

**University of Alberta**

**Luck and its Effects on Knowledge**

by

**Kenneth Alan Boyd**



A thesis submitted to the Faculty of Graduate Studies and Research  
in partial fulfillment of the requirements for the degree of

**Master of Arts**

Department of Philosophy

Edmonton, Alberta

Fall 2007



Library and  
Archives Canada

Bibliothèque et  
Archives Canada

Published Heritage  
Branch

Direction du  
Patrimoine de l'édition

395 Wellington Street  
Ottawa ON K1A 0N4  
Canada

395, rue Wellington  
Ottawa ON K1A 0N4  
Canada

*Your file* *Votre référence*

*ISBN: 978-0-494-33115-6*

*Our file* *Notre référence*

*ISBN: 978-0-494-33115-6*

#### NOTICE:

The author has granted a non-exclusive license allowing Library and Archives Canada to reproduce, publish, archive, preserve, conserve, communicate to the public by telecommunication or on the Internet, loan, distribute and sell theses worldwide, for commercial or non-commercial purposes, in microform, paper, electronic and/or any other formats.

The author retains copyright ownership and moral rights in this thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without the author's permission.

#### AVIS:

L'auteur a accordé une licence non exclusive permettant à la Bibliothèque et Archives Canada de reproduire, publier, archiver, sauvegarder, conserver, transmettre au public par télécommunication ou par l'Internet, prêter, distribuer et vendre des thèses partout dans le monde, à des fins commerciales ou autres, sur support microforme, papier, électronique et/ou autres formats.

L'auteur conserve la propriété du droit d'auteur et des droits moraux qui protègent cette thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation.

---

In compliance with the Canadian Privacy Act some supporting forms may have been removed from this thesis.

Conformément à la loi canadienne sur la protection de la vie privée, quelques formulaires secondaires ont été enlevés de cette thèse.

While these forms may be included in the document page count, their removal does not represent any loss of content from the thesis.

Bien que ces formulaires aient inclus dans la pagination, il n'y aura aucun contenu manquant.

  
**Canada**

## **Abstract**

This project is an examination of how luck has an effect on our ability to acquire knowledge, and how evidence can qualify a lucky belief as knowledge-worthy. Many contemporary theories make the concept of epistemic luck ubiquitous, and if luck is to be understood as predominantly malignant, then we are forced to accept that most of the beliefs we have are simply lucky to be true and not knowledge-worthy. It is my contention that through an accumulation of supporting evidence we are able to use luckily true beliefs responsibly in rational belief formation. In other words, although the effects of luck are able to prevent a belief from being known, epistemic luck is not as harmful as has traditionally been portrayed.

## **Acknowledgements**

I must acknowledge first and foremost Dr. Adam Morton, who went above and beyond the call of duty as a supervisor, and without whom this work would not have been possible. Also deserving more recognition than space allows are my fellow graduate students at the University of Alberta, especially Victor Kumar, Julian Fowke, Octavian Ion and Yual Chiek, for their support and friendship. I'd also like to thank Dr. Wesley Cooper, Dr. Bruce Hunter, and Dr. Alinda Friedman for their role as committee members, and, of course, my family, especially my sister Jessica, who has always believed in me, even when I haven't.

# **Luck and its Effects on Knowledge**

<b>Introduction</b>	<b>1</b>
<b>Chapter 1: Does Epistemic Luck Wear Off?</b>	<b>3</b>
<b>Chapter 2: Rejecting the Incompatibility Thesis: How Lucky Beliefs Can Lead to Knowledge</b>	<b>18</b>
<b>Chapter 3: Another Look at Epistemic Luck: Is Modal Rarity Sufficiently Luck-Determining?</b>	<b>32</b>
<b>Chapter 4: Barn Façades and Car Crashes</b>	<b>48</b>
<b>Chapter 5: The Effects of Evidence</b>	<b>60</b>
<b>Chapter 6: Lotteries Don't Fade</b>	<b>73</b>
<b>Final Thoughts: Revisiting the Incompatibility Thesis</b>	<b>82</b>
<b>Bibliography</b>	<b>84</b>

## Detailed Table of Contents

<b>Introduction</b>	1
<b>Chapter 1: Does Epistemic Luck Wear Off?</b>	3
• <i>A Working Conception of Luck</i>	3
• <i>The Stopped Clock Example</i>	6
• <i>Luck Wearing Off: The Job Interview Examples</i>	7
• <i>The Principle of Confirmation</i>	10
<b>Chapter 2: Rejecting the Incompatibility Thesis: How Lucky Beliefs Can Lead to Knowledge</b>	18
• <i>Setting the Stage: The Identical Twins</i>	19
• <i>Judy, Trudy, and Fading</i>	26
• <i>The Battle of Hastings</i>	28
• <i>Another Look at “Knowledge-Harmful” and “Knowledge-Harmless”</i>	31
<b>Chapter 3: Another Look at Epistemic Luck: Is Modal Rarity Sufficiently Luck-Determining?</b>	32
• <i>Narrowing the Scope</i>	35
• <i>Reflective Epistemic Luck</i>	38
• <i>Angst and Value</i>	43
<b>Chapter 4: Barn Façades and Car Crashes</b>	48
• <i>Car Crashes</i>	48
• <i>After the Barn Façade</i>	50
• <i>Digging Deeper: Relevance as an Internal or External Phenomenon</i>	56
<b>Chapter 5: The Effects of Evidence</b>	60
• <i>Fading and Propinquity</i>	60
• <i>Activation Levels</i>	63
• <i>Implicit and Explicit Evidence</i>	67
• <i>Unfortunate Consequences?</i>	71
<b>Chapter 6: Lotteries Don’t Fade</b>	73
• <i>Outlining the Problem</i>	73
• <i>Transient and Permanent Alternatives</i>	76
<b>Final Thoughts: Revisiting the Incompatibility Thesis</b>	82
<b>Bibliography</b>	84

## List of Figures

Figures a – c: The Sheep Spotter	14
Figure 1: Implicit evidence provides support for a knowledge-worthy true belief	22
Figure 2: Insufficient implicit evidence does not provide enough support for a knowledge-worthy true belief	23
Figure 3: Jane forms her initial belief concerning the time	27
Figure 4: Jane forms a later belief, after confirming evidence	27
Figure 5: Nader’s belief based upon evidence given by both his teacher and his father.	30
Figure 6: Figure 6: Nader’s belief based solely on the testimony of his father	30
Figure 7: Considering all possible worlds, those in which John’s house is on fire (F) is minute compared to those in which it is not on fire (NF). Thus in this situation, John’s belief that his house is on fire, B(F), is only lucky to be true and as a result not knowledge- worthy.	37
Figure 8: After new evidence is presented, the grayed circle represents the scope of possible worlds now under consideration. At this scope, B(F) becomes knowledge-worthy.	37
Figure 9: Formation of the belief of p: there is a real barn, from the viewpoint of the agent, before learning of his presence in barn façade country.	52
Figure 10: From the viewpoint of the informed agent, the salience of barn façades makes the original belief true only as a matter of luck.	52

Figure 11: The belief is knowledge-worthy when the background assumptions (grayed) and their knowledge-preventing effects (grayed arrows) are safely ignored.	62
Figure 12: The possibility of being in a car accident is made salient and knowledge-harmful by the experience of a car accident.	63
Figure 13: Unconscious-implicit evidence of a car accident being unlikely defeats the knowledge defeater P(will be in a car accident).	69
Figure 14: Potential defeaters (gray) for the belief in p: the coin will come down heads on the next toss, as modeled on previous examples presented thus far.	78
Figure 15: Potential defeaters (gray) for the belief in p: I will lose the lottery, as modeled on previous examples presented thus far.	79



## **Introduction**

“Epistemic luck permeates the human condition whether for good or ill.”

- Linda Zagzebski

We feel that we know a lot. Unfortunately many modern conceptions of knowledge make knowing a lot overly difficult (as the skeptics believe), uncontrollably transient (as implied by John Hawthorne), subject to the whim of suggestion or seemingly unrelated circumstances (as Duncan Pritchard’s conception of luck seems to claim) or, in general, requiring abilities far beyond the capacities of most rational agents. These flaws are representative of invariant theories of knowledge in general, that is, theories that make it a definite fact about a person whether she knows a particular proposition regardless of context, for while we feel we know a lot, additional intuitions, helped along by Edmund Gettier, tell us that we have to work for the knowledge to which we feel we are entitled. After Gettier came the introduction of the notion of epistemic luck to the general philosophical consciousness and the immediate consequence that having a justified, true belief is not itself sufficient for knowledge. As a result, a great deal of the epistemological work in the aftermath has been aimed at attempting to find the elusive “fourth ingredient” that makes it perfectly clear which beliefs are knowledge-worthy, and which are not, reformulating the existing recipe in hope of more intuitive results, or rejecting the notion of a knowledge recipe altogether.

In this investigation I am not so interested in completing the knowledge formula as I am in examining just how knowledge-harmful the notion of luck really is. As agents with only a limited mental and epistemological capacity, we are often lucky, or flat out

wrong, when forming beliefs; nevertheless, we often consider ourselves bearers and sources of knowledge-worthy beliefs. We often base beliefs upon sketchy or incomplete evidence, yet often we feel that we still know, or are, at least, capable of knowing in a reliable manner. With this general feeling in mind, I explore in the first chapter a phenomenon I refer to as “fading”, where if we happen to have a belief that is true simply as a matter of luck, the belief can become knowledge-worthy as a result of the accumulation of supporting evidence; in essence, the effects of luck “fade out”, eventually resulting in the once malignant effect becoming benign. I then expand upon the notion of evidential support in chapter two and show how false or luckily-true beliefs can provide support for knowledge-worthy beliefs, albeit in somewhat unorthodox ways. These first chapters serve as an attempt to reconcile our ability to be successful epistemic agents with our often haphazard belief-forming processes: we possess many knowledge-worthy beliefs, although they may be knowledge-worthy in a way unknown to us. In chapter three I look at the fading phenomenon and its consequences when applied to modern conceptions of epistemic luck, specifically Pritchard’s and Hawthorne’s, and conclude that it may be our conception of luck that is preventing more knowledge than it should. In chapter four I reject the knowledge-harmfulness of internalist conceptions of luck as a natural consequence of the first half of the investigation. Chapter five is concerned with fleshing out the nature of the fading phenomenon and the epistemic effects of different kinds of evidence with a specifically externalist’s view of knowledge in mind. In the final chapter, I apply the fading paradigm to a classic epistemological puzzle, the lottery problem, and present a simple proposal: one cannot know that they are going to lose the lottery because the luck involved in lotteries simply does not fade.

## Chapter 1: Does Epistemic Luck Wear Off?

The infamous “Gettier examples” have resulted in the widespread belief that knowledge is incompatible with a certain variety of luck; specifically, if beliefs one has are true in virtue of pure happenstance as opposed to any epistemic work of ones own doing then we are intuitively disinclined to count those beliefs as knowledge. Duncan Pritchard has called this “veritic luck”, simply the notion that “it is a matter of luck that the agent’s belief is true” (Pritchard 2005a, 146), and employs a number of examples to show why veritic luck and knowledge do not get along. But while immediate inferences based on veritically lucky evidence no doubt appear to be knowledge unworthy, will any and every further inference based on the same evidence be subject to the same knowledge-preventing effect? I believe that that answer is “no”, that, in a way, epistemic luck *wears off*, allowing us to form knowledge-worthy beliefs that are based in part on veritically lucky premises. Bertrand Russell’s “stopped clock” Gettier-style example will serve as the centerpiece of this chapter, and I will attempt to show via my own series of examples that intuitively, the knowledge preventing factor of luck can and does “wear off.” I will refer to the phenomenon whereby luck has an increasingly minor affect on knowledge claims as “fading”, a term chosen for its suggestion of something becoming less prominent over time, further implying its ability to disappear entirely.

### A Working Conception of Luck

I begin with a conception of what it means for something to be epistemically “lucky”. Although not perfect<sup>1</sup>, I first appeal to Duncan Pritchard’s characterization in

---

<sup>1</sup> Luck, its recent definitions and the flaws therein are all topics that I will return to in further chapters. For now, working with this conception will at least allow me to set the stage for discussion.

his recent book “Epistemic Luck” as a starting point. Pritchard puts forth two necessary conditions for something to count as lucky:

(L1) If an event is lucky, then it is an event that occurs in the actual world but which does not occur in a wide class of the nearest possible worlds where the relevant initial conditions are the same as in the actual world. (Pritchard 2005a, 128)

(L2) If an event is lucky, then it is an event that is significant to the agent concerned (or would be significant, were the agent to be availed of the relevant facts). (Pritchard 2005a, 132)

It is (L1) that is of the greatest concern here<sup>2</sup>, but it is necessary to add a third premise, one put forth by Wayne Riggs, who finds (L1) and (L2) insufficient for reasons that will become clear in a moment:

(L3) If an event E is lucky for [an agent] S, then S was not sufficiently responsible for bringing about E. (Riggs, 19)

(L1), (L2) and (L3) combined supply a working conception of luck that I will appeal to at least temporarily. They also supply an easy test to determine whether something is, in fact, lucky<sup>3</sup>, by meeting or not meeting the criteria. Notice, however, that while as a totality the criteria are sufficient, (L1) with (L2) or (L2) with (L3) are also sufficient pairs. For instance, it is not always necessary for (L3) to be met, as is illustrated in the following example:

---

<sup>2</sup> (L2), the condition that human affairs are necessarily involved when considering luck, will play very little role in the discussion from here on out. Although Pritchard himself gives it considerable attention, I consider it more of a caveat (although a necessary one) than anything else.

<sup>3</sup> Note that we are distinguishing between “luck” and “non-luck” or “luck-neutrality”, and *not* “lucky” meaning “good luck” with its antonym “bad luck.” I will take this for granted throughout.

*The Terrible Basketball Player:* Steve is a terrible basketball player. In a game he takes a last-second shot from half-court, and the ball happens to go in, winning the game for his side.

(L1) is met, as Steve's lack of skill at basketball means that in almost every nearby possible world his last-second shot *does not* go in. (L2) is met as the situation directly concerns Steve, his teammates and his (no doubt furious) opponents. (L3) is not met, however, since Steve is indeed sufficiently responsible for the successful basket; after all, he took the shot.

Similarly, (L1) does not necessarily need to be met for an event to count as lucky, as is illustrated in the following précis of an example put forth by Riggs:

*The Skilled Basketball Player:* A very skilled basketball player always plays excellently in practice but, because of the pressure, plays terribly during actual games; in fact, he has not made a single basket all season. During one particular game, due to forces outside of his understanding, he calms down momentarily, takes a shot, and it goes in. (modified from Riggs, 18)

Here we are inclined to say that the skilled basketball player *was not* lucky in making his shot (that is we would think the basketball player was sufficiently responsible for his success), but (L1) dictates that since the basket would not have been made in many nearby possible worlds (because his nerves almost always prevent his natural talent from having an affect), it should count as lucky. As is shown in *The Skilled Basketball Player*, (L1) has a tendency to give false positives when determining luck, and therefore is over-permissive and not a sufficient condition. Thus while in any given situation it is unclear how many of the given postulates need to be met for an event to count as "lucky", it is

safe to say, at least for now, that if both conditions (L1) and (L3) are violated, then that event is definitively *not lucky* (note that (L2), pertinence to human affairs, is obviously incapable of determining luck by itself). This is enough to begin the investigation.

### **The Stopped Clock Example**

I will begin by presenting Pritchard's slightly modified version of an example originally put forth by Bertrand Russell:

*The Stopped Clock:* Our hero comes downstairs every morning around the same time and looks at the time on the old clock in her hall. The clock tells her that it is 8:22, and as a result, our agent forms the belief that this is the time. Furthermore, the clock is right, because it *is* 8:22 a.m. Nevertheless, unbeknownst to our protagonist, the clock has in fact broken down (it broke down 12 hours before).

