standard would produce too many failures to detect aircraft that are in fact in the airspace, i.e., false negatives.

System theorists distinguished 'signal' from 'noise'. 'Signal' refers to radar data that originate with those objects that the radar operator is trying to detect ('target objects'), particularly aircraft. So, signal is a sign of a target object like aircraft. In contrast, 'noise' refers to radar data that does not originate with objects that the radar operator is trying to detect. (The terms 'signal' and 'noise' are treated as mass nouns because the data to which they refer is generally produced in continuous streams.)²⁸ According to SDT, data is flagged as signal on the basis of whether its strength meets a threshold level of strength; and the question for radar theorists is what that threshold level should be.

SDT is based on the principle that the optimal detection threshold for data strength reflects the relative importance to the radar operator of the following two contradictory objectives: (1) obtaining a high probability that the data that has been flagged as signal actually is signal (i.e., actually is a sign of a target object), and (2) obtaining a low probability of failing to flag as signal that data that is in fact signal. For example, if, in a given set of circumstances, objective (1) is

²⁷ For example, see Marcum (1947).

²⁸ Epistemologists might think of noise as being analogous to the evidence in Gettier-Style Cases, which does not have the right kind of connection with the world condition the subject is aiming to identify.

significantly more important than objective (2), then detection thresholds should be set high; and *vice versa*. (Examples follow in a moment.)

SDT recognizes that the relative importance of these two contradictory objectives depends entirely on the radar operator's practical circumstances. For example, radar operators who aim to detect approaching enemy fighter planes will likely tend to find objective (2) significantly more important than objective (1). It would tend to be important that they achieve a very low probability of failing to detect an incoming fighter; and perhaps *not* terribly important that they achieve a very high probability that the detections they make are true. Making a few false calls would *not* be too high a price to pay in order to ensure that they do not miss any calls they should have made. So, they will set a relatively low detection threshold. *In contrast*, radar operators who are monitoring incursions into their airspace by a friendly neighbour may well find objective (1) more important than objective (2). Say, that the negative repercussions of falsely accusing their neighbour of making the incursion is high. Here, it may be more important to obtain a very high probability that that which is flagged as an aircraft actually is one, than to obtain a very low probability of failing to detect an aircraft. Under such circumstances, a high detection threshold may be set.

So, SDT suggests that, in principle, the operator's optimal detection threshold is determined on the basis of how the operator would answer the following question:

Operator's Question: What is the higher priority for you at this time:

Objective A: increasing the probability that detections are in fact target objects (by raising the system's detection threshold), or

Objective B: decreasing the probability of failing to detect target objects (by lowering the system's detection threshold)?

The underlying thesis may be expressed as follows: one moves toward the optimal detection threshold in the circumstances by determining whether it is more important to (1) decrease the risk of being overly trusting (by raising standards) or (2) decrease the risk of being overly sceptical (by lowering standards). Consider the extremes. If (1) is consistently deemed to be more important than (2), operators would end up being sceptical even with strong evidence (a situation that is similar to that described in §3.4.2). If (2) is consistently deemed to be more important than (1), operators would end up being trusting even with weak evidence (a situation that is similar to that described in §3.4.1).

4.2.2 Factors that Determine Optimal Detection Thresholds

SDT suggests circumstantial factors that operators (e.g., radar system operators) would do well to take into account when they answer the Operator's Question. These factors determine optimal detection thresholds for the systems they operate, under the circumstances they encounter. (These factors will be incorporated into my working hypothesis of the KIP's operating principles; and further on we will see that it is plausible that analogous factors determine the Epistemic Standards that we seem to set.)

First, operators do well to take into account the consequences of (1) true positives (i.e., flagging data as signal when it is in fact signal) relative to false negatives (i.e., failing to flag data as signal when it is in fact signal), and (2) false positives (i.e., flagging data as signal when it is in fact noise) relative to true negatives (i.e., not flagging data as signal when it is in fact noise). In other words, they do well to assess how the consequences of true positives (relative to false negatives) compare with the consequences of false positives (relative to true negatives). The early SDT theorist J. Marcum gestures toward this thesis:

What about the operator watching a cathode ray tube—what are his criteria for calling “signal”? . . . If he is told that he will be subject to severe penalties if he calls a false alarm (calls a signal when it subsequently turns out that there was none), then he will be very cautious. If a doubtful pip appears on the screen, he will use discretion and say nothing. This means that under these conditions . . . the probability of detecting a target at a given range is decreased. (1947, p.10)

Marcum is suggesting that when operators answer the Operator's Question they do well to take into account the cost of false positives. If that cost is substantial (e.g., there are penalties for calling a false alarm, or negative repercussions of falsely accusing a neighbour of making an incursion into one's airspace), then an operator may do well to shift her priority toward Objective A, increasing the probability that detections are in fact target objects by raising the system's detection threshold. Elsewhere, Marcum suggests the converse: if the benefit of a true positive is substantial (e.g., there is a substantial benefit in being alerted to an incursion by an enemy fighter), then an operator may do well to shift her priority toward Objective B, decreasing the probability of failing to detect target objects by lowering the system's detection threshold. So, we have two fundamental factors: the cost of false positives and the benefit of true positives. (I set aside, in my limited discussion of SDT, cases in which there is a *cost*

associated with a true positive, and cases in which there is a *benefit* associated with a false positive.)

Consider the situation of an operator of a radar system that is used to prevent mid-air collisions at a busy airport. The benefit of those true positives, which will prevent a collision, is extremely high; while the cost of false positives is low, i.e., the cost of having an aircraft make only a small (completely unnecessary) course change. So, all other things being equal, the operator of such a system may place a priority on Objective B, decreasing the probability of failing to detect target objects (by lowering the system's detection threshold). The operator will prefer to (a) risk being overly trusting of her evidence suggesting the presence of an aircraft in the airspace than to (b) risk being overly sceptical of it.

SDT suggests, in particular, that operators' optimal detection threshold correlates with a value equal to the cost of a false positive *divided by* the benefit of a true positive:

Correlation 1:

optimal threshold	correlates with	$\frac{\text{cost of a false positive}}{\text{benefit of a true positive}}$
----------------------	--------------------	---

So, as the cost of false positives increases, detection thresholds increase; and as the benefit of true positives increases, detection thresholds decrease, e.g., at the busy airport where the benefit of a true positive for a collision is extremely high and the cost of a false positive is not high, the optimal detection threshold may be low.²⁹

I suggested at the outset that when operators answer the Operator's Question they *also* do well to take into account the consequences of *false negatives* (i.e., not detecting target objects that are, in fact, present) and of *true negatives* (i.e., not detecting target objects when there are no target objects). However, my interest here in SDT is strictly limited to those features that may be applied to the knowledge cases that epistemologists have addressed, and to other cases I have found worthwhile looking at in this investigation. In almost all of these particular cases, these additional factors do not seem to be relevant, in the sense that they do not seem to play a role in them. In these particular cases, the consequences of false negatives and of true negatives are so small or insignificant that they are not an appreciable practical consideration. Even so, they *could*, in principle, be highly significant in a range of knowledge cases. Let us a look at this a little more closely.

²⁹ For additional discussion, see Godfrey-Smith (1991, pp.232-237).

Both false negatives and true negatives involve not responding as if a target object is present, i.e., not responding as if the questioned world condition exists. There are two general kinds of consequences that may follow from such a non-response. The first involves either failing to obtain the benefit of a true positive or avoiding the cost of a false positive. However, these consequences are already fully accounted for when the benefit of true positives and the cost of false positives is factored into answers to the Operator's Question. So, we do not take this kind of false negative and true negative consequence into account when we answer the Operator's Question; because, if we did, we would be taking it into account twice. However, the second kind of consequence should be taken into consideration when we answer the Operator's Question. A paradigmatic consequence of the second kind involves the *conservation of a limited resource* that would have been expended by responding: As a result of not responding, a limited resource will be available to be *better* deployed in some other way. A second paradigmatic consequence of the second kind arises when a non-response would be met with a *sanction, punishment or penalty*. In these cases, such a non-response constitutes a failure to meet a requirement, obligation, duty, or constitutes breaking a law. Even so, as I suggest, while consequences of the second kind should, in principle, be taken into account when we answer the Operator's Question, they are generally not significant in the knowledge cases that epistemologists investigate. So, for the sake of simplicity, I set aside consideration of both false negatives and true negatives.

*

Say there are two radar systems that have very different traffic histories—one located in Toronto, Ontario, and the other at an outpost in the Canadian Northwest Territories. Aircraft have very often entered the airspace of the former system, call it the 'high traffic system'; while they have only infrequently entered the airspace of the other, the 'low traffic system'. Accordingly, the *prior probability* of the presence of aircraft in the airspace of the high traffic system is significantly higher than it is in the airspace of the low traffic system. As a result, all other things being equal, the historic ratio of signal to noise at the high traffic facility is higher than that at the low traffic facility—because the higher traffic leads to the production of more signal, while the quantity of noise would generally be more or less the same in both systems. So, SDT assumes that, all other things being equal, a given radar dataset received by the high traffic system is more likely to be signal than is the same dataset received by the low traffic system.

If, under these circumstances, the detection thresholds of both the high traffic and low traffic systems were the same, then, all other things being equal, the ratio of true positives to false positives achieved by the high traffic facility would be higher than that achieved by the low traffic facility. Under these circumstances, we may expect the operator of the low traffic system to tend to place a higher priority on Objective A—increasing the probability that detections are in fact target objects (by raising the system’s detection threshold)—than would the operator of the high traffic system. Raising the detection threshold in the low traffic system will tend to improve its ratio of true positives to false positives.

So, according to SDT, when operators answer the Operator’s Question they do well to take into account the prior probability of the presence of target objects. This factor is often referred to as the prior probability of signal. It is negatively correlated with the optimal detection threshold: the lower the prior probability of signal, the higher is the optimal threshold (e.g., the low prior probability of signal that the low traffic system must deal with puts upward pressure on its detection threshold, which in turn improves its ratio of true positives to false positives).

Further, according to SDT, when operators answer the Operator’s Question they also do well to take into account the prior probability that a given dataset is noise. The latter is a measure that is a probability complement of the prior probability of signal, i.e., they add up to 1. In particular, SDT suggests that that operators’ optimal detection threshold correlates with a value equal to the prior probability of noise *divided by* the prior probability of signal:

Correlation 2:

$$\text{optimal threshold} \quad \text{correlates with} \quad \frac{\text{Prior Pr(noise)}}{\text{Prior Pr(signal)}}$$

I will suggest an intuitive way of seeing why it is necessary to take both the prior probability of signal, and its complement, the prior probability of noise, into account. However, I will address one matter first.

We may notionally combine the two ratios shown in Correlation 1 and Correlation 2 in order to obtain a formulation that includes all the factors that I have described to this point.

Correlation 3:

$$\text{optimal threshold} \quad \text{correlates with} \quad \frac{\text{cost of a false positive}}{\text{benefit of a true positive}} \times \frac{\text{Prior Pr(noise)}}{\text{Prior Pr(signal)}}$$

Correlation 3 suggests that in order to set the optimal threshold (that strikes the optimal balance between the risk of being overly sceptical and the risk of being overly trusting), radar operators may modulate their detection threshold in a way that correlates (positively) with the product of the two ratios.

Let the variable ‘SDT Index’ equal the product of the two ratios above:

$$\text{SDT Index} = \frac{\text{cost of a false positive}}{\text{benefit of a true positive}} \times \frac{\text{Prior Pr(noise)}}{\text{Prior Pr(signal)}}$$

Thus, the optimal threshold correlates with the SDT Index. As the latter moves up and down, so too should the optimal detection threshold.

Now, let me suggest a way to see why it is necessary to take both Prior Pr(signal) and its complement, Prior Pr(noise), into account. View the two numerators of the SDT Index as being associated; and view the two denominators as being separately associated as follows:

$$\text{SDT Index} = \frac{\text{cost of a false positive} \times \text{Prior Pr(noise)}}{\text{benefit of a true positive} \times \text{Prior Pr(signal)}}$$

Seeing the two numerators multiplied together gives one the sense that the cost of a false positive is appropriately weighted by the prior probability that a dataset is misleading; and seeing the two denominators multiplied together gives one the sense that the benefit of a true positive is appropriately weighted by the prior probability that a dataset is actually indicative of the facts. Thus, the product of the two numerators may be viewed as something like the expected cost that will be incurred due to false positives (based solely on Prior Pr(noise)).³⁰ Similarly, the product of the two denominators may be viewed as something like the expected benefit that will be produced due to true positives (based solely on Prior Pr(signal)). So, the SDT Index as a whole may be thought of as *the expected cost of a false positive relative to the expected benefit of a true positive* (based solely on the Prior Pr(noise) and Prior Pr(signal)). Such a conception of the SDT Index has a well-known counterpart in the academic and professional business world, in the concept of the ‘risk/reward environment’—which refers to something like the expected cost of a bad investment decision relative to the expected benefit of a good investment decision. In what is sometimes referred to as a ‘high risk/reward environment’, it

³⁰ Less elliptically, the product may be viewed as something like the projected average cost due to false positives prior to the Present Evidence (incurred in a large number of iterations in which the system flags a target object and gets that number of false positives suggested by the prior probability value).

seems that the expected cost of a bad investment decision is high relative to the expected benefit of a good investment decision. So, we are advised to adopt a cautious investment strategy in such an environment. In a 'low risk/reward environment', the expected cost of a bad investment decision is low relative to the expected benefit of a good investment decision. We may adopt a less cautious investment strategy in such an environment. In fact, further on, we will find it helpful to view the SDT Index in much the same way when we use it to predict or explain the Epistemic Standards that we seem to set in knowledge cases.


4.3 How Signal Detection Theory Applies to the KIP

4.3.1 Introduction

Biologists use SDT to explain detection threshold levels that various organisms depend on when they identify objects, events, or environmental conditions on the basis of equivocal or deceptive evidence. These sorts of explanations have been made mainly in relation to *innate detection thresholds*: It seems that nature often selected for the best balance between the risk of being overly sceptical and the risk of being overly trusting. For example, SDT explains the innate detection threshold level exhibited by female birds when they identify male mating signals (Price, 2013), the detection threshold level exhibited by bees when they identify floral displays (Leonard et al., 2011), and the innate detection threshold level exhibited by various organisms that are involved in triggering the adoption of different Phenotypes (Moran, 1992). Recall the earlier example, *Hieracium umbellatum*, which grows bushy, broad leaves on rocky, sea-side cliffs, and grows prostrate with narrow leaves on sand dunes. The evidence that these plants gather about soil conditions is equivocal. SDT explains their detection threshold level against which this evidence is measured in the process of identifying soil conditions. This in turn determines which Phenotype they adopt.

However, SDT may also be used to explain detection threshold levels *that individuals set* on the basis of the world conditions they encounter. D. Harvell (1986) discovered that bryozoans—an aquatic, colonial, moss animal—set a detection threshold on such a basis. These creatures generally adopt a non-spined Phenotype under 'normal' conditions. However, when they detect sea slugs in their locality, they may adopt a spined Phenotype. Even so, there are significant costs these creatures incur growing those spines; and it is worthwhile doing so only if they are necessary in order to protect against the slugs (see Figure 4-5).

Figure 4-5: Bryozoan spines³¹



Bryozoans detect the slugs by sensing the concentration of a water-borne chemical. The higher the concentration, the more likely it is that they are close by. They adopt the spined Phenotype when they sense that this concentration meets or exceeds a detection threshold. However, bryozoans are able to *modify* the level of this threshold. In effect, they can make themselves more or less ‘sceptical’ about nearby slugs. In particular, Harvell found that the threshold that bryozoans adopt when they live in small colonies is higher than it is when they live in large ones (p.816). He explains that the cost of producing spines is proportionately greater in smaller colonies because the size of the colony periphery is proportionately greater. So, the cost of a false positive is higher when they live in smaller colonies, and, as a result, they tend to do better being more sceptical about slugs being in their vicinity under these circumstances.

Godfrey-Smith suggests that the detection thresholds that bryozoans set may take standard SDT factors into account (1998, pp.207-254). Further, he shows in principle that bryozoans may respond to the equivocal evidence they gather (i.e., chemical concentrations) in a way that tends to advance their Inclusive Fitness.

The question is: Do SDT principles apply in the human case? In answering this question, we bear in mind that these principles might apply to either (a) certain innate detection thresholds, which thresholds are not generally dependent on environmental conditions, and (b) detection thresholds that individuals seem to set on the basis of the world conditions they encounter. Innate detection thresholds seem to be mainly exhibited in our adjudications of Elementary Propositions (e.g., something is quickly approaching, or there is a snake in the grass); and often govern reflex, autonomic, and a range of stimulus-response behaviours—including startle

³¹ Copied from D. Harvell (1986, p.812).

responses and defensive physical responses. (Even so, it appears that there are important exceptions.³²) In contrast, the detection thresholds that individuals seem to set on the basis of the world conditions they encounter are mainly exhibited in our adjudications of Advanced Propositions (e.g., my bank will be open this Saturday); and seem to mainly govern the deliberate and calculated performance of actions and formation of attitudes. These thresholds seem to be reflected in the Epistemic Standards we seem to set, and relate to the operation of the KIP. So, the question on which I am focused here is, more precisely: Do SDT principles (which are reflected in the SDT Index) explain the appearance of Epistemic Standards that move up and down?

4.3.2 Direct Evidence

4.3.2.1 Four Traits

I have outlined indirect or circumstantial evidence that Godfrey-Smith has presented for the hypothesis that SDT describes our belief formation practices: He has shown (1) that a belief

³² Some evolutionary psychologists suggest that, under certain world conditions, innate detection thresholds in the form of innate Epistemic Standards are employed in relation to certain kinds of Advanced Propositions. They suggest that our distant ancestors may have faced protracted selection pressure for low Epistemic Standards in relation to these particular propositions. Their theory, called Error Management Theory (*EMT*), “proposes that psychological mechanisms are designed to be predictably biased when the costs of false-positive and false-negative errors were asymmetrical over evolutionary history” (Haselton and Nettle, 2006, p.81).

Fleeing from an area that contains no predator may be inconvenient but is much less costly than the failure to detect a predator that really is close by. [Error management theory] predicts that an optimal decision rule [for the behaviours] will minimize not the crude rate of error, but the net effect of error on fitness. (Haselton et al., 2005, p.731)

Haselton and Nettle explain further:

For ancestral humans, the [benefit of true positives follow from responding to] aggressive intentions on the part of others, whereas the [cost of false positives of] over-inferring aggressiveness was low, especially for members of competing coalitions. This asymmetry did not characterize inferences about in-group members, in which costly within-coalition conflict would have resulted from unwarranted inferences of hostility or aggressiveness. (Haselton and Nettle, 2006, p.53)

So, certain of the biases that have been identified in the experimental social psychology literature may be explained as following from innate predispositions to adopt low Epistemic Standards for certain propositions. Prominently, such propositions feature negative conclusions about members of different racial or ethnic groups, or members of an ‘out-group’. These include negative conclusions pertaining to their generosity and kindness (Brewer, 1979) and to their level of hostility and violence (Quillian and Pager, 2001).

formation device that operated under SDT principles would be favoured by natural selection; (2) that SDT principles may apply to the operation of the device that determines which Phenotype bryozoans adopt; and (3) that belief formation devices could *in theory* be governed by SDT principles. However, I believe that a great deal of direct evidence supports the hypothesis that SDT principles are reflected in the production of Knowledge Intuitions, in the operating principles of the KIP. I outline some of that evidence in this section.

I organize my presentation of this evidence around the plausibility that we possess four particular traits. I distinguish the four traits only for the purpose of facilitating an orderly presentation of evidence. In fact, it is our possession of just one of those traits, Trait 3, that is the heart of the matter: Our possessing that trait alone would constitute a strong indication that the KIP follows SDT principles. However, Trait 3 is a broadly defined trait that, in fact, includes Traits 1 and 2, i.e., if we possess Trait 3, then we must also possess Traits 1 and 2. So, I present evidence for Traits 1 and 2 separately, but ultimately in support of Trait 3. Moreover, it seems that Trait 3 could not have evolved unless we possessed Trait 4. So, I present evidence for Trait 4 separately, but again ultimately in support of Trait 3.

Even so, all the evidence that I present in relation to these traits is meant only to show that it is plausible that the KIP follows SDT principles: I do not suggest that the evidence is strong enough to secure the thesis. However, in Chapters 5 and 6, I present a number of diverse knowledge cases that further test the thesis; and, in Chapter 7, I show that the developed theory allows me to explain basic differences in the variety of theoretical views that are expressed in the epistemological literature, which still further tests the thesis. These are the traits that I will argue we possess:

Trait 1: We assess evidential strength with a degree of precision and a frequency of being correct that are sufficient for the application of SDT principles.

Trait 2: We assess whether our evidence meets our Epistemic Standard in a way that depends on the context.

Trait 3: We seem to set Epistemic Standards that correlate (positively) with the SDT Index.

Trait 4: We do not possess a cognitive device that can trump the KIP in relation to the control of calculated deliberate behaviour.

First, it is certainly plausible that we possess Trait 1. We assess the strength of our evidence with a degree of precision and a frequency of being correct that is sufficient for the application of SDT principles. (To be clear, by ‘strength of our evidence’, I refer to its ‘*predictive validity*’, i.e., its ability to predict that which it is evidence for.) It is easy to reach this conclusion because the minimum degree of precision and frequency of being correct called for by SDT are very low.

Regarding precision: SDT principles could in fact apply to a system that was capable of distinguishing only between two probability ranges in relation to the existence of a world condition, e.g., low probability and high probability. In some contexts, the system’s detection threshold could be set so that a world condition is flagged (as existing) only if there is a high probability that it exists; while in other contexts, its detection threshold could be set so that the world condition is flagged if there is either a low or a high probability that it exists. Even so, human probability assessments often seem to be considerably more fine-grained. A physician may assess that patient X has a 1 in 10 chance of having ailment Y—suggesting an ability to distinguish between 10 probability ranges.

Regarding frequency of being correct: In order for SDT principles to apply to a system, it must decide which of the probability ranges that it has distinguished applies to its evidence, and it must do so with some minimal frequency of being correct. For example, say that Harry calls Sam, Sam does not pick up the phone, and Harry wonders whether this constitutes sufficient evidence that he can take it that Sam is not home. Harry distinguishes three probability ranges in relation to the strength of his evidence, ‘low’, ‘medium’ and ‘high’; and decides that the strength of his evidence is ‘high’. The fact that Sam did not pick up the phone confers a ‘high’ probability on the proposition p that Sam is not home. An adjudication of p that is based on SDT principles would require that Sam be correct that the evidence confers a ‘high’ probability on p .

In principle, SDT principles could apply to systems even if the frequency with which they correctly identify the probability range that applies to its evidence is low—provided that the frequency is better than chance. Even so, *the ability of SDT principles to lead to behaviour that tends to advance individuals’ interests correlates with this frequency of correctness*. So, if Harry is often incorrect, the ability of SDT principles to advance his interests will be diminished.

However, if we often distinguish only a small number of probability ranges, say 3, 4 or 5 probability ranges (as Harry does in the last case), then the frequency with which we correctly decide which of those ranges applies would seem to be reasonably high. Moreover, if in fact the

frequency with which we correctly assess which of these probability ranges applies were particularly low, then it would seem that we would not have been as *successful* a species as we have been.

So, I conclude that it is plausible that we possess Trait 1. Considerably more discussion is required to show that it is plausible that we possess Traits 2, 3 and 4. I address them in the following three sections.

4.3.2.2 Trait 2

Is it plausible that we possess Trait 2, according to which we assess whether our evidence meets our Epistemic Standard in a way that depends on the context (so that the production of Knowledge Intuitions depends on context)? (Further on, in my discussion of Trait 3, I examine the separate matter of the particular practical contextual factors that correlate with the Epistemic Standard that we seem to set.) SDT suggests that such a state of affairs is critical to achieving, in the circumstances, an optimal balance between the risk of being overly sceptical and the risk of being overly trusting.

I will make an assumption that is analogous to one standardly made by those who employ SDT when they design radar systems or by biologists who employ SDT to understand how an organism identifies environmental conditions on the basis of equivocal or deceptive evidence: A system's or organism's assessments of whether the evidence meets a detection threshold are made dependent on world conditions by *modulating its detection threshold according to those world conditions*—and not by modulating its assessments of the predictive validity of the evidence. A radar system may be made responsive to world conditions by modulating its detection threshold—and not by modulating its assessment of the predictive validity of its input data. So, SDT principles are reflected, and only reflected, in the modulation of detection thresholds. I make the analogous assumption that if humans assess whether their evidence meets their Epistemic Standard in a way that depends on the context, then they do so by modulating their Epistemic Standard according to the context. Thus, we may reduce the principal question that we address in this section to *whether it is plausible that our Epistemic Standard depends on context*.

Some support for the hypothesis that our Epistemic Standard depends on context is found in the substantial body of Contextualist and Subject Sensitive Invariantist literature (which I briefly

discussed in §2.3). That literature is largely based on cases in which the use of ‘know’ seems to depend on context.

Additional support for the hypothesis that the Epistemic Standards that we seem to set depend on context follows from recognizing that the process of adjudicating an Advanced Proposition often involves adjudications of additional Advanced Propositions that are ancillary to the principal Advanced Proposition. In particular, we often need to adjudicate ancillary Advanced Propositions in order to assess the strength of the evidence for the principal Advanced Proposition. For example, the committee’s adjudication of the principal proposition that the pipeline development would produce a net benefit to the community depends on their assessment of the quality of an environmental impact study, which involves an adjudication of the ancillary proposition that the study is reliable. Further, adjudications of such ancillary propositions are based on ancillary evidence, which may run for or against those ancillary propositions; and these adjudications determine whether the ancillary evidence is strong enough to meet an Epistemic Standard that applies to the ancillary proposition. So, we may further support the hypothesis that the Epistemic Standards that we seem to set depend on context by identifying support for the following hypothesis: A context-dependent Epistemic Standard is used in our adjudications of the *ancillary Advanced Propositions* that are required to assess the evidence for principal Advanced Propositions. In other words, *we may further support the main hypothesis we address here by finding that our assessments of the strength of evidence (as being sufficient or not) depend on context.*

Indeed, it seems that assessments of evidential strength depend on context. For example, contrary evidence may be dismissed or viewed as being anomalous in one context, but not in another; or contrary evidence may be assessed as being weaker in one context than in another. The committee may seem to have treated a high quality environmental impact study as being more important in one situation, than it did in another. The judge may have treated the partner’s testimony as being more important in one trial than he did in another.

This picture is supported by social psychology experiments that find that we are ‘apt to accept [empirical evidence that confirms favoured hypotheses] at face value while subjecting “disconfirming” evidence to critical evaluation’ (Lord et al., 1979, p.2098). It looks as if the Epistemic Standard that we set in relation to ancillary propositions about the validity of evidence depends on whether the evidence confirms or disconfirms favoured hypotheses. Disconfirming evidence can even strengthen the intensity of a favoured hypothesis (Batson, 1975). For example,

each side of an international conflict often believes that the other side is up to no good. An offer of concessions may be interpreted as a sign of weakness or trickery rather than as evidence against the favoured belief. Likewise, if the evidence is mixed, one side may attend only to the evidence that supports its [favoured] belief, so that the belief is strengthened by this part of the evidence but not weakened by unfavourable evidence. (Baron, 2008, pp.208-209)

More generally, a wide range of social psychology experiments reveal a marked tendency to assess evidence in a way that supports *already held* hypotheses; a phenomenon referred to as either ‘confirmation bias’ or ‘myside bias’:

Instead of a naïve scientist entering the environment in search of the truth, we find the rather unflattering picture of a charlatan trying to make the data come out in a manner most advantageous to his or her already held theories. (Fiske and Taylor, 1984, p.88)

(Further on (in §5.3), I advance an explanation for such variation in our assessment of the strength of evidence that is consistent with my general theory of Knowledge Intuitions.)

In summary, the support for the hypothesis that we possess Trait 2— that our Epistemic Standard depends on context—comes from Contextualist and Subject Sensitive Invariantist literature, from *prima facie* evidence that our assessments of evidential strength depend on context, and from social psychology experiments that suggest that our assessments of evidential strength depend on context. So, it seems plausible that we possess Trait 2.

4.3.2.3 Trait 3

Is it plausible that we possess Trait 3, according to which we seem to set Epistemic Standards that correlate (positively) with the SDT Index? Recall that index:

$$\text{SDT Index} = \frac{\text{cost of a false positive}}{\text{benefit of a true positive}} \times \frac{\text{Prior Pr(noise)}}{\text{Prior Pr(signal)}}$$

We may determine whether it is plausible that the Epistemic Standards that we seem to set correlate with the SDT Index by determining whether it is plausible that those standards

correlate positively with each of the two variables that are in the numerator, and correlate negatively with each of the two variables that are in the denominator. For example, one part of our investigation can be to see whether the Epistemic Standards that we seem to set correlate negatively with the benefit of a true positive. In effect, the SDT Index makes four separate predictions, one for each of the variables. However, before looking at the variables in turn, we need to look briefly at a certain general matter.

When biologists apply SDT, the cost of false positives and the benefit of true positives are standardly assessed in relation to Inclusive Fitness;³³ and it may well be that innate human detection thresholds (which thresholds are not generally dependent on environmental conditions, and which thresholds generally govern reflex, autonomic, and a range of stimulus-response behaviours) reflect the cost of false positives and the benefit of true positives in relation to our ancestors' Inclusive Fitness. However, it seems that *the Epistemic Standards* that individuals set (on the basis of the world conditions they encounter) may *not* reflect costs and benefits that those individuals measure in relation to Inclusive Fitness. Particularly, learning may well play a role in how individuals assess the cost of a false positive and the benefit of a true positive and we may learn to assess them in a way that does not correlate with Inclusive Fitness. However, SDT allows that these variables may be assessed in relation to *whatever it is that the operator values*, i.e., whatever the operator takes to be positive, good or beneficial. For example, say that an operator of a radar installation valued not being interrupted during her afternoon nap. For that operator, the cost of a false positive *during naptime* is much higher than it is at other times. As a result, during naptime, the SDT Index and optimal detection threshold (which correlates with the index) is higher than usual. The operator may manually raise the detection threshold of her system at naptime and thereby decrease the probability of being woken up. So, my working hypothesis is that *we assess the cost of a false positive and the benefit of a true positive on the basis of our values*, which may or may not help us to enhance our Inclusive Fitness. (It seems that many individuals have values that are not conducive to enhancing their Inclusive Fitness. It is another matter entirely as to whether our values are *generally* conducive to enhancing Inclusive Fitness. However, it would be hard to see how we would have done well as we have as a species if this were not so.)

Let us now see whether it is plausible that the Epistemic Standards that we seem to set correlate in the right way with each of the four constituent variables of the SDT Index.

³³ More precisely, biologists standardly look at proxies for Inclusive Fitness such as the ability to defend against predators or the caloric intake necessary to maintain a mechanism.

1. Cost of False Positives

It does indeed seem that the Epistemic Standards that we seem to set correlate positively with the cost of a false positive, consistent with what the SDT Index predicts. As this cost increases, the Epistemic Standard that we seem to set does too. This is precisely what is revealed in the knowledge attribution cases that Contextualists and Subject Sensitive Invariantists use to support their theories. (In Contextualist cases, attributors focus on the costs relative to their own practical interests; whereas, in Subject Sensitive Invariantist cases, attributors focus on the costs relative to the subject's practical interests.) For example, in Stanley's Low Stakes Case and Stanley's High Stakes Case, the cost of a false positive—i.e., the cost of assuming that the bank is open on Saturday when it is not—is higher in the High Stakes case than it is in the Low Stakes case; and so too is the Epistemic Standard that we seem to set.³⁴

In addition, there is some social psychology research that suggests that, in circumstances in which the cost of a false positive is relatively low, the Epistemic Standard that we seem to adopt tends also to be low: 'people appear to be sensitive to the context in which they make their predictions, and tend to exhibit greater optimism . . . when they expect that the consequences of being inaccurate will be less severe' (Armor and Taylor, 2002, p.339).

2. Benefit of True Positives

Let us see whether the Epistemic Standards that we seem to set correlate *negatively* with the benefit of a true positive; whether, for example, high benefits correlate with low Epistemic Standards. If so, the situation would be analogous to a radar operator setting a low detection threshold in order to avoid failing to detect approaching enemy fighters and to avoid failing to obtain the benefit that such a detection may produce.

Such a correlation appears to be revealed very often when evidence of a threat emerges, and the agent is in the position of being able to defend herself effectively against that threat. So, if the threat is real, the benefit of a true positive is the benefit produced by defensive actions. For

³⁴ A number of contrasting paired cases presented in the literature reveal this correlation. See, for example, Stewart Cohen's airport case (1999, p.58), Jeremy Fantl & Matthew McGrath's train cases (2002, pp.67-68), and Chandra Sripada & Jason Stanley's allergy cases (2012, pp.11-12).

example, there may be a substantial benefit in taking oneself as knowing that the email that includes a solicitation is a scam, and deleting it. This puts downward pressure on our Epistemic Standard we set when we consider whether the email is a scam.

The situation predicts a tendency to over-detect dangers, to be overly cautious. James Sage predicts that such over detection may be a trait that is common to a range of species: 'cautious belief-forming processes (say, those based on weak inductive generalizations) can systematically generate false beliefs, but still be [adaptive]' (2004, pp.102-103). We may tend to believe 'that all spotted mushrooms are poisonous (because consumption of spotted mushrooms was once followed by illness)' (p.103). Similarly, the situation predicts inappropriate defensiveness or hypervigilance. For example, under ordinary circumstances, working alongside someone who has HIV should be perfectly safe, given the way in which the HIV virus is known to be transmitted. However, it is widely reported that many people exercise excessive caution in relation to working alongside someone who has that condition. Again, if we would *tend to* benefit by behaving as if a threat exists, then downward pressure on the Epistemic Standard for the proposition that the threat exists is produced.

Additionally, a substantial experimental literature seems to indicate a negative correlation between the benefit of a true positive and the Epistemic Standard that seems to be set in relation to a range of positive propositions about the self. Taylor and Brown outline some of the benefits of taking oneself to know such propositions as being true:

overly positive self-evaluations, exaggerated perceptions of control or mastery, and unrealistic optimism . . . appear to promote . . . the ability to care about others, the ability to be happy or contented, and the ability to engage in productive and creative work. (Taylor and Brown, 1988, p.193)

Further, these positive attitudes allow many individuals to perform at higher levels:

It is widely believed that misjudgement produces dysfunction. Certainly, gross miscalculation can create problems. However, optimistic self-appraisals of capability that are not unduly disparate from what is possible can be advantageous, whereas veridical judgments can be self-limiting. When people err in their self-appraisals, they tend to overestimate their capabilities. This is a benefit rather than a cognitive failing to be eradicated. If self-efficacy beliefs always reflected only what people could do

routinely, they would rarely fail but they would not mount the extra effort needed to surpass their ordinary performances. (Bandura, 1989, p.1177)

It seems plausible that many of the benefits of the positive self-assessments indicated in the literature would only materialize if those assessments turned out to be not far from the truth. For example, perceptions of control and estimates of capabilities would only be beneficial if they are not far from the truth, not 'unduly disparate' from the truth as Bandura puts it. So, the benefit of a true positive, or of a close-to-true positive, in relation to positive self-appraisals is the benefit produced by the various actions implied above, e.g., engaging in creative work and mounting extra effort. The social psychology research suggests that this seems to lead to a lowering of Epistemic Standards. For example, that research indicates that members of different groups tend to hold such optimistic beliefs—including entrepreneurs, physicians, clinical psychologists, lawyers, negotiators, engineers and security analysts.^{35 36}

It seems that epistemologists have paid very little attention to the apparent effect of the benefit of a true positive on Epistemic Standards. However, the subject seems to have come up in other areas of philosophy, particularly in studies of self-deception, wishful thinking and motivated irrationality.