This is clearly a case where our agent's belief is only luckily true, since in most of the nearest possible worlds where the agent forms her belief about this matter and the relevant initial conditions for this event are the same as in the actual world...her belief will be false. (Pritchard 2005a, 148)

The intuitive conclusion is that the agent, call her Jane, does not have knowledge of the time: the fact that she has a true belief is dependent upon the lucky coincidence that her clock happened to stop functioning exactly 12 hours beforehand. According to our working conception of luck, in many nearby possible worlds the clock stops at many different times and thus leads Jane to form a false belief (satisfying (L1)); the incident is significant for Jane (satisfying (L2)); and the clock's stopping was an occurrence outside of Jane's control (satisfying (L3)). Thus, when all tolled, it seems that the event in

question is definitively lucky, and that Jane's belief, although true, does not count as knowledge. So far, so good.

### **Luck Wearing Off: The Job Interview Examples**

The "Stopped Clock" example emphasizes our intuitions that beliefs based on premises that happen to be true purely as a matter of luck prevent those beliefs from counting as knowledge. These beliefs, in Pritchard's terminology, are veritically lucky; the question that I wish to raise is whether *all* inferences based on veritically lucky premises lead us to the same conclusion as the "Stopped Clock" example, that is that they do not intuitively count as being knowledge-worthy. I will illustrate my concern via the following examples:

*The Job Interview:* Jane has recently landed an interview at a large, prestigious company. Her father, very proud of his daughter, decides to give her a family heirloom, a pocket-watch, for good luck. Jane is told that the watch, if taken care of, is an extremely reliable timepiece. Let us assume that this is the case. Now, the watch has not been used in a long time, so it displays some arbitrary time from when it was last left to run out, and all parties are aware that it needs to be reset. Jane wakes up the morning of her interview and, remembering her present, decides to set her watch according to the clock in her front hall. The hall clock, however, has experienced a Russell-type stoppage as outlined above. Jane uses the information given by the veritically luckily correct clock to set her watch, and winds it, at which point it begins to function flawlessly. During breakfast, Jane looks at her watch and forms a belief that her interview begins in half-an-hour.

She leaves, travels via her preferred mode of transportation, and arrives at her interview right on time.

In this scenario, we are intuitively liable to label Jane's belief that her interview starts in half-an-hour as not knowledge-worthy. Note that as it stands now, "The Job Interview" is only slightly different from the "Stopped Clock" example, the singular difference being that the pocket-watch has been inserted as a sort of "middle-man" between Jane and her belief. This does not seem to affect our intuitions thus far, though: we still think her belief has come out to be true by virtue of happenstance, and indeed it seems as though our luck criteria (L1) to (L3) are again all met. For now, nothing has changed. Consider, then:

*The Job Interview II:* Jane's interview goes smashingly well: she is offered a corner office, a secretary, the works, and she starts the next day<sup>4</sup>. Having always been a little superstitious, Jane attributes some of her roaring success to the "lucky charm" that is her pocket-watch. As a result, Jane becomes enamoured with the timepiece, and takes meticulous care of it, making sure that it is wound constantly. Furthermore, Jane trusts no time-telling device other than her own, and thus her sole source for information about the time is her watch. She looks at her watch to make sure she's at the bus stop at the right time and, sure enough, she is. Her favorite television program comes on at a certain time in the evening and she checks her watch to make sure she catches it; sure enough, she does.

Now imagine that Jane forms beliefs about each of the events outlined in the above example. Before she leaves for the bus station, she believes "my bus will be here in five

---

<sup>4</sup> Say the interview was conducted on a Monday and that Jane then starts on Tuesday, to avoid any potential over-the-weekend-next-business-day confusion.



minutes.” Similarly, before she sits down in front of the television she believes “my favorite show is about to start.” Her beliefs are true (her bus *does* come in five minutes, her favorite show *is* about to start) but they are nevertheless based upon the same lucky piece of evidence given initially by the faulty front hall clock: Jane’s pocket-watch, only as a matter of luck, happens to tell the right time. The difference is that in the cases where she’s watching television or getting the bus back from work, I believe we are far more inclined to say that Jane’s beliefs count as knowledge.

What is to account for the change in our intuitions? No doubt the following possibility comes to mind almost immediately: for Jane, getting to her interview on time is extremely important, whereas missing the bus or her favorite television show is no big deal. In other words, in “The Job Interview” the stakes are high, whereas in “The Job Interview II” the stakes are relatively low. There is an epistemological intuition that the lower the stakes, the more likely we are to ascribe knowledge<sup>5</sup>, but I feel that the stakes are not, at least in the scenarios I have presented, relevant. For example, consider the following:

*The Meeting:* After working at her new job for a month, Jane has a very important meeting with some investors one day at 3 p.m. At 2:55 p.m. Jane looks at her watch and tells her coworker: “the meeting will start in five minutes, so head down to the boardroom and I’ll meet you there.”

Again, here it seems as though we would intuitively ascribe knowledge to Jane in this case, i.e. that she knows that the meeting will start in five minutes. There are two important factors to consider: first, the stakes are high (it is a *very* important meeting!),

---

<sup>5</sup> See, most recently, Jason Stanley, *Knowledge and Practical Interests* (2005). Talk of stakes and practical interests will also receive more detailed treatment in later chapters.

and second, the time on Jane's watch is *still* correct in virtue of being based upon lucky evidence (remember, Jane looks only at her watch to tell the time, she takes extremely good care of it, it has not stopped running consistently since its first setting, etc.).

Considering the stakes, we may be inclined to say that in a situation as important as "The Meeting", faultily set timepieces are salient possibilities that should be considered. But notice that this seems to be an issue that does not affect the present examples: if the stakes are so high as to warrant the double-checking and history of the setting of a timepiece, then the veritic luck should no longer be the guiding basis of our knowledge ascriptions, but is rather playing a peripheral role, the central issue being the high stakes of the meeting. Thus I do not think that it is the stakes involved that are most important at this point<sup>6</sup>; rather it seems that at the time of the meeting, if not well beforehand, the luck involved in the initial watch setting no longer seems to matter. This raises the question: how do the effects of epistemic luck wear off?

### **The Principle of Confirmation**

The potential solution I wish to explore explaining why we are intuitively inclined to count Jane's beliefs at the time of "The Job Interview II" as knowledge-worthy is that the main source of her beliefs, i.e. her pocket-watch, has been confirmed by numerous pieces of evidence as capable of providing true beliefs. An explanation by appeal to confirmation might proceed as follows: despite the fact that Jane looks only at her pocket-watch to tell the time, the fact that she is consistently correct (i.e. constantly arriving on time for scheduled events) gives confirming evidence that her watch is an accurate and reliable timepiece. At the time of "The Meeting" it is no longer the case that

---

<sup>6</sup> I am not denying that stakes play a role in our epistemic intuitions, rather that an appeal to stakes often does not provide a satisfactory solution to problems of epistemic luck.

Jane's belief is based solely upon a faulty piece of evidence (i.e. the faulty front hall clock), but rather that this evidence is supplanted by numerous confirmations. Luck has been replaced as the source of Jane's belief by evidence that we believe she is epistemically responsible for obtaining.

It is now evident why "The Job Interview" example was constructed such that Jane looks at no timepiece other than her pocket-watch: had she looked at the clock-tower in the centre of town, for instance, (which, for the sake of argument, tells the correct time in a way that is not veritically lucky) she would have "confirmed" the accuracy of her watch; if this were to happen we would be much more inclined to attribute knowledge concerning time-based claims to Jane right off the bat. Specifically, we would be more inclined to attribute her with knowledge concerning how much time she had to get to her interview. The explanation for this relies on the following principle:

*Principle of Confirmation:* Beliefs that just-so-happen-to-be-true and are thus not knowledge conducive can be made to be knowledge conducive by confirming evidence.

The Principle of Confirmation is itself a special case of a more general theory:

*The Fading Phenomenon:* Beliefs that are true as a matter of luck and are knowledge-harmful can become benign through a number of processes of which the agent has varying degrees of conscious awareness.

The Fading Phenomenon is left intentionally vague; indeed the first three chapters of the investigation are aimed at fleshing out this principle, namely by expounding upon the aforementioned belief-forming processes that are able to make knowledge-harmful luck benign. This leads me to reject a form of the incompatibility thesis, namely one that

states that luck and knowledge are entirely incompatible. I say a “form” of the thesis since the traditional form of the incompatibility thesis has involved knowledge and luck *simpliciter*. As expressed by Andrew Latus,

Luck and knowledge are not incompatible at all...Imagine a person who usually works in a windowless library carrel, but one day, on a whim, decides to work at a table near a large window. It begins to rain outside. She notices this and so comes to believe that it is raining. Surely, she knows it is raining, but it is just a matter of (circumstantial) luck that she is in a position to know this. (Latus, 164-165).

Thus, as Latus calls it, “circumstantial luck” and knowledge are not incompatible.

Nevertheless, there does seem to be a category of luck, as Pritchard has dubbed it the aforementioned “veritic” luck, which is separate from the knowledge-harmless doxastic luck (two categories that I will explore in further detail in later chapters). The new form of the incompatibility thesis then becomes that knowledge and veritic luck are incompatible entities, yet as I have begun to show in this chapter, this, too, is a theory that I reject. In Jane’s case it is the correctness of her claims about the time that becomes less and less knowledge-harmful as the amount of evidence she gathers increases.

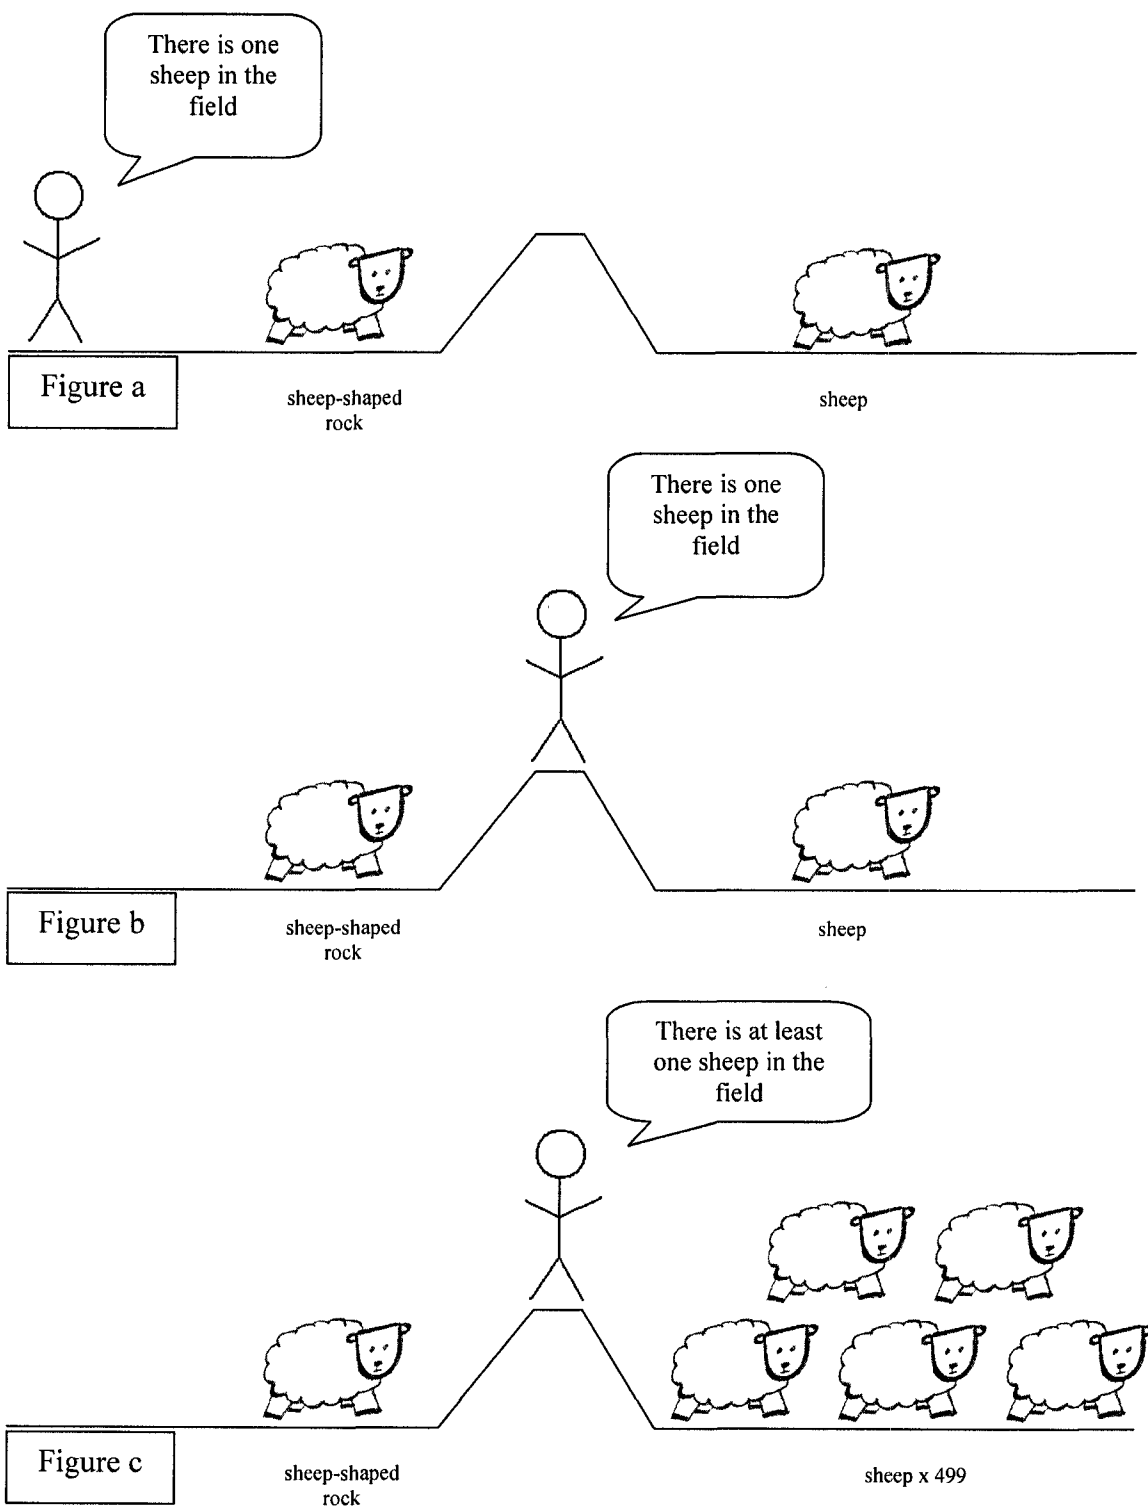
Consider another Gettier-style example, which I have borrowed from the general epistemic consciousness and will call “The Sheep Spotter”: John claims to see a sheep in a field, and from this information forms the belief: “there is at least one sheep in the field.” In actuality, unbeknownst to John, what he thinks is a sheep is really a cleverly painted sheep-shaped rock. His belief is true, however, as there *is* a *real* sheep standing just over a hill, out of sight of John but still within the confines of the field (see Figure a).

Straightforward Gettier intuitions tell us that John's belief does not count as knowledge. If, however, John *confirms* his evidence, that is he takes a walk over the hill and sees the real sheep, now it seems that his initial belief may indeed count as knowledge. Of course, a revised belief such as "there are at least *two* sheep in the field" would still be false, as John's evidence would then be based upon one real sheep and one sheep-shaped rock. Further, his confirmation needs to be of the right kind: if he saw yet another sheep-shaped rock, again mistaking it for the real thing, he would again be without knowledge.