3. Prior Probability of Signal

Recall the example I offered earlier that I suggest shows why radar operators would tend to do well to negatively correlate the detection thresholds of their systems with the prior probability of signal. Aircraft have very often entered the airspace of a radar facility located in Toronto, Ontario, the high traffic facility; and aircraft have only infrequently entered the airspace of another radar facility located at an outpost in the Canadian Northwest Territories, the low traffic facility. The prior probability of signal is high at the former facility and low at the latter facility. As a result, the historic ratio of signal to noise at the high traffic facility is higher than it is at the low traffic facility. So, all other things being equal, if the high traffic and the low traffic

³⁵ A comprehensive listing of these studies is found in Griffin and Tversky (2002, p.230).

³⁶ Even so, Julie Norem suggests that anxious personality types lower their Epistemic Standard, not in relation to optimistic propositions about the self, but rather in relation to pessimistic propositions about the self. This 'pessimism is a strategy that helps us to work through our anxious thoughts rather than denying them, so that we may achieve our goals' (2001, p.3). Her research shows a marked tendency for these individuals to take themselves as knowing these pessimistic propositions as being true; and they seem to lower their Epistemic Standards in these cases. However, I set analysis of these findings aside.

systems both used the same detection threshold, then the ratio of true positives to false positives achieved at the high traffic facility would be higher than that achieved at the low traffic facility; and the operator of the low traffic system would tend to do better, tend to improve its ratio of true positives to false positives, with a higher detection threshold than would the operator of the high traffic system. My suggestion is that the situation in relation to the setting of Epistemic Standards is analogous to this situation. We tend to do well if our Epistemic Standards correlate negatively with what, for us, is analogous to the prior probability of signal. What is, for us, analogous to the prior probability of signal?

The expression 'prior probability of signal' may be viewed as the prior probability of the presence of target objects. At a radar installation, it is the prior probability of aircraft being in the system's airspace. This notion may be transposed to analyses of knowledge cases by generalizing from it: The prior probability of signal concerns the prior probability that questioned world conditions are present; or, alternatively, the prior probability that a questioned Advanced Proposition is true. Let us here focus on this prior probability that a questioned Advanced Proposition is true. I will refer to it as 'Prior $\Pr(p=\text{TRUE})$ '. The overarching question for us is whether the Epistemic Standards that we seem to set correlate negatively with $\Pr(p=\text{TRUE})$.

However, I suggest that Prior $\Pr(p=\text{TRUE})$ be viewed only as an analytical construct: This measure is not a particularly faithful representation of assessments that we ordinarily make—though it may help us to analyse knowledge cases. In particular, it does not seem plausible that the production of ordinary Knowledge Intuitions depends on our making actual prior probability calculations—coming up, for example, with a number between zero and one that reflects the number of times the proposition might be true out of one hundred iterations of the present circumstances. However, it *does* seem plausible that we keep track of the 'natural frequency' of events; and that the Epistemic Standards that we seem to set tend to be negatively correlated with such a frequency. So, Prior $\Pr(p=\text{TRUE})$ is an analytical construct that gets at these natural frequencies.

The natural frequency of events is a frequency that we experience or envision. For example, the natural frequency with which I receive mail delivery (i.e., old-fashioned paper mail) is once a day, and the natural frequency with which I am disturbed by Fido's barking is three or four times a day. This sense we have of the natural frequency of events need not be terribly fine grained. For example, my sense of the frequency with which the word of politicians pronounced in a hotly contested race is kept when they are in office is that this *quite often* does not occur.

Likewise, my sense of the frequency with which cars in my neighbourhood remain where they are parked overnight (because they are not stolen) is that this *almost always* occurs.

A series of cognitive psychology experiments over the last several decades reveal that we are good at reasoning with natural frequencies—while we are *not* characteristically good at reasoning with probabilities. For example, in one study, subjects were apt to provide correct responses to questions based on the statement that ten out of every 1000 forty-year-old women participating in routine screening have breast cancer. It seems that they were readily able to picture 1000 participants in these screenings, and ten of them having the ailment. However, subjects were *not* good at answering questions based on the statement that the probability of breast cancer is one percent for forty-year-old women participating in routine screening (Gigerenzer and Hoffrage, 1995, cited in McCarley and Benjamin, 2013, p.469).³⁷

As I suggested, it seems that our Epistemic Standards in relation to questioned propositions tend to be negatively correlated with such natural frequencies (and I support this hypothesis in the pages and chapters that follow). Accordingly, when I use the defined term, $\text{Pr}(p=\text{TRUE})$, in my analysis of knowledge cases, I will be referring to the subject's sense of certain natural frequencies (which I look at more closely in a moment).

To see the effect of our assessments of the Prior $\text{Pr}(p=\text{TRUE})$ on Epistemic Standards, it is helpful to first distinguish two basic kinds of evidence that we may use when we adjudicate propositions: (1) the evidence that gives us a sense of Prior $\text{Pr}(p=\text{TRUE})$, which is broadly-based or historical in nature, and (2) evidence that is directly connected to the specific proposition and, very often, recently gathered. We may not be consciously aware that our adjudications are partly based on the former type of evidence, whereas the latter type is always prominent in the context. (I do not view these evidence categories as being hard and fast, but rather as family resemblance descriptions of evidence types that will facilitate my analysis.) I refer to the latter type, the type (2) evidence, as “Present Evidence”.

In knowledge cases, type (1) evidence, and type (2) evidence or Present Evidence, both confer a probability on the questioned proposition p ; and the sum of the two probabilities equals the total probability conferred on p . For example, the sense a doctor has of the total probability

³⁷ Jason S. McCarley and Aaron Benjamin review the history of this research (2013, pp.468-470). They suggest that a number of studies that focused only on reasoning with probabilities or relative frequencies, and not on natural frequencies, *incorrectly* concluded that we are not good at Bayesian reasoning which takes prior probabilities into account.

conferred on the proposition p that his patient has diabetes, is based on both (a) type (1) evidence, the historical incidence of diabetes relative to that segment of the population that has the patient's medical profile, and which gives the doctor a sense of Prior $\Pr(p=\text{TRUE})$; and (b) the Present Evidence, the patient's present symptoms and lab results.

It is evident that we would do well to negatively correlate our Epistemic Standards with Prior $\Pr(p=\text{TRUE})$ —simply because the weaker our historical evidence is (i.e., the lower our sense of the Prior $\Pr(p=\text{TRUE})$), the stronger our Present Evidence will need to be in order to fully compensate for that weakness, and so the higher our Epistemic Standards will need to be. If the doctor has the sense that Prior $\Pr(p=\text{TRUE})$ is very low (i.e., the natural frequency with which individuals of the patient's medical profile contract diabetes is very low), then he would do well to set a very high Epistemic Standard, i.e., he would do well to require very strong Present Evidence (e.g., highly consistent symptomology and compelling lab results) to compensate for the weakness in his historical evidentiary support for p .

It is intuitively plausible that the Epistemic Standards that we seem to set are negatively correlated with Prior $\Pr(p=\text{TRUE})$. For example, in this last case in which the doctor has the sense that Prior $\Pr(p=\text{TRUE})$ is very low, it is plausible that the doctor would have set a very high Epistemic Standard, i.e., he would have required very strong Present Evidence in order to remove any doubt that his patient has the ailment. On the other hand, say the doctor had another patient who might also have contracted diabetes. However, this patient's medical profile included several strong correlates with diabetes—so that the doctor had the sense that Prior $\Pr(p=\text{TRUE})$ is high. It seems plausible that the doctor would have set only a moderate Epistemic Standard, i.e., only moderately strong Present Evidence would have been sufficient to remove doubt that the patient has the ailment.

*

If, as I suggested, our sense of $\Pr(p=\text{TRUE})$ informs our adjudications of everyday Advanced Propositions that arise in quick succession, then clearly that sense does not generally come from statistical studies of the kind that the doctor relies on. As I proposed earlier, that sense of $\Pr(p=\text{TRUE})$ generally comes from our sense of certain natural frequencies. I will demonstrate this using Stanley's Low Stakes Case.

In this case, Hannah sets an Epistemic Standard in relation to the proposition p that the bank will be open the coming Saturday. This standard is low and positively correlated with the low

cost of a false positive in the case. I suggest additionally that the standard is negatively correlated with her sense of $\Pr(p=\text{TRUE})$. For example, it may have been that Hannah had the sense that $\Pr(p=\text{TRUE})$ was relatively high, i.e., that the prior probability of the bank being open on Saturday is high; and, if so, this would have put additional downward pressure on her Epistemic Standard. Let us look at how she might have gotten such a sense of $\Pr(p=\text{TRUE})$.

Hannah might have reflected on the frequency with which propositions that are of the same kind as p are true. For example, she might have recognized that (i) two months ago, the proposition p_1 , that the competing nearby bank A will be open the coming Saturday, would have been true, (ii) one month ago, the proposition p_2 , that nearby bank B will be open the coming Saturday, would have been true, and (iii) two weeks ago, the proposition p_3 , that nearby bank C will be open the coming Saturday, would have been true. Thus, propositions that are of the same kind as p are frequently true. Having a sense of this may constitute Hannah's sense that $\Pr(p=\text{TRUE})$ is high, i.e., that the prior probability of the bank being open on Saturday is high (which would have put downward pressure on her Epistemic Standard in the case).

SDT predicts that if, instead, Hannah had the sense that the $\Pr(p=\text{TRUE})$ was low, then this would have put upward pressure on her Epistemic Standard, and potentially raised it to a level that her Present Evidence would not have been strong enough to meet. Under such circumstances, she would not have taken herself as knowing that the bank would be open—even though the stakes are low. In fact, precisely such an affect is suggested in Keith DeRose's Contextualist theory (which I outlined in §2.3).

While DeRose argues that context determines the range of relevant alternatives, *his examples* seem to indicate that it is, *in particular*, the subject's sense of $\Pr(p=\text{TRUE})$ that determines the range of relevant alternatives. He seems to suggest, for example, that if the subject has the sense that $\Pr(p=\text{TRUE})$ is low, then the range of alternatives may be extended, effectively raising the subject's Epistemic Standard. Referring to his low stakes bank case (which I reproduced in §2.3), he makes the following suggestion:

If very many nearby banks have discontinued their Saturday hours in the last two weeks, then it seems that my original claim to know may well have been false, and if I admit that I did not know after this surprising fact about local banks is called to my attention, I will be taking back and contradicting my earlier claim to have known. (1992, p.921)

If DeRose learns that very many nearby banks have discontinued their Saturday hours, then he will have the sense that Prior $\Pr(p=\text{TRUE})$ is lower than what he assumed; and he will, as a result, require stronger Present Evidence (which in the case consists only of his having seen that the bank was open two weeks before on a Saturday) in order to eliminate doubt. His Epistemic Standard is effectively lifted, and his Present Evidence fails to meet it. So, he fails to take himself as knowing that the bank will be open on Saturday.

Additionally, the hypothesis that Epistemic Standards negatively correlate with our sense of $\Pr(p=\text{TRUE})$ is broadly consistent with a hypothesis that Stewart Cohen defends. As I pointed out (in §2.3), he suggests that an alternative possibility becomes relevant and needs to be ruled out (in order for the subject to have knowledge) if the subject's *internal statistical evidence* indicates that the probability of the alternative is high (1988, p.102). Notice that if a subject's internal statistical evidence indicates that the probability of an alternative is high, then that subject will have the sense that $\Pr(p=\text{TRUE})$ is low; and if the latter is low, then SDT predicts that the subject's Epistemic Standard will be high. This is a state of affairs that is broadly consistent with the subject finding that she has an additional alternative possibility that is relevant and that needs to be ruled out.

4. Prior Probability of Noise

Let us see whether the Epistemic Standards that we seem to set correlate *positively* with what is, for us, analogous to the prior probability of noise. Earlier I transposed the notion of the prior probability of signal to knowledge cases by generalizing from the former to the notion of the Prior Probability that a questioned Advanced Proposition is true. Further, I suggested that I would use the term, $\Pr(p=\text{TRUE})$, to refer to the sense we get of the natural frequency with which propositions of the kind that are in question are true. So too, let us transpose the notion of the prior probability of noise to knowledge cases by generalizing from it to the notion of the Prior Probability that a questioned Advanced Proposition is false; and let us use the term, $\Pr(p=\text{FALSE})$, to refer to the sense we get of the natural frequency with which propositions of the kind that are in question are false.

Because Prior $\Pr(p=\text{TRUE})$ and Prior $\Pr(p=\text{FALSE})$ are probability complements that add up to 1, when one increases, the other decreases. So, if it is plausible that the Epistemic Standards that we seem to set are negatively correlated with Prior $\Pr(p=\text{TRUE})$ (and I believe that I have shown that it is so), we may predict that the Epistemic Standards that we seem to set are positively

correlated with Prior $\Pr(p=\text{FALSE})$. As this probability increases, our Epistemic Standard should increase.

*

To summarize, it is plausible that the Epistemic Standards that we seem to set correlate in the right way with each of the four constituent variables of the SDT Index. Therefore, it seems plausible that the Epistemic Standards that we seem to set correlate positively with the SDT Index, and that we possess Trait 3.

4.3.2.4 Trait 4

Is it plausible that we possess Trait 4, according to which we do not possess a cognitive device that can trump the KIP in relation to the control of calculated deliberate behaviour? Before answering this question, note the critical connection between this trait and Trait 3 above (according to which the Epistemic Standards that we seem to set correlate positively with the SDT Index). It would seem that Trait 3 could not have evolved if we had not already possessed Trait 4. Trait 3 allows us to tend to intuit knowing that which we would tend to do well to intuit knowing—because the KIP will have sought to achieve an appropriate balance between the risk of being overly sceptical and the risk of being overly trusting in the circumstances. However, if the KIP and the Knowledge Intuitions it produces could be trumped by some other cognitive device in relation to the control of calculated deliberate behaviour (particularly if it could be trumped by a device that is not governed by SDT principles) then the adaptive pressure for Trait 3 could be significantly diminished, if not eliminated. (So, it may well be that Traits 3 and 4 are best viewed as a single more inclusive trait. However, as I suggested earlier, I separated the traits for presentational reasons.)

A look at the plausibility of the hypothesis that we possess Trait 4 reveals a wide range of apparent counterexamples. These suggest that behaviour may be controlled by a cognitive device that is not the KIP. I have in mind all those cases in which the Present Evidence fails to meet the Epistemic Standard that appears to be set, but in which the subject decides to take a chance by acting or forming an attitude as if the questioned proposition is true anyway. The KIP seems to advise us to not act as if the questioned proposition is true; but we do so anyway. For example, I do not know that my house will catch fire, and yet I buy fire insurance; and I do not know that I will win the lottery, and yet I buy a lottery ticket. The behaviour that is featured in

such cases generally depend on a conscious, rational cognitive device. (We will see that this cognitive device does not appear to be governed by SDT principles: It does not seem to depend on whether the strength of the Present Evidence meets an Epistemic Standard that correlates with the SDT Index.)

However, a close look at these cases reveals that the conscious, rational cognitive device to which I refer nonetheless depends on Knowledge Intuitions. These Knowledge Intuitions arise in these cases not in relation to the main questioned proposition, but rather in relation to certain other propositions. As a result of this, control of behaviour by the KIP is maintained. In this section, I discuss these cases and defend this last claim.

Given that all these apparent counterexamples involve an agent taking some sort of chance—insofar as they have failed to take themselves as knowing to be true the main questioned proposition—it would seem that the counterexamples would all have one or the other of the following two features:

- (1) The agent's behaviour is precautionary in nature, performed in order to avoid a potential harm. I call these 'Insurance Cases'.
- (2) The agent's behaviour constitutes a gamble taken in the hope of obtaining a benefit. I call these 'Gambling Cases'.

In Insurance Cases, agents assume or conclude that they would be harmed by the occurrence of an event. They question whether the harmful event will occur, and find that they cannot rule out the possibility that it will. So, they take an appropriate precaution, performing a particular action or forming a particular attitude. They act though they do not intuit knowing that the event will occur. For example, I assume that it would be bad for me if my house was damaged or destroyed by fire, and I cannot rule out the possibility that it will. So, *I buy fire insurance even though I do not intuit knowing that my house will catch fire.* Or, I assume that my clothes would get uncomfortably wet if it rained, and I cannot rule out the possibility of rain. So, I take an umbrella even though I do not intuit knowing that it will rain. Or, I conclude that my stock market investment would be adversely impacted by a market crash, and I cannot rule out that possibility. So, I sell off my stocks even though I do not intuit knowing that the market will crash.

In Gambling Cases, agents assume or conclude that the performance of a particular action or forming a particular attitude could produce some *benefit*, and they are not able to rule out the possibility that the action would not produce this benefit. So, they perform that action or form that attitude with the hope that the benefit would be produced, even though they do not know whether it will. For example, I assume that buying a lottery ticket could lead to a win, and I am not able to rule out the possibility that it would not. So, *I buy a lottery ticket even though I do not intuit knowing that I will win*. Or, I know that the restaurant is either on Oak Street or Pine Street, but do not know which. I assume that turning down Oak Street could lead directly to the restaurant, and am not able to rule out the possibility that it would not. So, I will turn down Oak Street even though I do not intuit knowing that the restaurant is on Oak Street. If I do not find the restaurant on Oak Street, then I will try Pine Street.

Before looking more closely at this, notice that F1 predicts that we would *not* take ourselves as knowing the main questioned proposition in either Insurance or Gambling Cases. For example, F1 predicts that in a typical case in which an individual buys fire insurance that individual would not take herself as knowing that her house will catch fire; and in a typical case in which an individual buys a lottery ticket that individual would not take herself as knowing that she will win. To see this, recognize that, under F1, when one takes oneself as knowing that *p*, behaving as if *p* has no limitations or restrictions (beside physical limitations). For example, if I take myself as knowing and have no doubt at all that the man will attempt to kill me, my behaviour may well breach ordinary boundaries. (This may be viewed as the reason that, in some cases, the elimination of all doubt requires very strong evidence.) So, according to F1, if the agent in Insurance and Gambling Cases took herself as knowing that *p*, then their actions and attitudes as if *p* would be unlimited or unrestricted. If, in fact, I took myself as knowing, as having no doubt at all, that my house will catch fire, I would do much more than buy fire insurance. I may take any kind of reasonable and practical preventative action; if I was not aware of any such action, then I may make inquiries about this; and if, finally, I concluded that there were no preventive actions that I could perform, I would move out of my house. Similarly, if, in fact, I took myself as knowing, as having no doubt at all, that I would win the lottery, I would do much more than buy a lottery ticket. I might also buy a yacht. Finally, F1 suggests that for Knowledge Intuitions to be produced one must have assessed that (unrestricted) behaving as if *p* would tend to produce better consequences than would (unrestricted) not behaving as if *p*. However, given the kinds of unrestricted behaviour that may follow from behaving as if *p* in Insurance and Gambling Cases (e.g., moving out of the house, buying a yacht), and given the strength of the Present Evidence standardly featured in these cases, we do *not* assess that (unrestricted) behaving as if *p* will tend to produce better consequences than that of (unrestricted) not

behaving as if p . So, F1 predicts that we would not take ourselves as knowing. In a typical lottery case, we assess that unrestricted behaving as if the lottery ticket is a winner would not tend to produce better consequences than that of unrestricted not behaving as if it is. So, F1 predicts that the individual would not take herself as knowing that the ticket is a winner.

Let us now look more closely at Insurance and Gambling Cases. In both kinds of cases, the agent's Present Evidence is *not* strong enough to produce Knowledge Intuitions in relation to the main questioned proposition, e.g., my house will catch fire, or I will win the lottery. Even so, agents have a strategy for dealing with such situations: They (1) take themselves as knowing that the main questioned proposition, p , is a *practical possibility*, and then (2) *adjudicate a different proposition* specifically about whether it would be either *prudent* or *worthwhile* to behave in some limited way as if p , e.g., buy fire insurance (but do not sell the house), or buy a lottery ticket (but do not buy a yacht).³⁸ For example, if I recognize that I am not going to be able to take myself as knowing whether or not it will rain, then I may (1) take myself as knowing that it *might* rain (i.e., that rain is a practical possibility) and (2) adjudicate the separate proposition that it would be *prudent* to behave in some limited way as if p , e.g., take an umbrella with me (but not change my plans altogether). If I realize that I am not going to be able to take myself as knowing that the stock market will crash, then I may (1) take myself as knowing that it *might* crash, and (2) adjudicate the separate proposition that it would be *prudent* to behave in some limited way as if it will crash, e.g., sell my stocks (but not bet against or 'short' the stock market). If I realize that I am not going to be able to take myself as knowing that I will win the lottery, then I may (1) take myself as knowing that I *might* win, and (2) adjudicate the separate proposition that it would be *worthwhile* to behave in some limited way as if I will win, e.g., buy a lottery ticket.³⁹

Critically, in both Insurance and Gambling Cases, all other things being equal, deliberate action still depends on having Knowledge Intuitions: All other things being equal, agents' actions and attitudes depend on their *taking themselves as knowing* (1) that the main questioned proposition is a practical possibility, and (2) that it is either prudent or worthwhile to perform

³⁸ It may be that my sense that it is worthwhile to buy a lottery ticket consists in my sense that the expected utility of buying the ticket is positive.

³⁹ If in a given case, subjects have taken themselves as knowing that the questioned proposition is a practical possibility but have not taken themselves as knowing to be true a proposition about the prudence or worthwhileness of a particular limited action, *then* they may try to identify and adjudicate an *alternative proposition* about the prudence or worthwhileness of a different limited action. For example, if I intuit knowing that it *might* rain, but do not take myself as knowing that it is prudent to take an umbrella, then I may adjudicate the alternative proposition that it would be prudent to wear my hooded coat. Then, if I intuit that I know that it is worthwhile wearing my hooded coat, all other things being equal, I will do so.

particular actions or to form particular attitudes. All other things being equal, my taking an umbrella depends on my taking myself as knowing (1) that it might rain, and (2) that it is *prudent* to take an umbrella. All other things being equal, my buying a lottery ticket depends on my taking myself as knowing that (1) I might win the lottery, and (2) that it is *worthwhile* to buy a ticket. So, although the actions that are performed and attitudes that are formed in these cases do not depend on intuitions of knowing the main questioned proposition; they do depend on knowing the *two secondary propositions to which I refer*. Further, it seems that if one does *not* take oneself as knowing *either* the main questioned proposition or the two secondary propositions, then, in the moment, one is *stuck and unable to act*.⁴⁰

*

It appears as though there is another class of counterexamples to the hypothesis that we possess Trait 4, suggesting that behaviour may be controlled by a cognitive device other than the KIP. These are ‘rational acceptance’ cases—in which an agent performs an action or forms an attitude on the basis of rationally accepting a proposition that they do not take themselves as knowing. Michael Bratman claims that such cases indicate ‘that the cognitive attitudes guiding practical reasoning and action go beyond our beliefs’ (1999, p.15); and, further, ‘we need to make room for both [belief and rational acceptance] in [a] model of practical reasoning’ (pp.20-21). Addressing these cases not only deals with the problem of apparent counterexamples to the hypothesis that we possess Trait 4; doing so will also help us deal with the objection (raised in §3.2.3) that what I am calling Knowledge Intuitions, and am analysing as having function F1 (and F2), would be better called, not that, but ‘intuitions of rational acceptability’.

⁴⁰ Compare the feelings of certainty or apprehensiveness of subjects in standard knowledge cases with those of subjects in Insurance and Gambling Cases. In the former cases, the subject takes herself as knowing the main questioned proposition, e.g., Hannah knowing that the bank is open on Saturday. In these cases, the subject feels, in the moment, *certain that one would do well to act or form an attitude as if p*, e.g., Hannah feels, in the moment, *certain that she will do well to wait until Saturday to deposit her cheque*. In Insurance and Gambling Cases, subjects do not take themselves as knowing the main questioned proposition, and, accordingly, *do not feel certain that they will do well when they act or form an attitude as if p*. If you do not take ourselves as knowing that the stock market will crash and nonetheless find that it would be prudent to sell your stocks, then you do not feel certain that you will do well when you sell your stocks. You may instead feel apprehensive. So, the mental states of subjects in Insurance and Gambling Cases are very different than those of subjects in standard knowledge cases.

Even so, when subjects intuit knowing the two secondary propositions that come up in Insurance and Gambling Cases, they feel certain about those two propositions. They feel certain that the main questioned proposition is a practical possibility and certain that it is prudent or worthwhile to perform the particular limited action or to form the particular attitude. If you do not take ourselves as knowing that the stock market will crash, then you may feel *certain that it is prudent* to sell your stocks.

Consider one kind of case that Bratman brings to our attention. Say that my very close friend, Tonya, has been accused of acting badly on a particular occasion and that I do *not* take myself as knowing that she *did not* do so. It seems at least possible that she did. The matter of her behaviour comes up in conversation with a third party. It is clear to me that if I say nothing about the behaviour, then the third party would infer that I believe that Tonya acted badly. So, on my reading of the social situation, I have two unpleasant alternative courses of action: either say nothing and thereby condemn my close friend by allowing the third party to make condemnatory inferences, or defend my close friend even though I do not take myself as knowing that she did not behave badly. So, I rationally accept that Tonya did not behave badly and defend her accordingly. Bratman suggests that, in such a case, rational acceptance assumes control of action, instead of knowledge.⁴¹ I do not know that my close friend did not behave badly, but I rationally accept that she did not do so, and my defence of her is based on rational acceptance, not a Knowledge Intuition.

However, look more closely. Let us make a plausible assumption about the case: My defence of Tonya that followed my rational acceptance (that she did not behave badly) was less vigorous than it would have been had it followed my taking myself as knowing (that she did not behave badly). Had it followed Knowledge Intuitions, my defence may have included a pronounced and sincere evincing of moral indignation about the suggestion that Tonya behaved badly, that she could have behaved that way—which was not included in the defence that followed rational acceptance. With this, we see that this rational acceptance case is a Gambling Case: (1) I intuited knowing that Tonya *might* not have behaved badly, i.e., I intuited knowing that the main questioned proposition is a practical possibility; and (2) I intuited knowing that it would be *worthwhile* to defend Tonya in some limited way, i.e., with limited vigour. I took a gamble on Tonya's innocence. Moreover, had I not taken myself as knowing these last two secondary propositions, I would not have defended Tonya at all. I find that all the cases that Bratman presents in his typology of rational acceptance cases (1999, pp.20-26) are either Insurance or Gambling Cases.

*

It seems that the mediation of behaviour in Insurance and Gambling Cases involves certain conscious states, and that these states depend on Knowledge Intuitions. I refer to the conscious

⁴¹ See also Jonathan Cohen (1989).

states that arise in Insurance Cases as *considered decisions to act cautiously*; and refer to the conscious states that arise in Gambling Cases as *considered decisions to act incautiously*. If I have decided to act cautiously, I may buy an insurance policy or take an umbrella along with me in case it starts to rain. If I have decided to act incautiously, I may buy a lottery ticket.

My use of these expressions, 'considered decisions to act cautiously' and 'considered decisions to act incautiously', is not meant to suggest that the decisions to which I refer are either (a) objectively cautious or incautious, or that (b) individuals take them to be excessively cautious or incautious. Rather, I want to suggest that, in the moment that such decisions are made, the agent is not certain that she will do well when she behaves accordingly; and I want to suggest that the agent has the sense that her behaviour is *appropriately* cautious or *appropriately* incautious. (This contrasts situations in which individuals take themselves as knowing that the main questioned proposition is true. If I take myself as knowing it will rain, and, on that basis, take an umbrella, then, at least in that moment, I am certain that I would do well to do so.)

It seems that the mental states that are often referred to as rational acceptances involve either considered decisions to act cautiously or considered decisions to act incautiously. For example, my rational acceptance that Tonya did not act badly involves a considered decision to act in a way that is appropriately incautious. Even so, it seems that rational acceptances are still *somehow* different than either the considered decisions to act cautiously in Insurance Cases or the considered decisions to act incautiously in Gambling Cases. While I may make a considered decision to act in a way that is appropriately incautious when I buy a lottery ticket, it would be odd to suggest that I have rationally accepted that I would win. (I set aside any further investigation into the relation between these mental states.)

However, let us look more carefully at the Knowledge Intuitions on which our considered decisions to act depend, i.e., the Knowledge Intuitions produced in Insurance and Gambling Cases. My decision to act cautiously and buy fire insurance depends on two prior Knowledge Intuitions:

 KI (1): The Knowledge Intuition that it is a practical possibility that my house will catch fire, and

 KI (2): The Knowledge Intuition that it would be prudent to behave in some limited way as if my house will catch fire, viz., prudent to buy insurance.

When I had KI (1), I started behaving as if that proposition is true. I did so by adopting an *ambivalent attitude* about whether my house would catch fire; and if I had not had KI (1), I would probably instead have had the definitive attitude that my house would not catch fire. Further, when I had KI (2), I started behaving as if that proposition is true. I did so by adopting a positive attitude about buying fire insurance; and if I had not had KI (2), I would probably instead have had a negative attitude about buying fire insurance. It seems clear that the considered decision to act cautiously that led to me buying fire insurance depended on these two attitudes, (1) my ambivalence about whether the house would catch fire, and (2) my positive attitude about buying fire insurance. An analogous analysis of the rational acceptance case may be made. The conclusion of that analysis is that the rational acceptance that led to me defending Tonya depended on (1) my ambivalence about whether she behaved badly, and (2) my positive attitude about defending her with a certain limited vigour. Thus, the critical control by Knowledge Intuitions and the KIP of behaviour is maintained in these cases.

Also, the foregoing constitutes an account of the difference between knowing and rationally accepting. In a word, Knowledge Intuitions permit unlimited actions and attitudes as if p , whereas rational acceptance (which must be supported by taking oneself as knowing to be true certain secondary propositions) permit certain limited actions and attitudes as if p .

*

In this section I brought forward important classes of counterexamples (the only ones I can think of) to the hypothesis that we possess Trait 4, according to which we do not possess a cognitive device that can trump the KIP in relation to the control of calculated deliberate behaviour. However, I found in these cases that control of behaviour is maintained by the KIP and the Knowledge Intuitions it produces. I take it from this that it remains plausible that we possess Trait 4.

4.3.2.5 Conclusion

In conclusion, it seems plausible that we possess the four traits that I identified, and this constitutes direct evidence that the operating principles of the KIP are SDT principles, evidence that the production of ordinary Knowledge Intuitions and our use of 'know' reflects SDT principles. Accordingly, my working hypothesis shall be that our Epistemic Standards correlate with the SDT Index. I add that the support for this hypothesis that comes from the direct

evidence above, will be accompanied by support from a wide range of consistent examples that I present in Chapters 5 and 6.

However, the term 'SDT Index' does not seem particularly fitting in discussions of everyday knowledge cases. It is unnatural, abstract, and does not seem particularly connected to the experience of taking oneself as knowing. Even so, recall (from §4.2.2) that a different way of conceptualizing the index emerges when we view the numerators as being connected, and the denominators as being separately connected. It may be conceptualized as the expected cost of a false positive relative to the expected benefit of a true positive, based on the prior evidence. This conception is analogous to the notion that is referred to as 'risk/reward environment'. I find that discussions of knowledge cases using this picture to be quite natural, and resonant with how we feel when we take ourselves as knowing. So, I will begin now to refer to the SDT Index as the 'Risk/Reward Quotient'. Thus,

$$\text{Risk/Reward Quotient} = \frac{\text{cost of a false positive} \times \text{Prior Pr}(p=\text{FALSE})}{\text{benefit of a true positive} \times \text{Prior Pr}(p=\text{TRUE})}$$

4.4 On Tracking Truth

If the Epistemic Standards that we seem to set correlate with the Risk/Reward Quotient; and if, in the moment, the Risk/Reward Quotient is low, then we may intuit knowing that p even if our evidence for p is weak. This prediction seems to go against the view suggested by some philosophers that natural selection favours reliable belief-formation processes. Richard Feldman explains the latter view.

[I]f a being has beliefs at all, it is better (that is, more conducive to survival) for it to have true beliefs than false beliefs. True beliefs about where one's food is are more helpful for finding food, and surviving, than are false beliefs. Similarly, true beliefs about where one's predators are and how to escape them are more survival enhancing than false beliefs about these matters. So, natural selection is likely to select for believers that have mostly true beliefs. The best way, perhaps the only way, for believers to have mostly true beliefs is for them to have reliable belief-forming mechanisms or strategies. Reliable mechanisms or strategies are ones that lead mostly to true beliefs. Hence, natural selection will select believers that have reliable belief-forming mechanisms. (1988, p.218)

Additionally, Christopher Stephens develops a detailed model that predicts that under what would seem to have been prevalent circumstances ‘non-reliable belief formation policies will get the organism into trouble’ (2001, p.183) and produce selection pressure for the production of reliable belief formation processes. Kim Sterelny (2003) similarly finds that in the human case, selection ‘will favour more reliable mechanisms over less reliable ones’ (2003, p.32) that track truth. (I do not suggest that either Feldman, Stephens or Sterelny conclude from this that true-belief producing Cognitive processes necessarily evolved.)⁴²

Even so, notice that neither F1 nor SDT suggest that the purpose of intuiting that one knows on the basis of weak evidence is to have false beliefs, because false beliefs are better. Instead, SDT suggests that Epistemic Standards may be set low for *the purpose of not failing to believe those propositions that advance our interests*, whether they are true or not. The situation is analogous to a radar operator setting a low detection threshold: She may do so in order to avoid failing to detect approaching enemy fighters (i.e., to avoid failing to obtain the benefit of a true positive).

Further, F1 and SDT suggest that even if, under some circumstances, we do not track truth particularly well; under *other* circumstances, particularly when the Risk/Reward Quotient is high, we do track truth well. Under the latter circumstances, the KIP will have to produce Knowledge Intuitions that are reliable. Truth will be ‘a fuel for success in dealing with the world’ (Godfrey-Smith, 1998, p.172).

It seems that the view that nature favours an *invariably reliable* belief-formation process fails to take into account the benefit that can follow from modulating evidential standards; while SDT fully takes such benefits into account. In fact, if SDT principles govern the KIP, then the hallmark of the KIP is that it actively adjusts Epistemic Standards up and down in different contexts.

⁴² Relatedly, correspondence theories of truth have been supported by naturalist philosophers, who suggest that correspondence may explain behavioural success. Godfrey-Smith provides a general discussion of this literature (Godfrey-Smith, 1998, Chapter 6).

Chapter 5: How Do We Do It?

In the last chapter, I found that it is plausible that SDT principles govern the production of Knowledge Intuitions and our use of ‘know’, so that the Epistemic Standards that we seem to set correlate with the Risk/Reward Quotient. In this chapter, I focus on practical matters—particularly on how we adjudicate propositions in a way that takes the Risk/Reward Quotient into account. Additionally, I look at special practical circumstances under which we strategically suspend Knowledge Intuitions, even though our Present Evidence meets our Epistemic Standard (in §5.2); how we adjudicate special propositions that are not associated with either a particular practical situation or any particular action or attitude (in §5.3); and how good we are at setting Epistemic Standards, and whether or not those standards are optimal, i.e., whether they lead to Knowledge Intuitions that tend to produce consequences that are the best that we are capable of achieving, given our circumstances, limitations, and so on (in §5.4).

5.1 The Basic Method

We may get a sense of how we go about adjudicating Advanced Propositions in a way that takes the Risk/Reward Quotient into account by looking at what it is like to have a sense of (a) the Risk/Reward Quotient, and (b) the strength of our Present Evidence and Epistemic Standard.

The Risk/Reward Quotient: Before looking at the Risk/Reward Quotient as a whole, let us look more closely at the sense we get of the Prior Probability that a questioned proposition is true, Prior $\Pr(p=TRUE)$, and also Prior $\Pr(p=FALSE)$. Earlier (in §4.3.2.3), I suggested that in my analyses, these measures would refer to a sense we have of a ‘natural frequency’ (the frequency with which experienced or envisioned events occur).⁴³ In particular, I indicated that Prior $\Pr(p=TRUE)$ refers to the sense we have of the natural frequency with which propositions of the kind in question are true (and that Prior $\Pr(p=FALSE)$ refers to the sense we have of the natural frequency with which propositions of the kind in question are false). In my earlier example, Hannah, in Stanley’s Low Stakes Case, might have recognized that (i) two months ago, the proposition p_1 , that the competing nearby bank A will be open the coming Saturday, would

⁴³ I pointed out (in §4.3.2.3) that a series of cognitive psychology experiments over the last several decades reveal that we are good at reasoning with natural frequencies while we are not characteristically good at reasoning with probabilities. We are also not good at reasoning with relative frequencies (see fn. 37).

have been true, (ii) one month ago, the proposition p_2 , that nearby bank B will be open the coming Saturday, would have been true, and (iii) two weeks ago, the proposition p_3 , that nearby bank C will be open the coming Saturday, would have been true. She might then have gotten the sense that propositions that are of the same kind as p are frequently true.

However, there are a vast number of everyday Advanced Propositions that we adjudicate; and it does not seem plausible that, for all of these, we would be able to get a sense of the natural frequency with which propositions of the kind that are in question are true. Even so, it seems that we manage under these circumstances by disregarding those features of propositions of the kind that are in question that are unlikely to be important. So, it seems more plausible to suggest that we may concern ourselves with the natural frequency with which *general propositions* of the kind that is in question are true. For example, Hannah's sense of the Prior Probability that her bank is open this Saturday may be derived, not so much from recognizing the situation in relation to the particular banks A, B and C. Instead, it may be derived from her getting the sense that banks generally are open on Saturdays (disregarding all the irrelevant particulars). So, it seems plausible that we arm ourselves with the ability to get a sense of Prior $\Pr(p=\text{TRUE})$ and Prior $\Pr(p=\text{FALSE})$ (which we use in our adjudications of everyday Advanced Propositions) by noticing and keeping track of the frequency of the appearance or occurrence of different *general kinds* of objects, relations, states and events—such as the frequency with which banks generally are open on Saturdays.

With this, let us consider the sense we may get of the Risk/Reward Quotient. Recall again the quotient:

$$\text{Risk/Reward Quotient} = \frac{\text{cost of a false positive} \times \text{Prior } \Pr(p=\text{FALSE})}{\text{benefit of a true positive} \times \text{Prior } \Pr(p=\text{TRUE})}$$

The sense we have of the Risk/Reward Quotient plausibly involves (a) modulating the negative valuation that constitutes our sense of the cost of a false positive in accordance with Prior $\Pr(p=\text{FALSE})$, (b) modulating the positive valuation that constitutes our sense of the benefit of a true positive in accordance with Prior $\Pr(p=\text{TRUE})$, and (c) weighing the numerator against the denominator.

However, on my account, we adjudicate Advanced Propositions in quick succession in order to flag the best course of action available to us from moment to moment. Is it plausible that we possess the cognitive speed and capacity to perform these multiple assessments in 'real time'? I

think so. First, notice that these assessments are not very different from those that are made when we determine the expected utility of actions or events—a measure widely used by economists and game theorists. These expected utility calculations seem to predict a wide range of human behaviours. Second, there are other domains in which, it seems, we take multiple factors into account almost instantly. For example, depth perception is produced by a process that takes many varied cues into account.⁴⁴ Third, it seems that infra-human species make analogous assessments. C. R. Gallistel offers an account of an experiment that suggests that ducks modulate their assessment of the benefit of an event in accordance with the frequency of its occurrence.

Every day two naturalists go out to a pond where some ducks are overwintering and station themselves about 30 yards apart. Each carries a sack of bread chunks. Each day a randomly chosen one of the naturalists throws a chunk every 5 seconds; the other throws every 10 seconds. After a few days' experience with this drill, the ducks divide themselves in proportion to the throwing rates; within 1 minute: after the onset of throwing, there are twice as many ducks in front of the naturalist that throws at twice the rate of the other. One day, however, the slower thrower throws chunks twice as big. At first the ducks distribute themselves two to one in favour of the faster thrower, but within 5 minutes they are divided fifty-fifty between the two "foraging patches." This behaviour is predicted by a model of an evolutionarily stable foraging strategy. The model assumes that ducks and other foraging animals can represent rates of return, the number of items per unit time multiplied by the average size of an item. (Gallistel, 1990, p.2)

The strength of the Present Evidence and Epistemic Standard: Let us now consider what it is like to have a sense of the strength of the Present Evidence and the Epistemic Standard. For example, in Stanley's High Stakes Case, what is it like for Hannah to have a sense of the strength of her evidence—that the bank was open two weeks before? And what is it like for her to have a sense of her Epistemic Standard?

We may measure the strength of Present Evidence and Epistemic Standards using probabilities. For example, the Present Evidence may confer a 95% probability of truth on p ; and find that the Epistemic Standard requires that the evidence confers a 98% probability of truth on p .

⁴⁴ Simon Grondin presents a summary of principle findings on depth perception cues (2016, pp.103-110).

However, as I pointed out earlier, we are, apparently, not good at reasoning with probabilities; whereas we are good at reasoning with natural frequencies. So, it may be that the sense we ordinarily have of both the strength of our Present Evidence and our Epistemic Standards consists of our sense of certain natural frequencies. Particularly, *we may represent the strength of our Present Evidence to ourselves as the frequency with which evidence like our Present Evidence correctly predicts propositions like the proposition in question.* For example, in Stanley's Low Stakes Case, Hannah may represent the strength of her Present Evidence (i.e., the bank having been open two weeks before on a Saturday) to herself as the frequency with which such evidence would correctly predict propositions of the sort that is in question. Say that she has noticed that since she switched to her present bank over a decade ago, the bank has only changed its operating hours five or six times. This would allow her to infer that evidence like her Present Evidence would *generally* correctly predict propositions like the questioned proposition, that the bank will be open this Saturday.

Notice that the yardstick we use to measure the strength of the Present Evidence should also be used to set Epistemic Standards. Epistemic Standards are merely thresholds that the strength of the Present Evidence must meet in order for Knowledge Intuitions to be produced. This state of affairs suggests that we may represent Epistemic Standards to ourselves as *the minimum natural frequency with which evidence like our Present Evidence correctly predicts propositions like the proposition in question.* For example, in Stanley's High Stakes Case, Hannah's sense of her Epistemic Standard may consist of her sense that she requires Present Evidence of a kind that would *practically never fail* to correctly predict proposition of the kind that is in question. If so, she may find that her Present Evidence, which she assesses would *generally* predict propositions such as that the bank will be open this Saturday does not meet this Epistemic Standard.

So, it would seem that in order to be armed with the ability to get a sense of the strength of the Present Evidence, we would need to *notice and keep track of the frequency with which different general kinds of objects, relations, states or events correctly predict other different general kinds of objects, relations, states or events.* For example, I have noticed that my friend *almost always* does what she says she will do. This arms me with the ability to get a sense of the strength of Present Evidence that consisted of my friend saying what she will do. Similarly, I have noticed that the word of politicians pronounced in a hotly contested race *frequently* fails to predict their policies once in office, that the conclusions of an outlier social psychology experiment *very often* proves to be incorrect; that circumstantial evidence *regularly* fails to predict the guilt of a trial defendant; that three consistent and independent eye-witness accounts *almost always*

describes what in fact occurred; and that a randomized controlled trial in relation to propositions about the efficacy of a new medical treatment is *almost always* correct.

In this emerging picture, it would seem that the way in which we set Epistemic Standards is straightforward. We first get a sense of the Risk/Reward Quotient in the circumstances and then associate by positive correlation this sense directly with an Epistemic Standard. For example, if one has the sense that the Risk/Reward Quotient is high, then one's Epistemic Standard may be high, say requiring that evidence like the Present Evidence *almost always* correctly predicts propositions like the questioned proposition; and if one has the sense that the Risk/Reward Quotient is low, then one's Epistemic Standard may be low, say requiring that evidence like the Present Evidence *very often* correctly predicts propositions like the questioned proposition.

In summary, I have outlined a general method by which Advanced Propositions may be adjudicated. I will refer to it as the 'Basic Method'.

- (1) We notice and keep track of the frequency of the appearance or occurrence of different general kinds of objects, relations, states or events. These frequencies are the basis for the sense we get of the Prior $\Pr(p=\text{TRUE})$ in particular instances. We may then get a sense of the Risk/Reward Quotient by modulating our sense of the cost of a false positive and the benefit of a true positive accordingly, and weigh the numerator of the Risk/Reward Quotient against its denominator.
- (2) We associate by positive correlation the Risk/Reward Quotient with an Epistemic Standard.
- (3) We notice and keep track of the frequency with which different general kinds of objects, relations, states or events correctly predict other different general kinds of objects, relations, states or events. These frequencies are the basis for the sense we get of the strength of the Present Evidence in particular instances.
- (4) We produce Knowledge Intuitions in relation to the questioned proposition if and only if the assessed strength of the Present Evidence meets our Epistemic Standard.

As I suggested, the sense we have of the frequencies that inform our sense of the Prior $\Pr(p=\text{TRUE})$ and inform our sense of the strength of the Present Evidence in particular instances are largely (if not entirely) learned. This state of affairs explains the practice of criticizing incorrect productions of Knowledge Intuitions that fail to take experienced events into account; and the practice of forgiving incorrect productions of Knowledge Intuitions that are caused by a lack of experience. For example, if Julie betrayed Florence and Florence has previously known Julie to betray others, then Harry may be warranted in blaming Florence for trusting Julie. On the other hand, if Julie betrayed Sally and Sally had never known Julie to betray others, then Harry would not be warranted in blaming Sally for trusting Julie.

5.2 The Adding Evidence Thesis

F1 predicts that under certain circumstances the Basic Method may be suspended. That is, it predicts that under certain circumstances the production of Knowledge Intuitions may be strategically barred or suspended—even if the Present Evidence meets the Epistemic Standard that is predicted by the Basic Method. Such a state of affairs may arise when we are presented with the opportunity to benefit by reducing the *risk* of a false positive—which we may do by gathering additional evidence, and thereby strengthening Present Evidence. However, for such a suspension to occur, the anticipated benefit of reducing false positive risk must outweigh the anticipated cost of gathering the additional evidence.

Recall F1:

F1: In a situation S , the function of intuiting that one knows that a proposition p is true is to indicate to oneself that, on the total evidence, behaving as if p in S will tend to produce better consequences than that of not behaving as if p in S .

F1 predicts that *if* behaving as if p (which entails not gathering additional evidence) would *not* tend to produce better consequences than not behaving as if p (which may include gathering additional evidence), *then* Knowledge Intuitions would not be produced. Moreover, F1 is fundamental to my theory; whereas SDT (and, by extension, the Basic Method) have the supporting role of suggesting operating principles of the KIP that would allow it to perform the functions described by F1 (and F2). However, SDT simply does not apply to the cases I outlined in which Knowledge Intuitions are suspended.

These special cases would arise only in virtue of our ability (a) to delay the production of Knowledge Intuitions, and delay performing an action or forming an attitude as if p (which would otherwise follow from the Basic Method), (b) to assess the value of reducing false positive risk by gathering additional evidence, (c) to assess the cost we are likely to incur in the process of gathering additional evidence, and (d) to determine whether the assessment in (b) outweighs the assessment in (c). SDT is not intended to apply to systems or organisms when they exercise such a set of abilities. The theory assumes a more static situation in which a system or an organism has certain fixed evidence, and must determine in the context whether that evidence is sufficient for it to behave as if the questioned world condition exists.

Consider an example based on Stanley's Low Stakes Case.

Case 5-1

It is Friday, and Hannah has just got off work with her paycheque in hand. As she starts her drive home, she reflects on a bill coming due. While she has enough funds in her bank account to cover it, the account's balance would drop below the minimum monthly balance and she would incur a small bank service charge. Even though the charge is small, she would like to avoid it. So, she plans to stop at the bank on the way home to deposit her paycheque.

Sometime later she arrives at the bank, and as she drives past it, she notices that the queues are very long. She then recalls that the bank was open two weeks before on a Saturday and decides to return to the bank tomorrow to make the deposit.

Suddenly, Hannah realizes that the bank app on her smartphone would indicate the bank's operating hours instantly. She reminds herself that banks do change their hours. She checks her app and it indicates that the bank is indeed open on Saturdays. Hannah feels fully reassured. She continues on her way home and plans to go to the bank on Saturday.

Hannah addresses the proposition p that the bank will be open on Saturday. Initially, her Present Evidence is limited to the bank having been open on a Saturday two weeks before; and, initially, that Evidence meets her Epistemic Standard. That initial standard is not particularly high because the Risk/Reward Quotient in her situation is low: If Hannah fails to deposit her cheque over the weekend, then she only faces having to pay a small bank charge.

However, even though the cost of a false positive for Hannah is low, she would nonetheless benefit by reducing her risk of a false positive; and she realizes that she has the opportunity to do so by checking her bank app. Moreover, Hannah realizes that the cost of checking her app is negligible. So, Hannah would plausibly find that the benefit of reducing false positive risk outweighs this cost. Accordingly, Hannah finds that it is worthwhile to check the app. She is happy to be reassured about the Saturday hours.

Generally, in these cases, subjects rightly assume that it is likely (but not certain) that the additional evidence they would gather would *strengthen or confirm* the evidence they have in hand—because the latter is strong enough to meet the Epistemic Standard that they associate with the Risk/Reward Quotient. For example, the evidence Hannah has prior to checking her bank app, i.e., the bank having been open two weeks before on a Saturday, gives her good reason to believe that the bank app will confirm p , that the bank will be open on Saturday. Even so, in these cases, subjects find that the benefit of reducing false positive risk outweighs the cost of gathering the additional evidence.

To be clear, the cost of gathering additional evidence in these cases may include multiple components. Particularly, it may include the subject's valuation of the energy and resources required to gather the additional evidence, and her valuation of lost opportunities to perform other actions while gathering the evidence. Such lost opportunities may include small things, such as having the opportunity to maintain the flow of a conversation (e.g., the cost of interrupting a flowing conversation to confirm a small point).

After Hannah realizes that it is worthwhile to gather additional evidence, she reminds herself that banks do change their hours. This reflects a suspension of her Knowledge Intuition about the bank being open, which she had up until that moment. It reflects the occurrence of a new doubt.⁴⁵

⁴⁵ It may be that Antonio Damasio's 'somatic-marker hypothesis' describes the process by which Knowledge Intuitions may thus be suspended. A somatic-marker is a sensation, like a gut-feeling, that accompanies anticipated consequences of actions. Damasio explains:

[a bad gut-feeling] forces attention on the negative outcome to which a given action may lead, and functions as an automated alarm signal which says: Beware of danger ahead if you choose the option which leads to this outcome. The signal may lead you to reject, immediately, the negative course of action and thus make you choose among other alternatives. (1994, p.173)

Perhaps, when one is struck by the sense that one would do well to *not* behave as if p and to gather additional evidence in relation to p —even if one's Present Evidence meets the Epistemic Standard associated with the Risk/Reward Quotient in the circumstances—the moment is accompanied by a somatic-marker that causes the suspension of the Knowledge Intuition.

Finally, after Hannah checks her app (gathers the additional evidence) and finds that it confirms p , her *total* Present Evidence more than meets her Epistemic Standard (that continues to reflect the unchanged Risk/Reward Quotient). She is fully reassured about the bank being open and her Knowledge Intuition that the bank will be open re-emerges.

The scenario I describe suggests an additional contextual factor in the production of Knowledge Intuitions.

Adding Evidence Thesis: One's Knowledge Intuitions are suspended in the moment (even if one's Present Evidence meets the Epistemic Standard predicted by the Basic Method) if one believes that (a) reducing false positive risk by gathering additional evidence is a practical possibility, and (b) the benefit of this reduction in false positive risk exceeds the cost of gathering the additional evidence.

When Knowledge Intuitions are suspended under the Adding Evidence Thesis, one's Epistemic Standard is effectively raised in the period during which additional evidence is gathered. So, the Adding Evidence Thesis is a *truth-conducive* contextual factor in the production of Knowledge Intuitions.

Notice too that the Adding Evidence Thesis suggests that Knowledge Intuitions may be suspended anytime the practical possibility of gathering additional evidence is identified; and also that identifying practical possibilities to gather evidence depends entirely on one's capacities, technologies, resources and strategies for doing so. In other words, one's ability to extract evidence from the world determines whether the Adding Evidence Thesis is relevant in particular cases. In Case 5-1, Hannah's ability to gather additional evidence depended on her having a smartphone with a bank app installed on it in her possession, and depended on her ability to use this technology.

The Adding Evidence Thesis resonates with Patrick Rysiew's suggestion that the function of knowledge attributions is to mark 'the opening and closing of specific lines of inquiry' (Rysiew, 2012, p.275). (Rysiew acknowledges that his approach is similar to that taken by Chris Kelp (Kelp, 2011) and Klemens Kappel (Kappel, 2010).) My thesis suggests that the production of Knowledge Intuitions and the use of 'know' depends on whether practical possibilities to gather

evidence have been identified; and so they depend on whether lines of inquiry are open or closed. Accordingly, they mark the openings and closings Rysiew refers to.

5.3 Basic Assumptions

We very often take ourselves as knowing to be true propositions that are not, in the moment, associated with a particular practical situation, nor a particular range of possible practical situations, nor any particular action or attitude, nor a particular range of possible actions or attitudes. For example, at this moment, the following two propositions meet such a description for me: Paris is the capital of France, and Albert Einstein proposed that $E=mc^2$. If the Basic Method is true, then the Epistemic Standards that we set under such circumstances would track our assessment of the Risk/Reward Quotient in the moment. But how are we to assess the Risk/Reward Quotient in relation to such propositions given that they are not directly associated with practical matters? I propose an answer to this question in this section.

It seems plausible that most (if not all) of the adjudications of propositions that we make—even those that are not directly associated with practical matters—serve a purpose. Further, it seems that very often the purpose of adjudicating propositions that are not directly associated with practical matters relates to the *potential* support the adjudications provide for our adjudications of other propositions, ones which *are associated* with practical matters. I propose that the situation is a little like that which Otto Neurath referred to in his famous boat metaphor.

Neurath likened the enterprise of building a body of knowledge using the scientific method to the task of sailors rebuilding their boat while at sea: They are not able to take the boat into dry dock, and must stand on one plank while repairing another (Neurath, 1932/1983, p.92). Quine suggested that the metaphor applies to philosophy:

The philosopher and the scientist are in the same boat. If we improve our understanding of ordinary talk of physical things, it will . . . be by clarifying the connections, causal or otherwise, between ordinary talk of physical things and various further matters which in turn we grasp with help of ordinary talk of physical things.
(Quine, 1960, p.3)

I am suggesting that the metaphor may be extended further to the adjudication of ordinary Advanced Propositions. Consider *first* that, for the agent, the ability of Present Evidence to predict a proposition generally depends on the agent taking herself as knowing certain ancillary propositions as being true. For example, say that I am considering driving to my friend's place to pay him a surprise visit. However, I want to first check to see if the storm has let up. So, I look out the window and see that the streets are very icy. On the basis of this Present Evidence, I take myself as knowing the proposition *p*: *it is dangerous to drive at this time*. On that basis, I decide to stay home. The Present Evidence (i.e., the streets being very icy) predicts *p* (i.e., that it is dangerous to drive at this time). However, my ability to make that prediction on the basis of that Present Evidence depends entirely on my taking myself as knowing *that, in general, ice is slippery*. I depend on a *general ancillary proposition* which we may view as a Neurathian plank I stand on when I adjudicate *p* on the basis of my Present Evidence.

Consider also that the Basic Method indicates that our adjudications of Advanced Propositions are based on intuitions of knowing (a) the frequency of the appearance or occurrence of different general kinds of objects, relations, states or events, and (b) the frequency with which different general kinds of objects, relations, states or events, predict other different general kinds of objects, relations, states or events. These propositions about frequencies are again general ancillary propositions that facilitate our being able to adjudicate the main questioned proposition. These too may be viewed as Neurathian planks. Consider some examples.

1. In Stanley's Low Stakes Case, Hannah's Knowledge Intuition about the bank being open on Saturday depends on the Epistemic Standard she sets, which (according to the Basic Method) in turn depends on her sense of the Prior Probability of the bank being open. In this example, Hannah's sense of this Prior Probability in turn depends on assumptions such as that banks in her area are usually open on Saturdays, and that disruptions of regular bank hours (due to power failures, currency shortages, labour disruptions or robberies) are extremely rare.

2. A doctor's Knowledge Intuition that his patient has type 2 diabetes depends on the Epistemic Standard he sets, which in turn depends on his sense of the Prior Probability of the patient having that ailment. This Prior Probability depends on an assumption that he makes about the frequency with which type 2 diabetes occurs in that segment of the population that has the patient's medical profile. For example, if the occurrence of the ailment is extremely rare, then the Epistemic Standard that the doctor would set would be extremely high.

3. Say that you are a pedestrian at a busy street corner waiting for an opportunity to cross safely. A vehicle stops opposite you; and you and the driver make eye contact. She clearly waves you on, suggesting that she will wait while you cross; and you intuit knowing the proposition p that *it is safe to cross the street under these circumstances*. Your Knowledge Intuition depends on the Epistemic Standard you set for p ; which, in turn, depends on your sense of the Prior Probability that p is true. This sense of the Prior Probability depends on such assumptions as (a) that the frequency with which drivers attempt to run over pedestrians is very low, (b) that the frequency with which drivers lose control of their vehicles due to medical conditions is very low, and (c) that the frequency with which cars lurch forward on their own due to a mechanical malfunction is very low. Your taking yourself as knowing to be true these ancillary propositions leads you to set a low Epistemic Standard in relation to p .

Let us assume that these sorts of ancillary propositions are derived from more general propositions. For example, the proposition, the frequency with which drivers attempt to run over pedestrians is very low, derives from the more general proposition, the frequency with which people in my area attempt to cause serious harm to others who mean them no harm is very low. The proposition, the frequency with which cars lurch forward on their own due to a mechanical malfunction is very low, derives from the more general proposition, the frequency with which cars have major mechanical malfunctions is very low.

I suggest that when we use such *general ancillary propositions* to facilitate our adjudications, we ordinarily believe without doubt in the moment that they are true: We stand on those planks that we feel safe standing on. I call such general ancillary propositions that are Advanced Propositions and that in the moment we take ourselves as knowing 'Basic Assumptions'.

I assume that we have previously positively adjudicated the Advanced Propositions that are the content of Basic Assumptions—whether we remember doing so or not. Further, I assume that when we make such an adjudication, we anticipate that it would serve a purpose—particularly the purpose of supporting future adjudications of other Advanced Propositions). I call propositions that constitute candidate content for a new Basic Assumption, 'Basic Propositions'. For example, I assume that at some point I adjudicated the Basic Proposition that the frequency with which people in my area attempt to cause serious harm to others who mean them no harm is very low (even though I do not remember doing so), and further that I knew that if I took myself as knowing this Basic Proposition to be true, then I would in effect be accepting a new Basic Assumption that would have the potential of being used in a range of future adjudications.

If we take ourselves as knowing that a Basic Proposition is true, that proposition will be the basis for a new Basic Assumption that may be used over and over again in future adjudications. So, acquiring a new Basic Assumption has the potential of being highly consequential; and some new Basic Assumptions may be extremely consequential, e.g., Basic Assumptions about the virtue and power of the scientific method and about the existence of a personal God.

If the Basic Method applies to our adjudications of Basic Propositions, then the Epistemic Standard that we set should correlate with the Risk/Reward Quotient in the situation. However, one may well not be able to *anticipate* the practical circumstances surrounding many, if not most, of the future adjudications in which a new Basic Assumption will play a role. So, how would one get a sense of the Risk/Reward Quotient under these circumstances?

It seems that we do get some reasonably clear sense of the general risk and the general reward that could go with adopting a new Basic Assumption. We may get a sense of this by first recognizing that the set of all of an individual's Basic Assumptions constitutes a basis for her ongoing adjudications of Advanced Propositions, and so a basis for the performance of deliberate actions and formation of attitudes. I will call an individual's set of Basic Assumptions her 'Worldview'.

Further, we might recognize that some Worldviews are better than others—insofar as they better facilitate successful ongoing adjudications of ordinary Advanced Propositions, and better facilitate adjudications of Advanced Propositions that are difficult to correctly adjudicate (say because they are highly abstract, or are not well supported by Present Evidence). In general, better Worldviews are better at helping us to advance our interests.

In this picture, it appears that the principal *risk* of incorporating a new Basic Assumption into one's Worldview is that it has a deleterious effect on that Worldview, i.e., it combines with already held views in a way that makes us worse adjudicators of new Advanced Propositions; and it appears that the principal *reward* of incorporating a new Basic Assumption into one's Worldview is that it has a helpful effect—by combining synergistically with already held views in a way that makes us better adjudicators of new Advanced Propositions.

Further, we may predict the extent to which a new Basic Assumption will be either deleterious or helpful on the basis of whether or not certain relations would obtain between that Basic Assumption and our already-held views. For example, Basic Assumptions that are *compatible*

with already held views are more likely to be synergistic and not deleterious than are incompatible Basic Assumptions—particularly when assessed against already held views that have in the past played an important role for us in making important and successful adjudications.

Second, Basic Assumptions that are *coherent* with already held views are more likely to be synergistic and not deleterious than are Basic Assumptions that do not cohere with already held views. The notion of coherence suggested by Laurence Bonjour seems to work here: Coherence between beliefs is afforded by ‘various sorts of inferential, evidential and explanatory relations’ (Bonjour, 1985, p.93). Again, this is particularly so when assessed against already held views that have played an important role in making important and successful adjudications.

Third, Basic Assumptions that are *complementary* with already held views are more likely to be synergistic and not deleterious than are Basic Assumptions that are not *complementary*. A Basic Assumption is complementary to the extent that it extends or fills in gaps in one’s Worldview—while not diminishing its coherence. Doing so may allow us to broaden the range of the Advanced Propositions that we are successfully able to adjudicate.

So, we may get a sense of the Risk/Reward Quotient when we are adjudicating Basic Propositions by weighing the three factors to which I referred—the compatibility, coherence, and complementarity of a Basic Proposition in relation to our already held views. If a Basic Proposition is compatible, coherent and complementary with them, then allowing it to be the content of a new Basic Assumption would likely not be risky, and may very well be helpful. So, the Risk/Reward Quotient in relation to that Basic Proposition would be low—and one may associate with it a low Epistemic Standard. Generally, the Risk/Reward Quotient is negatively correlated with the extent to which a Basic Proposition is compatible, coherent and complementary with already held views.

Experimental social psychology supports the picture that I have sketched. I refer to the research that deals with confirmation bias (which I discuss in §4.3.2.2). Recall Fiske & Taylor’s summary of the picture that emerges from this research:

Instead of a naïve scientist entering the environment in search of the truth, we find the rather unflattering picture of a charlatan trying to make the data come out in a manner most advantageous to his or her already held theories. (Fiske and Taylor, 1984, p.88)

However, even if confirmation bias is strong and our Epistemic Standards in relation to propositions that are compatible and coherent with already held views are low, my theory suggests that our Worldview may well develop in way that tends to make us better adjudicators of Advanced Propositions. It may well be that not every new Basic Assumption that is added to that Worldview is perfectly compatible and coherent, and may well be that we put a high priority on adding highly complementary Basic Assumptions, that increases the range of Advanced Propositions that we can successfully adjudicate. The situation would be like that which Neurath describes for the sciences in his final version (there were several) of his boat analogy:

Imagine sailors, who, far out at sea, transform the shape of their clumsy vessel from a more circular to a more fishlike one [which would presumably traverse the waters more smoothly and efficiently]. They make use of some drifting timber, besides the timber of the old structure, to modify the skeleton and the hull of their vessel. But they cannot put the ship in dock in order to start from scratch . . . A new ship grows out of the old one, step by step . . . The whole business will go on in a way that we cannot even anticipate today. That is our fate. (1944, p.47)

However, the prospect of ever-improving Worldviews is diminished to the extent that we assess the compatibility, coherence and complementarity of a Basic Proposition (for possible inclusion into our Worldview as a new Basic Assumption) in relation to only a small subset of Basic Assumptions. If so, we would tend to adopt new Basic Assumptions that are coherent with some already held views but not others. Such a process would lead to the creation of islands of coherent Basic Assumptions, so that the Basic Assumptions that constitute one island are not generally coherent with those of others. For reasons spelled out by coherentists, including the early work of C. I. Lewis (see §2.2), a failure to establish broad coherence increases the likelihood that individual islands of coherent Basic Assumptions contain significant falsehoods. So, the worry here is that while the process by which Worldviews develop generally serves our interests, those Worldviews may nonetheless contain caricature-like, narrowly coherent, but broadly incoherent elements.

5.4 Optimality

If indeed the Epistemic Standards that we seem to set are determined by the Basic Method under the Adding Evidence Thesis, are those standards optimal? That is, do they tend to produce consequences that are the best that we are capable of achieving, given our circumstances, limitations, and so on? Such a question is prompted because (1) SDT was developed largely for the purpose of determining optimal detection thresholds for systems and organisms, and (2) SDT is used by adaptationists to explain selection pressure for detection thresholds (which are exhibited in the response patterns of the organisms they investigate).⁴⁶ We find that there is reason to assume that our Epistemic Standards may very well not be optimal.

Recall my earlier distinction between mechanisms that govern autonomic and stimulus-response behaviours, and the mechanism that governs the performance of calculated deliberate actions and the formation of attitudes (in §4.3.1):

Mechanisms that govern autonomic and stimulus-response behaviours. These depend on innate detection thresholds that are more or less fixed for individuals. For example, visual data must be strong enough to meet an innate detection threshold in order for me to reflexively respond to an object that is fast approaching my head, e.g., by ducking or blinking. Call these ‘Type A’ mechanisms.

A mechanism that governs the performance of calculated deliberate actions and the formation of attitudes. This mechanism depends on the KIP which sets Epistemic Standards on the basis of the world conditions we encounter. Call this a ‘Type B’ mechanism.

Let us assume that much of the evolutionary history of Type A mechanisms predates much of the evolutionary history of the Type B mechanism. Such an assumption is plausible because it would seem that very many multi-cell organisms have Type A mechanisms. They are capable of responding to stimuli and employ an innate detection threshold. On the other hand, only few, relatively sophisticated, species possess a Type B mechanism. So, we may view the former as early stage systems, and that latter as a late stage system.

⁴⁶ The working hypothesis in these investigations is that these detection thresholds evolve to optimality. Earlier I offered these examples: (1) detection threshold exhibited by female birds for male mating signals (Price, 2013), (2) the detection threshold exhibited by bees for floral displays (Leonard et al., 2011) and (3) the detection threshold exhibited by the plant *Hieracium umbellatum* for soil types. The latter grows bushy, broad leaves on rocky, sea-side cliffs, and grows prostrate with narrow leaves on sand dunes.

Late stage adaptations sometimes evolve from earlier adaptations that had different functions. For example, the lungs in land animals evolved from the swim bladder in fish—a gas-filled sack used to control buoyancy. Andy Clark explains that ‘if evolution proceeds as a tinkerer, then each step in the evolutionary chain exploits a ‘net historical opportunity’ whose nature is determined by whatever materials happen to be available to adapt to a new requirement’ (1987, p.283). It may well be that some of the cognitive components of the Type B mechanism also evolved from earlier adaptations that had different functions. For example, the Basic Method suggests that we possess a device that keeps track of the frequency of the appearance or occurrence of different general kinds of objects, relations, states or events. Such a frequency of appearance or occurrence allows us to get a sense of the Prior $Pr(p=TRUE)$, which helps us assess the Risk/Reward Quotient and set Epistemic Standards in the circumstances. So, in the late-stage system, the frequency-tracking device’s function is to contribute one piece of contextually relevant information which we use in combination with several other pieces of information in order to set an Epistemic Standard. Calculated and deliberate action may or may not follow. However, this frequency-tracking device may have evolved from a much earlier device that had a different function. For example, it may have evolved from a device that facilitated automatic conditioned responses—like the device that mediates the Pavlovian responses exhibited today in many infra-human species, e.g., the device mediating Pavlov’s dog’s salivating when hearing a bell. The organism acquires a particular response to an unconditioned stimulus when the frequency of co-occurrence of that unconditioned stimuli with a conditioned stimulus meets a threshold. Such a function is very different than the function of the more evolved, modern frequency-tracking device used in the Type B mechanism.

Andy Clark calls adaptations that have this sort of evolutionary history, i.e., that involve gerrymandered design elements, ‘kludges’. The performance of kludges may well not be optimal. For example, when the swim bladder becomes the basis for a breathing device,

the initial elegance is lost. The new device, though functional, may be expected to embody design principles appropriate to swim bladders and in all likelihood inappropriate, from a pure design perspective, to lungs. At best, then, we may hope for elegant solutions to the earliest and most telling cognitive problems which faced our ancestors. In all other cases - including, sadly, most of those studied by contemporary cognitive science - we may expect to find cognitive kludges. (Clark, 1987, p.290)

This suggests that the Type B mechanism may not be optimal.

However, the Basic Method allows that our adjudications of Advanced Propositions may continuously improve. First, the Basic Method allows the possibility of improving our sense of the frequency of the appearance or occurrence of different general kinds of objects, relations, states or events. This, in turn, allows us to get better at predicting the value of Prior $Pr(p=TRUE)$, and so, better at assessing the Risk/Reward Quotient and setting advantageous Epistemic Standards. Second, our assessments of the strength of our Present Evidence may continuously improve: The Basic Method allows the possibility of improving our sense of the frequency with which different general kinds of objects, relations, states or events correctly predict other general kinds of objects, relations, states or events. Third, the Adding Evidence Thesis allows the possibility of improving our sense of the cost of gathering additional evidence. This in turn allows us to get better at determining whether we would do well to suspend Knowledge Intuitions and gather additional evidence.⁴⁷

⁴⁷ These conclusions are consistent with a suggestion Carolyn Price makes about the function of the systems that control learning. Roughly, that function is to improve our reasoning in the direction of a certain ideal (Price, 2002, p.613).

Chapter 6: Altruistic Knowledge Intuitions

6.1 Introduction

We found that an agent adjudicating a proposition may tend to produce consequences that she values by modulating her Epistemic Standard in accordance with the Risk/Reward Quotient; and, we assumed that the Risk/Reward Quotient reflects risks and rewards *for the adjudicator* and not someone else. That is, we assumed that the adjudicator assesses the benefit of a true positive and the cost of a false positive in relation to her own practical interests, and not in relation to those of others.

However, sometimes we anticipate that our adjudications of propositions will directly affect the practical interests of others; and sometimes, in these circumstances, we are focused on this effect, and not on our own immediate practical situation. Moreover, sometimes, not only are we concerned that those others are not harmed in any way as a result of our adjudications and actions, but, we are hopeful that our adjudications and actions have positive consequences for those others.

It would seem that under these circumstances, in order to tend to produce positive consequences for those others, our Epistemic Standards should modulate in accordance with a Risk/Reward Quotient that reflects the risks and rewards for those others whose practical interests we hope to advance. That is, we should assess the benefit of a true positive and the cost of a false positive, not in relation to our own practical interests, but in relation to those of others; and this seems to be precisely what we do. Moreover, I find that this practice is key to explaining the multiple impasses that we find in the epistemological literature—particularly, between Contextualists, Subject Sensitive Invariantists and Classic Invariantists—and key to explaining why we have two independent histories in epistemology (see Chapter 2). So, let us investigate the following thesis:

The Altruistic Assessment Thesis (AAT): We are sometimes focused on helping others; and when we are, our adjudications of relevant Advanced Propositions are based on our assessments of the benefit of true positives and the cost of false positives in relation to the practical interests of those others.