Is it enough, however, that John merely confirms his evidence and does not actively replace it? That is to ask, need he only form a belief about the real sheep, or need he also realize that he was duped initially? Consider the simple case: John sees a sheep-shaped rock, and then a real sheep, believing both of them to be real sheep. He then examines his evidence further and notices that he has been duped in the case of his first sighting, and that he was correct after all in his second sighting. In this scenario, it is clear that, after his confirmations are made, we should attribute John the knowledge that "there is at least one sheep in the field." The murky scenario is where John *does not* double-check the accuracy of his sheep sightings, but nevertheless forms the same belief (see Figure b). Are we wont to attribute knowledge to him in this case? Of this, I am uncertain.

But now consider the scenario where John obtains not *one* true belief about a real sheep in addition to his falsely held sheep-shaped rock belief, but instead he sees another *499* real sheep, and for each of them forms the true belief that "there is (yet another) sheep in the field" (see Figure c). What of John's initial knowledge claim in this

Figures for “The Sheep Spotter” Example



scenario? I think that here we do indeed want to ascribe the knowledge that “there is at least one sheep in the field” to John. In the case where John has as much knowledge-

worthy belief as he does veritically lucky belief (as in the one real sheep and one sheep-shaped rock scenario), it is unclear as to whether we would credit John with knowledge (perhaps it is here that factors such as high or low stakes help mould our intuitions). In the case where the veritically lucky belief is only one in five-hundred, however, intuitively it does not have the same knowledge-preventing affect it once had.

In Jane's case the principle of confirmation works implicitly<sup>7</sup>: each time Jane is on time for a meeting, catching the bus or watching her favorite television show, her belief as to what the time is at each of those moments is confirmed. Thus, it seems that out of our initial criteria for something to count as lucky, (L3) is violated: at the time of Jane's meeting, it seems that she is sufficiently responsible for bringing about her getting there on time, as she has implicitly confirmed that her watch is accurate. (L2) is, of course, met, but is not sufficient on its own to determine luck, so we may safely disregard it. What, then, of (L1)?

(L1), I believe, is also violated, but in order to see how, it is necessary to examine the criterion in detail, specifically the "relevant initial conditions" stated by Pritchard. My reasoning is as follows: returning to *The Job Interview* example, in the real world Jane forms a true belief about the time given by her front hall clock, but it is veritically lucky. In a nearby possibly world, however, she forms a false belief, since the front-hall clock stopped at a different time, say eleven and-a-half hours beforehand instead of twelve. As a result, in this possible world Jane arrives at her interview half-an-hour late, and therefore doesn't get the job. We can also assume that had Jane arrived early for her meeting (thereby exhausting the possibilities for nearby possible worlds: early, late, or on

---

<sup>7</sup> Another promissory note: "implicitly" in this context is espoused more fully in Chapter 5.

time<sup>8</sup>) she would also not get the job, for whatever reason. The relevant initial condition I appeal to, then, is that Jane is a rational agent, and thus if her job interview were to have gone poorly because of the time of her arrival, she would confirm the inaccuracy of her watch and revise her beliefs. Thus if we have the following relevant initial conditions present:

- Jane's front-hall clock is faulty;
- Jane sets her watch by the clock and does not initially confirm that setting;
- Jane's watch functions normally; and
- Jane is a rational agent;

then it seems that in nearby possible worlds, after an unsuccessful interview Jane *does* arrive at the bus stop on time (although now for a trip home, dejected, as opposed to elated) and she *does* catch her favorite television show on time, because she has changed the basis of her beliefs. Thus, it seems that a confirmation of evidence violates (L1) as well: in the real world as well as nearby possible worlds Jane is consistently on time after confirming or replacing the source of her beliefs. With (L1) and (L3) violated, it seems that Jane's beliefs no longer count as lucky, and thus they may be considered knowledge-worthy. The effects of luck have worn off.

If a strict incompatibility thesis stating that knowledge and veritic luck are incompatible is true, then it is under-described. It seems that when an agent forms a belief exclusively on the basis of a premise that is true only by happenstance then we do not intuitively want to ascribe knowledge to her. On the other hand, when a veritically luckily true premise acts as only one of many other premises on which a belief is based,

---

<sup>8</sup> I am, of course, ignoring possible world scenarios where Jane doesn't make it to the interview at all, for whatever reasons. I don't feel as though such concerns are relevant to the investigation at hand.



or when a large body of evidence contains a relatively small amount of veritically luckily true premises, it seems that what was once knowledge-demolishing has lost its efficacy. This chapter has served as an introduction to the fading phenomenon and one of its instantiations, the Principle of Confirmation, but there are still further ways in which the effects of luck can fade. I have also taken for granted a certain definition of luck, but there have already been conceptual snags presented that need to be dealt with. Finally, I have rejected an initial version of the incompatibility thesis, one that says that knowledge and veritic luck are incompatible *simpliciter*. It is these issues that are the focus of the chapter 2.

## Chapter 2: Rejecting the Incompatibility Thesis, and How Bad Beliefs Can Lead to Knowledge

As fallible agents, we often form beliefs that are not knowledge-worthy due to factors such as poor, misleading, or incomplete evidence. Yet despite our tendency to be epistemically irresponsible, we nevertheless feel we are more often than not capable of reliably producing knowledge-worthy beliefs. In this chapter I argue that fallible beliefs themselves are able to support knowledge claims, despite their shaky epistemological status, through two additional forms of fading: one where implications of already acquired evidence can provide implicit evidence used to support a weaker knowledge-claim, and another where the role of the spurious evidence can be seen as becoming relegated to a peripheral yet necessarily supportive role for the original knowledge-claim (i.e. it is not so much the content of the evidence as the existence of the evidence itself that is knowledge-conducive). I centre my investigation around two examples: first the Judy and Trudy example, involving a situation similar to the case of Jane and her pocket watch, where ones beliefs are true only as a matter of luck and the content of these beliefs are used to form additional beliefs. Second, the Battle of Hastings example, where ones beliefs are lucky to be true as the result of being formed via an unreliable method. In both cases, I claim, we are intuitively inclined to count the beliefs that are partially reliant upon epistemically poor evidence as knowledge-worthy, despite the agent's unawareness of, or lack of concern for the status of the evidence used to form those beliefs. The upshot of the investigation is that while agents may appeal to certain evidence as the sole basis of their belief, the knowledge-worthiness of those beliefs is actually rooted in additional implications that the agent has not consciously considered, and need not have consciously considered in order to be classified as knowledge. This also has

ramifications for how we can define epistemic luck in the first place, notably arguing against modal conceptions such as those put by Pritchard.

### **Setting the Stage: The Identical Twins Example**

I will begin by presenting a case where a belief is based upon previous beliefs that are themselves not knowledge-worthy due to the harmful effects of epistemic luck:

*The Identical Twins:* Eugene has invited a set of identical twins to his party, Judy and Trudy. The night of the party there are many people around, and Eugene is asked whether Judy and Trudy decided to show up. He looks to his left, and in one room sees Judy, who he mistakes as Trudy. He looks to his right, and in another room sees Trudy, who he mistakes as Judy. Eugene tells his friend: both Judy and Trudy are at the party, and his belief is correct.

Here Eugene initially formulates two beliefs: one, that Judy is at the party, and two, that Trudy is at the party. Each belief individually is only true by matter of luck: his belief that Judy is at the party is based on his actually seeing Trudy, but since Judy *is* at the party (just out of sight at the time of the belief formation) his belief is nonetheless true. The same reasoning applies to his belief about the presence of Trudy. Neither of Eugene's beliefs about the whereabouts of Judy or Trudy should count as knowledge, but it nevertheless seems that his belief that the *both of them* are at the party is knowledge worthy.

A conflict is brewing. Eugene would claim that he has formed his final belief that both twins are present in the following way:

$B(\text{Judy is at the party}) \ \& \ B(\text{Trudy is at the party}) \ \rightarrow \ B(\text{Judy and Trudy are at the party}),$

arriving at a knowledge-worthy belief (the belief that both Judy and Trudy are at the party) solely from premises that are not themselves knowledge-worthy. This is perhaps surprising as he has acquired no new evidence between the initial formation of his beliefs and his reaching his conclusion. Surely we cannot have a knowledge-worthy belief based exclusively upon evidence that is itself only true as a matter of luck? Perhaps even more surprising is that it seems a simple matter of deduction to get from:

Judy and Trudy are at the party

to

Judy is at the party

by conjunction elimination. Yet B(Judy is at the party) was one of the beliefs that was deemed to not be knowledge-worthy based on the evidence Eugene had received, namely because it is only lucky that it is true. Nevertheless, if we follow these steps it appears that the belief has become knowledge-worthy simply by applying a few logical progressions, but without the addition of any new evidence<sup>9</sup>. I do not wish to delve into the related debate of whether knowledge is closed under known entailment<sup>10</sup>, and in order to make the conflict more apparent I will take for granted that knowledge of a conjunction of propositions entails knowledge of each proposition individually<sup>11</sup>. I do wish to try and make clearer how knowledge-worthy beliefs can be based upon premises which are not themselves knowledge-worthy, maintaining that Eugene's belief that both

---

<sup>9</sup> Note the key qualifier, "has become *knowledge-worthy*", i.e. that is not to say that Eugene automatically knows the logical consequence (which would assume an unrealistically strong form of the "knowledge is closed under entailment" thesis) but that given his background beliefs, he is a position such that were he to consider the logical entailment he would know it.

<sup>10</sup> Although I think it highly likely that it is.

<sup>11</sup> Although there is, perhaps, room for debate here. Again, I will not concern myself with such matters.

twins are at the party is knowledge-worthy, and to qualify some of the seemingly nasty consequences. To try to illustrate this point further, consider a variant of the above:

*The Identical Twins, II:* Eugene is hosting the same party, and the same inquisitive friend again asks if Judy and Trudy are present. However, the party venue is so full that Eugene cannot see everyone simply by looking around, and instead has to navigate to different rooms. He enters one room and sees Judy, mistaking her to be Trudy. He then takes a convoluted course through another set of rooms where he again sees Judy (who has since moved from the punch to the hors d'oeuvres), believing this time that it is, in fact, Judy. In actuality, both Judy and Trudy are at the party, but Eugene has simply seen Judy twice, thinking her to be, on different occasions, both her and her sister. He returns to his friend and tells her that he believes that both Judy and Trudy are at the party, and his belief is correct.


In this situation Eugene's belief that Trudy is at the party is only luckily true, since he forms the belief that Trudy is at the party on the evidence of seeing Judy, whilst unbeknownst to him Trudy is in another room. Thus his belief that Trudy is at the party should not count as knowledge. Eugene's belief that Judy is at the party, however, is based on actually seeing Judy, and so might be a candidate for knowledge. However, in this situation it seems that he did not consider the relevant possible alternatives (i.e. that it could have been the other twin that he was seeing, as to him they are phenomenologically indistinguishable) and thus his belief should again not count as knowledge<sup>12</sup>. Again he makes the claim that he knows that both twins are at the party, and again his claim is

---

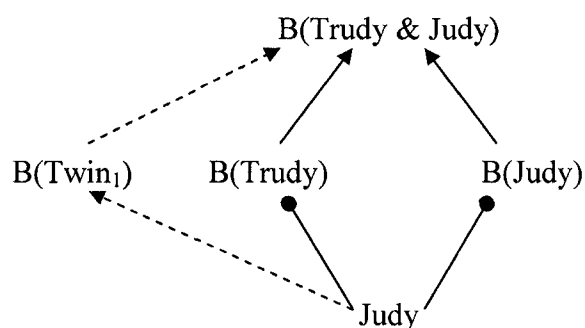
<sup>12</sup> If this is the case then it seems impossible for Eugene to be able to know that either Judy or Trudy is present in *any* situation, but this possibility does not affect the results of the investigation.

based upon two beliefs that are not knowledge-worthy, similar to the first identical twins scenario. Yet here it seems much more intuitive to say that the final belief that both Judy and Trudy are at the party is not knowledge-worthy.

This second example shows that it is not simply the possibility that “two epistemic wrongs make a right” in order to explain the knowledge-worthiness of the proposition in both examples. The solution I wish to put forth is that further implications (specifically implicit ones) of the evidence presented can eventually lead to knowledge-worthy conclusions that are at least partially based upon beliefs that are not themselves knowledge-worthy. Referring to the first twins example, Eugene first forms a luckily-true belief that Judy is at the party. His belief in Judy’s presence at the party also implies that one of the set of the twins is at the party, let us denote her  $Twin_1$ , and denote the belief in her presence as  $B(Twin_1)$ . He then forms a luckily-true belief that Trudy is at the party, this belief further implying that there is a second, different twin present; continuing the same syntax, denoting her as  $Twin_2$  and the belief in her presence as  $B(Twin_2)$ . These further implications give implicit evidence that can support Eugene’s knowledge claim that Judy and Trudy are both at the party (see Figure 1).



**Figure 1: Implicit evidence provides support for a knowledge-worthy true belief**



**Figure 2: Insufficient implicit evidence does not provide enough support for a knowledge-worthy true belief**

In the above diagrams, Judy and Trudy respectively provide different kinds of evidence: a dotted arrow represents implicit evidence resulting in a belief, be it formed implicitly or explicitly, whereas a blunt-headed arrow represents a Gettier-style belief formation, where the resulting belief is true, but only as a matter of luck. Solid lines represent a connection that an agent believes he is making from beliefs that he has previously formed, which may or may not represent a knowledge-worthy belief forming process.

Notice now the contrast with the second twins example, where the mistaken beliefs are formed based upon evidence given by only one twin: in this situation, the first belief again implies that there is one twin at the party (see Figure 2). But the second belief does not imply that there is another, different twin, as it is based upon evidence given by the same twin. In other words, there is only one source of additional implicit evidence, and thus the final knowledge-claim, that both twins are at the party, is ill-supported, and thus is cannot be considered knowledge-worthy.

We may then be led to the perhaps strange conclusion that Eugene can know that both twins are at the party, and that he can know that one of the twins is present and that another twin is present also, but can know neither that *j*: Judy is at the party nor *t*: Trudy

is at the party. Further examination, however, shows such kinds of situations to be relatively mundane. Consider again the case of John the sheep-spotter: John thinks he sees a sheep in a field, and thus forms the belief, “there is a sheep in this field.” What he is actually seeing, however, is a cleverly painted sheep-shaped rock; a real sheep is grazing over a hill and out of sight of John, thus making his belief true but not knowledge-worthy. We *would* nevertheless feel justified in ascribing to John the knowledge that *something* is in the field, even if he is mistaken as to what it is. If we were to *ask* John if there were something in the field, he would likely respond by saying, “yes, there is a sheep over there.” If we are aware of the true nature of the situation and inform John of his mistake, he could still claim that he at least knew there was *something* there, and he would be justified in making this claim, as weak as it is.

What, then, of the initial conflict, where we ascribe to Eugene the knowledge that both Judy and Trudy are at the party, yet deny him knowledge of the logical entailments *j* and *t*? I believe we *are* able to ascribe Eugene this further knowledge, but only as a restricted version: we should ascribe him knowledge that Judy *qua* Twin<sub>1</sub> is present, and also that Trudy *qua* Twin<sub>2</sub> is present. Thus knowledge under conjunction elimination holds, from B(Judy *qua* Twin<sub>1</sub> and Trudy *qua* Twin<sub>2</sub> are present) as a belief counting as knowledge to B(Judy *qua* Twin<sub>1</sub>) and B(Trudy *qua* Twin<sub>2</sub>) as also being knowledge-worthy.

Knowledge-worthy beliefs have been portrayed as requiring the support of a certain amount and quality of evidence, and if this support is too weak the construction collapses<sup>13</sup>. Beliefs which when initially formed and found not to be knowledge-worthy can be buttressed by additional evidence, this additional evidence being able to provide a

---

<sup>13</sup> Many thanks to Andrew Jeffery, who provided the architectural metaphor to which I appeal throughout.



strong enough support for the knowledge claim. I have argued earlier for the effects of explicit evidence, that is to say that a belief based on a poverty of evidence can be supplemented to the point of becoming knowledge-worthy by occurrences that confirm the belief or deny opposition to it. In the situation presented in this investigation the claim is perhaps stronger, that even without additional explicit evidence a belief can be knowledge-worthy based upon propositions supported by the evidence that are not necessarily given any credence or consideration by the belief-former. An immediate implication of this is that the agent may have a knowledge-worthy belief, but if asked to justify that belief would appeal to evidence that is in itself insufficient to support the knowledge claim. Thus Judy and Trudy situations make a case for an externalist treatment of epistemic luck, inasmuch as it ignores the internal reflections of the would-be knower<sup>14</sup>.

Two further consequences should be noted: first, it is not clear, at least from the examples that I have entertained, *how much* evidential support is needed for any given knowledge claim. Indeed it seems that there is no all-encompassing rule, as “heavier” knowledge claims, i.e. beliefs that claim more, will usually require more support to make the claim knowledge-worthy, where as “lighter” knowledge claims, i.e. such as the situation above where John the sheep-spotter merely claims to see *something*, might not require very much support at all. Second, although I am a proponent of an invariant theory of knowledge, nothing I have presented in this section prevents a contextualist treatment of the subject (although contextualists may deem the work thus far as superfluous) and is nicely compatible with theories of stakes and practical interests, such

---

<sup>14</sup> For a more detailed examination of this kind of luck, see Duncan Pritchard’s discussion of “reflective epistemic luck” in his book *Epistemic Luck*.

as those presented explicitly by Stanley and hinted at by Hawthorne<sup>15</sup>. For instance, one might see stakes as either “weighing down” or “lightening” a given knowledge claim, thus requiring more or less evidence to support it, respectively (it does seem, however, that at any given degree of importance in stakes theory, or at any given degree of knowledge-relation in contextualist theories, Judy and Trudy phenomena can still be present).

### **Judy, Trudy, and Fading**

Continuing the theme from the previous chapter, Judy and Trudy phenomena seem to present another type of case where the effects of epistemic luck, where once harmful, can become harmless, i.e. another instance of the fading phenomenon.

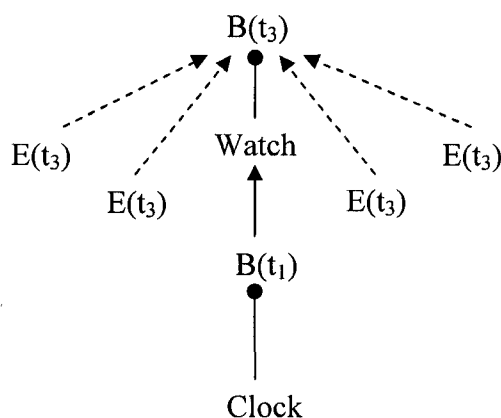
Returning briefly to the example of Jane and her pocket watch, a belief that she formed which was not knowledge-worthy due to factors of luck became knowledge worthy as her belief was given more and more support. Importing the Jane case into the diagrammatic syntax presented in this chapter, at the time of the initial watch-reading, Jane’s belief is formed in the following manner:

---

<sup>15</sup> See Stanley, *Knowledge and Practical Interests* and Hawthorne, *Knowledge and Lotteries* in particular for their respective takes.



**Figure 3**  
Jane forms her initial belief concerning the time



**Figure 4**  
Jane forms a later belief, after confirming evidence

Here it was determined that Jane's belief is not knowledge-worthy at the time she forms  $B(t_2)$  (as represented in Figure 3) but that at a later time, after numerous confirming evidence is presented (in her case, implicitly),  $B(t_3)$  is found to be knowledge-worthy. As seen in Figure 4, the further implicit evidence causes the once knowledge-preventing effect of the lucky-to-be-true belief (as represented by the blunt-headed arrow) to lose its harmfulness. While in the case of Jane the additional evidence originates from outside sources (i.e. the confirmation presented by showing up to meetings on time, catching her favorite television show, etc.), Eugene's "additional" evidence is derived from implications of evidence with which he has already been presented. Thus the fading phenomenon as experienced by Jane sees her maintaining the same knowledge claim (i.e. her belief about the time) but requiring additional, external evidence to effectively dilute the effects of epistemic luck (although this evidence may be implicit or explicit), whereas in the case of Eugene, the amount of evidence he has accumulated remains the same, but

the demands of his knowledge-claim are lessened, i.e. from B(Judy is at the party) to B(one of the twins is at the party) or B(Judy qua Twin<sub>1</sub> is at the party).

### **The Battle of Hastings**

The upshot thus far is that Judy and Trudy cases exhibit a type of epistemic fading where it is not required that *new* sources of evidence be present, but rather that considering the implications of already acquired evidence, combined with lessening the demands of the knowledge-claim<sup>16</sup> can also cause the effects of epistemic luck to wear off, which I will introduce as the following:

*The Principle of Lessening Demands:* Beliefs that are knowledge-harmfully lucky and are not sufficiently supportive of a further knowledge-claim can be sufficiently supportive of the same knowledge-claim if it is construed in a “less-demanding” way (where “less demanding” means requiring less evidential support).

The common thread running through these examples has been that an agent can form knowledge-worthy beliefs on the basis of evidence of which she is not consciously aware, with implications that she has not consciously taken into consideration. A consequence of this is that the agent may *appeal* to evidence as an explanation of her belief that is not wholly supportive of her knowledge claim, all the while her belief remaining knowledge-worthy, a topic I will examine in more detail in the upcoming chapters.

The cases thus far have provided examples of two different ways in which epistemic fading is able to take place. I believe there is also a third type, as exemplified in the following example:

---

<sup>16</sup> This talk should not be confused with talk of contextualism, although perhaps such an analogue could be made: if one considers the same evidence but claiming a less-demanding k-relation perhaps similar effects could be noted.


*The Battle of Hastings*: Nader has been raised not to believe any information his schoolteachers tell him unless it is confirmed by his father. One day in history class, Nader is taught that the Battle of Hastings took place in 1066, but refrains from believing that the statement is true. Returning home that afternoon, Nader asks his father when the Battle of Hastings took place; the father, however, has no idea, but not wanting to lose face in front of his son, picks a date at random, and happens to reply “1066”. Nader now forms a belief in the proposition BOH: “the Battle of Hastings took place in 1066”, and his belief is true<sup>17</sup>.

Here the evidence provided by the authority of Nader’s teacher, who, let us assume, is a trustworthy and reliable agent, provides strong enough evidence on its own to support belief in BOH; to satiate Nader’s unusual epistemological upbringing, however, he requires confirming evidence from his father to form the belief at all, which, in this case, is not itself knowledge-supportive. Notice here that the support given by his father’s evidence must support a belief in the teacher’s evidence, and cannot support a knowledge claim on its own. Thus Nader must believe the proposition provided by his teacher: it is not good enough to say that Nader *only* believes his father, as then his belief in BOH would be based upon evidence that exhibits Gettier-style instability. If the agent *actively ignores* sources of evidence, believing them to be inaccurate or unsupportive, then this evidence should not be considered as a source of evidence that can play a supporting role


---

<sup>17</sup> Many thanks to an anonymous audience member at the Western Canadian Philosophy Conference in 2006 who helped flesh out this example.

in a belief<sup>18</sup>. Thus, the structure of Figure 5 is sufficient for B(BOH) to be knowledge-worthy, whereas it is not in Figure 6:



**Figure 5: Nader's belief based upon evidence given by both his teacher and his father.**



**Figure 6: Nader's belief based solely on the testimony of his father**

In Figure 6 the evidence provided by the father exhibits Gettier-style instability and is thus, in itself, not supportive of knowledge. In Figure 5, however, the same evidence shows the same instability, but is nevertheless a crucial component in the formation of B(BOH).

The upshot of this discussion is that sources of evidence that are knowledge-harmful because of their epistemic-luck content can nevertheless be knowledge-conducive by filling a different kind of epistemological niche, in this case as a stepping-stone to belief formation. The situation involving Nader shows a form of fading where the epistemically-harmful belief can be seen as being relegated to a more structural role: whereas a cornerstone cannot by itself support the mass of a building, it is nevertheless a crucial component in its construction, and indeed in this case the belief could not be formed at all without it. Thus we have a third variety of fading:

<sup>18</sup> A difference should be noted between “implicit” evidence, i.e. a type of evidence that is obtained but not consciously recognized by an agent, and evidence that is actively ignored, the latter not being able to play a knowledge-supportive role in a belief formation.

*Principle of Auxiliary Support:* Beliefs or evidence that is not itself supportive of a further belief because of a knowledge-harmful variety of luck can nevertheless be knowledge-worthy by causing the agent to form the belief in the first place.

In other words, relegating a belief to the role of structural support in the belief-forming process makes benign the once knowledge-harmful luck component.

### **Another Look at the Nature of Luck: Refining Knowledge-Harmful and Knowledge-Harmless**

Throughout this and previous discussions I have been making liberal use of the terms “knowledge-harmful” and “knowledge-harmless”, terms that I borrowed from Pritchard who distinguished the two in order to put forth his modal conception of luck. “Harm” in this case means preventing ones true beliefs from counting as knowledge, and indeed the more interesting category of the two, the harmful luck, can be broken down into two further categories, what Pritchard calls “veritic luck” and “reflective epistemic luck.” Briefly, he defines the former as being a “matter of luck that the agent’s belief is true”, whereas the latter is defined as “given only what the agent is able to know by reflection alone, it is a matter of luck that her belief is true.” (Pritchard 2005a, 175) Of these two it is no doubt the latter that is most controversial, the upshot for Pritchard’s theory being that while Gettier-type examples force us to add at least some degree of externalism to our epistemological accounts (categorized by veritic luck), reflective epistemic luck also forces us to reexamine internalist conditions on knowledge. It seems, however, that the notion of fading that I have been extolling is incompatible with “reflective epistemic luck”, as I have attempted to show that agents are able to have perfectly knowledge-worthy beliefs even though their knowledge-worthiness is due to

factors that are outside of the agent's conscious considerations. In the sections following, I will re-examine Pritchard's conception of luck in light of the discussion of fading, and the plausibility of a modal conception in general given the results of the examples heretofore presented.

### Chapter 3: Another Look at Epistemic Luck - Is Modal Rarity Sufficiently Luck-Determining?

To reiterate, Pritchard's luck criteria are as follows: for an event to count as lucky, that event must (L1) occur in this world but not occur in many nearby possible worlds, where the relevant initial conditions are the same, and (L2) the event in question is pertinent to human affairs, or would be if that person were aware of all the factors that should concern him. Pritchard's example of a paradigm case of luck is winning the lottery: if one wins the lottery in this world, no doubt in many nearby possible worlds one does not actually win the lottery, and thus the event should be considered lucky<sup>19</sup>. Even with this seemingly straightforward example of luck, the notion of relevant initial conditions is tricky. As Pritchard notes,

[I]f one includes in the initial conditions for the event the demand that the balls fall into the lottery machine in a certain way, then one will no longer generate the desired result that the event is lucky...because the specification of the initial condition will *determine* the event in question across all nearby worlds. What we have in mind by this clause is thus some conception of the initial conditions which does not understand them in such a way that, individually or collectively, they determine the event in question. (Pritchard 2005a, 131)

---

<sup>19</sup> Lotteries are discussed in more detail in Chapter 6.



The problem is essentially this: say the winning numbers are determined by the fallings of ping-pong balls into a chute. If the system in which this occurs is *determined*, say by the complex interaction of a number of physical variables, then it seems likely that we would want to include the ping-pong ball position-determining factors as part of our relevant initial conditions. Doing so, however, results in a violation of (L1), since in the possible worlds where the positions of the ping-pong balls are determined in the same way, the same numbers will be drawn, and the agent will win in those worlds. Including these factors would therefore dictate that there is actually no luck involved in lotteries, a very counter-intuitive result<sup>20</sup>. Pritchard's first attempt at a solution is to essentially ignore the relevant initial conditions restraint, and consider the class of *all* possible worlds, his reasoning being that considering all possible worlds the agent's winning the lottery will occur in very few of them, and thus still count as lucky (the other possible worlds would also consist of many instances in which the agent does not play the lottery at all, say, thus making the occurrence especially rare).

Opening consideration to all possible worlds, however, cannot be in itself a guide to luck. Consider a possible revised condition in response:

(L1\*) An event is lucky if it occurs in this world but does not occur in many nearby possible worlds.

This construction, however, opens the doors for a whole host of counter-intuitive examples. For instance, consider the following:

*The Skilled Basketball Player, at School:* Michael Jordan doesn't usually pick up his son Michael Jr. from school, but his wife has been in a minor car accident and

---

<sup>20</sup> It would also be hard-pressed to show that there is any luck involved in *anything*, a position that is perhaps defensible, but is clearly not a position Pritchard would wish to take (for details on a possible defense, and subsequent rejection, see Richard Johns, "A Subjective Theory of Objective Chance").

thus cannot chauffer on this particular day. Walking up to the school a group of schoolchildren playing basketball recognize him and ask him to take a free-throw.

He does, and it goes in.

Now, considering the entire scope of possible worlds, Michael Jordan makes his basket in this world, but not in a wide variety of nearby possible worlds, as in those worlds he is not in the position to make the basket in the first place (his wife's accident would never have occurred, and thus he wouldn't go to school at all). So what is lucky according to (L1\*)? First, presumably Mrs. Jordan's accident itself. Second, for the schoolchildren to have seen Michael Jordan is no doubt lucky for them. But his making the basket is something we would intuitively not consider lucky, regardless of its modal rarity. Notice also that none of the factors leading up to Michael Jordan's shot-making actually *determine* whether the ball goes in the basket. In this situation, determining factors would perhaps be the force with which the shot was taken, the position of his arms, the wind intensity and direction, etc. Indeed we need not look to such convoluted examples to see why modal rarity, widely construed, is insufficient in determining a harmful variety of luck. Presumably many of the beliefs I acquire in an every-day sort of fashion I would not acquire in many nearby possible worlds, due to the sheer complexity of the factors that are involved in even the most mundane occurrences. Pritchard realizes that such types of "luck" are not knowledge-harmful, classifying them under "content epistemic luck" which he defines as it being "lucky that the proposition is true." (134) However, by potentially dropping the "relevant initial conditions" condition of (L1), Pritchard seems to equivocate what he wants to count as harmful luck with harmless luck.