Earlier, I introduced and defined the term ‘Altruistic Knowledge Intuitions’. That term may, alternatively, be viewed as the Knowledge Intuitions that arise in the situations to which the AAT refers. (In contrast, my defined term ‘Egoistic Knowledge Intuitions’ are the Knowledge Intuitions that follow from self-oriented assessments.)⁴⁸

*

My terminology suggests that I regard the cases to which I refer as cases of altruism. Let us be reminded of the point I made earlier (in §3.3) about the particular meaning of the term ‘altruism’ in the biological literature, which meaning I employ here. In *everyday* use, ‘altruism’ refers to behaviour that follows from an ultimate desire or motivation to help others, and not from a desire or motivation that is ultimately egoistic. However, in biology, ultimate desires and motivations are entirely irrelevant. The sole criterion for whether behaviour is altruistic is whether its direct (non-accidental) effect is to decrease the donor’s Inclusive Fitness and to increase the recipient’s. So, from a biological perspective, altruistic behaviour may well stem from egoistic ultimate desires or motivations, such as the desire to enhance one’s reputation or to create an indebtedness. As I suggested earlier, in order to keep my investigation consistent with the biological literature, I employ the biological meaning.

However, I find the application of the standard biological meaning problematic in one respect. Recall the suggestion I made earlier (in §4.3.2.3) that we do not assess the consequences of our behaviour directly on the basis of Inclusive Fitness, but rather on the basis of our values (i.e., whatever it is that is taken to be positive, good or beneficial). Moreover, it seems that in some cases, our values—which may, to a significant extent, be learned—do not seem to be compatible with Inclusive Fitness. This state of affairs potentially complicates the present

⁴⁸ I allow the possibility that we are sometimes *primarily* self-interested and *secondarily* other-interested, or *vice versa*. An analysis of such cases might proceed by focusing on the primary orientation. Further, it seems plausible that an individual may in some instances vacillate between orientations in relation to the same proposition and in the same general context. Certain self-ascriptions of knowledge may be anticipated to lead to both self-oriented actions and to other-oriented actions. However, I set aside consideration of such cases.

discussion. So, I will remedy the problem by making ‘altruism’ a defined term, and define it in relation to values. Let us say that an individual’s behaviour or thought (including an adjudication of an Advanced Proposition) is ‘Altruistic’ if and only if she values neutrally or negatively its direct consequences in relation to her own practical and self-oriented interests, and values positively its consequences in relation to the practical interests of those she intends or envisages helping.⁴⁹

There are, of course, many ways in which we may help others. However, most of the cases in the epistemological literature in which Altruism plays a role feature an agent who is focused on helping others *by means of communicating a proposition to those others*. I focus on just such cases.

Recall that I refer to these last agents as ‘Senders’ and those to whom their communications are directed as ‘Receivers’. Further, *in the cases on which I am focused*, Senders assume that their Receivers are very likely to take their communicated propositions as being true. As a result, Senders assume that Receivers are very likely to behave as if their communicated propositions are true and likely to benefit if they are true and be harmed if they are false. In these cases, if the Altruistic Assessment Thesis is true, Senders’ adjudications of their communicated propositions are based on their assessments of these Receivers’ benefits and harms.

Earlier (in §3.3), I proposed a primary function for Altruistic Knowledge Intuitions that arise in these communication cases, F2—which is an applied special instance of my primary function hypothesis F1. F2 suggests that the function of a Sender’s Altruistic Knowledge Intuition in relation to p is to indicate to herself that communicating that p will tend to help her Receiver(s) more than would not communicating that p .^{50 51}

⁴⁹ Notice that I have *not* appealed to the values of those who we intend or envisage helping. This is relevant when there is a marked difference between the two. For example, I have suggested that if I intend to help or envisage helping a masochist, I will assess the consequences of my behaviour, not on the basis of the masochist’s values, but rather on the basis of my own. This assumption simplifies the theory and the analysis; and I do not have an opinion as to how much data (if any) this move would potentially leave unexplained.

⁵⁰ Knowledge cases that involve communications do not necessarily involve Altruistic Knowledge Intuitions. An individual X who has an Egoistic Knowledge Intuition in relation to p may communicate p to an agent Y in order to either get Y to perform an action or form an attitude that X believes would advance her own interests, regardless of whether or not it is good for Y . Or, an individual who performs an action on the basis of an Egoistic Knowledge Intuition may merely report on having performed that action after the fact. I may say to myself ‘It is 5 pm. So, I’ll go home now’ and then later explain ‘I knew it was 5 pm. So, I went home’.

⁵¹ F2 suggests that if (a) Altruistic Knowledge Intuitions arise in these cases, then (b) we have the sense that communicating that p will tend to help our Receiver(s) more than would not communicating that p . However, it does *not* suggest that if (b), then (a). F2 allows, for example, that circumstances may arise

6.2 Human Altruism

Let us consider whether the AAT is plausible. On one hand, it is obvious that we sometimes focus on helping others, obvious that we are often successful in helping others, and obvious that we sometimes take the practical interests of others into account. On the other, it is not evident that our taking the practical interests of others into account extends to the unconscious operations of the KIP. That is, it is not at all clear that the production of Knowledge Intuitions sometimes depends on the agent's reading of the practical interests of others, as the AAT suggests. Even so, the thesis is supported by a range of ordinary knowledge cases with which it is consistent; and supported also by the biological literature. Let us first look at the latter.

Biological research, which I review in Appendix 3, suggests that we possess Altruistic traits. It suggests that we are innately Altruistic (under the assumption that individuals' values very broadly tend to support their respective Inclusive Fitness). Further, it suggests that not only do we attempt to help others, but *we actually often succeed in doing so*, and that benefits for others are often produced. In fact, the evolution of Altruistic traits depends on producing those benefits. But, how does our behaviour tend to produce positive consequences for others? As I suggested earlier, Signal Detection Theory offers an answer to the question, just as it did in relation to the production of positive consequences for ourselves, viz., when we intend to help or envisage helping others and we are adjudicating propositions that are anticipated to lead to behaviour that would affect those others, we assess the benefit of true positives and the cost of false positives in relation to the practical interests of those others.

Further, the biological research predicts the particular circumstances under which *ceteris paribus* Altruistic behaviour will arise—and predicts also that Altruistic behaviour is generally not likely to arise outside of them. It is thus plausible that the Altruistic Assessment Thesis will

under which a Sender has the sense that communicating what she takes to be a falsehood will tend to help her Receiver(s) more than would not communicating the presumed falsehood. I set aside cases of the latter kind. As I indicated earlier (in §3.5), in this investigation, I set aside cases in which the knowledge attributor is intentionally ironic or dishonest—even if they are doing so for the purpose of helping their Receivers.

ceteris paribus be applicable under these particular circumstances, though it may not be applicable outside of them.⁵² These are those circumstances, roughly described:

- (1) The individual intends or envisages communicating with closely-related kin—or a more distantly related individual, if the cost of preparing and sending a helpful communication is a small fraction of the latter individual’s expected benefit (see §A.3.1 in Appendix 3).
- (2) The individual intends or envisages communicating with someone with whom she has an ongoing relationship that involves reciprocal interactions, e.g., one is an employee, and the other her employer (assuming that the employee’s work in fact benefits the employer and the employer regularly reciprocates by paying the employee’s wages) (see §A.3.2).
- (3) The individual with whom one is intending or envisaging communicating is a ‘co-operator’ (i.e., someone who has a good record of reciprocating and helping others), and others in the community will learn about the communication producing reputational rewards for the altruist (see §A.3.3).
- (4) The individual has a co-operative character or personality, and it is strong enough to compel her to provide the level of co-operation called for in the circumstances (see §A.3.4).

Let us look at how Senders employ the Basic Method under the AAT. They first get a sense of the benefit for their Receivers of a true positive and the cost for their Receivers of a false positive. This enables them to get a sense of the Risk/Reward Quotient. They then associate the Risk/Reward Quotient with a particular Epistemic Standard, e.g., a high Risk/Reward Quotient is associated with a high Epistemic Standard. (I look at how Senders assess whether the conditions contained in the Adding Evidence Thesis have been met under the AAT in §6.3.) A low stakes example and high stakes example follow.⁵³

Case 6-1: Low Stakes

⁵² It may be that the Altruistic Assessment Thesis would not be relevant even under one of the described circumstances, if, in the moment, contrary egoistic or other motivations outweigh the individual’s Altruistic motivations.

⁵³ These are based on Stanley’s Low Stakes Case and Stanley’s High Stakes Case.

Sarah is driving her sister Hannah home on a Friday afternoon. Sarah asks Hannah if she would mind if they stopped at the bank on the way so that she can deposit her paycheque. Sarah explains that while it is not important that she do so as she has no impending bills, it would nonetheless be efficient. Coincidentally, Hannah also uses the same bank and knows that they will not have to go too far out of their way to get there. She tells Sarah that she does not mind stopping.

As they drive past the bank, they notice that the lines inside are very long. Hannah says, 'I know the bank will be open tomorrow, since I was there just two weeks ago on Saturday. You might prefer coming back and depositing your paycheque then'.

Case 6-2: High Stakes

Sarah is driving her sister Hannah home on a Friday afternoon. Sarah asks Hannah if she would mind if they stopped at the bank on the way so that she can deposit her paycheque. Sarah explains that she has a very important impending bill coming due, and not enough in her account to cover it. So, she has to make sure that she deposits her paycheque by Saturday. Coincidentally, Hannah also uses the same bank and knows that they will not have to go too far out of their way to get there. She tells Sarah that she does not mind stopping.

As they drive past the bank, they notice that the lines inside are very long. Hannah says, 'The bank was open just two weeks ago on Saturday. Mind you, it is very important that you deposit your cheque by Saturday, and banks do change their hours. I can't say I know that the bank will be open tomorrow. It'd be a good idea to deposit your cheque now'.

In both cases, Hannah is a Sender to her Receiver, her sister Sarah; in both cases, Hannah adjudicates the proposition p that the bank will be open on Saturday; and, in both cases, her Present Evidence consists of her memory of the bank being open two weeks before on a Saturday. Still, she plausibly takes herself as knowing that p , in the first case, and plausibly does not take herself as knowing in the second. Yet we know nothing about Hannah's practical interests. This seems to show that Hannah's Epistemic Standard in relation to p is determined by her sense of the cost of a false positive for Sarah, which is much higher in the second case than it is in the first.

Case 6-1 features a first-person knowledge attribution, and Case 6-2, a first-person knowledge denial. Even so, the Basic Method under the AAT may be employed in cases that feature second-person or third-person knowledge attributions or denials. Here are examples that are

modifications of Case 2-3 and Case 2-4—which cases I used in §2.3 to support Subject Sensitive Invariantism.

Case 6-3: Low Stakes

Frank knows that his brother Herman wants to get a sense of the American political scene just so that he can better understand and enjoy political newscasts. Herman tells Frank that he has learned from a newspaper article that Clinton is a Democrat. Frank replies, ‘Good, you know that Clinton is a Democrat’.

Case 6-4: High Stakes

Frank’s brother Herman is a Democratic Party strategist. Frank knows that Herman is worried that Hillary Clinton quietly switched party affiliations, from the Democratic Party to the Republican Party. Frank knows that this would be very bad news for Herman. Herman tells Frank that he has learned from a newspaper article that Clinton is a Democrat. Frank replies, ‘No, you do not know that Clinton is a Democrat on the basis of that article’.

Frank plausibly intuitively knows that his brother Herman knows in the low stakes case, and plausibly does not intuitively know that Herman knows in the high stakes case. Yet we know nothing about Frank’s practical interests. This seems to show that Frank’s Epistemic Standard is determined by his sense of the cost of a false positive for Herman, which is higher in the high stakes case than it is in the low stakes case.

6.2.1 An Aside on Subject Sensitive Invariantism and Contextualism

Case 6-3 and Case 6-4 support Subject Sensitive Invariantism (the view that the truth conditions for knowledge attributions depend on the practical interests of the subjects of those attributions). Frank’s knowledge attribution in the low stakes case, and his knowledge denial in the high stakes case, depend on the subject’s practical interests. Clearly, a wide range of cases that feature (a) second-person or third-person attributions or denials, and (b) the employment of the Basic Method under the AAT—would be consistent with SSI. These are, in particular, cases in which the attributor assesses the Risk/Reward Quotient in relation to the practical interests of the subject of her second-person or third-person knowledge attribution or denial. On the other hand, these cases are clearly not consistent with Contextualism (the view that the truth conditions for knowledge attributions depend on the attributor’s practical interests).

As I suggested earlier (in §2.3), cases that feature first-person attributions or denials under egoistic assessments of the practical circumstances support both SSI and Contextualism. As an egoistic assessment, the attributor's Epistemic Standard is determined by her own practical interests. However, the attributor and the subject of those attributions or denials are one and the same. So, it is also true that Epistemic Standards in these cases are determined by the subject's practical interests; and cases that feature first-person attributions or denials and the employment of the Basic Method under egoistic assessments (i.e., not under AAT) support both SSI and Contextualism.

Also, I suggested earlier (again in §2.3) that cases that feature second-person or third-person attributions or denials under egoistic assessments of the practical circumstances support Contextualism. The attributor's egoistic assessments make her adjudications dependent on her own practical interests. So, cases that feature second-person or third-person attributions or denials and the employment of the Basic Method under egoist assessments support Contextualism.

This situation suggests that both Subject Sensitive Invariantism and Contextualism are correct within the boundaries that I have suggested here; and that they are complementary in the sense that they make correct predictions in relation to different cases. However, they both overgeneralize by suggesting universal applicability of a theory that in fact applies to only a subset of cases.

6.2.2 An Aside on the Receiver's Point of View

I have examined the production of Knowledge Intuitions and the use of 'know' in communications cases from the point of view of the Sender. These cases may also be investigated from the point of view of the Receiver. In fact, many epistemologists focus on the latter, the Receiver's point of view. Case 6-5 below exemplifies cases of the kind that these epistemologists are interested in.

Case 6-5

Jonah is in London to attend an important meeting. He is on his way to the meeting, which starts very soon. So, he is in a great hurry to find the nearest tube station. He remembers that his colleague, Hayley, lives in the neighbourhood in which

he is presently situated and uses the tube to get to work. Jonah says to himself, 'Hayley will know how to get to the nearest station. I'll call her'. Jonah calls Hayley, and she tells him that he should turn left on Wellesley Road. Jonah thanks Hayley and says, 'I knew you would know'.

Jonah has set an Epistemic Standard for (a) Hayley knowing where the nearest station is, and for (b) himself knowing where it is. This Epistemic Standard seems to be determined by his own practical interests. Typically, in these cases, the Receiver's Epistemic Standard seems to be determined by their own practical interests.⁵⁴

So, these cases feature the Receiver's Egoistic Knowledge Intuitions; and I addressed cases of this kind in the previous chapter. However, certain investigations by others into such communication cases from the Receiver's point of view suggest that a fundamental function of 'know' is revealed in just such cases. So, their findings compete with my F1 (and F2, which is a special case of F1). So, let us look briefly at some of this literature.

Edward Craig refers to Receivers as 'inquirers', and suggests that the function of 'know' follows from their need to flag reliable sources of information (1999).⁵⁵ So, in Case 6-5, Jonah is an inquirer who uses 'know' to flag Hayley as a reliable source of information. Jonah says, 'Hayley will know where the nearest station is'.

Steven Reynolds suggests that the function of 'know' follows from Receivers having an interest in rewarding Senders. They do so for the purpose of getting Senders to communicate messages that are more reliable:

the aim of improving testimony explains why we have and use our concept of knowledge. If we were to introduce a term of praise for testimony, using it at first to praise testimony that apparently helped us in our practical projects, it would come to be used as we now use the word 'know'. (2002, p.139)

⁵⁴ Even so, Receivers may, in some cases, intend or envisage *resending* the messages they receive for the purpose of helping others. Such Receivers adjudicate the messages they receive as a Sender to third-party Receivers. I look at cases of this kind in §6.4.4.

⁵⁵ See §3.5 for a general discussion of Craig's book.

John Greco has a somewhat similar view. He suggests that the function of 'know' follows from Receivers intending to 'give the person [i.e., the Sender] credit for getting things right [for coming to true belief] . . . due to his own abilities, efforts and actions' (2003, p.111).

Let us compare. Craig's suggestion about the function of 'know' relates to the process by which Receivers flag reliable sources of information. Reynolds' and Greco's suggestions about the function of 'know' seem to relate largely to ensuring that the reputations of good Senders are promoted, and that their good work is acknowledged and rewarded. I suggest that these hypotheses about the function of 'know', if they are correct, are secondary or ancillary to the fundamental function of Knowledge Intuitions that my theory postulates: An individual's Egoistic Knowledge Intuitions identify Advanced Propositions that will be a good basis upon which she may navigate the world; and an individual's Altruistic Knowledge Intuitions identify Advanced Propositions that will be a good basis upon which *others* may navigate the world. The process of navigating the world (or helping others to do so) sometimes involves Craig's suggested flagging of reliable sources of information. However, as we have seen, it often involves very much more, e.g., flagging propositions about the existence of non-perceptible world conditions as ones which may be taken as being true. Further, Reynolds' and Greco's suggested praise-giving or credit-giving may be viewed as helping Receivers establish the circumstances under which it is plausible that the Altruistic Assessment Thesis will *ceteris paribus* be applicable, and Senders will be inclined to help Receivers. The praise-giving or credit-giving help produce rewards for Altruism. (Recall circumstance (3) under which the AAT is plausibly applicable: The individual with whom one is intending or envisaging communicating is a 'co-operator' and others in the community will learn about the communication producing reputational rewards for the altruist.) So, while 'know' may sometimes be used to praise or give credit to Senders, this is ancillary to the larger function of identifying Advanced Propositions that will be helpful to Receivers.

6.3 The Adding Evidence Thesis Holds for Senders

I have said that Senders set Epistemic Standards using the Basic Method under the AAT, i.e., they assess the cost of false positives and the benefit of true positives in relation to their Receivers' practical interests. I now add that the Adding Evidence Thesis may apply to Senders. That is, Senders assess whether the conditions contained in the Adding Evidence Thesis have been met; and if they have, their Knowledge Intuitions may be suspended, even if their Present

Evidence meets the Epistemic Standard indicated by the Basic Method under the AAT. In particular, their Knowledge Intuitions will be suspended if (a) they believe that reducing false positive risk for their Receivers by gathering additional evidence is a practical possibility, and (b) they value the benefit for their Receivers of such a reduction in risk more than they value the cost of gathering the additional evidence (e.g., saving energy and resources). Here is an example.

Case 6-6

Sarah is driving her colleague Hannah home on a Friday afternoon. Sarah asks Hannah if she would mind if they stopped at the bank on the way so that she can deposit her paycheque. Sarah explains that while it is not terribly important that she do so as she has no impending bills, it would nonetheless be efficient. Coincidentally, Hannah also uses the same bank and knows that they will not have to go far out of their way to get there. She tells Sarah that she does not at all mind stopping.

As they drive past the bank, they notice that the lines inside are very long. Realizing that it isn't very important to Sarah that her paycheque be deposited right away, Hannah says, 'I know the bank will be open tomorrow, since I was there just two weeks ago on Saturday. I truly do not mind waiting now. However, you might prefer coming back and depositing your paycheque tomorrow'.

Suddenly, Hannah realizes that the bank app on her smartphone would indicate the bank's operating hours instantly. She reminds herself that banks do change their hours. Hannah then says, 'Just to make sure that the bank's open, I'll check a bank app that I have on my phone. It'll just take a second'. She checks and finds that the bank is indeed scheduled to be open tomorrow. Hannah feels reassured.

Just as in Case 6-1, it seems plausible and natural that Hannah would initially take herself as knowing that the bank will be open on Saturday, given that it was open just two weeks before on a Saturday. Moreover, it seems that her Epistemic Standard in that moment tracks a Risk/Reward Quotient based on Sarah's practical situation, i.e., the AAT applies to Hannah's assessment of the Risk/Reward Quotient. However, a moment later, it seems that Hannah has determined that the benefit of reducing false positive risk for Sarah by checking her bank app outweighs the very small cost of doing so; and, consistent with the Adding Evidence Thesis, her Knowledge Intuitions are suspended while she checks her app. So, while Hannah is checking her app, her Epistemic Standard is raised so that her Present Evidence is not strong enough to meet it; and the suspension of her Knowledge Intuition is indicated by her reminding herself that banks change their hours.

*

Senders may find themselves in the position of being able to gather additional evidence and reducing false positive risk for their Receivers, while having their Receivers incur the cost of gathering that evidence. In these cases, Senders, as Altruists, take the cost of gathering evidence, which is incurred by their Receivers, into account when they assess whether it is worthwhile gathering the evidence. So, in these cases, the suspension of Senders' Knowledge Intuitions (under the Adding Evidence Thesis) depends on valuations of the benefit to the Receiver of reducing false positive risk, and on the cost incurred by the Receiver gathering additional evidence. The pair of cases that follow, taken together, demonstrate this principle.

Case 6-7

Jonah, a physician, has just received the results of tests he ran on his patient Peter and feels he now has more than enough evidence to diagnose condition *C*. In fact, Peter presents as a textbook case of *C*. The preferred treatment for *C* is highly effective, and produces only mild side effects. He will meet with Peter tomorrow to inform him.

Later that day Jonah learns of a test for *C* that has just become available for use, and that has very high predictive validity. Peter would have to pay only a very small fee for the test and Jonah could easily administer it in his office. Jonah feels that he would be happy to have the diagnosis confirmed, because he would prefer that Peter not have to incur unpleasant side effects (though they are mild) without obtaining any benefit. He lets Peter know that he would like to administer the additional test before settling on a treatment plan.

Case 6-8

Jonah, a physician, has just received the results of tests he ran on his patient Peter and feels he now has more than enough evidence to diagnose condition *C*. In fact, Peter presents as a textbook case of *C*. The preferred treatment for *C* is highly effective, and produces only mild side effects. He will meet with Peter tomorrow to inform him.

Later that day Jonah learns of a test for *C* that has just become available for use, and that has very high predictive validity. Unfortunately, it would take Peter a full day to get to the closest medical facility that could administer the test. Jonah knows that this would be very hard on Peter. Jonah reminds himself that he already has more than enough evidence to diagnose for *C*, and he is sure that the new test would only confirm his diagnosis. Jonah tells himself that he knows that Peter has *C*. Jonah does not

consider the matter further, and plans to meet with Peter tomorrow to talk about his diagnosis.

In both cases, Jonah considers the proposition p that Peter has condition C ; and Jonah's Present Evidence is constituted by tests that he has already administered. We assume, in both cases, that Jonah's Epistemic Standard for p tracks his assessment of the Risk/Reward Quotient under the AAT, i.e., Jonah's assessment of the cost of a false positive and the benefit of a true positive is based on the seriousness of condition C for Peter, Peter's prognosis, and the efficacy of the recommended treatment.

In Case 6-7, Jonah's Present Evidence initially meets his Epistemic Standard; and, initially, he takes himself as knowing that Peter has C . However, later in the day, these Knowledge Intuitions are suspended when he learns about a new test and assesses that it would be worthwhile to administer it—the cost of which, for Peter, would be very small. Jonah also learns of a new test in the second case. However, there, the cost of administering the test for Peter would be high, Jonah does not find it worthwhile to have Peter take it, and Jonah's initial Knowledge Intuitions are not suspended. So, it seems as if the suspension of Senders' Knowledge Intuitions under the Adding Evidence Thesis can depend on Senders' valuations of the cost that Receivers incur gathering additional evidence.⁵⁶

6.4 Communicating with Multiple Receivers

Let us look at cases in which a Sender intends or envisages communicating the same message to *a number of Receivers*. To start, we focus on cases in which the Sender assumes that the relevant practical circumstances of her multiple Receivers are similar, or in which the message may produce some minimum benefit for all. This is the state of affairs, for example, for a social worker who is about to advise a group of inmates prior to their release from prison to dress neatly and conservatively for job interviews; and it is the state of affairs for a lawyer who is about to explain to a group of CEOs of major corporations that a new federal regulation applies to their companies.

⁵⁶ This difference in Jonah's response in the second case is broadly consistent with the finding that the "display of cost information [for laboratory tests] . . . can lead to a modest reduction in ordering of laboratory tests" (Horn et al., 2014, p.708). However, the costs to which Horn's study refer are actually born by the American government's Medicaid program. Even so, it seems likely that the effects of cost information would be at least as great if they were born by the patient.

I suggest that, regardless of the number of Receivers for a given communication, Senders' Epistemic Standards are set using the Basic Method under the Adding Evidence Thesis. In §6.4.1, look at how the Basic Method is employed, and what it predicts about the effect that increasing the number of Receivers has on Epistemic Standards. Then, in §6.4.2, I look at what the Adding Evidence Thesis predicts about the effect that increasing the number of Receivers has on Epistemic Standards.

6.4.1 The Basic Method with Multiple Receivers

Senders may employ the Basic Method by first getting a sense of the Risk/Reward Quotient for their group of Receivers (which, in these cases, is a practical possibility because Senders assume that the relevant practical circumstances of their multiple Receivers are similar, or that their messages may produce some minimum benefit for all). For example, the Risk/Reward Quotient may be low for the social worker suggesting to the inmates that they would do well to dress neatly and conservatively for job interviews. The cost of a false positive is low, i.e., the cost to the inmates of dressing neatly and conservatively for job interviews, when doing so would in fact not be helpful, is low; while the benefit of a true positive may well be very high, i.e., dressing that way may significantly improve their chances of landing a job. The social worker may then associate the low Risk/Reward Quotient with a low Epistemic Standard: She may not require particularly strong Present Evidence in order to take herself as knowing that neat and conservative clothing will help. On the other hand, the Risk/Reward Quotient may be very high for the corporate lawyer telling the CEOs that a new federal regulation applies to their companies. The cost of a false positive for the CEOs may be high, i.e., the cost of *unnecessary* compliance may be high; while the benefit of a true positive, from the CEOs' perspective, may well be very low, or even so low as to be inappreciable, i.e., they will not be at all rewarded, honoured or congratulated for complying with a new federal regulation. The lawyer may then associate the high Risk/Reward Quotient with a high Epistemic Standard: She may require strong Present Evidence in order to take herself as knowing that the new regulation applies to the CEOs' companies.⁵⁷

⁵⁷ This is a case in which a significant cost of a false negative could arise: If the CEOs failed to take themselves as knowing that the new regulation applies, and failed to comply with the regulation, when in fact it does apply, then, if they get caught, their companies would be subject to federal fines. However, recall (from §4.2.2) that I would aim to simplify my investigation by setting aside false negative and true negatives. To simplify the analysis of the CEO case, let us assume that the fines would be very small relatively to corporate incomes, and that the CEOs believe that the likelihood of getting caught is low. So,

Further, the Basic Method seems to predict that under certain circumstances the Epistemic Standard that Senders set depends on the number of Receivers to whom they intend or envisage sending a message. (I set aside consideration of the Adding Evidence Thesis for the moment.) Consider, for example, the situation that the corporate lawyer faces as a Sender. As the number of Receiver-CEO's increases, the lawyer senses that the *total* cost of a false positive incurred by the multiple CEOs increases, whereas the *total* inappreciable benefit of a true positive gained by the multiple CEOs remains inappreciable. So, as the number of Receiver-CEO's increases, the Risk/Reward Quotient and the lawyer's Epistemic Standard are predicted to increase. Here, *the Sender's Epistemic Standards positively correlate with the number of Receivers*. (However, if one has assumed that the lawyer's Epistemic Standard was very high to begin with, it might be difficult to appreciate the increase.) The opposite effect is predicted when Senders get the sense that the cost of a false positive for Receivers is so low that it is inappreciable, whereas the benefit of a true positive for Receivers is appreciable: All other things being equal, Senders' Epistemic Standards are here predicted to *negatively correlate with the number of Receivers*. (These valuations need not correlate in a linear way with the number of Receivers. It may well be that we are, in general, increasingly insensitive to a given increase in the number of Receivers.)

Even so, the Basic Method predicts that in cases in which the Sender has a sense that both the benefit of a true positive and the cost of a false positive for Receivers are appreciable, Senders' Epistemic Standards are *not* affected by the number of Receivers. All other things being equal, as the number of Receivers multiplies, Senders may get the sense that the cost of a false positive, and the benefit of a true positive, for a typical Receiver, are both being multiplied by the same number, i.e., the number of Receivers. So, Senders may get the sense that the ratio expressed by the Risk/Reward Quotient remains unchanged, producing no change in pressure on Epistemic Standards. Let us look at an example.

A municipal official has received a report from the national weather service that a raging storm front has turned sharply and is moving directly toward her municipality. It is scheduled to land in approximately 12 hours. In the moment, the official will determine whether or not she takes herself as knowing the proposition p that an evacuation order is warranted. Moreover, in that

the expected cost of a false negative is inappreciably low. Even so, we are left with the question as to what would motivate the CEOs to comply with the new law if they took themselves as knowing that it applies. To answer, let us say that some of the CEOs would comply with the law out of principle, while others would comply because their employment contracts require them to do so.

moment, she is mindful that if she issues the evacuation order and the storm does not hit, then those residents who evacuated will have been greatly and unnecessarily inconvenienced, i.e., that the cost of a false positive for the townsfolk is significant. She is also mindful that if she does issue the order and the storm hits, then those residents who evacuated will have avoided three or four days of impassable roadways, possible power outages, and freezing temperatures, i.e., the benefit of a true positive for the townsfolk is also significant. Further, the *total* cost of a false positive produced for the entire community will be the product of the cost of a false positive for a typical resident multiplied by the number of effected residents; and the *total* benefit of a true positive produced for the entire community will be the product of the benefit of a true positive for a typical resident again multiplied by the number of effected residents. So, the official may be mindful that the costs or benefits may be multiplied thousands of times over. However, the ratio of the *total* cost of a false positive for the townsfolk to the *total* benefit of a true positive does not depend on the number of residents. The ratio would be the same whether there are 50 or 50,000 residents—because both the numerator and denominator of the ratio are multiplied by the number. So, the Basic Method predicts that the official's Epistemic Standard in relation to p would not depend on the population of her community. The Basic Method explains that for the official, all other things being equal, the ratio expressed by the Risk/Reward Quotient remains unchanged as the number of townsfolk increases; and, as a result, her Epistemic Standard in relation to p does not change. This prediction seems to be correct.

In summary, the Basic Method seems to correctly predict that the Epistemic Standard that Senders set depends on the number of Receivers to whom they intend or envisage sending a message if and only if either (1) the cost of a false positive for Receivers is appreciable, whereas the benefit of a true positive for them is so low that it is inappreciable, or (2) the cost of a false positive for Receivers is so low that it is inappreciable, whereas the benefit of a true positive for them is appreciable. Moreover, under (1), Epistemic Standards are positively correlated with the number of Receivers; while under (2), they are negatively correlated.

6.4.2 The Adding Evidence Thesis with Multiple Receivers

Recall that under the Adding Evidence Thesis, a Sender's Knowledge Intuition may be suspended if she assesses that (a) the benefit for her intended or envisaged Receiver(s) of reducing false positive risk (by gathering additional evidence), outweighs (b) the cost of gathering the additional evidence. The importance of the Adding Evidence Thesis is multiplied

when there are multiple receivers because (a) the total benefit produced for these Receivers (by reducing false positive risk) multiplies as their numbers multiply, whereas (b) the cost of gathering additional evidence is not affected by Receiver numbers. Senders may gather additional evidence just once, and then leverage that evidence many times over when they provide well supported messages to their many Receivers. Thus, the greater the number of Receivers, the greater is the prospect that (a) the total benefit of reducing false positive risk, will outweigh (b) the cost of gathering additional evidence—provided that the cost of a false positive for a typical Receiver is not inappreciable low. So, the greater the number of Receivers, the greater is the prospect that Senders' Knowledge Intuitions will be suspended while they gather additional evidence, again provided that the cost of a false positive for a typical Receiver is not inappreciable low. Thus, increasing the number of Receivers tends to put upward pressure on Senders' effective Epistemic Standards.⁵⁸

Further, the effect of increasing Receiver numbers on the total benefit of reducing false positive risk depends on how high the cost of a false positive is for a typical Receiver. For example, in the corporate lawyer case in which the cost of a false positive for a typical Receiver is high, increasing Receiver numbers would greatly increase the total benefit of reducing false positive risk. With high Receiver numbers, the corporate lawyer may find it to be worthwhile to gather additional evidence even if the cost of doing so is very high. However, even if the cost of a false positive for a typical Receiver is relatively low, increasing Receiver numbers sufficiently could make it worthwhile to gather additional evidence and lead to Senders' Knowledge Intuitions being suspended while additional evidence is gathered. Consider the situation in the social worker case. Recall, she communicates the proposition p that the inmates would do well to dress neatly and conservatively for job interviews. However, say that there is some small cost of a false positive for a typical inmate, a small cost for them behaving as if p when p is false. Dressing neatly and conservatively for job interviews would require such an inmate to buy clothes on a tight budget and make them feel a little awkward having to dress in a way they are not used to. The Adding Evidence Thesis suggests that the number of Receivers could increase to the point that the social worker's valuation of the total benefit or reducing false positive risk for the many inmates would outweigh her valuation of the cost she would incur gathering

⁵⁸ Notice that the positive correlation (between the number of Receiver and Epistemic Standards) which is predicted by the Adding Evidence Thesis cannot be mitigated by a negative correlation predicted by the Basic Method in multiple-Receiver cases. The former positive correlation only applies to cases in which the cost of a false positive for Receivers is not inappreciably low, whereas the latter negative correlation predicted by the Basic Method only applies to cases in which the cost of a false positive is inappreciably low (and the benefit of true positive for Receivers is appreciable).

additional evidence. If this occurred, then her Knowledge Intuitions would be suspended and her Epistemic Standard would effectively increase while she gathered that additional evidence.

This state of affairs prompts three *ceteris paribus* predictions. *First*, as Receiver numbers increase, it becomes increasingly likely that Senders' Knowledge Intuitions will be suspended under the Adding Evidence Thesis. The *second* prediction follows from considering the circumstances under which it would not be worthwhile gathering additional evidence. A Sender may find that in a particular context it is not worthwhile to gather additional evidence because either (a) the improvement that the additional evidence is anticipated to produce in the predictive validity of the total Present Evidence is too low, or (b) the cost of gathering the evidence is too high. Even so, if the number of Receivers increased sufficiently (producing a corresponding increase in the total benefit of reducing false positive risk for Receivers), then it may well be worthwhile to gather that same additional evidence. The second prediction comes into sight: As Receiver numbers increase, (i) the minimum improvement in predictive validity of the total Present Evidence required in order to make it worthwhile to gather additional evidence decreases, and (ii) the maximum cost that it is worthwhile incurring to gather additional evidence increases. *Third*, as Receiver numbers increase, increasingly unlikely alternative hypotheses may become relevant and needing to be ruled out in order for the Sender to take herself as knowing.

6.4.3 Alfred Hitchcock's *Rear Window*

We have looked at the effect of increasing Receiver numbers on Epistemic Standards which are suggested by the Basic Method and by the Adding Evidence Thesis. Let us now look at the picture that emerges when these effects are combined.

We found that the Basic Method predicts a positive correlation between Epistemic Standards and Receiver numbers when the cost of a false positive is appreciable and the benefit of a true positive is inappreciable; and we found that the Adding Evidence Thesis predicts a positive correlation between Epistemic Standards and Receiver numbers when the cost of a false positive is appreciable, regardless of whether the benefit of a true positive is inappreciable. Together, these predictions suggest the following:

Multiple Receivers Prediction: Senders' Epistemic Standards correlate positively with Receiver numbers—except when there is no appreciable false positive risk for Receivers.⁵⁹

If the Multiple Receivers Prediction is true, then it would seem that the potential could be created for rather strange epistemic conflicts and tensions. Senders may intuit knowing that p in the moment that they intend or envisage communicating to only a few Receivers, and find too that their Knowledge Intuitions are suspended in the moment that they intend or envisage communicating with many Receivers—even though the strength of their Present Evidence remains unchanged. It seems that this is the nature of the central dramatic tension in Alfred Hitchcock's movie *Rear Window*.

The movie's protagonist, Jeff, is laid up at home due to a serious injury. He spends his days in a wheelchair, in front of a window overlooking a courtyard and an apartment building. He sometimes has company. He is visited regularly by a personal care nurse and his girlfriend. Even so, he is alone enough that he has taken to spying on a neighbour who lives in an opposite apartment. After a number of days of this, his observations lead him to consider the proposition p that his neighbour has killed and dismembered his wife.