### Narrowing the Scope


Indeed this is a point that Pritchard wishes to acknowledge. Again considering *The Skilled Basketball Player, at School*, to avoid the problem of modal rarity we need to limit the scope of relevant initial conditions around Michael Jordan making the basket; Pritchard is initially concerned with “non-determining conditions”, but quickly focuses his attention on scope-determining factors. The direction of the investigation then becomes not simply the project of determining if something is modally rare enough, but from which modal standpoint is rarity luck-determining, and, more importantly, knowledge-harmful. Something at which Pritchard gestures is that when making knowledge attributions we need to consider the scope of possible worlds such that the agent forms his belief in the same way across those worlds. Throughout his book he returns to the example of “gullible John”, a character who will believe pretty much anything his (difficult to be considered) friends tell him. Regardless of the number of times he has been fooled in the past (and *knows* he has been fooled in the past), gullible John still forms beliefs on the basis of his friends’ very sketchy evidence. On one occasion John’s friends tell him that his house is on fire, and on the basis of this information John comes to believe it to be true. As it turns out, though, John’s house *is* on fire, and so his belief happens to be true<sup>21</sup>, although intuitively this belief is not knowledge-worthy. At this point in the example, we consider John not to have knowledge because his belief is true simply by matter of veritic luck: the cause of his belief (i.e. the evidence given by his friends) does not limit the possible worlds under consideration merely to those in which John’s house actually is on fire; indeed, the

---

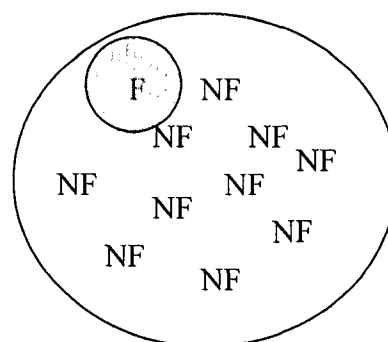
<sup>21</sup> Is it necessary that John’s friends also don’t know that his house is on fire? If they don’t, it seems straightforward that John doesn’t know, either. If they do, it seems comparable to a “barn façade country” example, where our intuitions should tell us that in that case he again is without knowledge.

possible worlds we need to consider are those in which John's house *is not* on fire, which are, presumably, the majority. If, however, John is motivated to go see if his house really is on fire, and then sees that it is the case, John's belief that his house is on fire is now based on the empirical evidence that he is directly experiencing. Thus at this point in the example, the scope of possible worlds to be considered is narrowed: John's belief is now based upon the evidence given to him such that his house is actually on fire. John's belief in this case is *not* veritically lucky, as in almost all nearby possible worlds where John forms his belief in the same way, he believes his house *is* on fire, his house is on fire, and there is an appropriate connection between his belief and reality.

This seems to be a paradigm case of the effects of luck wearing off under the Principle of Confirmation: John's belief that his house is on fire when he is standing in front of his engulfed home is based on simple empirical evidence. However, he would have never acquired that evidence if he had not believed something that was only luckily true, namely the bad evidence he received when told his house was on fire. Thus the belief that is true by luck is made knowledge-worthy by changing the luck in question from veritic luck to doxastic epistemic luck (138). The luck has worn off, by Pritchard's standards, due to a change of scope in our considerations of which possible worlds are relevant to consider (see Figures 7 and 8).



**Figure 7:** Considering all possible worlds, those in which John's house is on fire (F) is minute compared to those in which it is not on fire (NF). Thus in this situation, John's belief that his house is on fire, B(F), is only lucky to be true and as a result not knowledge-worthy.



**Figure 8:** After new evidence is presented, the grayed circle represents the scope of possible worlds now under consideration. At this scope, B(F) becomes knowledge-worthy.

Despite the conclusions just reached, there remains a nagging question: what degree of modal rarity is required for an event to count as lucky, after the appropriate narrowing of scope has taken place? Is it enough to say that the event occurs simply in a “minority” of worlds, or need we specify a (perhaps even a rough) modal frequency? Pritchard, somewhat surprisingly, does provide specifics: if the event occurs in less than or equal to 50 percent of the possible worlds, we are most likely to be inclined to count the occurrence of those events as lucky. He provides the following example: consider a game show contestant who is presented with a final choice between two possible answers. The contestant has not the slightest inkling as to which is the correct answer, so she guesses. To her delight, she gets the answer correct, winning fabulous prizes. This scenario, in the mind of Pritchard, is obviously one of luck.

He then asks us to consider a slightly modified version of the above, whereby the contestant has four final choices instead of two, where three of the choices are correct and one is incorrect (it's an easy game). Again she has no inkling as to which answers are

correct and which are incorrect, again guesses, and again picks one of the right answers. Pritchard believes that our intuitions tell us that in this case luck has had less of an influence, that the agent in the modified scenario is somehow not as lucky as the agent in the original scenario; as Pritchard himself says “I think it would be unlikely that we would put this down to luck since the odds were squarely in her favor” (2005a, 130). This conclusion sounds strange, as it seems to directly contradict lottery intuitions. Indeed if the game show contestant in the modified example forms a belief as to which answer is the correct one, her belief should count as a paradigm instance of veritic luck; however, its modal commonality at this point in Pritchard’s investigation prevents it from being classified as lucky at all. This seems to be a dead end for Pritchard.

### **Reflective Epistemic Luck**

Thus far it seems that issues I have presented which are concerned with fading are compatible with a modal conception of luck: harmful luck can become harmless as additional evidence refocuses our modal attention to just those possible worlds in which the occurrence of the event upon which our belief is based is no longer modally rare. Pritchard still maintains that there is a strict incompatibility between veritic luck and knowledge, however we can satisfy the demands of his theory by showing that, as I have previously stated, veritic luck can “turn into” one of the benign varieties of luck through a modal refocusing. Of course the problem of the degree of rarity required is still at large, but it is one that I will not attempt to solve<sup>22</sup>. Rather, I turn to Pritchard’s second type of knowledge-harmful luck, his reflective epistemic luck condition. He introduces this type of luck mostly to deal with chicken-sexer examples; briefly, an agent is able to determine

---

<sup>22</sup> It may, in fact, be a lost cause: it seems that answering the question “what degree of modal rarity is luck determining?” is as difficult as answering the question of what makes a belief lucky in the first place.

the sex of a chicken extremely reliably, yet either does not know how she is doing it, or believes she is doing it a certain way while she actually does it in a different way (say she believes the knowledge of the sex of the chicken is delivered to her by some visual and tactile clues, whereas in fact there is a subtle, consciously undetectable pheromone that gives her the answer). Pritchard wants to be able to cope with these examples; specifically he wants to deny the naïve chicken-sexer knowledge by saying that her beliefs are true only as a matter of luck<sup>23</sup>. On the veritic luck model, however, the modal results Pritchard seeks do not obtain: in nearby possible worlds, the chicken-sexer forms the same beliefs, and is continuously correct, at least in the way that possible worlds have been “ordered” thus far in the investigation (i.e. narrowed from the class of all possible worlds to those in which the agent forms her belief in the same way). Pritchard thus prescribes the following: when dealing with cases like the chicken-sexer, we need to consider what the agent is able to reflectively know, essentially how *they* would order nearby possible worlds. An *experienced* chicken sexer, who has excellent reflective evidence as to why she is accurate, would order nearby possible worlds in a way such that the abilities she has are shared throughout those worlds. In this case, those nearby possible worlds put the experienced chicken-sexer in a scenario where she is still able to reliably and knowingly sex chickens, and thus her beliefs should count as knowledge.

In the naïve chicken-sexer’s mind, however, she would order nearby possible worlds in a way such that those in which she still has her “abilities” of being able to see and feel the difference between the sexes of chickens are nearest. In the possible worlds now ordered, the naïve chicken-sexer does not (presumably) have the subtle pheromone

---

<sup>23</sup> It is not now, nor has it ever been crystal clear to me whether the naïve chicken-sexer should be ascribed or denied knowledge, but my reservation in coming to a decision does not affect the course of the investigation.

detecting ability in the majority of them, resulting in the conclusion that she does not have true beliefs in the majority of nearby possible worlds on this ordering, and thus does not have knowledge. There are further problems lurking here, amongst which: what if the chicken-sexer doesn't know in any way how she's successfully sexing chickens? Would she have a reflective ordering at all? What if the ability she uses is not responsible for the truth of her beliefs, but is nevertheless "tied up with" the "real" ability in some way, through either coincidental high modal correlation, or that the one ability requires the other to be performed? Consider the following scenario: while typing his essay, Dinu believes that the following process is involved: the sound of hitting a key on his keyboard summons an invisible fairy who then makes the appropriate letter appear on screen by magic. Of course, in actuality there is a process of electrical circuits being completed by the key press, but nevertheless in all possible worlds where he believes the fairy-causation theory he still manages to type successfully. Thus the condition of success in different possible worlds does not sufficiently diffuse the notion of epistemic luck: under Dinu's internal modal re-ordering, he will still be a successful typist

However, the reason why the reordering of modal possibilities in the naïve chicken-sexer example gives desirable results (i.e. shows that her beliefs are only true in a way that is knowledge-harmfully lucky) is that the agent's beliefs are based upon epistemically unstable evidence. As has been shown in the Judy and Trudy and Battle of Hastings examples, the report of the agent as to how her belief was formed can point to evidence that is not in itself knowledge-supportive, while implicit, unaccounted for evidence is doing the work. Under a reflective model, it seems that we must deny knowledge both to Eugene of the presence of the twins at his party, and to Nader as to the



date of the Battle of Hastings, yet this seems to be knowledge that we are intuitively wont to ascribe. If the intuitions as to the nature of the aforementioned examples are shared, a reflective modal notion of epistemic luck, then, does not seem to appropriately account for the effects of fading. Furthermore, a reflective model of luck does not seem to account for epistemic “grey-areas”, where the agent is partially cognitively responsible for their true beliefs. If one is only partially aware, or only partially correct as to the mechanisms the agent uses to form a belief, then under a reflective modal reordering it could end up that in possible worlds where the closeness of worlds is determined by a process determined by internal reflection, the agents’ beliefs could come out as lucky and thus not knowledge-worthy, despite that agents’ partial knowledge of the belief-forming mechanism. In short, reflective modal reordering potentially demands an unrealistically high level of self-conscious awareness of the processes that one is using to form her beliefs.

Pritchard does, however, add a carefully placed caveat: “the range of nearby possible worlds that are relevant will be restricted in terms of the way in which the agent *believes* (or would believe) she formed her belief in the actual world, rather than...in terms of the way she *in fact* formed her belief” (2005a, 175). The “or would believe” clause is explained in a footnote: “We need this qualification because the agent might be so unreflective as to have never actually thought about how it is that she is forming her belief in the target proposition” (2005a, 180). Thus we often do not remember the source of our beliefs, or sometimes even why we believe them in the first place, yet we often want to count many of these beliefs as knowledge-worthy. Is this a viable compromise when facing fading issues? Perhaps the modal ordering should be such that the closest

possible worlds are those in which I obtained my belief in the way I *actually* did, i.e. using an externalist's "actually" where implicit evidence and the non-considered implications of evidence are taken into account, as opposed to the way I *believe* that I did. If this is the case, then fading intuitions presented in the Judy and Trudy examples are satiated; unfortunately, under this restructuring the chicken-sexer example no longer becomes lucky, the phenomenon that Pritchard was trying to account for in the first place.

The upshot of the discussion in this chapter is that we require a reformulation of the incompatibility thesis with different varieties of luck in mind. Under a modal conception, we have seen that fading can occur if we describe the phenomenon as a change from knowledge-harmful to knowledge-harmless, but this process seems artificial in cases such as The Battle of Hastings and Judy and Trudy cases: after all, in these cases it nevertheless seems that the veritic luck involved is still, in some ways, present in the construction of the knowledge-worthy belief, although it is no longer knowledge-harmful in the same way. My generally externalist project also denies the incompatibility between reflective luck and knowledge, a topic I will return to in a later chapter. Nevertheless, the externalist's project has, in recent literature, been challenged by those who are more sympathetic to the internalist's case. Knowledge of the type I have described has been referred to as "brute", and it is has been claimed that the *value* of knowledge truly lies in our ability to be consciously aware of the correct and responsible methods by which we acquire it. It is to this issue that I turn now.

### Angst and Value

Linda Zagzebski states in *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*:

The value of the truth obtained by a reliable process in the absence of any conscious awareness of a connection between the behavior of the agent and the truth he thereby acquires is not better than the value of the lucky guess. (304)

Similarly, as Pritchard states:

[W]e don't just want agents to be forming beliefs in such a way that we can rely on the truth of those beliefs, we also want agents to be *cognitively responsible* for their beliefs, and this is only possible if they form beliefs in ways that are responsive to the reflectively accessible grounds that they have in favour of their beliefs. (2005a, 184)

Although Zagzebski's attack in particular is directed toward externalist treatments of chicken-sexer cases, Pritchard's criticism could be made regarding the nature of knowledge that I have examined thus far. A common reaction towards such complaints is to at least partially agree: knowledge-worthy beliefs where the agent is fully aware of all the steps and background beliefs taken into account to form the final belief may indeed be more desirable, or more valuable, than ones in which such an awareness is either limited or non-existent<sup>24</sup>. This preference, however, does not imply that beliefs formed with processes of the latter type cannot be classified as knowledge.

I tend to be sympathetic with this reaction, but with a caveat: the kind of "flawless" knowledge that Zagzebski finds almost exclusively worth pursuing is only

---

<sup>24</sup> There are, however, dissenting opinions here. See Hilary Kornblith, "Knowledge and its Place in Nature", pp.103-136, for an excellent series of accounts as to how reflection can cause more epistemic harm than good, although his arguments are mostly directed against coherentist theories.

rarely, if ever, found within your everyday belief-former. Indeed it is doubtful that ones knowledge-worthy beliefs meet the criteria of either internalists or externalists but not both, and there is rarely, if ever, a knowledge-worthy belief formed by limited, often irrational agents such ourselves such that every epistemologically relevant factor is taken into conscious consideration. Pritchard has referred to the kind of knowledge as that “which one can take full cognitive responsibility for” (2005a, 203) as involving a combined elimination of all luck involved, veritic and reflective. As I have attempted to show in this and previous sections, however, (and will explore in further detail in Chapter 4) internal reflections often prolong the knowledge-harmful effects of reflective luck long after we are willing to ascribe knowledge to the agent, and in addition reflective luck seems to fade and become salient as the result of factors wholly outside of epistemological concern, i.e. as the agent “happens to think about them”, as grudgingly expressed by Hawthorne.

So is knowledge that has resulted from a faded veritically or reflectively luckily true belief valuable in a sense similar to the one desired by Zagzebski, or is it, on some level, merely brute? Perhaps it lies somewhere in between. Of greatest importance in cases such as Jane and her watch is the extra evidence that she must be responsible for obtaining, even though she is not consciously aware of such evidence when she is acquiring it (at least inasmuch as it is pertinent to certain belief formations) or when making a knowledge claim. Indeed, here lies the crux of my opposition to certain forms of internalist constraints, namely that many, if not most knowledge-claims made in everyday life are based upon some amount of evidence that would be considered incomplete or inadequate from an internalist standpoint. As Michael Williams has noted,

Knowledge and justification always involve...an element of *epistemic luck*. A belief whose truth is *wholly* accidental cannot count as knowledge. But...getting things right is never wholly nonaccidental either (59).

Pritchard in particular reacts strongly against Willaimes' claim, saying that all knowledge must be completely luck-free, at least when involving knowledge-harmful luck, and posits that the true source of our epistemic worries is what he refers to as "epistemic angst":

[W]e are unavoidably subject to what I will call an *epistemic angst*, where *angst* is here understood as a general fear about the nature of our epistemic position which is not due to any *specific* empirical challenge to our putative knowledge. Instead, it is caused by the discovery, in the context of reflection, that the ultimate scope of our cognitive responsibility is severely restricted. More specifically, what we discover is that the kind of cognitive responsibility that we standardly attribute to ourselves is only possible given the correctness of a backdrop of anti-sceptical assumptions the truth of which we cannot be cognitively responsible for. Crucially, however, no plausible theory of knowledge (even an internalist theory of knowledge) can adequately allay the problem of epistemic *angst*, and this is the source of scepticism. Scepticism is therefore an *existential* problem, not in the sense that anyone actually endorses the sceptical conclusion, but in the sense that the source of scepticism lies in an essential feature of the human condition.

(Pritchard 2005b, 204)

Pritchard seems correct on at least one crucial point: if epistemic angst is an *existential* problem, any epistemic theory is going to be hard-pressed to allay those kinds of worries.

From an epistemic viewpoint, however, it seems that the most interesting question surrounding epistemic angst is whether it is knowledge-harmful, inasmuch as whether it is able, by itself, to prevent beliefs from becoming knowledge-worthy, and whether it is value-harmful, that is to say whether it causes knowledge-worthy beliefs to lose their value by making them merely brute. He continues,

Ultimately, what reasons we have for our everyday beliefs will be of a pragmatic, rather than an epistemic nature, in that it is only by setting aside certain kinds of error-possibilities that one can coherently engage in the practice of offering grounds in the first place...(Pritchard 2005b, 203)

Pritchard and others paint a picture of knowledge where agents, after careful reflection, realize that in any given knowledge-claim they are taking for granted a plethora of background assumptions for which they have no evidence one way or another, but that, in order to “get through the day”, as it were, these concerns are put aside, albeit perhaps reluctantly. If skeptical possibilities and potentialities of the influence of luck are ignored in order to form what we would typically consider a knowledge-worthy belief, they are done so grudgingly, that angst remaining an ever-present epistemic din that we must ignore if we are able to use the concept of knowledge at all.

I believe, however, that this picture is in many ways contrary to our natural epistemic inclinations: we often do not feel that ignoring knowledge-harmful possibilities is an act of pragmatic desperation, but is actually epistemically virtuous, given many common sorts of circumstances. The agent who is able to tell someone to “snap out it” when entertaining extremely unlikely or outright skeptical possibilities that are potentially knowledge-harmful is displaying a sensible epistemic character, not an

unfortunately necessary pragmatic vice. My line has been that it is precisely this ability to ignore factors that we consider inappropriate to entertain in belief formations that constitutes a good amount of our epistemically virtuous character. Of course, Pritchard's notion of angst is not specific to any one given knowledge claim, or based on any one skeptical possibility, but rather exists as a general epistemic ennui from which an overarching form of skepticism arises; but more interesting seem exactly those worries that *do* pertain to specific knowledge-claims, and are based upon evidence that may give credence to possibilities that were once thought to be merely skeptical and outside the concern of everyday knowledge-worthy beliefs. Thus although I disagree with Pritchard's overall conclusion, it nevertheless poses the question of when epistemic angst is warranted, and when it is knowledge-harmful. Along these lines, in the following section I will examine in more detail the correlation between hesitation to form beliefs because of certain types of evidence, and when such hesitation is epistemically virtuous and vicious.

## Chapter 4: Barn Façades and Car Crashes

The fading phenomenon has shown how beliefs that are not knowledge-worthy because of incomplete or inappropriate evidence can become knowledge-worthy as the result of the accumulation of supporting evidence, implicit or explicit, obvious or implied. The motivation for the investigation stemmed from the counter-intuitive results of theories like Pritchard's notion of reflective luck, which puts an emphasis on the internalized construction of beliefs by the would-be knower, and its subsequent incompatibility with knowledge-worthiness. I have proposed a different approach, namely that potentially knowledge-destroying possibilities can be safely ignored if they have sufficiently faded. It is not so much my concern that "sufficiently faded" is in need of qualification; rather I wish to examine which kinds of possibilities are subject to the fading phenomenon in general. I will begin by examining how experiencing unlikely phenomena can make their possibility temporarily and acutely salient, presenting a knowledge-harmful factor which is then susceptible to fading. On this account, it is not simply the agent's cognitive awareness that determines whether a given possibility can be ignored in an epistemically responsible manner, but rather a possibility becomes safely ignorable only after it has sufficiently faded.

### Car Crashes

We are often wont to ascribe knowledge-worthy status to "everyday" propositions such as "I will be alive in five minutes" or "I will be at the meeting in an hour."<sup>25</sup> At the same time we may recognize that there exist a number of possible events that whose

---

<sup>25</sup> Intuitions are again somewhat divided over so-called instances of "knowledge of the future", some saying that it is never possible, regardless of the example used. I use the cases presented in this section more out of convenience and succinctness than anything else, although counterfactuals or more convoluted statements regarding the present or the past could be substituted without loss of comprehensibility.



occurrence we cannot exclude, which might prevent our initial claim from becoming true: that I am otherwise perfectly healthy but am struck by a sudden and fatal heart attack in 3 minutes time, or that on my short drive from my home to my office I am involved in a car accident, or, potentially more unlikely still, I am struck and killed by an oblivious driver as a pedestrian. Yet if we make the knowledge claim in the first place, either we do not consciously consider these factors, or we feel, at least to some extent, that such factors are not important enough for us to retract the claim, or to prevent us from forming the belief in the first place. Furthermore, if it turns out that we *are* correct (i.e. that I *am* alive in five minutes, I *do* get to my meeting on time, etc.) then we do not typically think that we are correct *in spite of* those other factors. That is to say that we do not consider ourselves lucky that our belief is true, but rather that we are correct because of some virtuous epistemic reasoning (if we consciously think about our epistemic character at all). Anticipating discussion in a later section, I wish to focus on two factors surrounding these cases: first, to continue the example, that it is potentially the case that being killed in a car accident (in whichever scenario we choose) has a “greater chance” of occurring than winning the lottery, yet we are more likely to ignore the chance involved in the former than in the latter. This suggests a fundamental difference in kind between what I will henceforth refer to as “car crash cases” and “lottery cases.”<sup>26</sup> Second, that fatal car crashes, sudden and unexpected heart attacks, and many other unlikely events *do* occur, and we have experiences, with varying degrees of intensity and propinquity, of such events. When we experience such events, we typically re-evaluate the salience of such kinds of possibilities, and may similarly re-evaluate the processes we used to form beliefs that we once considered knowledge-worthy.

---

<sup>26</sup> So far, this is nothing new. But bear with me.

To outline the problem at hand, I first look to John Hawthorne, who, when presented with the task of determining which factors need to be considered when making knowledge claims or attributions, states that we need to consider those factors which are “epistemic possibilities for the agent”, appealing to the intuition that if said agent consciously believes that he is making an oversight (i.e. by ignoring alternatives of which he is consciously aware), then we are wont to count their beliefs as not being knowledge-worthy. This prescription, however, results in two undesirable consequences: first, that the fewer alternatives that the agent consciously considers, the more knowledge we are forced to ascribe to her, and second, that the relevance of potentially knowledge-defeating factors is subject to the agent’s entertainment of them. The line I wish to explore, then, is that relevant alternatives become less relevant (i.e. fade) the more evidence (either positive or negative) we have to rule them out. In the coming sections I will examine theories that contain internalist components to make their shortcomings more pronounced, and examine the different kinds of evidence we typically use to rule out alternatives, and the epistemic impact of various other types of evidence.

### **After the Barn Façade**

In order to frame the problem, I will begin by examining the well-known barn façade example, in which our hapless protagonist unknowingly drives into “barn façade country” and, as luck should have it, correctly identifies the only real barn for miles around. Typically the story ends here: our agent does not have knowledge because his belief just-so-happens to be true. Traditionally the reason why has been that the existence of fake barns is a relevant alternative that the agent has not sufficiently considered when making his knowledge-claim. Other explanations put forth in the

literature include: that the agent is an unreliable real barn detector, or that he does not track the truth in the situation properly (i.e. had he looked at a fake barn, so that the proposition  $p$  = “there is a real barn” would be false, the agent would still believe it<sup>27</sup>). Pritchard’s modal model seems to attempt to cover these intuitions, i.e. the agent’s correct indication of there being a real barn is a modally rare occurrence. As the type of luck covered by Pritchard’s modal theory and more traditional relevant alternatives theory is coextensional when it comes to barn façade cases, I feel it a safe practice to simply refer to such cases as relevant alternatives, and will leave it to the reader to choose their preferred theory of luck (at least for now).


Consider now an extension of the example, where we continue to follow the agent to a rest-stop just outside of barn façade country, at which an attendant asks if our protagonist enjoyed his drive past the many convincing barn façades for which the region is a popular tourist destination. It is safe to assume that our agent would be confused, and would most likely, if prompted, second-guess his previous claim to know the proposition  $p$ : there is a real barn. Whether he does so question his own ability is again perhaps not of the greatest concern; what *is* of concern is the rest of his drive, which will, for argument’s sake, take our agent past a number of real barns. Indeed now our agent would most likely be very wary of claiming to know “there is a real barn” as he has just recently been notified of the existence of barn façades in general, and more acutely that they exist in abundance, and very close by (after all, if someone has taken the time to erect all those façades, who is to say that there are not many more?). He has realized that either he is unsure of whether he made a mistake in his initial claim, or if he stubbornly insists that he

---


<sup>27</sup> There is much literature surrounding this account, namely the “truth-tracking” account put forth by Robert Nozick. For a recent defense of the truth-tracking account, see Sherrilyn Roush, *Tracking Truth: Knowledge, Evidence, and Science*, 2006.

continues to know his original proposition then he at least realizes that he could have made a mistake very easily, which will also most likely affect his belief formations in the near thereafter. Indeed, we need only be “very nearly bitten” in order to be “twice shy”, and often feel it prudent to be so, even in matters epistemic.

Our agent is now in possession of a relevant alternative that he may not have previously considered. Here the internalist constraint again rears its head: our agent has attached relevance to an alternative that he either had not previously considered, or considered the possibility thereof safely ignorable, thus preventing his previous belief from being knowledge-worthy, and preventing further belief-formations in the future:



**Figure 9: Formation of the belief of p: there is a real barn, from the viewpoint of the agent, before learning of his presence in barn façade country**



**Figure 10: From the viewpoint of the informed agent, the salience of barn façades makes the original belief true only as a matter of luck**

But we should not want to say that our agent’s barn identification skills have been completely epistemically neutralized, or if not neutralized then conclusively shown to be inadequate in certain contexts<sup>28</sup>. It is counter-intuitive to think that once we have made a mistake of the type presented in classic barn façade cases that we cannot make knowledge claims concerning similar situations in the future. In other words, we would

<sup>28</sup> As his phenomenological experience of both barns and barn façades are identical, say, from the driver’s seat of his car, we might conclude that from this context our agent’s abilities are inadequate; if he exited his vehicle and looked all around the structure, however, he would not be so fooled by the façades, and thus in this context his abilities would remain adequate.

expect the relevant alternative of barn façades to be knowledge-damaging to barn-related propositions, but presumably not permanently. Our epistemology should be able to handle such phenomena. The question becomes: how damaging are relevant alternatives, and for how long?

Hawthorne introduces a version of invariantism called “sensitive moderate invariantism”, a theory that focuses on the epistemic concerns of the potentially knowing subject, as opposed to the knowledge-ascriber. The “moderate” invariantist still claims that knowledge is not context-dependent, but in addition is willing to allow for a more “relaxed” view of knowledge, so that, say, one could well have knowledge that was in violation of the epistemic probability constraint. Left unmodified, this “simple” moderate invariantism suffers from a number of undesirable consequences, including denials of some forms of epistemic closure<sup>29</sup>, and tensions when considering practical reasoning<sup>30</sup>. In order to make this base theory more robust, Hawthorne takes a page from contextualism and includes factors such as the attention of the subject, their practical interests, and their perceived stakes of the situation into the knowledge formula. As a result, salience is essentially entirely determined by the interests of the subject, and concerns of the “practical environment” are given an active role in the knowledge formula.

Two problems present themselves. First, the “practical environment” condition is vague at best: “Insofar as it is unacceptable – and not merely because the content of the

---

<sup>29</sup> Specifically, multi-premise closure. For example, by allowing knowledge to be compatible with the epistemic possibility constraint, one could claim to have knowledge of some lottery propositions. Consider the case where there is a lottery with 1000 tickers. For each ticket individually, such a theory could allow the knowledge that that ticket will lose. As a result, one could claim for each of the first 999 tickets that they will lose, and thus be able to deduce that the one remaining ticket would be the winner. But this is absurd.

<sup>30</sup> I will not go into the gory details here. Instead, see Hawthorne (149).

belief is irrelevant to the issues at hand – to use a belief that  $p$  as a premise in practical reasoning on a certain occasion, the belief is not a piece of knowledge at that time” (176). To clarify this formulation somewhat, Hawthorne offers the following example: “Thus when offered a penny for my lottery ticket [which cost a dollar] it would be unacceptable to use the premise that I will lose the lottery as my grounds for making such a sale” (176). Granted that this scenario seems intuitive enough, but the example lies on the extreme end of the spectrum; the intuitions surrounding these types of cases become significantly less clear as the variables are changed even slightly. It is less clear, for example, whether I should accept 50 cents for my ticket instead of 1, less clear still whether I should accept just slightly less, or even slightly more than the initial dollar I spent.

The second consequence is that by putting salience much more firmly in the mind of the subject rather than the ascriber, epistemic anxiety rules supreme. While Hawthorne is correct to note that salience does not imply mere attention, he allows subjects to be saved from skeptical suggestions, but not from themselves. The skeptically minded subject no doubt takes his own skepticism seriously, but in doing so prevents his own knowledge to an intuitively unacceptable degree. Hawthorne’s solution is that epistemic possibilities for the agent “come and go” as we think or abstain from thinking about them (Hawthorne, 176): when we have the conscious wherewithal to consider them important, they are knowledge-harmful, but if they are not part of our conscious considerations then they do not factor into our knowledge equations.

There are two immediate consequences of this view as related to our extended example:

- 1) that if, at any point in the future, our agent should, for some reason or another, suddenly make the memory of his trip through barn façade country occurrent and decides that barn-existence claims are of a high epistemic priority, then he is, at least until he stops thinking about it, unable to make knowledge-worthy claims pertaining to the existence of barns<sup>31</sup>, and,
- 2) if our agent happens to be somewhat absent-minded (i.e. he quickly forgets about the existence of barn façades, for whatever reason), he would be able to make a knowledge-worthy claim as to the existence of a real barn *immediately* after leaving barn façade country.

I think that both of these consequences present hard bullets to bite, but perhaps more importantly, imply that Hawthorne's conception of knowledge is overly transient, and depends upon an equally transient degree of relevance in our relevant alternatives, one that seems to be affected exclusively by factors internal to the agent, i.e. his practical environmental concerns<sup>32</sup>. I have already attempted to defend a theory that takes a familiar externalist viewpoint that states that it is not strictly necessary for the agent to be consciously aware of alternatives for them to be knowledge-harmful. More importantly, the relevance of a relevant alternative additionally hinges on the actions of the agent, and similarly alternatives need not be harmful simply because they are taken seriously. Thus a relevant alternative becomes less relevant as the agent gathers positive or negative evidence, explicitly or implicitly, that either reinforces the ability of the agent as a

---

<sup>31</sup> Assuming, for example, that he is unaware of the degree of ubiquitousness of barn façades in general, or, say, the geographical boundaries of barn façade country.

<sup>32</sup> Note that Hawthorne himself recognizes some of the potentially counter-intuitive results of his theory. He notes that "As Timothy Williamson pointed out...this picture makes for some intuitively odd counterfactuals. There will be cases where someone does not know that *p*, but we can assert 'If the stakes had been lower, he would have known that *p*'... The question is whether, all things considered, it is a price worth paying." (177) It is my contention that the price is high enough to warrant a serious attempt at an alternative theory.

knower in certain circumstances, or indicates that the same type of relevant alternative will not again or is extremely unlikely to again occur.

The nature of this evidence is a topic I will cover shortly, but for now it is enough to notice that, following these lines as I have set them out, we are able to avoid another potentially counter-intuitive consequence of Hawthorne's theory, being that, roughly, the fewer alternatives we consider when making knowledge claims or ascriptions, the more likely we are to claim to know or to ascribe knowledge. Under the proposed restructuring, the theory would yield that the epistemological virtue of being a scrutinizing and thoughtful agent would remain virtuous (i.e. having a disposition to form knowledge-conducive beliefs), since said agent would be less likely to make knowledge claims or attributions when alternatives are relevant, while at the same time making the thoughtless agent epistemically vicious (i.e. prone to forming beliefs which are not knowledge-worthy) since they would tend to ignore such alternatives. The next step is to examine the nature of "relevance", which I will approach now.