Over the course of the movie, Jeff continuously gathers evidence for p . At early stages, his Present Evidence appears to be strong enough to report the matter to people who are close to him, but not strong enough to report the matter to the police. I suggest that calling the police constitutes, for Jeff, a communication to a large number of Receivers. It is, in effect, a call to many members of the police force, members of the judiciary, and, in a sense, a call to the broader community.

There are three discernible stages in the movie that correspond to three levels of strength of Jeff's Present Evidence. In the first stage, the strength of that evidence is such that he is comfortable asserting that p to his nurse and his girlfriend, but only to them. Here, he explains the situation to his girlfriend,

⁵⁹ There is a second general prediction that follows from combining these effects, which I do not focus on. We found that the Basic Method predicts a negative correlation between Epistemic Standards and Receiver numbers when the cost of a false positive is inappreciably low and the benefit of a true positive is appreciable; and found that the Adding Evidence Thesis does *not* predict that Epistemic Standards would be effected by Receiver numbers when the cost of a false positive is inappreciably low. So, the Adding Evidence Thesis does not predict either a mitigation or an amplification of the effect of Receiver numbers on Epistemic Standards predicted by the Basic Method in these cases.

“Why would a man leave his apartment three times on a rainy night with a suitcase, and come back three times? . . . Why didn't he go to work today? . . . What's interesting about a butcher knife and a small saw wrapped in newspaper? . . . Why hasn't he been in his wife's bedroom all day? . . . There's something terribly wrong . . . That'd be a terrible job to tackle. Just how would you start to cut up a human body?”. (Hitchcock, 1954, 37:00 - 39:00 minutes)

In this scene, Jeff does not appear to have any doubt at all that p . So, at this stage, Jeff seems to have the sense that his Present Evidence is strong enough to meet an Epistemic Standard that is appropriate for providing testimony to his nurse and girlfriend, but not strong enough to meet an Epistemic Standard that is appropriate for providing testimony to the police.

The second stage is reached after Jeff sees the neighbour handling a large trunk secured with rope. With this evidence, Jeff is comfortable providing testimony that expresses p , not only to his nurse and girlfriend, but also to his friend, Tom, a police detective. He calls him on the phone, and says, indignantly, ‘It is probably nothing important at all, just a little neighbourhood murder. That's all’ (42:00 minutes). The nurse overhears Jeff's phone conversation with Tom and asks ‘You called the police?’ Jeff replies, ‘It wasn't an official call. He's just a friend’ (48:00 minutes). Jeff later insists that Tom investigate without reporting the matter to police officials. Tom reassures Jeff: ‘I won't report this to the department . . . No sense in you getting a lot of ridiculous publicity’ (52:40 minutes).

Finally, in the third stage, Jeff has gathered enough additional evidence to be comfortable making an official call to the police.

So, the movie seems to support the Multiple Receivers Prediction, i.e., Jeff's Epistemic Standards seem to correlate positively with the number of Receivers that he intends or envisages communicating with. Moreover, much of the dramatic tension in the movie is created by the discordance between, on one hand, Jeff taking himself as knowing that the grisly murder has taken place when he is talking to his close confidants; and on the other, him finding that his Present Evidence is not strong enough that he can call the police. Near the end of the movie, this tension is released when Jeff finally has enough evidence to call the police.

Let us look at cases in which a Sender intends or envisages passing a message to a very large number of Receivers—hundreds, thousands or even millions. Messages may be passed orally or by the written word, may spread well beyond the community, and may be picked up in the distant future. I call such Senders, ‘Broadcasters’, and call their intended or envisaged communications ‘Broadcasts’.

Published philosophers are Broadcasters, i.e., they intend or envisage passing messages to hundreds or thousands of Receivers. Consider Descartes: In a dedicatory letter to the Sorbonne for his *Meditations on First Philosophy* he explains that the intended audience for his book includes all those who do not have faith that the

soul does not die with the body, and that God exists . . . [and one does well to direct such a book to these people because, for them, there is] practically no moral virtue, that they can be persuaded to adopt until these truths [about the soul and God] are proved to them by natural reason. (1996, p.3)

Descartes also intended that his *Meditations* serve as a stable foundation for science:

I realized that it was necessary . . . to demolish everything completely and start again right from the foundations if I wanted to establish anything at all in the sciences that was stable and likely to last. (p.12)

Like other Senders, Broadcasters are Altruistic. So, I suggest that the Altruistic Assessment Thesis is applicable in Broadcast cases. That is, Broadcasters’ adjudications of the propositions they communicate are based on their assessments of the benefit of true positives and the cost of false positives in relation to the practical interests of their Receivers. Even so, the biological research would seem to allow that this claim is plausible *only in relation to Broadcasts that are made under one (or more) of the four circumstances I described earlier* (in §6.2 and Appendix 3). Recall the four circumstances.

- (1) The individual intends or envisages communicating with closely-related kin—or a more distantly related individual, if the cost of preparing and sending a helpful communication is a small fraction of the latter individual’s expected benefit (see §A.3.1 in Appendix 3).

- (2) The individual intends or envisages communicating with someone with whom she has an ongoing relationship that involves reciprocal interactions, e.g., one is an employee, and the other her employer (assuming that the employee's work in fact benefits the employer and the employer regularly reciprocates by paying the employee's wages) (see §A.3.2).
- (3) The individual with whom one is intending or envisaging communicating is a 'co-operator' (i.e., someone who has a good record of reciprocating and helping others), and others in the community will learn about the communication producing reputational rewards for the altruist (see §A.3.3).
- (4) The individual has a co-operative character or personality, and it is strong enough to compel her to provide the level of co-operation called for in the circumstances (see §A.3.4).

It seems that Broadcasts would only rarely be made under either of the first two of these circumstances, while they may very frequently be made under the third or fourth. Regarding circumstance (1), it seems that, in most Broadcast cases, most Receivers would *not* be closely-related to the Broadcaster. Regarding circumstance (2), it is simply not a practical possibility for Broadcasters to have ongoing dyadic relations with significant numbers of the very large audience for their Broadcasts. However, regarding circumstance (3), it seems that prospective Broadcasters generally envisage that others in the community will learn that they have sent a helpful message to co-operators; and, regarding circumstance (4), it may well be that the character or personality of many or most prospective Broadcasters is co-operative enough that they feel compelled to incur the cost associated with Broadcasting certain messages regardless of the self-oriented benefits that follow. So, it does seem that the biological research would allow that it is plausible that the Altruistic Assessment Thesis is applicable in Broadcast cases.

*

The Multiple Receivers Prediction calls for a positive correlation between Senders' Epistemic Standards and Receiver numbers—except when the cost of a false positive is inappreciably low; and recall that much of this correlation has to do with the suspension of Knowledge Intuitions under the Adding Evidence Thesis. If this is correct, then in Broadcasting cases (which feature extremely high numbers of Receivers) Broadcasters set very high Epistemic Standards, and find that their Knowledge Intuitions are very often suspended while they gather additional

evidence—except when the false positive risk for Receivers is inappreciably low. This prediction is supported by the general view in the sciences that scientific conclusions, though they may be reliable, are nonetheless tentative. Further, if this correct, then Broadcasters tend to find the following:

- (1) It is worthwhile gathering additional evidence that only modestly increases the predictive validity of their total Present Evidence. This prediction is supported when multiple researchers undertake independent studies to confirm a given hypothesis, which is common in most evidence-based disciplines. For example, at least nine different kinds of experiments have been performed to confirm the value of Avogadro's Constant.
- (2) It is worthwhile incurring a high cost to gather additional evidence. This prediction is supported when researchers—including say historians, archaeologists or botanists—expend a great deal of time and energy, travel great distances, or risk life and limb gathering additional evidence. It is further supported by the existence of 'big science', such as the research that involves the Large Hadron Collider or one of the giant telescopes.

I pointed out earlier (in §5.2) that the Adding Evidence Thesis suggests that the possibility of suspending Knowledge Intuitions arises in each instance in which we identify the practical possibility of gathering additional evidence; and further that a subject's prospects for making those identifications depends entirely on her capacities, technologies, resources and strategies. So, it would seem that as our capacities and technologies improve, the two tendencies I outlined under (1) and (2) above would strengthen. This picture is particularly important in Broadcasting cases because of the multiplication of benefits that comes from sending well-supported messages to large numbers of Receivers.

*

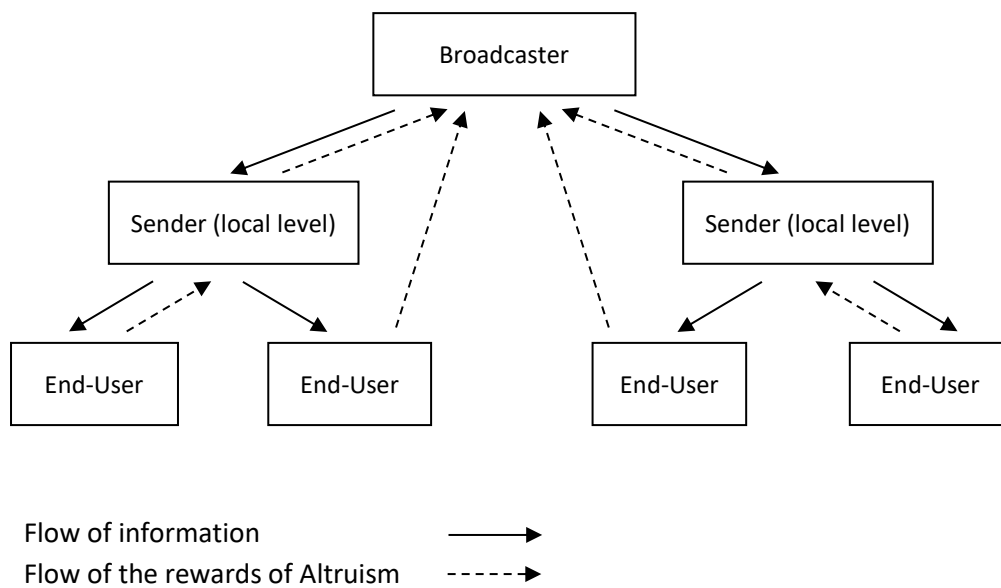
The fact that a given Broadcast messages may be sent to diverse Receivers whose practical circumstances may be very different, who live in different communities or in the distant future, or who have very different Worldviews, has implications for the kind of messages that Broadcasters do well to send. They should send messages that have a broad relevance, and that are of general importance. However, such messages may present a challenge for many *Receivers*, i.e., applying the general messages that they receive to their own particular

circumstances, or taking practical guidance from them. Some Receivers may not see the importance or relevance of such highly general messages. So, the generality of messages can constitute a significant barrier to producing benefits by Broadcasting.

Even so, this lost opportunity to produce benefits may in large part be recovered. Local ordinary Senders may be able to help relatively small numbers of local Receivers derive practical benefits from these messages. Moreover, these local Senders may do so under *any* of the four circumstances I outlined above (i.e., helping kin, advancing dyadic reciprocation, burnishing the Sender's reputation and exercising the Sender's helpful or co-operative character). As a result, a division of labour may arise between Broadcasters and local Senders that has to do with passing helpful messages. Both kinds of Senders would play an important role in what may be viewed as an information economy. They both facilitate the distribution, not of the community's physical wealth, but of its informational wealth; and that distribution would be controlled by the 'invisible hand' of human Altruism.

In this picture, Broadcasters provide helpful information to some number of local Senders as shown in Figure 6-1. These latter Senders may then use those Broadcast messages to fashion derivative messages that they send to local Receivers—which (to avoid confusion in this context) I call 'End-Users' (see Figure 6-1). These derivative messages are intended to be helpful, relevant and understandable to End-Users. (Further, local Senders may well base their local communications on the messages of several Broadcasters.)

Figure 6-1: Schematic flow of information and reciprocity in an information economy



Informational flows of this kind are prevalent in modern communities. A representative of the European Banking Authority (a Broadcaster) may depend on a large number of analysts, commentators and local authorities (local Senders) to assist them in putting their economic assessments to good use, and realizing the benefits that this information can produce. A public health researcher (a Broadcaster) may depend on a large number of dieticians (local Senders) to advise dietary managers in hospitals and schools. Generally, basic science depends on a wide network of local Senders to produce benefits for large numbers of End-Users.

*

The correlation between Broadcasters' Epistemic Standards and their Receiver numbers, together with the potentially great diversity among the End-Users of their messages, would seem to strain or even break the connection between Broadcasters and the practical interests of any particular End-User. Broadcasters do not need to get a sense of the cost of a false positive, the benefit of a true positive, or the Risk/Reward Quotient for *any particular* End-User, but do need to do so for the *group* of End-Users.

The picture that emerges here is that Broadcasters' Epistemic Standards are (a) very high, and (b) do not accord with Epistemic Standard Variabilism, i.e., those standards are not affected by

the interests of any particular End-User. The situation is consistent with the standard view of knowledge taken by First History theorists, i.e., Classic Invariantism. Moreover, I suggest that the reason that the present theory comes together with the First History view in this way is that First History theorists are, by and large, focused on Broadcasters' Knowledge Intuitions. In the following chapter, I look at this more carefully and suggest that these theorists' narrow focus on Broadcasters' Knowledge Intuitions facilitates their being able to play an important role in an information economy.⁶⁰

Before moving on, I would remind us that we have found in this chapter that Classic Invariantism, Contextualism and Subject Sensitive Invariantism are all correct within their own boundaries, and that these boundaries do not, by and large, overlap. (There is a general class of exceptions in relation to overlap: Cases that involve first-person attributions or denials, and the employment of the Basic Method under egoist assessments, support both Contextualism and SSI). They all make correct predictions in relation to different kinds of cases; and so, in this sense, they are all complementary (outside of the exception I just described). However, again, they overgeneralize by suggesting universal applicability of a theory that in fact applies to only a subset of cases.

⁶⁰ Edward Craig and I both purport to offer a naturalistic explanation for the appearance of high Epistemic Standards that do not accord with Epistemic Standard Variabilism. So, let us compare our conclusions.

Recall, from my earlier discussion of his theory (in §3.5), his suggestion that it is normal that our view in relation to epistemic matters 'will develop in the direction of objectivisation' (1999, p.84). That is, practical pressures normally cause us to move away from consideration of 'what any particular person wants at any particular time or place' (p.84) and toward consideration of our own future needs (p.84) and the needs of others (p.88). *Under objectivisation*, the function of knowledge is to flag

someone who is a good informant as to whether *p* whatever the particular circumstances of the inquirer, whatever rewards and penalties hang over him and whatever his attitude to them. That means someone with a very high degree of reliability. (p.91)

Moreover, he suggests that the alternative to objectivisation, in which individuals focus on the pure 'here and now for me as I am here and now' (p.83) may not actually ever arise; or, if it arises, it will be short-lived (p.84). So, Craig's theory, in its final form, is focused on, or limited to, a post-objectivisation situation.

Chapter 7: Why Epistemology Bifurcated

In Chapter 2, I suggested that the epistemological literature bifurcated over the last 100 years or so, and sketched two independent histories within epistemology that I called the First History (in §2.2) and the Second History (in §2.3). The theory of Knowledge Intuitions that I developed in Chapters 3, 4, 5 and 6 suggests an explanation of why the literature would have bifurcated in the way that it has. In the first section of this chapter, I present and defend that explanation. In the second section, I address the view suggested by many First History theorists that Second History theory is not of particular philosophical importance. If they are right, there would not be much point in worrying about this bifurcation: We ought to concern ourselves only with First History theory.

7.1 Why the Literature Bifurcated in the Way That It Has

To see how the theory of Knowledge Intuitions that I have laid out explains the bifurcation of the epistemological literature, one must first recognize that the epistemologists who contribute to the two histories are Broadcasters, i.e., Altruistic Senders to large numbers of diverse End-Users. They intend to help or envisage helping both readers of philosophy journals and others who may potentially benefit, directly or indirectly, from good epistemology. Further, it would seem that epistemological messages could be helpful to End-Users in two general ways—one of which may be associated with the First History, and the other with the Second History. Let us look at these two ways to help.

First, Broadcasters in epistemology could help large numbers of End-Users (indirectly) by helping Broadcasters *in other fields* achieve *their* aims of helping End-Users—particularly, those other Broadcasters who work in evidence-based disciplines like the sciences. Epistemologists who pursued such an aim would have a special role in the information economy of which they are a part. If my Figure 6-1 represented their information economy, then these epistemologists would hold a position that is on a higher level than that of the other Broadcasters represented in that figure. They would send messages that are primarily directed to those other Broadcasters. What sorts of messages would these be?

Earlier, (in §6.4.4) we found that, for Broadcasters, the key to producing valued benefits is to reduce false positive risk for End-Users. (However, we found that this generalization does not hold in cases in which the messages that are communicated are such that false positive risk for Receivers is inappreciably low.) The total potential harm associated with that risk is multiplied as the number of End-Users multiplies. This picture suggests that epistemologists could help other Broadcasters by *sending messages that help them address this problem of false positive risk—which they may do by helping other Broadcasters identify propositions that are unlikely to prove false*. I suggest that those epistemologists who send such messages are writing First History theory; and the general aim of helping other Broadcasters address the problem of false positive risk is what makes the First History a coherent research program within epistemology.

I suggest that this last general aim explains the importance in First History literature of the view that having sufficient evidence is no less than a responsibility, obligation or duty. Recall W. K. Clifford's extended metaphor about the ship-owner who sent out a ship filled with emigres (from §2.2):

he thought that perhaps he ought to have her thoroughly overhauled and refitted, even though this should put him at great expense. Before the ship sailed, however, he succeeded in overcoming these melancholy reflections . . . [and decided that he] would put his trust in Providence, . . . [with the result that the ship] went down in mid-ocean and told no tales' (1876, p.289).

Clifford makes vivid the multiplication of harms that may follow from sending false messages to large numbers of End-Users. Recall too (from §2.2) the importance in First History literature of the deontological conception of justification, which suggests that the very meaning of the term 'epistemic justification' should be rooted in this responsibility. Alvin Plantinga points out, '[o]riginally and at bottom, epistemic justification is deontological justification' (1993, p.14). Also recall (from §2.2) the suggestion in First History literature that we have a responsibility and duty to believe and to withhold belief in a way that accords with the evidence. William Alston explains that epistemic 'principles will *forbid* beliefs formed in such a way as to be likely to be false and either *permit* or *require* beliefs formed in such a way as to be likely to be true' (original italicization, 1988, pp.258-259). The primary motivation for such views is explained by the importance of reducing false positive risk for large numbers of End-Users.

The second general way in which epistemological messages could be helpful to End-Users is by shedding light on a basic aspect of the human condition and of human nature, shedding light on

the circumstances under which we take ourselves as knowing or of not knowing that particular world conditions obtain under ordinary circumstances. As I suggested earlier (in §1.2.2), this would in turn shed light on the ordinary use of 'know', the 'shiftiness' of 'know', as Matthew Chrisman puts it (2012, pp.121-122). From a different perspective, epistemologists could analyse, describe and explain the appearance of variation in Epistemic Standards. Many epistemologists have availed themselves of this opportunity; and I suggest that those who have write Second History theory. It is this aim that makes the Second History a coherent research program within epistemology. I add that, unlike First History epistemologists, those who write Second History theory hold a position in that information economy of which they are a part that is *on the same level* as that of other Broadcasters.

*

Let us now look at these two research programs more carefully. In particular, let us recognize that the two research programs are radically different: They call for (1) different conceptions of knowledge, (2) different views on the evidential standard required for knowledge, (3) different analytical focuses and (4) for knowledge theory that has different criteria for success. Let us look at these in turn.

1. Different conceptions of knowledge

Achieving the First History aim of helping Broadcasters identify propositions that are unlikely to prove false is facilitated by holding the view that knowing does not tolerate the possibility of falsity. Epistemologists assume that propositions that are truly known cannot later prove false. (I do not mean to suggest that the primary First History aim is necessarily the only reason that epistemologists hold this view—though I suggest that it is a significant factor in their holding the view.) This conception of knowledge is reflected, for example, in Bertrand Russell's suggestion that 'what does not go beyond our own personal sensible acquaintance must be for us the most certain' (1914, p.67). Russell suggests that Advanced Propositions are *inherently* susceptible to being proven false; and so, in general, knowledge claims in relation to Advanced Proposition should be made with an appropriate degree of caution. This view that knowledge does not tolerate falsity is also reflected in D. M. Armstrong's suggestion that the pursuit of knowledge leads to an indefinite regress of reasons (1973, pp.152-154). Knowledge depends on having a reason, *r*, for taking ourselves as knowing that *p*; having a further reason, *s*, for taking

ourselves as knowing that r has the appropriate relation to that which makes p true; having a further reason, t , to ensure that s qualifies as a good reason; and so on, indefinitely.

However, achieving the *Second History* aim of accounting for the ordinary use of 'know' forces epistemologists to address the fact that we sometimes take ourselves as knowing propositions that turn out to be false. Second History theorists, it would seem, would have to allow that propositions that are truly known can later prove false (though if they are, they would no longer be known). Indeed, many of these theorists make this allowance: They accommodate fallibilism—the view that it is possible to know without entailing evidence—in their theories of knowledge. Recall (from §2.3) Austin's early Second History step toward such an accommodation,

If we have made sure [that the bird is] a goldfinch, and a real goldfinch [by ruling out the relevant alternatives], and then in the future it does something outrageous (explodes, quotes Mrs. Woolf, or what not), we don't say we were wrong to say it was a goldfinch, we don't know what to say. Words literally fail us . . . [However, if we have ruled out the relevant alternatives, we cannot] be proved wrong, whatever happens . . . (1979, pp.88-89)

Austin's suggestion that, when we have ruled out all of the relevant alternatives, we cannot be proved wrong, seems to constitute a step toward the hypothesis that a knowledge attribution that follows from ruling out the relevant alternatives may be deemed to be true—even if it is *possible* that it will later prove to be false (and if it does prove to be false, would cease to be deemed to be true). So, knowing is fallible. Even so, it is Steward Cohen who brings fallibilism to the fore in the Second History with his paper "*How to Be a Fallibilist*" (1988) (which I discuss in §2.3). Here, he helps other Second History theorists accommodate fallibilism in their knowledge theory.

2. Different views on the evidential standard required for knowledge

I showed earlier (in §6.4.4) that my developed theory of Knowledge Intuitions predicts that Broadcasters in evidence-based disciplines, such as scientists, would tend to hold the view that the evidential standard for knowledge is very high and not dependent on the practical circumstances of any particular individual (when the cost of a false positive for End-Users is appreciable). Epistemologists are able to achieve the First History aim of helping these

Broadcasters by sharing this view, i.e., by being Classic Invariantists. However, most *Second History* epistemologists seem to have found that they would *not* be able to achieve their aim of accounting for the ordinary use of 'know' by holding this view. Classic Invariantism does not describe, analyse or explain the ordinary use of 'know'. Many Second History epistemologists have found that they could go some distance toward achieving their aims by maintaining less stringent views on epistemic justification, such as those defended by Contextualists, Subject Sensitive Invariantists and Moderate Invariantists. (However, as Broadcasters, Second History epistemologists may nonetheless hold the view that the evidential standard for knowledge in relation to the epistemology that they and others produce is very high and not dependent on the practical circumstances of any particular individuals.)

3. Different analytical focuses

Epistemologists are able to achieve the First History aim of helping Broadcasters identify propositions that are unlikely to prove false by focusing on the predictive validity of evidence and on how such evidence can turn out to be inadequate or faulty. Consider again H. H. Price's suggestion:

When I see a tomato there is much that I can doubt. I can doubt whether it is a tomato that I am seeing, and not a cleverly painted piece of wax. I can doubt whether there is any material thing there at all. Perhaps what I took to be a tomato was really a reflection . . . (1932, p.3)

Or, consider C. I. Lewis' suggestion that the 'congruence' or the 'coherence' of different pieces of evidence confer a much higher probability upon hypotheses than is generally assumed (1946, pp.344-346).

However, investigations into the predictive validity of evidence and on how such evidence can turn out to be inadequate or faulty do not shed light on the appearance of variation in our Epistemic Standards, or on the ordinary use of 'know'. Problems of evidence cut across contexts. So, Second History theorists do not standardly concern themselves with matters of evidence. Instead, they pursue their aims by investigating the practical circumstances that seem to influence Epistemic Standards.

4. Different criteria for successful knowledge theory

Generally, the principal criterion for successful theory is that it helps one to achieve one's research aims. So, the standard view among First History theorists is that a successful knowledge theory is one that may help us identify propositions that are unlikely to prove false; while the standard view among Second History theorists is that a successful knowledge theory predicts or explains the production of Knowledge Intuitions, or accounts for the use of 'know', in a wide range of cases.

*

Some of the theory that I included in my brief First History (in §2.2) may not seem to advance the First History aim of helping Broadcasters in other fields achieve their aim of identifying propositions that are unlikely to prove false. I refer to theory that suggests either that almost any proposition could prove false, or that we generally do not know when we have knowledge. Such theory may seem to suggest that Broadcasters should give up trying to identify propositions that are unlikely to prove false.

Cartesian epistemology which casts doubt on the veracity of sense perception and the existence of an outside world is an early example; and Cartesian concerns arise in the substantial literature on epistemic scepticism. Externalist theory may also be of concern in this regard. (A theory of justification is externalist if it includes, as a necessary condition for knowledge, a state of affairs that is beyond the potential knower's Cognitive perspective, e.g., the view that knowledge requires that a belief be the product of a reliable belief formation mechanism.) Externalism suggests that knowers may *generally* not know that they know. (I discuss externalism and opposition to it in §3.2).

Even so, this literature may well be helpful to Broadcasters in other fields. For example, as Tim Maudlin points out, it may inform scientific investigations:

[p]hilosophical scepticism focuses attention on the conceptual weak points in [scientific] theories and in arguments. It encourages exploration of alternative explanations and new theoretical approaches. (2015)

So, scientists and other Broadcasters may well find such epistemology helpful, even if they do not accept its radical conclusions. Moreover, scepticism, externalism and such suggestions as

Bertrand Russell's that 'what does not go beyond our own personal sensible acquaintance must be for us the most certain' (1914, p.67) may bear on Broadcasters' assumptions and attitudes about the sufficiency of evidence, and the ways in which evidence can fail (i.e., with any inductive inference). In this way, the literature may influence evidence gathering and experimental design, and it may support or encourage views such as that scientific conclusions are only tentative even if they are extremely reliable, or the Popperian view that we should value both the imaginative development and patient falsification of scientific theory.

7.2 The Validity of Second History Theory

As I suggested earlier, many First History theorists do not seem to think that Second History theory is of particular philosophical importance. If they are right, there would not be much point in worrying about the bifurcation in the literature that I have described. We ought to concern ourselves only with First History theory. So, let us look at First History criticism of the underlying basis for Second History theory to see if there is something there.

In *earlier* First History theory, the view that Second History theory is not philosophically important followed from the assumption that evidential standards for taking oneself as knowing, under ordinary circumstances, are *not* set in accordance with principles that matter. C. I. Lewis explains and defends this assumption even before concerted work on Second History theory began. It is a pre-emptive attack on Second History theory:

[Epistemological theory] cannot be merely a verbally more precise rendering of common-sense, nor a direct generalization from actual practice. Though it rises from what is implicit in experience, its procedure must be critical, not descriptive. So far as it is to be of use, it must assume the function of sharpening and correcting an interpretation which has already entered into the fabric of that experience which is its datum. Logical principles aim to replace . . . our naïve morality, and metaphysics, our unreflective ontological judgments. (1929, p.19)

Lewis' general observations either suggest or imply that the principles that govern the use of 'know' and the production of ordinary Knowledge Intuitions are simply not philosophically important. Such a view is perfectly understandable if either,

- (1) the theorist has not yet identified a principle that matters reflected in the ordinary use of 'know' or the Epistemic Standards that we seem to set (such as that they track the cost of a false positive), or
- (2) the theorist has focused *exclusively* on examples in which everyday Epistemic Standards seem to have been set in accordance with principles that do not matter much, i.e., that have little predictive or explanatory value.

David Lewis and Peter Unger, in their early writings, express views that are like C. I. Lewis', and do so under condition (1) above. (Both philosophers, in much later writings, adopt very different views on this subject.⁶¹) Both Lewis and Unger explain the appearance of Epistemic Standards that move up and down as merely the consequence of the imprecise use of 'know', and the cause of this lack of precision is not a matter of epistemological import. Lewis suggests that the ordinary use of 'know' may be imprecise just as the use of geographical shape adjectives, such as 'hexagonal', may be imprecise. He uses 'France is hexagonal' (1979, p.352) as an example. Unger (1975) suggests that 'know' belongs to the class of 'absolute terms'—a class that includes 'flat', 'dry', and 'empty'—and, like other absolute terms, is often, or generally, used imprecisely. We may say that the table is flat. However, if the table was truly flat, then no other object could be flatter. I note too that consistent with these suggestions, both Lewis and Unger, in those writings, suggest that true knowledge depends on entailing evidence.⁶² So, at the time that these early papers were written, neither Lewis nor Unger seem to have yet identified an important principle reflected in the use of 'know'.

Keith Lehrer expresses views that are like C. I. Lewis', and does so under condition (2) above. In particular, he focuses exclusively on unrepresentative examples that seem to reflect principles (in relation to ordinary Epistemic Standards) that have little predictive or explanatory value. In particular, he focuses on (a) an example of wishful thinking that invites belief with poor evidence, and, oppositely, on (b) an example that features an unusual psychological impediment to belief:

⁶¹ See, for example, Lewis (1996) and Unger (1986).

⁶² Both Lewis' and Unger's analyses fail, I suggest, because they both depend on the term 'know' being used in ordinary contexts in a way that would make it a gradable expression. The lack of precision in the examples they present is afforded by gradability: there are degrees to which an object may be hexagonal, degrees to which an object may be flat, degrees to which an object can be empty. However, Jason Stanley, has shown that 'know' does not come in 'varying degrees of strength' (2005, p.36).

One person may find, for example, that he believes that someone is tenderly concerned about his welfare, but, looking at the evidence, concludes that this is probably not true. He wishes it to be so with such fervour that he cannot help but believe it nonetheless. Similarly, there may be something that is so distasteful for a second person to believe that he cannot do so, even though the person becomes aware that it is evidently true. In the quest for truth, the first person might refuse to assent to what he believes, and the second might assent to what he does not believe . . .

. . . we find ourselves with a basket of apples . . . and we face the problem of sorting through the collection to decide which ones are good to eat and which are not . . . we sort through our beliefs to decide which ones should receive our assent and which ones not. (1979, pp.66-67)

In *later* First History theory, the basis for the view that Second History theory is not philosophically important is fundamentally different than it is in early First History theory. This change seems to correspond with the growing acceptance in the twentieth century that a particular cognitive process, or a particular cognitive trait or mental state, may well have an important function, even it is not obvious what that function might be. This point of view suggests that *if*, in fact, everyday standards of evidence move up and down, then that movement might well be governed by principles that matter. So, it seems that as the new point of view took hold, the basis of the First History dismissal of Second History theory transitioned from (a) the view that in those ordinary cases in which evidential standards seem to move up or down evidential standards are not set in accordance with principles that matter, to (b) the view that standards of evidence do not, in fact, move up and down. Epistemic Standard Variabilism is false. Thus, the debate between First History and Second History theorists *reduces to whether the thesis of Epistemic Standard Variabilism is true*. (Except that Moderate Invariantism, which I view as Second History theory and which I discuss briefly in §2.3, is immune from attacks from First History theorists on the basis of the falsity of Epistemic Standard Variabilism.)

Where are we, very roughly, in relation to this debate? On one hand, the most important evidence in support of Epistemic Standard Variabilism to be produced so far comes from the low stakes and high stakes cases presented by Contextualists and Subject Sensitive Invariantists (such as Stanley's Low Stakes Case and Stanley's High Stakes Case⁶³). In these paired cases, it seems that a change of context causes a change in Epistemic Standard. Additionally, some

⁶³ See fn. 34 for a listing of other often-discussed paired cases.

Contextualists suggest that cases in which sceptical alternatives are presented, such as ‘sceptical paradoxes’, constitute evidence for Epistemic Standard Variabilism. Consider this example: John says, ‘That animal is a zebra’; Mary replies, ‘Perhaps it’s a cleverly painted mule’; and John responds, ‘You may be right. I guess I don’t know that it’s a zebra’. Contextualists explain that the appearance of a higher Epistemic Standard in John’s last response is a consequence of a shift in the context of utterance, which shift was caused by Mary’s introduction of a sceptical alternative. On the other hand, First History theorists have presented a number of varied arguments against Epistemic Standard Variabilism.⁶⁴ Even so, it seems that there are a relatively small number of core arguments amongst these. In the pages that follow, I look at what appear to be these core arguments, and argue that they fail.

7.2.1 Reflective Access to the Process that Modulates Epistemic Standards

First History theorists have argued that if Epistemic Standards are context-dependent, then their being so should be apparent, and it is not. Earl Conee writes,

[h]ow can it be that fluent English speakers are well enough attuned to [the contextual factors that contextualists claim determine the meaning of ‘know’] while none noticed that ‘knows’ is context-sensitive until recently, and some are unable to recognize this by semantic reflection . . . ? (2005, p.55)

Conee suggests that there exists an intrinsic link between (a) the ability to move Epistemic Standards up and down in accordance with the practical circumstances, and (b) having reflective access to those processes. (Perhaps he believes that reflective access is needed to move Epistemic Standards up and down in accordance with the practical circumstances.) Such an inference seems to have no basis. Much of our Cognitive processing is unconscious. Moving Epistemic Standards up and down on the basis of practical circumstances may well also be an unconscious process. Therefore, contra Conee, being ‘attuned’ to contextual factors does not indicate that we should have ‘noticed’ Epistemic Standard Variabilism, i.e., we may well be attuned via unconscious processes. Moreover, we often fail to have a clear understanding of

⁶⁴ I set aside here explanations that later First History theorists and Moderate Invariantists have offered for the appearance of Epistemic Standard Variabilism; which explanations do not depend on Epistemic Standards moving up and down. These explanations are not arguments against Epistemic Standard Variabilism *per se*, but rather constitute competing accounts of the data.

why we doubt or fail to doubt some given proposition given the evidence. Such a phenomenon suggests the workings of an unconscious process that governs whether we take ourselves as knowing.

7.2.2 Interpretations of Acontextual Knowledge Attributions

Epistemologists often present knowledge attribution examples *without enough information* about the practical circumstances for us to get a sense of the Risk/Reward Quotient or whether the Adding Evidence Thesis is applicable. I describe such examples as being ‘acontextual’. One might think that if Epistemic Standard Variabilism is true (particularly if our Epistemic Standards depend on the practical circumstances), then we should simply not be able to adjudicate these acontextual knowledge attributions: If we do not know anything about the practical circumstances, we should not be able to set an Epistemic Standard, and therefore not have a basis upon which to adjudicate those attributions. However, we do in fact somehow make sense of these attributions. Even so, a closer look suggests that while we may well not need context to make sense of acontextual knowledge attributions, our ability to do so does not disprove Epistemic Standard Variabilism. To see this, consider these statements which are uncontroversially context-dependent and that are presented without any context:

- (1) Simon only drinks local wines.
- (2) Simon is short.
- (3) Simon lives in an old house.

One may make sense of (1) by assuming that Simon only drinks wines that are produced in the *country* in which he lives—even though, in some contexts, it means that he only drinks wine from the particular *region* in which he lives. One may make sense of (2) by assuming that Simon is short relative to his demographic—even though, in some contexts, it means that Simon is short relative to professional American basketball players. People who live on the Canadian West Coast may make sense of (3) by assuming that the house that Simon lives in is somewhere between 40 and 70 years old—even though Europeans may make sense of the statement by assuming that Simon’s house is much older.

These last suggested interpretations were all based on a *best bet about the nature of the context*. For example, one may interpret (2) on the basis of a best bet that the sentence would have been presented in a discussion about Simon’s height relative to his demographic. The fact

that we are able to interpret these statements when they are presented outside of a context does not disprove that they are context-dependent.