### **Digging Deeper: Relevance as an Internal or External Phenomenon**

There is another case to consider in the extended barn façade example: the mundane situation where our agent makes a true claim as to the existence of a barn and is denied knowledge because of the luck involved, but who never learns of his circumstances, i.e. that he had, in fact, driven through barn façade country at all. In this situation the agent does not consider the presence of barn façades as relevant alternatives, and thus is likely to make similar claims to knowledge in the future, if he is so inclined. Thus the tricky scenario is where the agent, ignorant of his knowledge-unworthy claim and brush with making a false claim, finds himself afterwards in a normal scenario, that



is where no Gettier-factors are hiding around the bend. To continue the theme, say our agent drives out of barn façade country and into the regular countryside, where there are no barn façades but instead a multitude of real barns. Our agent now claims knowledge of the same proposition  $p$ : there is a real barn, not for an instant taking into account the possibility of deception, for, from his perspective, he has no reason to believe he could be so deceived.

In the first of the presented situations where the agent is not ignorant, we might be inclined to say that he is being epistemically sloppy should he claim  $p$ ; after all, he has recently learned of his close brush with deception, and therefore should be more careful in the future before making such claims. Thus in the scenario where our agent is aware of a relevant possible alternative, yet that alternative is not actualized (i.e. in the regular, real-barn spotted countryside), our wariness to ascribe knowledge stems at least partially from the agent's internal constraint, and we consider our wariness to be justified. Where the agent is ignorant, however, and finds himself in the same scenario where no potentially relevant alternative is actualized, nothing in our theories of luck tell us that he should not have knowledge in this situation: to cite a few such theories, indicating a barn where the context contains no fake barns is a modally common situation, the agent exists in the right causal connection with the object of his proposition, his continued correctness makes him a reliable barn-spotter, etc. Thus, according to the popular theories surrounding epistemic luck, the ignorant agent's claim should be knowledge-worthy, while the informed agent's claim to the same proposition should *not* be considered knowledge-worthy. This is the dilemma I presented earlier: we *should* be hesitant to ascribe more knowledge to the ignorant agent as opposed to the informed one, but our

current theories give us no reason to prevent the ignorant agent from knowing, and in particular give us no reason to think that the informed agent is epistemically better than the ignorant agent.

Furthermore, while Hawthorne in particular is adamant that we need not take any and every possibility seriously simply by virtue of its mention (Hawthorne, 64), his theory does not prevent the particularly neurotic agent from seriously considering any number of miniscule possibilities, and thus by his own accord prevents himself from knowing very much at all. To continue Hawthorne's example (Hawthorne, 64), after having finished watching *The Matrix*, if our friend is seriously considering the possibility that he is the more technologically upgraded brain-in-a-vat as portrayed in the film, then we should be able to tell him to "snap out of it": even though he may be taking it seriously, we typically think that brain-in-vat scenarios are not possibilities that should be taken seriously in general<sup>33</sup>. Thus, there seems to be more to the story than just what the agent is seriously considering according to internal reflection, but rather what he *should* be considering. Consider, then:

*Principle of Faded Alternatives:* Alternatives are relevant and thus potentially knowledge-preventing if their significance has not adequately faded.

In other words, an alternative remains relevant until additional evidence is acquired such that it loses its relevance in a manner similar to cases outlined earlier. Note that according to the above principle, it is not the agent's internal reflection that is knowledge-harmful (either in the sense of preventing previous beliefs from being knowledge-worthy or preventing the formation of additional beliefs), but rather the agent

---

<sup>33</sup> One may again appeal to a sort of epistemic angst in these kinds of scenarios; although, to reiterate, it is unclear how knowledge-harmful epistemic angst really is.

is epistemically virtuous if she considers relevant those possibilities which have not adequately faded, and is epistemically vicious if she does not; and in this latter case, should her belief be true, will be true only as a matter of luck, and not knowledge-worthy.

This principle naturally leads to another ingredient in the epistemic luck recipe:

(L4) A belief held by an agent S is lucky if the relevance of possible alternatives for S have not sufficiently faded in significance.

This condition, of course, is only concerned with the epistemic side of luck, as opposed to criteria such as those found in Pritchard's theory which attempt to cover the notion of luck more generally<sup>34</sup>. Adding (L4), however, seems to account for at least some of the loose ends found in Pritchard's recipe: in particular, with (L4) as an additional premise, the difference between knowledge-harmful and knowledge-harmless luck becomes much more overt when dealing with a modal conception. I also believe that (L4) gives many of the same results provided by Pritchard's notion of reflective epistemic luck, but without the unfortunate consequences of reflection preventing more knowledge than it should.

This chapter began with a brief look at a classic paradox of epistemic reasoning, namely that it does not seem that we always base our beliefs on the basis of the probability of the occurrence of potentially defeating alternatives. Thus there is a conflict between which alternatives we perhaps think we *should* consider more strongly and the alternatives we *actually* consider. The lottery case is a prime example: we would most likely attribute a higher probability to being in a car accident than winning the lottery, yet we are very wary to make knowledge-claims concerning matters of the latter, and much more liberal in matters of the former. Solutions as to why this is so have ranged from

---

<sup>34</sup> Although as I have attempted to show previously, I do not believe that they handle epistemic issues in an entirely appropriate way

conceding that the typical agent is, in actuality, a generally poor reasoner to the stakes and relevance theories I have been dealing with throughout. As I have alluded, I do not generally consider the former to be true, and the latter contains a number of flaws that I have been highlighting at various points in the last four chapters. Seeking a middle ground, I now look in more detail to the specifics of the fading phenomenon.

## Chapter 5: The Effects of Evidence

### **Fading and Propinquity**

As I have already alluded, there exists a correlation, at least upon internal reflection, between the nature and propinquity of an experience and its epistemic impact. For instance, if one makes the normally knowledge-worthy claim that they will be at a very important meeting on time by means of car travel, being in a car accident not only makes that belief no longer knowledge-worthy (since the original claim then turns out false) but will also make the previously ignored alternative a relevant factor in future knowledge-claims and belief formations. Thus when asked if the agent will be at a particular place and time reachable by car in the future, the possibility of being in an accident may remain a relevant alternative, preventing the formation of a belief. If the agent is temporarily absent-minded, or stubborn, she may choose to form such a belief regardless; however, if her previous experiences are once again made occurrent, say by suggestion or recollection, she may end up re-evaluating the structure of her belief and retract the claim. Thus having the experience of an alternative once considered irrelevant can make that alternative knowledge-harmful, assuming, at least for the time being, that the agent's reflection is in some ways knowledge-affecting.

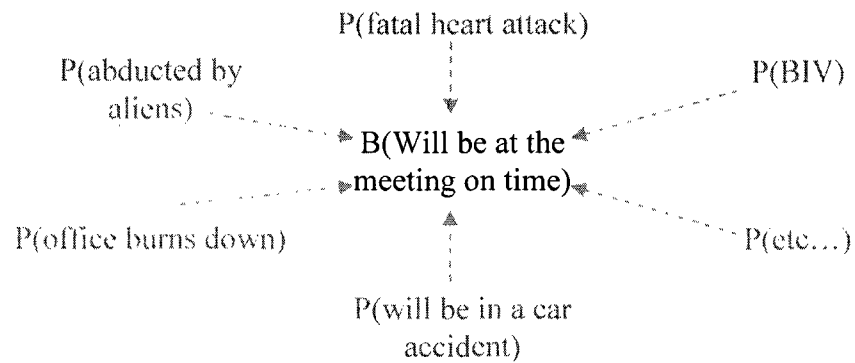
Additionally, it seems that the greater the “intensity” and propinquity of the experience, the more likely the agent is to let that alternative be knowledge-harmful in the future. This is an effect often found in so-called “everyday” cases of knowledge claims and attributions, and I believe it is prudent to examine the intuitions behind those cases, and use them to help further a theory of knowledge that is intuitively valuable. An experience of a possible alternative first-hand will likely be more knowledge-harmful or knowledge-harmful for a longer period of time, than those that are heard about second hand or through some more distant source. Hawthorne mentions a similar phenomenon briefly, in discussion of “availability heuristics”:

[I]n many cases, our estimation of the likelihood of an event is affected by the ease with which we can recall or imagine it. So, for example, when a certain scenario is made vivid, the perceived risk of that scenario may rise dramatically (Hawthorne 2004: 164).<sup>35</sup>

To continue the example, being in a car accident is typically more likely to affect beliefs that would usually take such possibilities for granted than merely hearing about a similar accident on the radio; although in the case of the latter the alternative is still potentially made salient, but perhaps only for a short period of time (potentially a longer period for the particularly neurotic agent). Consider, then, the following status of an agent’s beliefs at the time of forming the belief B(will be at the meeting on time):

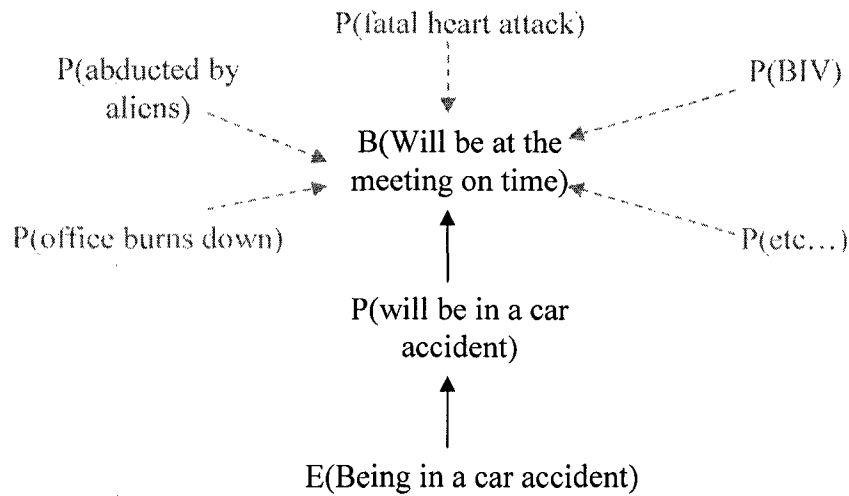
---

<sup>35</sup> Hawthorne takes his information from studies by Slovic, Fischhoff and Lichtenstein in their 1982 study “Fact Versus Fear: Understanding Perceived Risk” (in Kahneman, Daniel and Amos Tversky, *Judgment Under Certainty: Heuristics and Biases*. Cambridge: Cambridge University Press). The more detailed epistemological consequences of uses of the availability heuristic are outside of the scope of the current investigation (but are nevertheless terribly interesting).



**Figure 11: The belief is knowledge-worthy when the background assumptions (grayed) and their knowledge-preventing effects (grayed arrows) are safely ignored**

Figure 11 shows a belief with a standard complement of background assumptions taken for granted as not knowledge-harmful, either consciously or unconsciously (shown in gray). If the agent has an experience of a car accident, with some degree of intensity  $i$  and propinquity  $\rho$ , we see an effect on beliefs as shown in Figure 12, which shows the highlighted alternative, made salient by the evidential support, and displays a line of reasoning that might be possessed by an agent from an internal reflective position (note that it shows a specific instance of a possibility becoming salient, that is one that is based upon direct evidential support, as opposed to being recalled from memory).



**Figure 12: The possibility of being in a car accident is made salient and knowledge-harmful by the experience of a car accident**

While  $P(\text{being in a car accident})$  is salient, it presents a knowledge-harmful variety of luck, inasmuch as it presents a relevant alternative that cannot be internally disregarded by the agent, and thus if the agent still chooses to form  $B(\text{will be at the meeting on time})$ , and it turns out to be true, it will be so only as a matter of luck. Thus whereas in previous knowledge-claims and belief formations, such as those that are represented by Figure 11,  $P(\text{being in a car accident})$  would have been ignored<sup>36</sup>; it cannot, however, be so ignored at the time of Figure 12. The upshot is that if a background assumption that an agent has taken for granted enters the consciousness of said agent, then she should consider that kind of event, at least for the time being, a relevant alternative. While suggestion does not necessarily imply salience, experience does (although to varying degrees).

### Activation Levels

Again from an internal reflective viewpoint, different possibilities require different amounts of evidence for the agent to consider them relevant, what I will refer to

<sup>36</sup> For the sake of argument, I will assume that it was ignored in an epistemically virtuous manner.

as the “activation level” of a given possibility. For example, it might be that hearing of a car accident on the news may make such a possibility temporarily salient for one agent while it may not for another. Similarly, the *duration* of this salience is affected by similar factors, assuming the belief structure presented in Figure 12 is not ever-present, i.e. that an agent with such a belief structure will eventually be able to make knowledge-worthy claims concerning propositions concerning making it to work on time<sup>37</sup>. So far this is similar to the pictures sketched by Hawthorne and Pritchard inasmuch as factors internal to the agent determine knowledge-conduciveness or harmfulness; for now I have merely added that these factors are often made salient by the effects of a particular intensity and propinquity of experience. It also assumes, at least for the time being, that if an agent consciously believes that she does not know a proposition, then she does not know that proposition, regardless of the amount of evidence in its favor. In other words, under this and many internalist models, the agent has epistemic veto power.

Consider, for ease of future reference, the following formalizations: first, that the salience of a potential alternative varies directly with the amount of evidence accumulated in support of that possibility:

Salience of possibility  $\propto$  amount and strength of evidence for that possibility

The knowledge-affecting power of evidence for a possibility can be further expressed by a number of factors, some of which might include: intensity, propinquity, and trustworthiness of the source of the evidence<sup>38</sup>. Thus,

Knowledge-affecting power  $\propto$  propinquity, intensity, and trustworthiness

---

<sup>37</sup> By “duration” I do not strictly mean “amount of time” but rather that a possibility with higher epistemic duration affects a larger number of beliefs in the future for the agent.

<sup>38</sup> This is perhaps an incomplete list, and the terms themselves are under-defined. I use these formulae as a skeletal model as opposed to a fully fleshed-out theory.



Formalized further, for any given possible alternative  $P(x)$ , the salience of that possibility,  $P(x)_s$  is equal to the amount of accumulated evidence,  $E(x)$  in its favor, or:

$$P(x)_s = \Sigma E(x)$$

While the nature of that evidence is affected by the propinquity, intensity, and trustworthiness of the source, or approximately:

$$E(x) = \rho_x i_x t_x^{39}$$

The next question logically becomes: how large does  $P(x)_s$  need to be in order for  $P(x)$  to become a relevant alternative in the agent's belief formation? I have suggested that each possibility requires a certain quality or quantity of evidence to make it relevant, an "activation level" which I will designate  $P(x)_a$ . Thus, as a general rule, if  $P(x)_s > P(x)_a$ , then  $P(x)$  is a relevant alternative *for the agent* and thus is potentially knowledge-harmful. Jason Stanley, in *Knowledge and Practical Interests*, seems to endorse a similar type of model, where the "activation levels" as I have called them are represented by their stakes-theory analogue expressed by whether the proposition under scrutiny is a "serious practical question" for the agent. Stanley's example involves two agents who both claim to know that a given bank is open on a Saturday: in the situation where the bank being open or closed has serious no practical consequences (i.e. the "low stakes" case), the agent is able to know with less evidence than the agent for whom the bank being open or closed is of considerable concern (i.e. the "high stakes" case) (Stanley, 96-7). Applying

---

<sup>39</sup> Missing from these formulae are, of course, any sort of "units" or detail as to how to quantify any individual variable. Nor do I intend to involved myself in such issues; the formulae presented in this and other sections should be treated as illustrative, as opposed to mathematically viable.

these cases to the current issues, the low-stakes agent has the activation levels of possible defeaters set high, whereas the high-stakes agent has them set relatively low<sup>40</sup>.