So too, it seems like many acontextual knowledge attributions are interpreted on the basis of a best bet about the nature of the context. For example, I interpret the acontextual knowledge attribution, 'I know that the animal is a zebra', on the basis of a best bet that the attribution is uttered by an individual visiting a zoo with a friend or family member in the course of a casual conversation. As a result, I assume that the Risk/Reward Quotient and Epistemic Standard that applies to the example are *not* particularly high.

However, it seems that some of us interpret some acontextual knowledge examples, not on the basis of a best bet about the nature of the context; but, rather, on the basis of an internalized theory of justification. These attributions are interpreted in a way that is consistent with such an internalized theory; and we see this frequently in the examples that are presented in First History literature.

We find, in particular, that First History epistemologists interpret acontextual knowledge attributions in a way that is consistent with the justification theories they defend. For example, early and mid-twentieth century First History theorists—such as Bertrand Russell, H. H. Price, C. I. Lewis and D. M. Armstrong⁶⁵—standardly defend the view that we know *very little if any* of what we think we know; and, consistent with this view, they find that acontextual knowledge attributions reflect an extremely high evidential standard. For example, H. H. Price suggests that the statement 'I know that that is a tomato' reflects an extremely high standard; and that, as a result, the statement is probably false. In contrast, Earl Conee, who holds the more moderate view that the 'most plausible unvarying standard for truth is very high, but not unreachably high' (2005, p.52) finds that acontextual knowledge attributions reflect a more moderate evidential standard. However, regardless of the particular theory of justification one has internalized, if one consistently interprets acontextual knowledge attributions in a way that is consistent with that theory, then it will seem as if Epistemic Standard Variabilism is false—even though one needs only a proper context to reveal it.

*

⁶⁵ I discuss their theories in §2.2.

Herman Cappelen and Ernie Lepore (2003) present what may be the most elaborate attack on Epistemic Standard Variabilism in the literature, which they do via an attack on Contextualism. That attack is based on three purported tests for context sensitivity of the truth conditions of knowledge attributions. In their ‘test (i)’ and ‘test (ii)’, they look to see whether the truth conditions of a given knowledge attribution change when it is made *under different circumstances*. However, Cappelen & Lepore do not say a word about the *practical features* of those circumstances, i.e., features that are needed to get a sense of the Risk/Reward Quotient. In the end, they make the mistake of endeavouring to test for context sensitivity using acontextual attributions; and as a result, they get negative findings for context sensitivity which findings are false.

Let us look carefully at one of these two tests, test (ii). (We do not need to look at both because they share the same shortcoming.) Here is a schematic version of it:

Test (ii): Initially, agent A, under unknown practical circumstances C_1 , attributes knowledge to agent B. Agent A says ‘B knows when the train leaves’. Then, an observer, agent D, under unknown practical circumstances C_2 , reports that A has endorsed B as a good informant with respect to when the train leaves. D says, ‘A said that B knows when the train leaves’.

Thus, the attribution, ‘B knows when the train leaves’ has been transplanted from A’s unknown practical circumstances, C_1 , to D’s unknown practical circumstances, C_2 . Cappelen & Lepore suggest that if such an operation is unproblematic, then ‘know’ must not be context-dependent; and indeed the operation is not problematic. So, this test produced negative findings for context sensitivity (2003, pp.34-35). The theorists further suggest that the evidential standard connoted by the knowledge attribution is the same in both C_1 and C_2 .

Even so, as I suggested, the test is simply not sensitive to context sensitivity. The reader may set an Epistemic Standard for the attribution, ‘B knows when the train leaves’, in both C_1 and C_2 , in a way that is either consistent with a best bet about the practical circumstances in C_1 and C_2 , or consistent with internalized justification theory. In either case and as a result, the reader’s sense of the Risk/Reward Quotient and Epistemic Standard in both C_1 and C_2 is likely to be the same. In order to effectively test for context sensitivity, we would need to know a little more about agent D’s and agent A’s practical circumstances, and those circumstances would need to be significantly different (in the right way), similar to what is presented in the paired cases used to support Contextualism or SSI (see §2.3).

(Cappelen & Lepore's investigation includes a third test, 'test (iii)'. In that test, Cappelen & Lepore do present knowledge attributions in practical contexts, and one may indeed have a sense of the Risk/Reward Quotient in those contexts. So there is no call to discuss it in this particular section. Even so, Cappelen & Lepore's report negative findings for Epistemic Standard Variabilism in the test, and so I discuss it elsewhere, in Appendix 4. I conclude that while the test is valid, Cappelen & Lepore have not conducted it properly. If conducted properly, it is positive for Epistemic Standard Variabilism.)

7.2.3 An Observer Effect

Another line of attack against Epistemic Standard Variabilism is based on the effect of critically examining Knowledge Intuitions in the moment that they arise. Such critical examination often produces doubt. It is argued that this doubt is a reflection of the doubter's true invariant evidential standard for knowledge. So, critical examination reveals our true conception of knowledge. Earl Conee explains,

In ordinary contexts, when nothing much turns on it, people will claim knowledge, and attribute knowledge to themselves and others, in belief and in speech . . . [However,] if asked whether some proposition to which knowledge is ascribed on some such basis is really known, or truly known, or really and truly known, fluent speakers have a strong inclination to doubt or deny that it is. Only the most conspicuous facts of current perception, the clearest memories, triple-checked calculations, and the like will often pass some such 'really and truly' test. The present view has it that the answer to this 'really and truly' question reveals what a speaker judges to be knowledge when she is trying her best to apply her best thought as to what is the actual standard for the truth of a 'knowledge' attribution. (2005, p.52)

The questioning that Conee references, whether the proposition is really and truly known, is an example of the critical examination to which I refer. This seems to produce doubt. However, there is reason to believe that this effect of critical examination is *not* revelatory in the way that Conee and others suggest. Critical examination produces an 'observer effect' (a notion often used in the physical and social sciences) in which the examination alters the examined object. When we check tire pressure, we release air and thereby lower the pressure. So too, when we critically examine Knowledge Intuitions, we may produce doubt. David Lewis makes a similar

point—suggesting that examinations make knowledge elusive: ‘Examine it, and straightaway it vanishes’ (1996, p.560). So, I maintain, critical examinations of knowledge intuitions produce tainted data—data that has no value to researchers who are trying to understand the production of Knowledge Intuitions.

I can suggest two causes of such an observer effect. *First*, the effect is the result of the knowledge attributor inferring that the critical examiner has a reason for scrutinizing her attribution, and that reason weighs against knowing. Such an inference would be consistent with the Gricean hypothesis that proper conversational contributions are relevant to the purpose of the conversation (Grice, 1989). The attributor takes it that her attribution is being critically examined *because* the examiner has assessed the likelihood of the questioned proposition being true as being lower than has the attributor. The examiner may (a) possess contrary Present Evidence, (b) assess the Prior Probability that the proposition is true as being lower than she does, or (c) assess the predictive validity of the Present Evidence as being lower than she does. So, the question of whether one really and truly knows may suggest that there exists reason to doubt. J. L. Austin rightly challenges the conversational appropriateness of such critical examinations,

The doubt or question ‘But is it a *real* one?’ has always (*must* have) a special basis, there must be some ‘reason for suggesting’ that it isn’t real, in the sense of some specific way, or limited number of specific ways, in which it is suggested that this experience or item may be phony. (1979, p.87)

A *second cause* of an observer effect may be that critical examination suggests to the knower that the examiner is aware of potential costs of a false positive of which the knower is unaware. For example, I might worry that an examiner is asking me if I really and truly know because he is worried about a grave consequence of a false positive for me. If so, I may raise my assessment of the Risk/Reward Quotient in the situation, and associate with it a higher Epistemic Standard—which standard my evidence fails to meet.

Under either of these causes, the knower in these cases treats the critical examination itself as new Present Evidence that weighs against the proposition in question. Further, the knower’s assessment of the strength of this new evidence may depend on either (a) how good an adjudicator of the questioned proposition the knower believes the examiner would be, or (b) whether the knower believes that the examiner has information about the consequences of a

false positive that she does not. Nonetheless, this new contrary evidence, the critical examination, may well be strong enough to produce doubt.

The notion that critical examination constitutes new contradictory Present Evidence may be used to explain a certain often-made observation that informs some versions of Relevant Alternative Theory. It is observed that the mere act of advancing an alternative hypothesis makes the hypothesis relevant in the context. According to RA Theory, once an alternative is relevant, knowledge depends on ruling it out.⁶⁶ I am suggesting that when an agent who takes herself as knowing that p is presented with the alternative hypothesis q , she may infer that q would not have been presented unless the presenter had good reason for doing so—particularly, she may infer either that the presenter has evidence which she does not that supports q , or that the presenter's evidence for q is stronger than her own. Thus, the presentation of alternative hypotheses introduces an observer effect—though it is a different kind of situation than that produced by critical examinations of knowledge attributions. So, from the knower's point of view, the mere presentation of alternative q constitutes evidence that weighs against p , and changes the nature of the case so that additional Present Evidence is needed to rule out q . (Notice again that it is plausible that the presentation of an alternative does not have the effect of revealing an otherwise hidden conception of knowledge.)

7.2.4 Basic Assumptions About Knowledge

In each of the last three sections I alluded to an incorrect basic assumption that seems to be held by many First History theorists:

- (1) If our Epistemic Standard accords with Epistemic Standard Variabilism, then this should be evident, i.e., Epistemic Standard Variabilism depends on our having reflective access to the process that modulates Epistemic Standards (in §7.2.1).
- (2) The meaning of 'know' is reflected in our interpretation of acontextual examples (in §7.2.2).
- (3) We should be able to critically examine Knowledge Intuitions or suggest alternative hypotheses without producing an observer effect (in §7.2.3).

⁶⁶ See, for example, Steward Cohen (1988, p.106)

Different combinations of these basic assumptions are reflected in the First History literature that argues against Epistemic Standard Variabilism. In an interesting attack on Epistemic Standard Variabilism, Timothy Williamson makes assumptions (1) and (2), at the same time that he explicitly identifies assumption (3) as one that he will avoid making, just in case it is false. Williamson writes,

John: I know that this is a zebra.

Mary: How do you know that it isn't a mule cleverly painted to look like a zebra?

John: Hmm, for all I know it is a painted mule. So I was wrong. I don't know that it is a zebra after all.

. . . a contextualist can reconcile [John's] first and last remarks by postulating that Mary's question causes the word 'know' to shift its reference [consistent with Epistemic Standard Variabilism]. But [this] move [doesn't make] sense of John's admission 'So I was wrong', for [the move involves] explaining how John was not wrong in his first speech. More subtly, in context his final words 'after all' imply an admission of error. (2005, pp.220-221)

Williamson acknowledges the possibility that assumption (3) may be incorrect by allowing the possibility that Mary's question changes the context of attribution, and in the process alters the truth conditions of John's initial knowledge attribution. Even so, Williamson suggests that when John implies that his initial knowledge attribution was made in error—call this 'John's *mea culpa*'—John reveals what he judges to be knowledge, contradicts a contextualist account of the case, and contradicts Epistemic Standard Variabilism. I suggest that this conclusion is incorrect and depends on basic assumptions (1) and (2).

First, notice that assumption (2) is reflected in Williamson not saying a word about the practical circumstances under which John initially takes himself as knowing that the animal is a zebra. I suggested earlier that we interpret such attributions in a way that accords either with (a) our best bet about the nature of the practical circumstances, or (b) internalized justification theory. It seems plausible that most of us would interpret John's first comment in accordance with a best bet about the circumstances. Perhaps, John was engaged in casual conversation while visiting the zoo, in a context in which the Risk/Reward Quotient is not extremely high.

Second, assumption (1) (that Epistemic Standard Variabilism requires reflective access to the process by which we set Epistemic Standards) seems to be reflected in Williamson taking John's *mea culpa* seriously, suggesting that the *mea culpa* reflects the meaning of 'know'. To be clear, it certainly seems plausible that, following Mary's question, John may think and say that he made an error in suggesting that he knows that the animal is a zebra. However, I suggest that John's *mea culpa* may be used as evidence against Epistemic Standard Variabilism *only if assumption (1) is true*, only if John has reflective access. Consider that if John does not have reflective access, then Epistemic Standard Variabilism could be true and John's *mea culpa* could be the result of his not having reflective access. We may test the plausibility of this suggestion with the help of a thought experiment. We may see how the conversation between John and Mary would plausibly have gone in a nearby world in which Epistemic Standard Variabilism is true and in which *he has* reflective access. If John would plausibly *not* have admitted error, then, on *our* world, John's *mea culpa* could well have been the result of his not having reflective access—and not the result of Epistemic Standard Variabilism being false as Williamson suggests.

To be clear, in the world in which the following experiment takes place, *everyone* has full reflective access to the process by which they modulate Epistemic Standards. Also, everyone expresses what is foremost on their mind, and they do so in a straightforward manner, without irony or deceit. Further, let us assume that John and Mary are engaged in ordinary conversation, they are not epistemologists, and that the exchange between them takes place while they are casually strolling through a zoo.

John: I know that this is a zebra.

Mary: How do you know that it isn't a mule cleverly painted to look like a zebra?

John: Are you feeling alright? Here we are together at the zoo, clearly in a context in which the Risk/Reward Quotient is just moderate. We should both associate my last knowledge attribution with an Epistemic Standard that would be met by evidence that has just moderately high predictive validity—which allows errors somewhere perhaps on the order of one in every one hundred trials. However, the frequency with which we might observe an animal that looks exactly like a zebra but that is in fact a painted mule would seem to be a very small fraction of this last frequency—perhaps only one in many millions. So, my evidence meets my properly modulated Epistemic Standard, and *that is how I know that it isn't a painted mule*. Additionally, not only does your question suggest that you have detached yourself from the present practical circumstances; but the Epistemic Standard implied in your question is higher than any I have

encountered in my life, one that is even higher than a scientific standard. I think you must be making a joke.

It seems plausible that on this nearby world John would *not* have assessed his initial knowledge attribution as having been made in error. So, John's *mea culpa* in our world may well have been the result of reflective access being false, and not, as Williamson suggests, the result of Epistemic Standard Variabilism being false. I suggest that we recognize that when John makes his *mea culpa* on our world, he does not know better, and it should not be taken seriously.

7.2.5 Three Experimental Philosophy Studies

Three X-phi studies seem to present experimental evidence that contradicts Epistemic Standard Variabilism: Buckwalter (2010), May et al. (2010), and Feltz and Zarpentine (2010). However, all three studies are based on experiments (each study includes multiple experiments) that all somehow have the same design flaw—with the exception of one of the experiments in one of the studies, which I say a few words about further on. This design flaw invalidates all the experiments that have them.

The design flaw to which I refer produces an observer effect—though a different kind of observer effect than the others I described earlier. The effect is produced by subjects in knowledge cases seeming to set an Epistemic Standard that is lower than what the reader would otherwise expect.

All three X-phi studies test for the extent to which experiment participants, the 'respondents', agree or disagree with the subjects of low stakes and high stakes knowledge cases. The low stakes cases that are presented to respondents are just like the low stakes cases that Contextualists and Subject Sensitive Invariantists present. In fact, the low stakes cases that are used in some of the experiments are taken verbatim from the literature. However, in all three studies, the structure of the high stakes cases presented to respondents is markedly different than those used in the high stakes cases standardly presented in the literature. High stakes cases in the literature present high stakes contexts that seem to elevate Epistemic Standards—which the subjects' evidence fails to meet. So, these subjects deny that they know. (See, for example, Stanley's High Stakes Case.) In contrast, the high stakes cases in the three studies present high stakes contexts that *do not* cause the subjects described in those cases to have higher Epistemic Standards than they might otherwise; and the subjects' evidence, in those

cases, meets their Epistemic Standard. So, the subjects, in both the low stakes and high stakes cases, in all the experiments, in all three studies (again, except one that I say a few words about further on) *take themselves as knowing* to be true the propositions in question. All these experiments then test for the extent to which respondents agree or disagree with those subjects.

So, respondents read cases in which the subjects, in both the low stakes and high stakes cases, claim to know. The researchers suggest that if Epistemic Standard Variabilism is true, then significantly more respondents should agree or strongly agree with the subjects in the low stakes cases than with the subjects in the high stakes cases; or, alternatively, more respondents should disagree or strongly disagree with the subjects in the high stakes cases, than with the subjects in the low stakes cases. However, the principal finding in all three studies is that there is a strong *tendency among respondents to agree with the subjects in both the low stakes and high stakes cases*. This is supposed to indicate that Epistemic Standard Variabilism is false.

However, let us consider the effect on respondents that might be produced by the subjects in the high stakes cases taking themselves as knowing. To get a sense of this, consider that the respondents would likely assume that the subjects in these cases base their Epistemic Standards on information that the respondents do not have. The subjects in the cases are like strangers to the respondents, and the communities in which they live are like communities they do not know. As a result, the respondents' sense of the Risk/Reward Quotient in the cases would seem to depend heavily on best bets. Consider, for example, Stanley's High Stakes Case. The only information that we get about the extent to which Hannah would disvalue failing to deposit the cheque by Saturday is that she has an impending bill coming due, and that it is 'very important' that the cheque be deposited. So, our sense of the cost of a false positive for Hannah is based on a best bet on *exactly* how important it is to her to deposit the cheque. Moreover, we do not get any information about the frequency with which, in Hannah's community, (a) banks are closed on Saturdays, (b) branches are being permanently shut down, and (c) banks alter their operating hours. So, our sense of the Prior Probability of Hannah's bank being open on Saturday is also a best bet.

As a result, respondents should have the sense that the subjects in the cases have a far better sense of the Risk/Reward Quotient than they do. The subjects have a precise sense of the extent of the cost of a false positive; and, they have a sense of the Prior Probability of the questioned proposition being true that is based on accumulated experiences in their own communities. The respondents may feel that they do not have the epistemic right to disagree or

to strongly disagree when the subjects in the high stakes cases attribute knowledge. From a different perspective, subjects taking themselves as knowing, in these high stakes cases, constitutes, for respondents, evidence that weighs against the appropriateness of elevated Epistemic Standards.

However, even if we should not expect respondents to disagree or strongly disagree with the subjects in these high stakes cases (for the reasons I describe); we may well expect that they would *not agree as strongly* with the subjects in the high stakes cases, as they would with the subjects in the low stakes cases. This is precisely what all three studies reveal. It is, in fact, one of the principal findings in the studies. May et al. conclude that while their 'data indicate that . . . raising stakes [does not move] most people from attributing knowledge to denying it' (2010, p.265). However, 'practical interests did . . . affect the . . . level of [respondents'] confidence in the truth of [the subject's knowledge attributions]' (p.270). For example, many respondents may indicate that they 'strongly agree' with the subjects in the low stakes cases, and indicate only that they 'agree' with the subjects in the high stakes cases.

There is one experiment, in one of the studies, that is not described by any of foregoing. The Feltz and Zarpentine (2010) investigation includes four separate experiments. The second through the fourth all include modified high stakes cases of the kind I describe. However, the first experiment does not do so. In fact, the first experiment uses Stanley's Low Stakes Case and Stanley's High Stakes Case verbatim; and they find that in this experiment, but only in this one, that their test results 'offer support for the thesis that practical facts do have some effect on ordinary ascriptions of knowledge' (p.689).

We find in these studies that when an observer effect is included in the analysis of the data, all three studies seem to favour Epistemic Standard Variabilism instead of going against it.

7.2.6 Conclusion

I suggested that the debate between First History and Second History theorists about the validity of Second History theory reduced to whether Epistemic Standard Variabilism is true; and I have found here that core arguments against Epistemic Standard Variabilism fail. So, in my view, the validity of Second History theory remains plausible.

Chapter 8: Gettier-Style Cases

In Gettier-Style Cases, such as Russell's Stopped Clock, the subject has a true belief supported by ordinary justification (see §2.2). However, the evidence that constitutes the justification is faulty, e.g., in Russell's Stopped Clock, the clock is stopped. Further, we do not have the intuition that the subjects in these cases have knowledge—even though they have justified true belief. Some First History theorists find that because these cases feature *ordinary* justification the failure to intuit that the subject has knowledge is appropriate and unsurprising. Recall D.M. Armstrong's observation,

But because possession of such [justification] could not constitute possession of *knowledge* I should have thought it obvious that they are too weak to serve as suitable grounds. It is not surprising, therefore, that Gettier is able to construct examples where a true belief is justified in an ordinary sense of 'justified', but the true belief is clearly not a case of knowledge. (1973, pp.152-153)

My theory suggests that such a response from First History theorists is perfectly appropriate: Their principal aim is to help Broadcasters, i.e., help Altruistic Senders who intend or envisage communicating consequential messages to a large number of Receivers (discussed in §7.1). So, ordinary justification may well not meet their Epistemic Standard.

Even so, under ordinary circumstances (i.e., when we are not in a First History theorist's or Broadcaster's frame of mind), we still do not intuit that the subjects in these cases have knowledge; and my developed theory may seem to suggest that we would. The subjects' Present Evidence would seem to meet the Epistemic Standard that might be set in the context. However, a closer look reveals that my theory *does* predict that we would ordinarily not intuit that the subject in these cases know. Moreover, it explains why this is so. I look at this in this short chapter.

Let me start by suggesting that it is plausible that Knowledge Intuitions are produced by our KIP both in real-life situations and when we read knowledge cases. It seems that when we read knowledge cases, we are able to engage the KIP by either putting ourselves in the knowledge attributor's shoes, or by envisaging ourselves helping the subject (say by sending a helpful message). If we engage the KIP using the former method, then Egoistic Knowledge Intuitions

may be produced; and if by the latter, then Altruistic Knowledge Intuitions may be produced. For example, when we read Stanley's Low Stakes Case we may put ourselves in Hannah's shoes in order to engage the KIP and to produce Knowledge Intuitions. However, it seems that we generally engage the KIP in Gettier-Style Cases by using the latter method, by envisaging ourselves helping the subject. We may help by either confirming or disconfirming the subject's belief. Even so, whether we engage the KIP using the former or the latter method, we employ the Basic Method and determine whether the Adding Evidence Thesis is relevant (under either egoistic or Altruistic assessments) in order to govern the production of Knowledge Intuitions.

Consider the role of the Adding Evidence Thesis in Gettier-Style Cases. For that thesis to be relevant to a reading of a Gettier-Style Case, the reader would need to get a sense of whether gathering additional evidence is a practical possibility, and whether the benefit of a reduction of false positive risk that would follow from gathering that evidence exceeds the cost of gathering it. However, Gettier-Style Cases do not standardly tell us enough about the background circumstances to do this. So, I will set aside the Adding Evidence Thesis and focus only on the Basic Method.

Employing the Basic Method in a reading of a Gettier-Style Case requires the reader to get a sense of the Risk/Reward Quotient. However, the questioned propositions in the Gettier-Style Cases presented in the literature are standardly presented *without enough information* about the practical circumstances for us to get a sense of the Risk/Reward Quotient. Earlier, I described presentations of this kind as being acontextual. (For further discussion of the acontextual presentation of propositions see §7.2.2.) For example, in Russell's Stopped Clock, we do not know why the man needs to know what time it is, or how important it is to him to know that. So, following our ordinary response to acontextually presented examples, our intuitive response to Gettier-Style Cases is based on a best bet about the nature of the context. For example, I assume that the man, in Russell's Stopped Clock, is in an ordinary context, and that the stakes for the man are not terribly high. Moreover, I suggest that the best bets we make when we read Gettier-Style Cases *in general* are that the stakes are not terribly high. (Earlier, in footnote 7, I explain why the Gettier-Style Cases that are presented in the literature do not generally feature high-standard contexts, such as scientific contexts.) So, I suggest that we ordinarily get the sense that the Risk/Reward Quotient in these cases is not terribly high. Following the Basic Method, the reader associates this sense with an Epistemic Standard that is also not terribly high. Finally, under the Basic Method, the reader will assess the strength of the Present Evidence; and if it meets her Epistemic Standard, then Knowledge Intuitions will be produced.

However, as if by divine intervention, the reader of these cases learns that the subject's Present Evidence is faulty. In fact, we learn that the evidence has *no predictive validity whatsoever*. It has no ability at all to predict that which the evidence is purportedly evidence for. (Recall from §4.3.2.1 that my theory is based on the notion that it is the predictive validity of evidence that determines its strength.) In Russell's Stopped Clock, the evidence consists in the time told by a clock that the reader learns has stopped.⁶⁷ Accordingly, the reader accepts that the Present Evidence has no predictive validity. As a result, the reader assesses that the Present Evidence fails to meet her Epistemic Standard, and, following the Basic Method, a Knowledge Intuition is not produced.

In summary, in these cases, we envisage helping the subject; we employ the Basic Method in order to determine whether we would confirm or disconfirm the subject's belief; we make a best bet that the Risk/Reward Quotient is not terribly high; we set an Epistemic Standard that is also not terribly high; we learn that the Present Evidence has no predictive validity and fails to meet our standard; our KIP fails to produce Knowledge Intuitions; and we conclude that we would not confirm the subject's belief.

However, even if the Basic Method explains our response to Gettier-Style Cases, one may worry that the Basic Method is, in fact, more flexible than what I have suggested. Perhaps, it does not need to take the predictive validity of the Present Evidence into account in these cases. After all, the reader learns that the questioned proposition is true (e.g., the time the clock tells is correct) and understands that, under most ordinary circumstances, the Present Evidence would have sufficient predictive validity to meet the Epistemic Standard that is set. Why must we

⁶⁷ This evidence does not have any predictive validity relative to the time of day—because the time told by the stopped clock does not have a causal connection with the actual time of day. Compare this with the counterfactual situation in which the man looks at a working clock. Here, working mechanisms cause the time told by the working clock to track the actual time of day, i.e., to track the daily axial rotation of the earth. So, if the mechanism is, in fact, working reliably, then the axial rotation may be viewed as a causal factor of both the time told by the working clock and the actual time of day. However, under Russell's Stopped Clock, such a causal relation is absent.

Generally, evidence *E* may confer a probability on *p* in virtue of the fact that either (1) *E* indicates the existence of a cause of the state or event described by *p*, (2) *E* indicates the existence of an effect of the state or event described by *p*, or (3) *E* indicates the existence of an effect of a hypothesized causal state or event and *p* describes a different effect of the same hypothesized causal state or event. I note that under possibility (2), when *E* indicates the existence of an effect of the state or event described by *p*, the inductive relation between *E* and *p* may follow from an *abductive* process, i.e., *p* may be the best explanation of the evidence. For example, smoke may confer a high probability on the truth of the proposition that there is a fire.

worry that the Present Evidence fails to have any predictive validity in the extraordinary circumstances that are presented in these cases?

To answer this last question, let us recognize a basic principle of Signal Detection Theory: Systems that respond to equivocal or misleading data in a way that tends to advance their operator's interests will of necessity assess the predictive validity of their evidence. Radar systems that failed to take into account the predictive validity of its evidence (i.e., inputted radar data) could get the system's operators into a great deal of trouble. So too, producing Knowledge Intuitions without taking into account the predictive validity of evidence could be very bad for us. Accordingly, if my story about the evolution by natural selection of the KIP is on the right track, then it would be safe to assume that taking the predictive validity of evidence into account is an innate trait associated with the KIP. In addition to which, I showed (in §4.3.2.1) that it is plausible that we possess just such a trait, i.e., Trait 1, according to which we assess evidential strength with a degree of precision and a frequency of being correct that are sufficient for the application of SDT principles.⁶⁸ From a different perspective, it seems plausible that there would have been sustained selection pressure for this sort of inflexibility in the operation of the KIP.

*

Paul Mosher suggests that 'our understanding [of] exactly what propositional knowledge is essentially involves our having a Gettier-resistant analysis of such knowledge' (2010, p.397). Even so, Gettier-Style Cases do not seem to fit comfortably in any theory of knowledge that has a wide base of support. Nonetheless, not only is my theory of Knowledge Intuitions Gettier-resistant, the theory explains our intuitions in these cases.

Moreover, these cases may be viewed as adding a measure of support to my thesis that the ordinary use of 'know' in knowledge attributions and denials generally depends on the production of Knowledge Intuitions (outlined in §1.2.2)—which thesis contrasts the standard assumption that the use of 'know' reflects a particular concept, definition, or theory of knowledge. Having described a plausible process by which Knowledge Intuitions would be produced (i.e., the process that employs the Basic Method and the Adding Evidence Thesis) we find that it correctly predicts how we use 'know' in Gettier-Style Cases; whereas the theory that

⁶⁸ Nonetheless, if behaving as if p will tend to produce better consequences than that of not behaving as if p even if not- p , then our Epistemic Standard may be so low that Present Evidence that has no predictive validity may meet that standard and Knowledge Intuitions may be produced.

equates knowledge with true belief plus ordinary justification fails to correctly predict our use of 'know' in these cases. Further, while externalistic formulations (e.g., knowledge is true belief plus reliable belief formation) do predict our use of 'know' in these cases, they are otherwise not widely accepted, either by First History theorists (see §2.3) or Second History theorists.

Chapter 9: Conclusion

9.1 Summary

As I suggested at the outset, my aim here was to address the bad state of the epistemological enterprise, in which we find a Gordian knot of competing and incompatible positions with respect to the basic problems of knowledge. I promised to cut the knot with a general theory of the origin, function and triggering-conditions of Knowledge Intuitions.

In Chapter 1, I focused on outlining a conception of Knowledge Intuitions that is suggested by a body of *prima facie* evidence. First, I suggested that Knowledge Intuitions, which are produced in the moment and often in quick succession, reflect our adjudications of Advanced Propositions about world conditions. They help us to picture the world decisively and unambiguously—as being filled with distinct, solid, and weighty objects, filled with people who have good or bad intentions, filled with systems and organizations that may work well or badly, and filled with events that are causally connected though those connections may be non-perceptible. Second, I suggested that the word ‘know’ gives voice to these natural Knowledge Intuitions. Third, I suggested that these adjudications of Advanced Propositions, and the picture of the world that we get from them, constitute a basis for performing deliberate actions and forming attitudes.

In Chapters 3 through 6, I developed a naturalized hypothesis of the function of Knowledge Intuitions, F1; and then developed, on the basis of that hypothesis, a general theory of Knowledge Intuitions that was able to explain why, given the evidence, we have or fail to have those intuitions, and to identify their triggering-conditions. My hypothesis and the general theory were developed on a foundation of (a) biological theory about the function of Cognition (and particularly Godfrey-Smith’s Environmental Complexity Thesis), (b) Signal Detection Theory, and (c) biological theory about altruism. This led to the following simple general prediction: *The Epistemic Standard that we seem to set depends upon (a) the Risk/Reward Quotient, assessed either Altruistically (under the AAT) or egoistically, and (b) whether the conditions contained in the Adding Evidence Thesis have been met, which assessment may also be made either Altruistically or egoistically.* Moreover, because ‘know’ gives voice to those intuitions, this last general prediction applies to the use of ‘know’; it predicts the use of ‘know’.

This last general prediction seemed to hold with all the knowledge cases that I encountered in the literature—along with a good number of non-standard examples. In particular, the general prediction proved consistent with cases standardly used to support Contextualism, Subject Sensitive Invariantism, and Classic Invariantism. Additionally, I found that Gettier-Style Cases (Chapter 8), sceptical challenges (and paradoxes) (§7.2.4), and rational acceptance cases (§4.3.2.4) are all consistent with my developed theory. In fact, not only are Gettier-Style Cases consistent with the theory, the theory explains our intuitions in those cases.

In Chapter 2, I offered two very short histories of epistemology that suggest that the epistemological literature has bifurcated over the last century —creating two separate, distinct, individually coherent and mutually incompatible histories within epistemology, i.e., the First History and Second History. Then, in Chapter 7, I explained the situation on the basis of my developed theory. It seems that much epistemological literature belongs to either one or the other of two separate and distinct research programs; and the research programs powerfully influence the theory that falls under them. They suggest appropriate (a) conceptions of knowledge, (particularly whether or not propositions that are truly known can later prove false), (b) evidential standards required for knowledge, (c) focuses of analysis (particularly, whether to focus on how evidence can turn out to be inadequate or faulty, or focus on the practical circumstances that seem to influence Epistemic Standards), and (d) criteria for successful knowledge theory. I found that the research program that produced the First History influences Classic Invariantist theory and a diverse range of compatible theories; whereas the research program that produced the Second History influences Contextualist, Subject Sensitive Invariantist and Moderate Invariantist theory, and other theories.

Further, the developed theory suggests that under different circumstances, Knowledge Intuitions have different triggering-conditions. I was able to understand the state of the field on the basis of this fact. For any one of several basic positions in epistemology that have standardly been viewed as being incompatible, there is a set of circumstances under which the position works, and other circumstances under which it does not.

Position 1: The meaning of ‘know’ is *not* revealed in ordinary contexts, but it is in high-standard contexts. This position works when we intend or envisage Broadcasting messages that carry appreciable false positive risk for Receivers. (Recall that I defined Broadcasters as Altruistic individuals who intend or envisage passing messages to large numbers of Receivers.) (See §6.4.4 in particular.)

Position 2: The meaning of 'know' is revealed in ordinary contexts, but not in high-standard contexts. This position works in most every circumstance except when we intend or envisage Broadcasting messages that carry appreciable false positive risk for Receivers.

Position 3: Epistemic Standards are influenced by the practical interests of knowledge attributors, but not the subjects of those attributions. This position works when attributors 'egoistically' assess the Risk/Reward Quotient and whether the conditions contained in the Adding Evidence Thesis have been met. (See Chapter 5 in particular.)

Position 4: Epistemic Standards are influenced by the practical interests of the subjects of knowledge attributions, but not the attributors'. This position very often works when attributors, in non-Broadcast situations, Altruistically assess the Risk/Reward Quotient and whether the conditions contained in the Adding Evidence Thesis have been met; and also works in cases that feature first-person attributions or denials, and in which attributors 'egoistically' assess the Risk/Reward Quotient and whether the conditions contained in the Adding Evidence Thesis have been met. (See Chapter 6, except §6.4.4.)

Position 5: The Epistemic Standard is not influenced by the practical interests of any particular individual at all. This position works when we intend or envisage Broadcasting messages that carry appreciable false positive risk for Receivers. (See §6.4.4.)

The picture that emerges from my developed theory suggests that epistemology seems to be in a bad state mainly because its practitioners overgeneralize as a matter of course. They claim universal applicability of theory that in fact applies to only a subset (though usually an important subset) of knowledge cases.

9.2 A Brief Comparison with Moderate Invariantism

I pointed out earlier (at the end of §2.3) that Moderate Invariantism is an approach that succeeds in accounting for our use of 'know' in cases that are used to support Contextualism, and also cases that are used to support Subject Sensitive Invariantism. Moreover, it is Second History theory, as is my theory of Knowledge Intuitions. The theories share the basic characteristics of Second History theory (which I described in §7.1), e.g., instead of focusing on

problems of evidence, they focus on the appearance of variation in evidential standards. So, the approach competes with my theory.

I am not able to offer decisive arguments against Moderate Invariantism. However, I can point to two states of affairs that weigh in favour of my theory. The first is that my theory seems to have greater explanatory and predictive adequacy than does Moderate Invariantism. First, it seems to better both predict and explain the evidential standards that we seem to set: (a) it identifies contextual factors that seem to have a role in determining evidential standards that are not identified by Moderate Invariantism (e.g., the benefit of a true positive, and whether reducing false positive risk by gathering additional evidence is a practical possibility), and (b) it explains that we take these contextual factors into account in an effort to identify when behaving as if p will tend to have more highly valued consequences than would not behaving as if p . Second, it explains the impasse between Contextualists and Subject Sensitive Invariantists. I will not review the account of the impasse given in §6.2.1 here, save to point out that while the latter aims to accommodate cases that feature Altruistic Knowledge Intuitions, the former are more focused on cases that feature Egoistic Knowledge Intuitions. Third, it explains why we have two independent and strikingly different histories in epistemology, each of which reflects a different basic aim, conception of knowledge, focus of analysis, and criterion for success (see Chapter 7). Fourth, it explains our responses to Gettier-Style Cases, which are (surprisingly) not consistent with the view that knowledge is justified, true belief (see Chapter 8). Fifth, it explains the calculated deliberate performance of actions and formation of attitudes—as either following from knowing the principal questioned proposition or from knowing certain secondary propositions (see §4.3.2.4).

The second state of affairs that I would point to is that Moderate Invariantism depends on a move that contradicts an entrenched principle in epistemology, and which my theory does not. (The principle could of course be wrong. However, an awful lot of theory has been based on it, and it seems at least somewhat unlikely that it is wrong.) Moderate Invariantism seems to either set aside or reject the notion that Knowledge Intuitions are a good indicator of knowledge, or of the meaning of 'know'. In DeRose's high stakes bank case (which is presented and discussed in §2.3), when DeRose asserts 'Ok, I don't know that the bank is open on Saturday', it would seem that, in the moment, he would not have intuited knowing that the bank is open, i.e., he would have had some doubt about the matter. However, Moderate Invariantists tell us that DeRose nonetheless knows; and that a claim, by DeRose, to know in this instance, though it may constitute a conversational impropriety, would nonetheless be literally correct. So, it seems that DeRose would not intuit knowing; yet he does know and, if he

made the claim to know, that claim would be correct. In contrast, my theory supports the view that Knowledge Intuitions are a good indicator of knowledge, or of the meaning of 'know'. Thus, I accept that it is legitimate for Contextualists, Subject Sensitive Invariantists and Classic Invariantists to allow their Knowledge Intuitions to inform the development of their theories (except that I suggest that they should not claim universal applicability of theory that in fact applies to only a subset of knowledge cases).

9.3 Cutting the Knot

The general theory of Knowledge Intuitions I have laid out presents a framework of sorts for epistemological investigation. The framework is described by the different circumstances under which Knowledge Intuitions are triggered by different conditions—such circumstances as whether one is focused on one's own practical interests or focused on helping others, and whether one is envisaging communicating to a large audience. The Gordian knot of competing and incompatible positions in epistemology may be cut by viewing those positions through the lens of this framework (as I have in my comments on Positions 1 through 5 above). Doing so, eliminates much of the incompatibility of such positions by limiting their scope in accordance with the framework.

Viewing the field of epistemology through the lens of this framework has additional significant advantages. Doing so allows us to see and appreciate a far greater richness in the field of epistemology than we do otherwise. If one assumes that all but one of a group of competing theories is incorrect, when in fact several of those theories shed light on a range of cases, one will have failed to appreciate the insight offered by those theories that one has dismissed, and one may well also dismiss valuable observations and findings on which those theories are based. Additionally, viewing the field of epistemology through the lens of this framework adds depth because it connects the literature to a picture of the role of knowing in life, and also to a picture of the dissemination of important messages in communities through information economies.

Appendix 1: The Extent to Which the Environmental Complexity Thesis Overgeneralizes

In Chapter 3, I suggested that the primary Teleonomic Function of the KIP depends on the primary Teleonomic Function of Cognition—just as the function of depth perception depends on the general function of physical responses to physical circumstances. I then based my hypothesis of the function of the KIP, F1, on Godfrey-Smith’s hypothesis of the function of Cognition, the ECT, and also my corollary of the ECT, ECT’, which applies in the human case:

ECT: ‘[t]he function of cognition is to enable the individual to deal with environmental complexity’ (Godfrey-Smith, 1998, p.3).

ECT’: The function of Cognition is to create or select specialized behaviours that tend to be more efficacious and efficient in the different world conditions that the individual encounters than would general-purpose behaviours.

However, some theorists worry that the ECT overgeneralizes—largely because human Cognition takes in a wide range of Cognitive traits that are assumed to have different evolutionary histories; and the validity of the inferences I made from the ECT to F1 is undermined to the extent that the ECT overgeneralizes. So, let us see whether it is plausible that the ECT holds over, at least, a very large part of Cognition—so that my inferences may be viewed as being worthy of further investigation.

The task at hand may be reduced to seeing whether it is plausible that nearly all of our Cognitive traits ultimately help us to deal with Environmental Complexity. This may be further reduced to seeing whether nearly all of our Cognitive traits ultimately facilitate behavioural flexibility. However, as I pointed out earlier, we may assume that Cognitive traits that have *only an instrumental function* (i.e., a function that pertains strictly to how the trait is actually utilized) are by-products of traits that have a Teleonomic Function, that evolved by natural selection. So, we may even further reduce the task to *seeing whether it is plausible that nearly all of those Cognitive traits that have a Teleonomic Function facilitate behavioural flexibility*.

*

Karen Neander finds that the ECT does *not* overgeneralize. On my reading, she argues that the ECT points to something like *a necessary ingredient, which becomes a common denominator, in the creation of adaptive pressure for Cognitive traits*:

The ECT concerns the *primary* teleonomic function of cognition: as Godfrey-Smith otherwise expresses it, it purports to sketch an explanation of why we have cognition *at all*. This is to be contrasted with an explanation of *why*, given that we have cognition, we have cognition with this or that particular feature. So—and it is important to note this—the ECT is therefore consistent with many features of cognition being otherwise explained; perhaps . . . serving some subsidiary teleonomic function. (1997, p.569)

And on the matter of these subsidiary functions, she writes,

[T]he ECT seems to exclude very little in the way of an adaptational explanation of cognition: just about anything would qualify. If cognition was selected in one case for foraging for food, and in another case for choosing a mate, and in another case for evading predators, these could all be seen as enabling the agent to deal with environmental complexity. If all the world was food we wouldn't need to tell food from non-food; if all potential mates were the same, there would be nothing to choose between them; and if predators were all entirely predictable, one wouldn't need to be so alert. (p.574)

This seems correct. Nonetheless, it is not *obvious* that this necessary ingredient, i.e., enabling behaviour flexibility, had a role in the creation of adaptive pressure for nearly all of our Cognitive traits. Godfrey-Smith is the first to express the very worry:

[Cognition] might be a conglomerate of traits, located in a tangled nexus of evolutionary forces, different in every case . . . [Even so, this book will treat] a basic cognitive tool-kit as a unit for adaptationist analysis. This assumption about the teleonomic unity of the tool-kit is not regarded as obvious or necessarily true, but as an assumption which opens up a field of possibilities worthy of philosophical exploration. (1998, pp.23-24)

Godfrey-Smith then allows that the ECT does *not* explain Cognition associated with (1) the production of conscious qualitative experiences or qualia, and it does not explain (2) dreams (pp.22-23); and there are other Cognitive traits that have been suggested as not being covered by the ECT. D. M. Walsh finds that '[m]any thoroughly advantageous features of cognition, like the innate knowledge of grammatical structure, may be beneficial precisely because they facilitate inflexibility' (1997, p.616) which contradicts ECT (and particularly ECT'). Also, Kim Sterelny suggests that some traits should be excluded from the ECT. He points out that cognition is an extremely broadly defined trait; and, he argues, the more broadly defined the trait, the less likely we are to being able to construct a plausible history of it, or 'phylogeny'. He writes, 'The more broadly we define swimming, the less likely it has an adaptive explanation, and the harder it will be to use phylogeny to test our adaptive hypothesis' (1997, p.560).

However, Godfrey-Smith's, Walsh's and Sterelny's reservations seem misplaced. Look at Sterelny's first. Sterelny's worry may well be based on a valid general concern: Adaptationist accounts may well tend to get into trouble when the investigated object is a very broadly defined trait. However, under my reading (and, it seems, Neander's), Godfrey-Smith's project is a special case. His aim is not to develop a proper phylogeny of Cognition *per se*. Instead, he is focused on identifying a *necessary ingredient, or a common denominator, in the selection pressure for a wide range of cognitive traits*. This is a far more limited project than that which Sterelny envisions.

Godfrey-Smith's and Walsh's reservations also seem misplaced. They seem not to have considered that many Cognitive traits are well understood as *component* traits of an overarching *composite* trait; and that while some of these component traits may not facilitate behavioural flexibility, the composite trait under which they fall may well do so. I explain.

Certain traits are composite traits, constituted by component traits. For example, many birds have the composite trait of an orangey, or carotenoid-based, plumage coloration, which has at least four distinct component traits: pigment elaboration, patch area, pigment symmetry, and patch area symmetry. Human Cognition itself is a composite trait that includes, as component traits, attention, memory, knowledge, problem-solving, beliefs, appraisals, interpretations, representations and expectations (Gruszka et al., 2010). However, these component traits may themselves be composite traits and have their own component traits. For example, the component trait, problem-solving, is itself a composite trait that includes component traits of abstraction, analogy, brainstorming and lateral thinking. Similarly, a component trait of the

composite trait, depth perception, is the ability to determine the extent to which the eyes converge when they focus on objects.

Critically, component traits do not need to facilitate behavioural flexibility in order to come under the ECT provided that the composite traits under which they fall do so. *The Teleonomic Function of the composite trait produces adaptive pressure for its component traits in virtue of the latter supporting the former.* For example, it is plausible that the composite trait, depth perception, produced adaptive pressure for the component trait, the ability to determine the extent to which the eyes converge when they focus, in virtue of the latter supporting the former. Accordingly, if depth perception can be explained by the ECT, as a trait that facilitates behavioural flexibility, then the ability to determine the extent to which the eyes converge when they focus, can also be explained by the ECT.

Consider again Walsh's suggestion that 'innate knowledge of grammatical structure may be beneficial precisely because [it facilitates] inflexibility' (1997, p.616). However, innate knowledge of grammatical structure would seem to be a component trait of the composite trait, the ability to communicate propositions about world conditions. Moreover, this composite trait facilitates behavioural flexibility, and so falls under the ECT. Accordingly, the component trait too, innate knowledge of grammatical structure, falls under the ECT.

Consider Godfrey-Smith's suggestion that the production of qualia and the production of dreams is not covered by the ECT. These also are plausibly component traits. They seem to facilitate the composite trait of problem-solving, which in turn facilitates behavioural flexibility. Accordingly, these component traits too, fall under the ECT.

These various considerations suggest that the ECT does not greatly overgeneralize; and it is plausible that nearly all of our Cognitive traits that have a Teleonomic Function can be explained as facilitating behavioural flexibility, and fall under the ECT.

Appendix 2: Jason Stanley's Bank Cases

Two cases presented by Jason Stanley (Stanley, 2005, p.4) follow:

Low Stakes

Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paycheques. It is not important that they do so, as they have no impending bills. But as they drive past the bank, they notice that the lines inside are very long, as they often are on Friday afternoons. Realizing that it isn't very important that their paycheques are deposited right away, Hannah says, 'I know the bank will be open tomorrow, since I was there just two weeks ago on Saturday morning. So we can deposit our paycheques tomorrow morning'.

High Stakes

Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paycheques. Since they have an impending bill coming due, and very little in their account, it is very important that they deposit their paycheques by Saturday. Hannah notes that she was at the bank two weeks before on Saturday morning, and it was open. But, as Sarah points out, banks do change their hours. Hannah says, 'I guess you're right. I don't know that the bank will be open tomorrow'.

Appendix 3: Biological Altruism

I suggested that we adjudicate Advanced Propositions using the Basic Method under the Adding Evidence Thesis; and the Basic Method and Adding Evidence Thesis call for assessments of the benefit of true positives and the cost of false positives. These assessments are often made in relation to our own direct practical interests. Even so, the biological research supports the thesis that under certain circumstances these assessments may be made in relation to the practical interests of others—which thesis I call the Altruistic Assessment Thesis (§6.2). The research (a) predicts circumstances under which Altruistic behaviour is evolutionarily stable (i.e., it supports Inclusive Fitness under the varying environmental and social conditions individuals encounter), (b) finds a correspondence between these predicted circumstances and the circumstances under which Altruistic behaviour is observed, and (c) hypothesizes that this Altruistic behaviour is regulated by a Cognitive mechanism that evolved by natural selection. Each of the following four sections discusses an area in which such a correspondence between predicted and observed Altruistic behaviour has been identified. So, each of these sections predicts particular circumstances under which Altruistic behaviour may arise; and thus indicates the circumstances under which it is plausible that the Altruistic Assessment Thesis is applicable.

A.3.1 Helping Kin

According to kin-selection theory, adaptive pressure is produced for traits that cause individuals to help closely-related kin to survive and reproduce. The genes that produce these traits would tend to spread within the community. This suggests that the Altruistic Assessment Thesis may well be applicable when individuals are communicating with closely-related kin.

However, kin-selection theory postulates that the maximum biological cost c that a donor is willing to incur, in order to produce a biological benefit b for a recipient, is directly proportional to the genetic relatedness between donor and recipient. So, genetic relatedness determines the predicted maximum c/b ratio reflected in Altruistic interactions; and the maximum c/b ratio reflected in Altruistic interactions between distantly related individuals is predicted to be very low.

Further, the relatedness between individuals may be precisely quantified, e.g., the coefficient of relatedness r between brothers is .5, and that between cousins is .125. So, the maximum cost C that an individual is willing to incur to produce benefit b for someone else may be precisely quantified: $rb > c$. Individuals would incur one unit of cost to help his brother only if that brother gets two or more units of benefit. This inequality was proposed by W. D. Hamilton (1964) and is generally called 'Hamilton's principle'. J. B. S. Haldane famously demonstrated the principle in a verbal response to a question about whether he would sacrifice his life to save a drowning brother. He would not, though he 'would save two brothers or eight cousins' (cited in McElreath and Boyd, 2007, p.82).

However, it seems that a very large proportion of human interactions is with strangers (Seabright, 2004). So, kin-selection theory may seem to predict that Altruistic human communication would be highly limited, and produced mainly between close-related kin. Even so, Mary Jane West-Aberhard identifies an important range of cases in which kin-selection theory predicts Altruistic behaviour would be directed toward non-closely-related beneficiaries:

The probability of altruism is increased if the beneficiary stands to gain a great deal (e.g., in emergencies), if the cost is low (e.g., if the altruist is . . . in control of an abundant resource), or both; and if the donor is particularly efficient at giving aid or if the beneficiary is particularly efficient at using it, as in the case of the specialized workers and queens of social insects, or if both situations obtain. (1975, p.1)

Many communication cases feature these elements that increase the probability of Altruism. Communicated messages have the potential of being highly beneficial to Receivers, whereas the cost incurred sending those messages may be very low. So this suggests that kin selection theory may apply to wider range of human communication cases than what one might otherwise surmise.

A.3.2 Direct Reciprocation

Direct reciprocation theory relates to donor-recipient interactions in which, as a result of those interactions, (1) the loss in Inclusive Fitness for the donor is less than the gain in Inclusive Fitness for the recipient, (2) there is a good likelihood that the recipient will later reciprocate in an interaction in which, again, the loss for the helper is less than the gain for the one being

helped. Under these circumstances, when a donor helps and the recipient later reciprocates, both parties come out ahead in a 'non-zero-sum game'. Donors seem to adopt the motto, 'I scratch your back, you scratch mine'. To be clear, under direct reciprocation theory, individuals will tend *not* to help or to become a donor, *unless* they assess that their prospective recipient (a) would likely reciprocate if the opportunity arose, and (b) has the resources or the prospect of obtaining the resources that are necessary in order to reciprocate.

The question arises: How do donors assess whether their potential recipients would indeed reciprocate if the opportunity arose? Biologists find that donors of several different species depend on the recipient's history of reciprocating. The term 'co-operator' is used to refer to an individual who has a good history (and also to refer to an individual who has a good history of being a donor). Those who have a bad history in relation to either reciprocating or being a donor are referred to as 'defectors' or 'free-riders'. So, potential donors adopt a simple rule: help if the individual is a co-operator, and do not help if she is a defector.

A reciprocation strategy based on this last rule is referred to as 'Tit-for-tat'; and such a strategy appears to be employed by a number of species—including vampire bats, African vervet monkeys, and more than forty-five species of fish who offer cleaning services to larger fish species in coral reefs.⁶⁹ Moreover, it seems that humans employ Tit-for-tat:

a broad array of experimental and neuroscientific evidence has accumulated over the last two decades supporting the hypothesis that our species' decision-making architecture includes both cognitive and motivational specializations whose design features are specifically tailored to enable gains through direct reciprocity (e.g., detection of defectors and punitive sentiment toward defectors). (Delton et al., 2011, p. 13335)

Even so, we are far from being infallible in our identifications of co-operators and defectors. Fortunately, the risk of error may be significantly reduced by forming twosomes whose members may receive ongoing benefits from ongoing reciprocal interactions. Members of these twosomes seem to adopt the motto, 'I will scratch your back from time to time, and expect that you will scratch mine from time to time'. The research indicates that the behavioural traits that regulate such dyadic relations are evolutionarily stable. Moreover, by the lights of some

⁶⁹ Matt Ridley offers an overview of the interesting history of discovery and investigation in relation to Tit-for-tat (1997, pp.60-78).

researchers, '[t]he most important condition necessary for the evolution of direct reciprocity is that interactions between pairs of agents be sufficiently repeated' (Delton et al., 2011, p.13335). This finding suggests that the Altruistic Assessment Thesis may well be applicable when Senders and their intended or envisaged Receivers have an ongoing relationship that involves reciprocal interactions e.g., one is an employee, and the other her employer (assuming that the employee's work in fact benefits the employer and the employer regularly reciprocates by paying the employee's wages).

A.3.3 Reputation-based Reciprocity

According to the direct reciprocation theory discussed in the previous section, community members tend to prefer helping co-operators to helping defectors. As a result, all other being equal, co-operators should tend to be more successful at eliciting help than defectors. Such a state of affairs suggests that (a) community members keep track of who the co-operators and defectors are, (b) an individual's reputation as either a co-operator or a defector may be a significant factor in their Inclusive Fitness, and, most important to the present discussion, (c) *selection pressure is produced for Cognitive traits that mediate reputation-building behaviour.*

The reputation-building behaviour that biologists standardly suggest evolves primarily consists in the individual bringing it to pass that others in the community learn that she behaves like a co-operator, and do not learn that she behaves like a free-rider. This state of affairs suggests a modification to the direct reciprocation rule that I referred to in the last section, viz., help if the individual is a co-operator, and not if she is a defector. The modification is as follows:

Reputation-Building Rule: Help if the individual is a co-operator and others in the community will find out.⁷⁰

Reputation-Building Rule is predicted to be an evolutionarily stable strategy; and this situation suggests that the Altruistic Assessment Thesis may well be applicable when the Sender's intended or envisaged Receiver is a co-operator, and others in the community will learn about the communication.

⁷⁰ The research in relation to this strategy is reviewed in Nowak and Sigmund (2005).

The extent of the selection pressure for Cognitive traits that mediate reputation-building behaviour depends on the benefits of having a good reputation. Richard Alexander argues that these benefits may extend beyond those that are produced by being a preferred partner for reciprocal interactions. He suggests that,

the beneficent individual may [also] be rewarded with direct compensation from all or part of the group (such as with money or a medal or social elevation as a hero) which, in turn, increases his likelihood of (and that of his relatives) receiving additional perquisites; [and] the beneficent individual may be rewarded by simply having the success of the group within which he behaved beneficently contribute to the success of his own descendants and collateral relatives. (1987, p.94)

Moreover, the benefit of helping co-operators when others in the community will find out extends to being able to avoid punishment. It seems that some community members are willing to incur a cost to directly punish those who fail to help co-operators. Ernest Fehr and Simon Gächter find that,

there is indeed a widespread willingness of the co-operators to punish [those who fail to help co-operators]. Our results indicate that this holds true even if punishment is costly and does not provide any material benefit for the punisher. In addition, we provide evidence that [those who fail to help co-operators] are punished the more heavily the more they deviate from the cooperation levels of the co-operators. [Those who potentially fail to help co-operators], therefore, can avoid or at least reduce punishment by increasing their cooperation levels. (2000, p.980)

The effect of this willingness on the part of co-operators to punish those who fail to help co-operators is to further increase selection pressure for Cognitive traits that mediate reputation-building behaviour—that mediate, in particular, following the Reputation-Building Rule.

A.3.4 Commitment Theory

The recognition of the possibility that co-operators could change their stripes, i.e., become defectors, has the potential of undermining (a) direct reciprocity, including that exhibited in ongoing dyadic relationships (described in A.3.2), and (b) the benefit of having a reputation as a

co-operator (described in A.3.3). Under both situations, potential co-operators may come to doubt whether their help will be reciprocated; and they may, as a result, refuse to help others. Under such circumstances, the non-zero-sum benefits that are produced by reciprocation may be lost.

Commitment theory postulates that individuals have a means by which they may reduce such doubt, and may in the process mitigate the potential loss of benefits caused by that doubt: They may reduce the doubt that others may have about them by (a) allowing themselves to be forced in some way to co-operate in the future, and (b) ensuring that others find about this.

We may commit ourselves by employing a 'commitment device' that effectively forces certain future co-operative behaviour. Ulysses, for example, in Homer's *Odyssey* used a commitment device in the form of a pact that he made with the crew of his ship. The Sirens' beautiful voices lured sailors in passing ships to crash on their island reefs. As Ulysses' ship approached the Sirens' island, Ulysses wanted to hear the Sirens' song without endangering himself or his men. So, as their ship approached the island, he ordered his sailors to tie him to the ship's mast and to put wax in their ears so that they would not hear the Sirens' song. Ulysses' future behaviour would thus be controlled.

Commitment theory postulates that an individual's personality or character functions as a commitment device inasmuch as it may effectively compel future behaviour. An individual who has a helpful character may feel that they must be helpful. So, the inferences we make about the future co-operative behaviour of community members depends both on their reputation as a co-operator (which is based on their history of co-operation and defection), and on their demonstrations of personality or character.

In paradigm cases, displays of character are made by behaving in a way that is consistent with one's character precisely when doing so is costly. (According to Robert Frank, this principle applies not only to positive reputational characters, such as honesty or trustworthiness, but also to negative ones, such as the character of retaliating against hostile actions. If one is to demonstrate that one is a retaliator by character, it is not enough to retaliate when doing so would prove immediately advantageous. One must retaliate when it is costly to do so (1988, pp.71-95).) So, we reduce the potential for others to doubt our future co-operation by

demonstrating that we are a co-operator by character, which we do by co-operating when it is costly to do so.⁷¹

The foregoing suggests that selection pressure would be produced for a Cognitive device that would mediate a co-operative character. Moreover, according to Mary West-Aberhard, such selection pressure may be especially strong. She postulates the existence of potent 'social selection' that emerges out of competitions between conspecifics. This is a 'competition in which an individual must win in contests or comparisons with conspecific rivals in order to gain access to some resource, including (under sexual selection) mates' (1983, p.155). Social selection leads to the evolution of character traits: 'Characters used in social competition can undergo particularly rapid and divergent evolution' (155).

It is plausible that our ancestors competed to attract partners with whom they could have reciprocal interactions, friends, allies, and other kinds of partners. These competitions would have been based on having a co-operative character. Thus, strong selection pressure for having a co-operative character may have been exerted.⁷² These findings suggest that the Altruistic Assessment Thesis may well be applicable when individuals' co-operative character is strong enough to compel them to provide the level of co-operation called for in the circumstances.

⁷¹ Even so, it seems we sometimes fail to distinguish (a) individuals who co-operate in order to create a good reputation, but who may well defect if they believed that the defection would go undiscovered, from (b) those who have a co-operative character, and who, as a result, may well *not* defect even if they believed that the defection would go undiscovered.

⁷² Robert Frank suggests that learning plays an important role in having a co-operative character: 'indoctrination and practice are required for [a range of character traits] to emerge' (1988, p.18). He argues further that such practice sometimes results in Altruistic behaviour with non-kin, in one-shot interactions, even if it is unlikely to enhance one's reputation, e.g., tipping a restaurant server while on holidays.

Appendix 4: Cappelen & Lepore's 'Test (iii)'

Herman Cappelen and Ernie Lepore (2003) perform three tests for the context sensitivity of the truth conditions of knowledge attributions; and report negative findings in all three. However, I showed earlier (in §7.2.2) that the first two are based on interpretations of acontextual examples, which invalidates them. The third test, 'test (iii)', in my view, is valid. However, Cappelen & Lepore do not conduct it properly. The problem follows from using unrepresentative examples (i.e., they make an inductive sampling error). If conducted properly, the test is positive for context sensitivity.

As I suggest, test (iii) is based on a valid principle, which Cappelen & Lepore demonstrate using the following example:

Suppose we know there are two contexts in which 'Yesterday John left' and 'Yesterday Bill left' are true respectively (though we don't know the time of these contexts). It doesn't follow that there is a context in which 'Yesterday John and Bill left' is true.
(p.35)

For example, it may be that yesterday the statement 'Yesterday John left' was true; it may be that today the statement 'Yesterday Bill left' is true; and it does not follow from this that there is a context in which the collective description, 'Yesterday John and Bill left' is true. The failure of the collective description to follow indicates that the two initial statements are context sensitive. This 'is because the semantic value of [the collective description] is determined in one context, and we have no guarantee that that semantic value . . . "captures" . . . the semantic values of [the initial two statements] in those contexts of utterance' (p.35).

Cappelen & Lepore then propose a test for context sensitivity of knowledge attributions based on this principle. I paraphrase a formulation of the test as follows:

Test for C-D: Suppose we know that there are two practical contexts in which 'S₁ knows that *p*' and 'S₂ knows that *p*' are true respectively. If it follows from this that there exists a context in which the collective description 'S₁ and S₂ know that *p*' is true, then we may infer that the two initial knowledge attributions are not context sensitivity.

They find that *it does follow that the collective description is true*, and conclude that knowledge attributions are not context sensitive. Even so, I find this conclusion to be false.

It is possible to demonstrate that their test is *in principle* able to report positive findings for context sensitivity of the truth conditions of knowledge attributions. To see this notice first that there may, *in general*, be *only one context* in which a given context sensitive expression is true. For example, it may well be that ‘Yesterday John left’ is true in *only one* context. Say, that John had spent his entire life up until yesterday in a prison. Yesterday he left; and died the moment he left. Here, it may well emerge that there is only one context under which the expression is true. Now, say that there is only one practical context in which ‘ S_1 knows that p ’ is true and only one practical context in which ‘ S_2 knows that p ’ is true; and these two contexts are different. Call them contexts C_1 and C_2 respectively. Further, the Epistemic Standard that S_1 sets in C_1 is lower than that which she sets in any other context; and the Epistemic Standard that S_2 sets in C_2 is lower than that which she sets in any other context. Moreover, say that S_1 ’s evidence just meets the extremely low Epistemic Standard in C_1 , and S_2 ’s evidence just meets the extremely low Epistemic Standard in C_2 ; and neither S_1 ’s evidence nor S_2 ’s evidence meets the Epistemic Standard that they set respectively in any other context. (Also, S_1 ’s evidence does not meet the Epistemic Standard that she would set in C_2 , and S_2 ’s evidence does not meet the Epistemic Standard that she would set in C_1 .) Under these circumstances, C_1 is the only practical context in which ‘ S_1 knows that p ’ is true, and C_2 is the only practical context in which ‘ S_2 knows that p ’ is true. It is evident that it does *not* follow from this that there exists a context in which the collective description ‘ S_1 and S_2 know that p ’ is true. This seems to suggest that test (iii) may in principle report positive findings for context sensitivity.

As I suggested, Cappelen & Lepore’s false negative finding is the result of their using *unrepresentative examples*. They present two such examples. Each *example* consists of (a) two or more cases in which a subject attributes knowledge in relation to a particular proposition, and (b) a collective description. Here is one of their examples—which includes three cases and a collective description:

- A knows that penguins eat fish.
- B knows that penguins eat fish.
- C knows that penguins eat fish.
- A, B and C know that penguins eat fish.

Cappelen & Lepore flesh out the various cases presented in their two examples by providing a practical context for each subject. One may indeed have a sense of the Risk/Reward Quotient in all the cases included in the two examples. Also, they describe all their subjects' evidence.

They then ask us to assess the truth of the collective description in the two examples, and in the process revise their criterion for context sensitivity: 'If . . . you . . . think it follows that [the collective description is true], then you have evidence 'know' isn't context sensitive' (p.36). This is not quite the same as the question that the test calls for—which is something like, 'Does it follow that there is a context in which the collective description is true?' Also, notice that they ask us to adjudicate the collective description acontextually (see §7.2.2). I suggest that the epistemologists who read their journal article and adjudicate the acontextual collective descriptions in their examples will do so using an Epistemic Standard that is based on an internalized theory of justification.⁷³ Let us say that the Epistemic Standard that they set will be based on internalized theory that is consistent with Earl Conee's suggestion that the 'unvarying standard for truth is very high, but not unreachably high' (2005, p.52).

I do not suggest that Cappelen & Lepore's revised criterion for context sensitivity that I just described in itself violates the integrity of the test. However, I do suggest that they have failed to conduct the test properly insofar as their examples include knowledge cases that are *not* representative of the kinds of cases that Second History theorists focus on. In particular, they give their subjects in these cases evidence that is arguably the best evidence humans are capable of possessing. As a result, it seems to the reader that the subjects would know in pretty much any context. So, it is highly likely that we will be inclined to answer in the affirmative the question that reflects the revised criterion, 'If . . . you . . . think it follows that [the collective description is true], then you have evidence 'know' isn't context sensitive' (p.36).

In one of their examples, the knowledge cases involve attributions that have the form 'Subject A knows that object X is red'. The collective description has the form 'Subjects A, B, etc. all see a red object'. The evidence held by all of the subjects consists of seeing, presumably under good light, that the object is red (p.36). This evidence is about as good as it gets. So, when the reader assesses whether it follows that the collective description is true, she is bound to find that all

⁷³ I discuss such a possibility in §7.2.2. I add that I do not believe that the assessments I refer to here in test (iii) are made on the basis of a best bet, primarily because the initial knowledge attributions that Cappelen & Lepore use in each of their examples are presented in very different practical contexts, suggesting very different Risk/Reward Quotients. As a result, there does not seem to be any one particular best bet suggested.

the subjects have evidence that meet her Conee-like Epistemic Standard. The collective description seems to the reader to be true.

Their other example, the one that involves penguins above, is more complicated. The example includes three cases, one involves the subject A, another B, and the third C. Look at the case involving B first. Cappelen & Lepore tell us that the knowledge attribution, 'B knows that penguins eat fish', is made by B's biology professor. So, the reader would likely assume that B's evidence meets the biology professor's Epistemic Standard. We are told further that Subject C possesses evidence that is so good that it succeeds 'in responding to scepticism' (p.36). So, both B's and C's evidence should easily meet the reader's Conee-like Epistemic Standard; and it should seem to the reader that these two subjects know in pretty much any context.

Before turning to Subject A, notice that if the reader did not already take herself as knowing the proposition p that penguins eat fish, then, after learning that both B and C have either extremely good or *ideal* evidence in favour of p , she too would take herself as knowing that p in pretty much any context.

Turn to Subject A. About A's situation, we are told only that 'A is a five-year-old, and . . . his (non-philosophical mother) . . . truly attributes knowledge to him' (p.35). It would seem to be a good bet that A's mother is a responsible educator for A. So, Cappelen & Lepore have created a situation in which (i) the reader takes herself as knowing that p , (ii) the reader understands that A's mother finds that A knows that p , and (iii) A's mother is likely to be a responsible educator for A. It would seem that under these circumstances, the reader is likely to take it that A's evidence would meet their own Conee-like Epistemic Standard.

So, the reader it likely to take it that all the subjects have evidence that meets their own Conee-like Epistemic Standard; and, so, take it that the collective description is true.

Thus, while 'know' seems to have failed Cappelen & Lepore's revised criterion for context sensitivity in both of their examples (i.e., the reader is likely to think that it follows that the collective description is true in both examples); 'know' fails only because the authors have used unrepresentative examples. Cappelen & Lepore give their subjects evidence that meet an Epistemic Standard that the reader is likely to set given the way in which the examples are presented. For example, the reader would not find that 'know' fails their test if they had presented some cases in which their subjects' evidence only just met a relatively low Epistemic Standard; and if, following a reading of these cases, the reader was likely to have set an

Epistemic Standard when adjudicating the collective descriptions that would not be met by those subjects' evidence. The reader would not have found that the collective description is true, and 'know' would not have failed Cappelen & Lepore's revised criterion. Also, under these circumstances, the reader would not have found that Cappelen & Lepore's initial criterion for context sensitivity was met, i.e., that it follows that there is a context in which the collective description is true. Moreover, I demonstrated earlier that their test is *in principle* able to report positive findings for context sensitivity of the truth conditions of knowledge attributions (which demonstration used knowledge attributions that are true in only one context).

Appendix 5: Defined Terms

Title case terms and capitalized acronyms have the meanings indicated here.

Adding Evidence Thesis *Introduced in §5.2* The thesis that one's Knowledge Intuitions are suspended in the moment (even if one's Present Evidence meets the Epistemic Standard predicted by the Basic Method) if one believes that (a) reducing false positive risk by gathering additional evidence is a practical possibility, and (b) the benefit of this reduction in false positive risk exceeds the cost of gathering the additional evidence.

Advanced Proposition *Introduced in §3.1* Distinguished from Elementary Propositions, Advanced Propositions (a) identify *ambiguous* or *non-perceptible* objects, states or events, (b) locate those objects, states or events in space or time, (c) attribute *ambiguous* or *non-perceptible* properties to objects, states or events, (d) attribute *causes* to states and events, or (e) make *predictions* about future states and events. (This is not intended to be a comprehensive list.) For example, the following statements express Advanced Propositions: (1) the sun will rise in the East tomorrow, (2) my bank will be open this Saturday, (3) she is a loyal employee, (4) my car (which I cannot see at the moment) is parked outside on the street, (5) that car over there is heavy, (6) he has the public's best interests at heart, (7) the bridge is safe, (8) the bird is a gadwall, (9) anthropogenic global warming is real, (10) the politician's policies directly led to the problems to which you referred, (11) the accused is guilty, (12) the watch works fine, (13) that is Jonah's sweater, and (14) Paris is the capital of France.

Altruism (and cognate adjectives and adverbs) *Explained in §6.1* An individual's behaviour or thought (including an adjudication of an Advanced Proposition) is 'Altruistic' if and only if she values neutrally or negatively its direct consequences in relation to her own practical and self-oriented interests, and values positively its consequences in relation to the practical interests of those she intends or envisages helping.

Altruistic Assessment Thesis (AAT) *Introduced at the beginning of Chapter 6* The thesis that we are sometimes focused on helping others; and when we are, our adjudications of relevant Advanced Propositions are based on our assessments of the benefit of true positives and the cost of false positives in relation to the practical interests of those others.

Altruistic Knowledge Intuitions, Egoistic Knowledge Intuitions *Introduced in §3.3. See also §6.2* All Knowledge Intuitions emerge from our valuations of the consequences of performing actions or forming attitudes. Egoistic Knowledge Intuitions emerge from those valuations that are oriented toward our own practical interests; Altruistic Knowledge Intuitions emerge from valuations that are oriented toward the practical interests of individuals who in the moment we intend to help or envisage helping.

Basic Assumption *Introduced in §5.3* A belief that is a constituent of our Worldview and which facilitates the adjudication of Advanced Propositions by giving Present Evidence a basis upon which it may secure predictive validity.

Basic Method *Introduced in §5.1* A procedure by which Advanced Propositions may be adjudicated, leading to the production of either Egoistic Knowledge Intuitions or Altruistic Knowledge Intuitions.

- (1) We notice and keep track of the frequency of the appearance or occurrence of different general kinds of objects, relations, states or events. These frequencies are the basis for the sense we get of the Prior $\Pr(p=TRUE)$ in particular instances. We may then get a sense of the Risk/Reward Quotient by modulating our sense of the cost of a false positive and the benefit of a true positive accordingly, and weigh the numerator of the Risk/Reward Quotient against its denominator.
- (2) We associate by positive correlation the Risk/Reward Quotient with an Epistemic Standard.
- (3) We notice and keep track of the frequency with which different general kinds of objects, relations, states or events correctly predict other different general kinds of objects, relations, states or events. These frequencies are the basis for the sense we get of the strength of the Present Evidence in particular instances.
- (4) We produce Knowledge Intuitions in relation to the questioned proposition if and only if the assessed strength of the Present Evidence meets our Epistemic Standard.

Basic Proposition *Introduced in §5.3* Propositions that are identified as candidates for the content of new Basic Assumptions.

Broadcaster *Introduced in §6.4.4* Senders who intend or envisage passing messages to hundreds, thousands or millions of Receivers. Their messages may be passed orally or by the written word across communities or down from one generation to the next.

Classic Invariantism The view that the evidential standard for knowledge is very high and not dependent on context—though the standard is not so high that knowledge is not a practical possibility.

Cognition, Cognitive *Introduced in §3.2.2* '[T]hose higher mental processes in humans and animals, such as the formation of associations, concept formation and insight, whose existence can only be inferred and not directly observed' (Henderson's Dictionary of Biology, 14th edn, 2008, p.131).

Contextualism *Discussed in §2.3* The view that the truth conditions for knowledge attributions vary across contexts in accordance with the attributor's practical interests.

ECT, ECT' See ENVIRONMENTAL COMPLEXITY THESIS

Egoistic Knowledge Intuitions See ALTRUISTIC KNOWLEDGE INTUITIONS.

Elementary Propositions *Introduced in §3.1* Propositions that (a) identify unambiguous and perceptible objects, states, events or conjunctions of states or events, (b) locate those objects or events in space and time, or that (c) attribute unambiguous and perceptible properties to objects, states or events. (I do not suggest that this is a comprehensive list.) For example, the following statements (made under good observational conditions) express Elementary Propositions: (1) the object is a bird, (2) the cat is on the mat, (3) the tomato is red, and (4) the red billiard ball hit the black one.

End-User *Introduced in §6.4.4* The notion of the End-User arises out of a suggested division of labour with respect to distributions of helpful information. Broadcasters may send messages to a number of Senders at a local level who use those message to fashion derivative messages that are helpful, pertinent, and understandable to local Receivers, who I call 'End-Users'. The latter are thus able to directly benefit from Broadcasters' messages.

Environmental Complexity *Introduced in §3.2* Environmental Complexity is an environmental parameter that is relative to a particular organism. It has to do with whether an organism's Inclusive Fitness is affected differently when it performs particular behaviours or adopts particular Phenotypes under either different environmental conditions that it encounters (i.e., under a type of environmental complexity called 'environmental diversity'), or at different times (i.e., under a type of environmental complexity called 'environmental variability'). If an

organism's behaviour or adopted Phenotype positively affects its Inclusive Fitness under some of environmental conditions it encounters, but not others, or at some times, but not others, its environment is viewed as being complex. For example, the environment for the plant *Hieracium umbellatum* is viewed as complex insofar as the plant benefits significantly by (i) adopting its bushy, broad leaf Phenotype when it grows on rocky, sea-side cliffs, and (ii) adopting its narrow leaf Phenotype when it grows on sand dunes.

Environmental Complexity Thesis (ECT) *Introduced in §3.2.3* Godfrey Smith's hypothesis of the function of Cognition: '[t]he function of cognition is to enable the individual to deal with environmental complexity' (1998, p.3). In §3.2.3, I propose a corollary of the ECT that is more informative and particularly applicable in the human case:

ECT': The function of Cognition is to create or select specialized behaviours that tend to be more efficacious and efficient in the different world conditions that the individual encounters than would general-purpose behaviours.

Epistemic Standard *Introduced in §1.2.3* The standard that one's Present Evidence must meet in relation to a proposition p in order to produce a Knowledge Intuition in relation to p . In §4.3.2.3 I show that Epistemic Standards correlate (positively) with the SDT Index. Further, because I adopt the working hypothesis that the word 'know' gives voice to our Knowledge Intuitions (in §1.2.3), 'Epistemic Standard' also refers to the evidential standard that we seem to apply when we use the word 'know'.

Epistemic Standard Variabilism *Discussed in §7.2* The thesis that the Epistemic Standard reflected in knowledge attributions vary across contexts in accordance with practical circumstances. The thesis is maintained in theories that suggest that it is the practical circumstances of the speaker that determines the Epistemic Standard, e.g., Contextualism; theories that suggest that it is the practical circumstances of the subject that determines the Epistemic Standard, e.g., Subject Sensitive Invariantism; and the theory presented in this thesis that it is the practical circumstances of either the attributor (when Egoistic Knowledge Intuitions are produced) or those of the individual(s) who the attributor intends to help or envisages helping (when Altruistic Knowledge Intuitions are produced) that determines the Epistemic Standard.

F1 *Introduced in §3.2* My primary hypothesis of the function of the intuition of having propositional knowledge in particular cases: In a situation S , the function of intuiting that one

knows that a proposition p is true is to indicate to oneself that, on the total evidence, behaving as if p in S will tend to produce better consequences than that of not behaving as if p in S .

F2 *Introduced in §3.2* My discussion of communication cases is facilitated by identifying a special case of F1 that is specific to the Altruistic Knowledge Intuitions produced in these cases: The function of a Sender's Altruistic Knowledge Intuition in relation to p is to indicate to herself that communicating that p will tend to help her Receiver(s) more than would not communicating that p .

First History, Second History *See §2.2 and Chapter 7* Two separate, distinct and individually coherent histories within epistemology. Epistemological theory builds largely on earlier theoretical developments within the history of which it forms a part. The theory that makes up the First History may be directly or indirectly associated with the aim of helping those who broadcast messages widely identify propositions that are unlikely to prove false. The theory that makes up the Second History may be directly or indirectly associated with the aim of shedding light on the circumstances under which we take ourselves as knowing or not knowing that particular world conditions obtain, and also to shed light on the ordinary use of 'know'.

Gambling Case, Insurance Case *Introduced in §4.3.2.4* In both Gambling Cases and Insurance Cases, agents perform an action or form an attitude even though they do not have enough evidence to take themselves as knowing the main questioned proposition.

In Gambling Cases, agents assume or conclude that the performance of a particular action or forming a particular attitude could produce some *benefit*, and they are not able to rule out the possibility that the action would not produce this benefit. So, they perform that action or form that attitude with the hope that the benefit would be produced, even though they do not know whether it will. For example, I assume that buying a lottery ticket could lead to a win, and I am not able to rule out the possibility that it would not. So, *I buy a lottery ticket even though I do not intuit knowing that I will win.*

In Insurance Cases, agents assume or conclude that they would be harmed by the occurrence of an event. They question whether the harmful event will occur, and find that they cannot rule out the possibility that it will. So, they take an appropriate precaution, performing a particular action or forming a particular attitude. They act though they do not intuit knowing that the event will occur. For example, I assume that it would be bad for me if my house was damaged or destroyed by fire, and I cannot rule out the possibility that it will. So, *I buy fire insurance even though I do not intuit knowing that my house will catch fire.*

Gettier-Style Case See RUSSELL'S STOPPED CLOCK

Inclusive Fitness An organism's ability to pass its genes on to subsequent generations, which correlates with its ability and that of its closest relatives to survive and reproduce.

Insurance Case See GAMBLING CASE

Knowledge Intuitions *Introduced in §1.1* When we have the intuition of knowing that a world condition exists, in that moment we believe without any doubt that it exists and are certain that we would do well to act or form an attitude as if it exists. I argue in §1.2.2 that the ordinary use of 'know' in statements having the form 'I know that p ' give voice to Knowledge Intuitions. These mental states are produced by the Knowledge Intuition Producer or KIP.

Knowledge Intuition Producer (KIP) *Introduced in §3.1* The mechanism that adjudicates Advanced Propositions and produces Knowledge Intuitions in relation to the propositions it adjudicates. It determines whether the strength of the Present Evidence meets an Epistemic Standard, and does so in a way that allows it to fulfil its function, whatever that function may be.

Moderate Invariantism *Discussed in §2.3* The view that the evidential standard for knowledge is invariant across contexts and low enough that we know *most* of what we think we know.


Multiple Receivers Prediction *Introduced in §6.4.3* Senders' Epistemic Standards correlate positively with Receiver numbers—except when there is no appreciable false positive risk for Receivers.

Operator's Question *Introduced in §4.2.1* SDT suggests that, in principle, the optimal detection threshold (e.g., in relation to the strength of the data required to flag a target object), evidentiary threshold or Epistemic Standard may be set on the basis of how the system operator answers the following question: What is the higher priority for me at this time, increasing the probability that detections are in fact target objects (by raising the system's detection threshold), or decreasing the probability of failing to detect target objects (by lowering the system's detection threshold)?

Phenotype *Introduced in §4.3.1* The physical characteristics of an organism which is determined by the interaction of its genotype and its environment.

Present Evidence *Introduced in §4.3.2.3* A type of evidence that is distinguished from the evidence that gives us a sense of Prior $\Pr(p=TRUE)$. Present Evidence is directly connected to the specific proposition and recently gathered, whereas the latter type of evidence is broadly-based or historical in nature; and while Present Evidence is always prominent in the context, we may not be consciously aware that our adjudications are partly based on the latter type of evidence. For example, in Stanley's Low Stakes Case and Stanley's High Stakes Case, Hannah's having seen that the bank was open on Saturday two weeks before constitutes her Present Evidence, whereas the evidence that gives her a sense of the Prior Probability that the bank is open on Saturday is not mentioned. (I do not view these evidence categories as being hard and fast, but rather as family resemblance descriptions of evidence types.) Knowledge Intuitions are produced when one's Present Evidence meets one's Epistemic Standard.

Principal Challenge of the KIP *Introduced in §4.1* In order to infer the operating principles of the KIP, I hypothesize its principle biological challenge as follows: (1) in response to equivocal or deceptive evidence, (2) the KIP produces Knowledge Intuitions, (3) that guide our responses to world conditions in a way that (4) has a tendency to produce consequences that we value. This challenge may be represented as follows:



Prior $\Pr(p=TRUE)$, Prior $\Pr(p=FALSE)$ *Introduced in §4.3.2.3* These constituent variables of the Risk/Reward Quotient are analytical constructs that are analogous to the notions used in SDT of prior probability of signal and of prior probability of noise respectively. Prior $\Pr(p=TRUE)$ refers to the 'natural frequency' (i.e., the experienced or envisioned frequency) with which propositions of the kind that are in question are true; and Prior $\Pr(p=FALSE)$ is the former's probability complement, i.e., the two probabilities add to unity.

Prior Probability *Introduced in §4.3.2.3* Prior Probability refers to the subjective sense of the probability conferred on a questioned proposition by all of one's evidence except one's Present

Evidence. It seems that much of our ordinary sense of Prior Probabilities is derived from our sense of the frequency of world states or events based on accumulated experience.

Relevant Alternative Theory (RA Theory) *Introduced in §2.3* Relevant alternative theory holds that ruling out all the relevant alternatives to a questioned hypothesis h is a sufficient condition for knowing that h . One need not rule out irrelevant alternatives, particularly sceptical alternatives. RA Theory thus allows the possibility that there is much that we know.

Receiver *Introduced in §3.3* Individuals to whom intended or envisaged communications of Senders are directed.

Risk/Reward Quotient *Introduced in §4.2.2* An alternative conceptualization of the SDT Index. One may view the two numerators in the SDT Index as being associated; and, separately, view its two denominators as being associated. The product of the two numerators may be viewed as something like the expected cost that will be incurred due to false positives (based solely on Prior Pr(FALSE)). Similarly, the product of the two denominators may be viewed as something like the expected benefit that will be produced due to true positives (based solely on Prior Pr(TRUE)). So, the SDT Index as a whole may be thought of as *the expected cost of a false positive relative to the expected benefit of a true positive* (based solely on the Prior Pr(FALSE) and Prior Pr(TRUE)). Such a conception of the SDT Index has a well-known counterpart in the academic and professional business world, in the concept of the ‘risk/reward environment’—which refers to something like the expected cost of a bad investment decision relative to the expected benefit of a good investment decision.

$$\text{Risk/Reward Quotient} = \frac{\text{cost of a false positive} \times \text{Prior Pr}(p=\text{FALSE})}{\text{benefit of a true positive} \times \text{Prior Pr}(p=\text{TRUE})}$$

Russell’s Stopped Clock *Introduced in §2.2* Bertrand Russell (2009) presents a counterexample to the thesis that knowledge is true belief supported by ordinary justification: ‘There is the man who looks at a clock which is not going, though he thinks it is, and who happens to look at it the moment when it is right; this man acquires a true belief as to the time of day, but cannot be said to have knowledge’ (2009, p.140). The man’s belief is based on seeing a clock clearly, which is ordinarily viewed as sufficient justification for knowing the time; and his belief is true. Still, he does not have knowledge. Russell thus identifies a class of ordinary cases in which the evidence fails to have the right kind of relation with that which makes the belief true—so that the belief is true only by accident. Cases of this kind are

standardly associated with cases presented by Edmund Gettier (1963). So, I refer to them as Gettier-Style Cases.

SDT Index *Introduced in §4.2.2* The SDT Index is a value that positively correlates with THE optimal detection threshold in a given context:

$$\text{SDT Index} = \frac{\text{cost of a false positive}}{\text{benefit of a true positive}} \times \frac{\text{Prior Pr(noise)}}{\text{Prior Pr(signal)}}$$

The optimal balance between the risk of being overly sceptical and the risk of being overly trusting is achieved by setting detection thresholds that correlate with the SDT Index.

The SDT Index may be applied to the operation of the KIP. We substitute the notion of an optimal detection threshold with the notion of an optimal Epistemic Standard. We may in principle establish an optimal balance between the risk of being overly sceptical and the risk of being overly trusting by adjusting our Epistemic Standards in a way that correlates with the SDT Index. A version of the SDT Index that is appropriate in relation to the KIP is as follows.

$$\text{SDT Index} = \frac{\text{cost of a false positive}}{\text{benefit of a true positive}} \times \frac{\text{Prior Pr(FALSE)}}{\text{Prior Pr(TRUE)}}$$

See also RISK/REWARD QUOTIENT.

Second History See FIRST HISTORY.

Sender *Introduced in §6.2* An individual who intends or envisages communicating with the aim of helping others. They intend or envisage their communication to be Altruistic.

Signal Detection Theory (SDT) *Introduced in §4.1* Certain man-made systems (such as radar systems) and many organisms have the challenge of determining how to respond to evidence about world conditions that is very often equivocal or deceptive, and have the challenge of doing so in such a way as to tend to best advance the system operator's or organism's interests. Signal Detection Theory suggests a general approach to addressing this challenge: strategically adjust the detection threshold or standard of evidence. The threshold or standard may move up and down under different circumstances. Further, SDT predicts the level of the detection threshold or standard of evidence that optimally advances the operator's or organism's interests. In principle, an optimal detection threshold is one that best balances the risk of being

overly sceptical against the risk of being overly trusting under the circumstances. See also SDT Index. (For a full discussion see §4.2.)

Stanley's Low Stakes Case, Stanley's High Stakes Case See Appendix 2.

Subject Sensitive Invariantism (SSI) *Discussed in §2.3* The view that the truth conditions for knowledge attributions vary across contexts in accordance with the practical interests of the subjects of those attributions.

Teleonomic Function *Introduced in §3.2* A function of a trait exhibited by an organism which explains why that trait evolved and persisted over the course of the organism's subsequent evolutionary history.

Trait 1, Trait 2, Trait 3, Trait 4 *Introduced in §4.3* The plausibility that we possess these four traits constitutes (relatively) direct evidence that the production of ordinary Knowledge Intuitions reflects SDT principles and that we do not possess a cognitive device that can trump the KIP in relation to the control of calculated deliberate behaviour.

Trait 1: We assess evidential strength with a degree of precision and a frequency of being correct that are sufficient for the application of SDT principles.

Trait 2: We assess whether our evidence meets our Epistemic Standard in a way that depends on the context.

Trait 3: We seem to set Epistemic Standards that correlate (positively) with the SDT Index.

Trait 4: We do not possess a cognitive device that can trump the KIP in relation to the control of calculated deliberate behaviour.

Worldview *Introduced in §5.3* An individual's Worldview consists of all her Basic Assumptions. This constitutes a basis for ongoing adjudications of Advanced Propositions, and so a basis for performing deliberate actions and forming attitudes.

Works Cited

- ALEXANDER, R. D. 1987. *The Biology of Moral Systems*, Hawthorne, New York, Aldine De Gruyter.
- ALSTON, W. 2001. Internalism and Externalism in Epistemology. In: KORNBLITH, H. (ed.) *Epistemology: Internalism and Externalism*. Malden, Massachusetts: Blackwell Publishers Inc.
- ALSTON, W. P. 1988. The Deontological Conception of Epistemic Justification. *Philosophical Perspectives*, 2, 257-299.
- ARMOR, D. A. & TAYLOR, S. E. 2002. When Predictions Fail: The Dilemma of Unrealistic Optimism. In: GILOVICH, T., GRIFFIN, D. W. & KAHNEMAN, D. (eds.) *Heuristics and Biases: The Psychology of Intuitive Judgment*. Cambridge: Cambridge University Press.
- ARMSTRONG, D. M. 1973. *Belief, Truth and Knowledge*, Cambridge, Cambridge University Press.
- AUSTIN, J. L. 1979. Other Minds. In: URMSON, J. O. & WARNOCK, G. J. (eds.) *Philosophical Papers*. 3rd Edition ed.: Oxford University Press.
- AYER, A. J. 1956. *The Problem of Knowledge*, London, MacMillan.
- BACH, K. 2005. The Emperor's New 'Knows'. In: PREYER, G. & PETER, G. (eds.) *Contextualism in Philosophy: Knowledge, Meaning, and Truth*.
- BANDURA, A. 1989. Human agency in social cognitive theory. *American Psychologist*, 44, 1175-1184.
- BARON, J. 2008. *Thinking and Deciding*, New York, NY, Cambridge University Press.
- BATSON, C. D. 1975. Rational processing or rationalization? The effect of disconfirming information on a stated religious belief. *Journal of Personality and Social Psychology*, 32, 176-184.
- BONJOUR, L. 1985. *The Structure of Empirical Knowledge*, Cambridge, Massachusetts, Harvard University Press.
- BRATMAN, M. 1999. *Faces of intention*, Cambridge, Cambridge.
- BREWER, M. B. 1979. In-group bias in the minimal intergroup situation: A cognitive-motivational analysis. *Psychological Bulletin*, 86, 307-324.
- BROWN, J. 2006. Contextualism and Warranted Assertibility Manoeuvres. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 5, 130, 407-435.
- BUCKWALTER, W. 2010. Knowledge isn't closed on Saturday: A study in ordinary language. *Review of Philosophy and Psychology*, 1, 395-406.
- CAPPELEN, H. & LEPORE, E. 2003. Context Shifting Arguments. *Philosophical Perspectives*, 17, 25-50.
- CHISHOLM, R. M. 1957. *Perceiving: A Philosophical Study*, 1957.
- CHRISMAN, M. 2012. Epistemic Expressivism. *Philosophy Compass*, 7, 118-126.
- CLARK, A. 1987. The Kluge in the Machine*. *Mind & Language*, 2, 277-300.
- CLIFFORD, W. K. 1876. The Ethics of Belief. *The Contemporary Review*, 29, 289-309.
- COHEN, L. J. 1989. Belief and Acceptance. *Mind*, 98, 367-389.
- COHEN, S. 1988. How to be a Fallibilist. *Philosophical Perspectives*, 581-605.
- COHEN, S. 1999. Contextualism, Skepticism, and the Structure of Reasons. *Nous Supplement: Philosophical Perspectives*, 13, 57-89.
- CONEE, E. 2005. Contextualism Contested. In: STEUP, M. & SOSA, E. (eds.) *Contemporary Debates in Epistemology*. Malden, MA: Blackwell Publishing.
- CONEE, E. & FELDMAN, R. 2004. Evidentialism. In: CONEE, E. & FELDMAN, R. (eds.) *Evidentialism: Essays in Epistemology*. Oxford: Oxford University Press.
- CRAIG, E. 1999. *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*, Oxford, Oxford University Press.
- DAMASIO, A. R. 1994. *Descartes' Error: Emotion, Reason, and the Human Brain*, New York, New York, Avon Books.

- DELTON, A. W., KRASNOW, M. M., COSMIDES, L. & TOOBY, J. 2011. Evolution of direct reciprocity under uncertainty can explain human generosity in one-shot encounters. *Proceedings of the National Academy of Sciences*, 108, 13335-13340.
- DENNETT, D. C. 1995. *Darwin's Dangerous Idea*, New York, Touchstone.
- DENNETT, D. C. 1996. *Kinds of Minds: Toward an Understanding of Consciousness*, New York, Basic Books.
- DEROSE, K. 1992. Contextualism and Knowledge Attributions. *Philosophy and Phenomenological Research*, 52, 913-929.
- DESCARTES, R. 1996. Meditations on First Philosophy. In: COTTINGHAM, J. (ed.) *Meditations on First Philosophy*. Cambridge: Cambridge University Press.
- DOUGHERTY, T. 2011. Introduction. In: DOUGHERTY, T. (ed.) *Evidentialism and its Discontents*. Oxford: Oxford University Press.
- DRETSKE, F. 1981. The Pragmatic Dimension of Knowledge. *Philosophical Studies*, 40, 363-378.
- ELGIN, C. 2007. Understanding and the Facts. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 132, 33-42
- FANTL, J. & MCGRATH, M. 2002. Evidence, Pragmatics, and Justification. *The Philosophical Review*, 111, 67-94.
- FEHR, E. & GÄCHTER, S. 2000. Cooperation and Punishment in Public Goods Experiments. American Economic Association.
- FELDMAN, R. 1988. Rationality, Reliability and Natural Selection. *Philosophy of Science*, 55, 218-227.
- FELTZ, A. & ZARPENTINE, C. 2010. Do You Know More When It Matters Less? *Philosophical Psychology*, 23, 683-706.
- FISHER, J. C. 2006. *Pragmatic Conceptual Analysis*. Doctor of Philosophy, University of Arizona.
- FISKE, S. T. & TAYLOR, S. E. 1984. *Social Cognition*, Reading, MA, Addison-Wesley.
- FRANK, R. H. 1988. *Passions Within Reason*, New York, W. W. Norton & Company.
- FRIGG, R. & HARTMANN, S. 2012. Models in Science. *The Stanford Encyclopedia of Philosophy* [Online]. Available: <http://plato.stanford.edu/archives/fall2012/entries/models-science/>.
- GALLISTEL, C. R. 1990. *The Organization of Learning*, Cambridge, Massachusetts, MIT Press.
- GETTIER, E. L. 1963. Is Justified True Belief Knowledge? *Analysis*, 23, 121-23.
- GODFREY-SMITH, P. 1991. Signal, Decision, Action. *Journal of Philosophy*, 88, 709-722.
- GODFREY-SMITH, P. 1994. A Modern History Theory of Functions. *Nous*, 28, 344-362.
- GODFREY-SMITH, P. 1998. *Complexity and the Function of Mind in Nature*, Cambridge, U.K., Cambridge University Press.
- GOERTZEL, T. 1992. *Turncoats and True Believers: The Dynamics of Political Belief and Disillusionment*, Buffalo, N.Y., Prometheus Books.
- GOLDMAN, A. I. 1976. Discrimination and Perceptual Knowledge. *Journal of Philosophy*, 18, 771-791.
- GOLDMAN, A. I. 1999. Internalism Exposed. *The Journal of Philosophy*, 96, 271-293.
- GOULD, S. J. & LEWONTIN, R. C. 1979. The Spandrels of San Marco and the Panglossian Paradigm: A Critique of the Adaptationist Programme. *Proceedings of the Royal Society of London B: Biological Sciences*, 205, 581-598.
- GOULD, S. J. & VRBA, E. S. 1982. Exaptation-A Missing Term in the Science of Form. *Paleobiology*, 8, 4-15.
- GRECO, J. 2003. Knowledge as Credit for True Belief. In: DEPAUL, M. & ZAGZEBSKI, L. (eds.) *Intellectual Virtue: Perspectives from Ethics and Epistemology*. Oxford: Oxford University Press.
- GRICE, H. P. 1989. Logic and Conversation. *Studies in the Way of Words*. Cambridge, Mass.: Harvard University Press.
- GRIFFIN, D. W. & TVERSKY, A. 2002. The Weighing of Evidence and the Determinants of Confidence. In: GILOVICH, T., GRIFFIN, D. W. & KAHNEMAN, D. (eds.) *Heuristics and Biases*. New York: Cambridge University Press.

- GRONDIN, S. 2016. *Psychology of Perception*, Switzerland, Springer International Publishing.
- GRUSZKA, A., MATTHEWS, G. & SZYMURA, B. 2010. *Handbook of individual differences in cognition: Attention, memory, and executive control.*, New York, Springer Verlag.
- HAMBOURGER, R. 1987. Justified Assertion and the Relativity of Knowledge. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 51, 241-269.
- HAMILTON, W. D. 1964. The genetical evolution of social behaviour. I, II. *Journal of Theoretical Biology*, 7, 1–52.
- HARVELL, C. D. 1986. The Ecology and Evolution of Inducible Defenses in a Marine Bryozoan: Cues, Costs, and Consequences. *The American Naturalist*, 128, 810-823.
- HASELTON, M. G. & NETTLE, D. 2006. The paranoid optimist: An integrative evolutionary model of cognitive biases. *Personality and Social Psychology Review*, 10, 47-66.
- HASELTON, M. G., NETTLE, D. & ANDREWS, P. W. 2005. The Evolution of Cognitive Bias. In: BUSS, D. M. (ed.) *The handbook of evolutionary psychology*. Hoboken, NJ US: John Wiley & Sons Inc.
- HASLANGER, S. 1999. What knowledge is and what it ought to be: Feminist values and normative epistemology. *Philosophical Perspectives*, 13, 459-480.
- HAWTHORNE, J. & STANLEY, J. 2008. Knowledge and Action. (Series: Epistemic Norms, Part Two). *Journal of Philosophy*. O, 105, 571-590.
- HENDERSON'S DICTIONARY OF BIOLOGY 2008. Cognition. In: LAWRENCE, E. (ed.) *Henderson's Dictionary of Biology*. 14th ed. Essex, England: Pearson Education Limited.
- Rear Window*, 1954. Directed by HITCHCOCK, A. Universal City, CA: Paramount Pictures.
- HORN, D. M., KOPLAN, K. E., SENESE, M. D., ORAV, E. J. & SEQUIST, T. D. 2014. The Impact of Cost Displays on Primary Care Physician Laboratory Test Ordering. *Journal of General Internal Medicine*, 29, 708-714.
- JACKSON, R. E. & CORMACK, L. K. 2007. Evolved Navigation Theory and the Descent Illusion. *Perception & Psychophysics*, 353-362.
- JAMES, W. 1966. *The Sentiment of Rationality. Essays in Pragmatism*. New York: Hafner Publishing Company.
- KAPPEL, K. 2010. On Saying that Someone Knows: Themes from Craig. In: HADDOCK, A., MILLAR, A. & PRITCHARD, D. (eds.) *Social Epistemology*. New York: Oxford University Press.
- KELP, C. 2011. What Is the Point of "Knowledge" Anyway? *Episteme*, 8, 53-66.
- KORNBLITH, H. 1983. Justified Belief and Epistemically Responsible Action. *The Philosophical Review*, 92, 33-48.
- LEHRER, K. 1971. Why Not Scepticism? *Philosophical Forum*, 71, 283-298.
- LEHRER, K. 1979. The Gettier Problem and the Analysis of Knowledge. In: PAPPAS, G. S. (ed.) *Justification and Knowledge*. London, England: D. Reidel Publishing Company.
- LEONARD, A. S., DORNHAUS, A. & PAPA, D. R. 2011. Flowers help bees cope with uncertainty: signal detection and the function of floral complexity. *Journal of Experimental Biology*, 214, 113-121.
- LEWIS, C. I. 1929. *Mind and the World Order: Outline of a Theory of Knowledge*, New York, Charles Scribner's Sons.
- LEWIS, C. I. 1946. *An Analysis of Knowledge and Valuation*, La Salle, Illinois, The Open Court Publishing Company.
- LEWIS, D. 1979. Scorekeeping in a Language Game. *Journal of Philosophical Logic*. AG, 79, 339-359.
- LEWIS, D. 1996. Elusive Knowledge. *Australasian Journal of Philosophy*, 74, 549-567.
- LORD, C. G., ROSS, L. & LEPPER, M. R. 1979. Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology*, 37, 2098-2109.
- MACLAURIN, J. 1998. How to defeat complexity. *Australasian Journal of Philosophy*, 76, 491-501.

- MARCUM, J. I. 1947. A Statistical Theory of Target Detection by Pulsed Radar. *The Research Memorandum*, 90.
- MAUDLIN, T. 2015. *Why Physics Needs Philosophy* [Online]. Nova. Available: <http://www.pbs.org/wgbh/nova/blogs/physics/2015/04/physics-needs-philosophy/> [Accessed].
- MAY, J., SINNOTT-ARMSTRONG, W., HULL, J. G. & ZIMMERMAN, A. 2010. Practical interests, relevant alternatives, and knowledge attributions: An empirical study. *Review of Philosophy and Psychology*, 1, 265-273.
- MCCARLEY, J. S. & BENJAMIN, A. S. 2013. Bayesian and Signal Detection Models. In: LEE, J. D. & KIRLIK, A. (eds.) *The Oxford Handbook of Cognitive Engineering*. Oxford: Oxford University Press.
- MCELREATH, R. & BOYD, R. 2007. *Mathematical Models of Social Evolution: A guide for the Perplexed*, Chicago, the University of Chicago Press.
- MOORE, G. E. 1959. Four Forms of Scepticism. *Philosophical Papers*. London: George Allen & Unwin.
- MORAN, N. A. 1992. The evolutionary maintenance of alternative phenotypes. *American Naturalist*, 139, 971-989.
- MOSER, P. 2010. Gettier Problem. In: DANCY, J., SOSA, E. & STEUP, M. (eds.) *A Companion to Epistemology, Second Edition*. Malden, MA: Wiley-Blackwell.
- NEANDER, K. 1997. The Function of Cognition: Godfrey-Smith's Environmental Complexity Thesis. *Biology and Philosophy*, 12, 567-580.
- NEURATH, O. V. 1932/1983. In: COHEN, R. S. & NEURATH, M. (eds.) *Philosophical Papers 1913-1946*. Dordrecht: Reidel.
- NEURATH, O. V. 1944. Foundations of the Social Sciences. In: NEURATH, O. V., CARNAP, R. & MORRIS, C. (eds.) *International Encyclopedia of Unified Science*. Chicago: University of Chicago Press.
- NOREM, J. K. 2001. *The Positive Power of Negative Thinking: Using Defensive Pessimism to Manage Anxiety and Perform at Your Peak*, Cambridge, MA, Basic Books.
- NOWAK, M. A. & SIGMUND, K. 2005. Evolution of indirect reciprocity. *Nature*, 437, 1291-1298.
- NOZICK, R. 1981. *Philosophical Explanations*, 1981.
- PITTENDRIGH, C. S. 1958. Adaptation, natural selection, and behavior. In: SIMPSON, A. R. A. G. (ed.) *Behavior and Evolution*. New Haven: Yale University Press.
- PLANTINGA, A. 1990. Justification in the 20th Century. *Philosophy and Phenomenological Research*, 1, 45-71.
- PLANTINGA, A. 1993. *Warrant: The Current Debate*, New York, Oxford University Press.
- POLLOCK, J. L. 1986. *Contemporary Theories of Knowledge*, Totowa, NJ, Rowman and Littlefield.
- PRICE, C. 2002. Rationality, Biology and Optimality*. *Biology and Philosophy*, 17, 613-634.
- PRICE, H. H. 1932. *Perception*, London, Methuen & Co. Ltd.
- PRICE, J. J. 2013. Why is birdsong so repetitive? Signal detection and the evolution of avian singing modes. *Behaviour*, 150, 995-1013.
- QUILLIAN, L. & PAGER, D. 2001. Black Neighbors, Higher Crime? The Role of Racial Stereotypes in Evaluations of Neighborhood Crime. *American Journal of Sociology*, 107, 717-767.
- QUINE, W. V. O. 1960. *Word and Object*, THE MASSACHUSETTS INSTITUTE OF TECHNOLOGY.
- REYNOLDS, S. L. 2002. Testimony, Knowledge, and Epistemic Goals. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 110, 139-161.
- RIDLEY, M. 1997. *The Origins of Virtue: Human Instincts and the Evolution of Cooperation*, New York, Viking Penguin.
- RUSSELL, B. 1914. *Our Knowledge of the External World*, London, The Open Court Publishing Company.
- RUSSELL, B. 2009. *Human Knowledge: It's Scope and Limits*, New York, Routledge.
- RYSIEW, P. 2001. The Context-Sensitivity of Knowledge Attributions. *Nous*, 35, 477-514.
- RYSIEW, P. 2005. Contesting Contextualism. *Grazer Philosophische Studien*, 69, 51-69.

- RYSIEW, P. 2012. Epistemic Scorekeeping. In: BROWN, J. & GERKEN, M. (eds.) *Knowledge Ascriptions*. Oxford: Oxford University Press.
- SAGE, J. 2004. Truth-Reliability and the Evolution of Human Cognitive Faculties. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 117, 95-106.
- SEABRIGHT, P. 2004. *The Company of Strangers: A Natural History of Economic Life*, Princeton, Princeton University Press.
- SOBER, E. 1997. Is the Mind an Adaptation for Coping with Environmental Complexity? *Biology and Philosophy*, 12, 539-550.
- SRIPADA, C. S. & STANLEY, J. 2012. Empirical Tests of Interest-Relative Invariantism. *Episteme*, 9, 3-26.
- STANLEY, J. 2005. *Knowledge and Practical Interests*, New York and Oxford, Oxford University Press.
- STEFANUCCI, J. K. & PROFFITT, D. R. 2009. The roles of altitude and fear in the perception of heights. *Journal of Experimental Psychology: Human Perception and Performance*, 424-438.
- STEPHENS, C. L. 2001. When is it Selectively Advantageous to Have True Beliefs? Sandwiching the Better Safe than Sorry Argument. *Philosophical Studies*, 105, 161-189.
- STERELNY, K. 1997. Where Does Thinking Come From? A Commentary on Peter Godfrey-Smith's Complexity and the Function of Mind In Nature. *Biology and Philosophy*, 12, 551-566.
- STERELNY, K. 2003. *Thought in a Hostile World: The Evolution of Human Cognition*, Malden, MA, Blackwell Publishing.
- STINE, G. C. 1976. Skepticism, Relevant Alternatives, and Deductive Closure. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*. AP, 76, 249-260.
- STROUD, B. 2001. Understanding Human Knowledge in General. In: KORNBLITH, H. (ed.) *Epistemology: Internalism and Externalism*. Oxford: Blackwell Publishers.
- TAYLOR, S. E. & BROWN, J. D. 1988. Illusion and well-being: A social psychological perspective on mental health. *Psychological Bulletin*, 103, 193-210.
- UNGER, P. 1975. *Ignorance: A Case for Scepticism*, London, Oxford University Press.
- UNGER, P. 1986. The Cone Model of Knowledge. *Philosophical Topics*, 14, 125-178.
- VAGNONI, E., LOURENCO, S. F. & LONGO, M. R. 2012. Threat modulates perception of looming visual stimuli. *Current Biology*, 22, R826-R827.
- WALSH, D. M. 1997. Peter Godfrey-Smith Complexity and the Function of Mind in Nature. *British Journal for the Philosophy of Science*, 48, 613-617.
- WEST-EBERHARD, M. J. 1975. The Evolution of Social Behavior by Kin Selection. *The Quarterly Review of Biology*, 50, 1-33.
- WEST-EBERHARD, M. J. 1983. Sexual Selection, Social Competition, and Speciation. *The Quarterly Review of Biology*, 58, 155-183.
- WILLIAMS, G. C. 1966. *Adaptation and Natural Selection*, Princeton, Princeton University Press.
- WILLIAMSON, T. 2005. Contextualism, Subject-Sensitive Invariantism and Knowledge of Knowledge. *The Philosophical Quarterly*, 55, 213-235.