The flaw in this construction is that it puts too much control in the reflective capabilities of the would-be knower. The exclusively internal reading is, at the same time, too strong and too weak of a model: too strong, as the particular epistemic views of the individual agent can prevent much more knowledge than we think is typically prudent (i.e. as exemplified by the aforementioned “neurotic” agent, the solipsistic philosopher or the naïve freshman after having read the *Mediations*), and too weak, as it allows neglectful agents to exercise their own kind of epistemic viciousness by ignoring the effects of particular experience, or internally setting the activation levels for alternatives very high (e.g., they let very little “bother them”, epistemically). Again, using Stanley’s theory as a model, being in a high or low stakes situation seems to be based at least in part (if not entirely) upon the concerns of the agent, which leads to the perhaps strange conclusion that the fewer high-stakes situations you put yourself in, or if you generally deal with unimportant epistemic issues, then you are more capable of possessing knowledge. Following the internalist project, we will no doubt be presented with cases where we want to ascribe knowledge in the former case although the agent denies it<sup>41</sup>, and prevent knowledge in the latter case although the agent claims it.

Of course, as has been discussed earlier, the luck involved in certain types of knowledge claims can *wear off*: my line is that in cases such as those presented in Figure 12, the presence of the relevant alternative can become less relevant, and knowledge-

---

<sup>40</sup> It is worth mentioning that Stanley is putting forth an “interest-relative *invariantism*.” There are perhaps a good number of possible connections between Stanley’s theory and the one I have presented thus far, but I will make such connections explicit only when they help to clarify issues I am raising.

<sup>41</sup> Intuitions are varied as to whether agents possess the aforementioned “epistemic veto power”. My intuition is that they do not.

claims can once again be made once the alternative has ceased to become knowledge-harmful. Again, it is not so much my concern to examine specific types of cases to determine precisely how long it takes for different kinds of alternatives to become less relevant, but rather that the phenomenon, first of all, relies on evidential support, second, makes a demarcation between being epistemically virtuous by considering those factors which have not sufficiently faded, and being epistemically vicious by ignoring factors which have not sufficiently faded, and third, occurs regardless of the agent's consciously awareness or lack thereof. As a brief sketch, say that an epistemic alternative becomes able to be ignored as the salience of that alternative drops below its activation level, i.e. when  $P(x)_s < P(x)_a$ . For this to occur, evidence that gives credence *against* the alternative, i.e.  $P(\sim x)$ , needs to be accumulated by the agent. The fading that occurs in such cases as that shown in Figure 12 shows evidence that goes *against* the possibility made salient, "defeaters of defeaters", if you will. Of course, the evidence for  $P(\sim x)$  will most likely be of the form of what I have been calling "implicit evidence." It is to the nature of this evidence that I turn now.

### **Implicit and Explicit Evidence**

The ideas I have presented thus far have pivoted around evidence, specifically the accumulation of evidence that is able to make knowledge-harmful luck harmless by confirmation, a situation discussed in the case of Jane and her watch, as well as that evidence that is able to defeat potential knowledge-defeaters so that the agent may once again safely take for granted a number of background possibilities, as discussed in the car accident case. The types of evidence I have mentioned have been "explicit" and "implicit", terms that are in need of clarification. Explicit evidence consists in the

agent's consciously aware accumulation of information with which she is then able to modify her beliefs. For instance, in the case of the car accident above, the explicit evidence that an agent receives in the form of either being in a car accident herself or hearing about one on the radio are both kinds of explicit evidence, in her case ones that make background defeaters salient and prevent belief formation or the knowledge-worthiness of other beliefs.

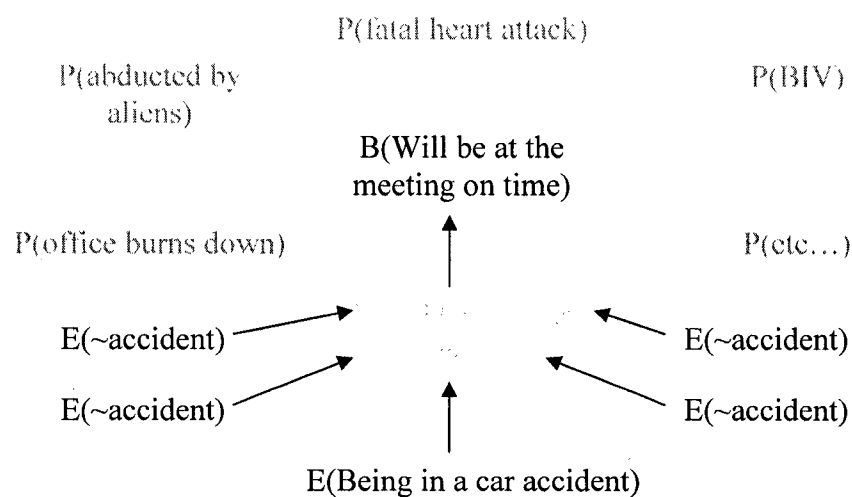
Implicit evidence comes in two varieties: first, evidence that is consciously accumulated by an agent, but affects her beliefs in a way that she is either unaware of entirely, or affects her in a way that is different from how she believes he is being affected. For instance, in the case of Jane and her watch, consistently catching her bus and favorite television show on time presents a kind of implicit evidence: although she is consciously aware of catching the bus, she is not consciously aware that such an activity also implies that her source for information about the time is correct. This kind of implicit evidence is also at work in the Judy and Trudy case, where the evidence acquired by the party host logically implies additional evidence of which he is consciously aware, but which he does not actively believes plays a role in the formation of his knowledge-worthy beliefs<sup>42</sup>. The second type of implicit evidence involves the *absence* of certain explicit evidence, inasmuch as an event *not* occurring can be good evidence for the non-occurrence of that event. This type of evidence can be consciously observed and experienced by an agent and used to form beliefs, what I will call conscious-implicit evidence, or accumulated without conscious awareness, what I will call unconscious-

---

<sup>42</sup> While reading this section, the following questions no doubt come to mind: which is epistemically "stronger", explicit or implicit evidence? Should we strive to acquire one kind rather than the other? No doubt there are others. My initial reaction to these two questions is that explicit evidence would seem to be capable of providing a greater epistemic impact overall, and that it should be given some kind of preference over its implicit counterpart, but I will not defend these views here.

implicit evidence. In the case of the former, a belief in the non-occurrence of an event can become more justified given the continued evidence of its not occurring: for example, the continued non-existence of a black swan can give credence to the belief that all swans are white (unless, of course, a black swan is discovered), and the evidence is conscious inasmuch as the White Swan Society is consciously concerned with it.

It is the unconscious-implicit evidence that I will focus on in this section. In the case of the agent who forms a belief similar to that displayed in Figure 12, the defeater of having an experience of a car accident can itself be defeated by an accumulation of evidence. The most common source of this evidence, I propose, is this unconscious-implicit evidence; that is to say, in this example specifically, a continued stream of evidence that involves the agent driving safely provides good evidence that a car accident is unlikely. Thus if our agent is herself involved in an accident while driving, the more safe trips she makes to her office or elsewhere give unconscious-implicit evidence that can, in sufficient quantity, defeat the knowledge-preventing defeater of the possibility of being in an accident. Hence Figure 13 shows a potential next-step in the epistemic process from Figure 12:



**Figure 13: Unconscious-implicit evidence of a car accident being unlikely defeats the knowledge defeater  $P(\text{will be in a car accident})$**

After enough evidence,  $E(\sim\text{accident})$ , is collected,  $P(\text{will be in a car accident})$  ceases to be salient, and retakes its former position as an ignored background assumption. A number of consequences and clarifications need to be made: first, not every defeater is capable of being defeated by the use of unconscious-implicit evidence, a topic I will discuss shortly in the section on lotteries. Second, it is again unclear *exactly* how much evidence is needed to defeat the defeater  $P(\text{will be in a car accident})$ ; to reiterate above, as a guideline  $P(x)$  ceases to become a relevant alternative once  $P(x)_s < P(x)_a$ , but these terms simply present the general case; presumably values will differ given different kinds of events. Third, this model implies that it is not solely within the agent's capacity to recall potential defeaters and by doing so making them knowledge-harmful: as the agent additionally possesses the unconscious-implicit defeaters of  $P(\text{will be in a car crash})$ , said defeater loses its knowledge destroying power.

Again this implies an externalist's reading: although the agent may herself be perturbed by the possibility of any number of potential defeaters, the implicit evidence she has received makes her belief knowledge-worthy. At least, that is, until it is reinforced by similar kinds of evidence. Which leads to the final upshot, that the ability to form knowledge-worthy true beliefs is in a continuous state of flux, where defeaters are given credence until they are themselves defeated by evidence to their contrary, at which point new evidence could be obtained, strengthening the initial defeater, at which point the whole process may repeat. This "flux", however, is evidentially based, and

thus, I believe, is more strongly grounded than a theory where the whim of the agent is able to sufficiently affect her knowledge-acquiring power<sup>43</sup>.

### **Unfortunate Consequences?**

The theory presented in this chapter, as it stands without further qualification, results in a number of undesirable consequences; specifically, it seems to give credence to certain fallacies of reasoning. Consider the following situation: Rosencrantz and Guildenstern are betting on the toss of a series of fair coins, a result of heads meaning Rosencrantz wins the coin, one of tails meaning the coin goes to Guildenstern. As it turns out, however, there has been an incredibly long run of heads, and Rosencrantz is becoming very confident that the next outcome will follow the same pattern. Applying the coin tossing case to the theory presented in this chapter potentially poses a problem: it seems to imply that, given that the possibility of the coin coming down tails is a relevant alternative for it coming down heads, then the explicit evidence of the coin coming down heads multiple times should be enough to defeat this potential defeater (say, for emphasis, that the string of heads is particularly long). If we model the coin-tossing case on the car-accident case shown in Figures 12 and 13, it was claimed that with the frequency of accidents becoming increasingly rare one is able to responsibly base beliefs upon the assumption that they will not be in a car accident in a given situation, so what is to prevent Rosencrantz from eventually being able to responsibly base a belief on the assumption that the next toss of a coin will come down heads, since it has happened so often?

---

<sup>43</sup> This theory is also not intended to cover merely “extraordinary” cases of belief formation and evidence acquisition; on the contrary, it aims to capture something about the nature of “mundane” or “everyday” knowledge-worthy belief formation.

Another problem has been lurking: I have been thus far referring to relevant possible alternatives in terms of the probability of their occurring, but have said nothing about whether knowledge is compatible with a certain degree of probability or if complete certainty is required for a belief-structure to be knowledge-worthy. Indeed there are two intuitions surrounding this notion that need to be addressed, which are summarized by Hawthorne:

*Epistemic Possibility Constraint*: if the epistemic probability for [agent] S that  $p$  is not zero, then S does not know that not- $p$ . (Hawthorne, 112)

Combine the above with the following:

*The Objective Chance Principle*: If at  $t$ , S knows that there is a nonzero objective chance that  $p$  at  $t$  (where  $p$  supervenes on the intrinsic facts about the future relative to  $t$ ), there is a nonzero epistemic probability for S that  $p$ . (Hawthorne, 92)

Which results in:

*The Chance-Knowledge Principle*: If at  $t$ , S knows that there is a nonzero objective chance that  $p$  at  $t$  (where  $p$  supervenes on the intrinsic facts about the future relative to  $t$ ), then, at  $t$ , S does not know that not- $p$ . (Hawthorne, 93)

And although not explicitly discussed by Hawthorne, the *Chance-Knowledge Principle* naturally leads into the following:

*The Lucky Chance Principle*: If at  $t$ , S knows that there is a nonzero objective chance that  $p$  at  $t$ , and claims to know not- $p$ , then if not- $p$  is true, S's belief is true only as a matter of luck, and thus is not knowledge-worthy.



The problem is that both the *Epistemic Possibility Constraint* and the *Objective Chance Principle* are intuitively viable, and since the *Chance-Knowledge* and *Lucky-Chance* principles follow directly from them, what is to prevent us from thinking that most, if not all of our beliefs are knowledge-unworthy due to a pervasive aspect of luck? It is to these issues that I turn in the final chapter.

## Chapter 6: Lotteries Don't Fade

### Outlining the Problem

With this in mind we now turn our attention to the lottery problem and the commonly held intuition that, despite the odds being squarely in our favor, we cannot know that we are going to lose the lottery. Typically, explanations as to why this is so have rested on notions of probability, that is that the chance of winning prevents any possible knowledge of losing. The immediate problem with this solution is that probability presumably plays a role in many mundane events of which we would still intuitively wish to ascribe knowledge, so we somehow need to construe the probability involved in lotteries as particularly important, or salient, if we wish to maintain this so-called "lottery intuition." John Greco presents a simple solution: "Why does the element of chance become salient in the lottery case? I would suggest that the very idea of a lottery has the idea of chance built right into it" (Greco, 124). This, however, is somewhat uninformative: why should a lottery have chance "built right into it" any more so than car accidents, any skeptical possibility, or really any possibility at all? I seek a more informative answer.

Indeed with discussion of this problem, we are brought full circle to one of the motivating causes for this investigation in its entirety: the conflict between the epistemic

possibility constraint and the objective chance principle that together seem to imply that most, if not all, of our beliefs are true due to some amount of luck. I have attempted to show that the possibilities of relevant alternatives are able to be made less and less salient given certain amounts and types of evidence, but applying this theory to lottery cases seems to give the counter-intuitive result that the fewer lotteries that are won by an agent the more likely the agent is able to know that she will lose the next one. After all, if, as I have suggested, driving safely provides evidence that is able to diminish the salience of the possibility of being in a car accident, why shouldn't a series of lottery losses diminish the salience of the possibility of winning<sup>44</sup>? At the end of his investigation, Hawthorne struggles with a number of potential solutions to a similar problem involving the unfortunate consequence of accepting both the epistemic possibility constraint and the objective chance principle, namely that almost all "everyday" kinds of knowledge are forced to be deemed non-knowledge-worthy (I will present the solutions in reverse order):

[One] option is one that cleaves to the Epistemic Possibility Constraint: when one knows  $p$ , there is a zero epistemic chance that not- $p$ . In that case, glosses of the practical environment idea in terms of epistemic chance will be circular and uninformative...In a deliberative environment where one ought to use  $p$  as a premise, one knows  $p$ . When one knows  $p$ , the epistemic chance for not- $p$  is zero.

---

<sup>44</sup> If ones intuitions are that one *can* know that one will lose a lottery, perhaps the theory presented up until this point will satiate those intuitions. Anticipating terms that I will introduce below, my initial formulation of a "solution" to the lottery problem involved a direct application of the fading principle, but appealed to the notion that we are presented with a constant stream of evidence (or that such evidence is in existence) that subverts ones ability to know that they will lose the lottery, i.e. that the possibility of winning a lottery was a transient alternative, but was nevertheless constantly being reinforced, and thus was permanently knowledge-harmful (although perhaps in a somewhat unorthodox way). I found this solution undesirable as it seems to condone the epistemic acts of the gambler in the gambler's fallacy, a topic I discuss in more detail below.

When the epistemic chance of not- $p$  is nonzero, one shouldn't use  $p$  as a premise in practical reasoning. (Hawthorne, 178)

I agree with Hawthorne that not only is this solution circular and uninformative, it seems to enforce a strict, statistical certainty to any knowledge-worthy claims, whether it be that complete statistical certainty is *required* for knowledge, or that knowledge itself implies that one is completely certain. As the latter implication would no doubt be denied by the majority of agents upon being presented with any skeptical possibility, perhaps, then, the difficulties lie in the epistemic possibility constraint itself. This leads Hawthorne to his second potential solution:

One approach gives up the Epistemic Possibility Constraint, allowing that knowledge that  $p$  is compatible with a small epistemic chance that not- $p$ . There is a natural way to spell out the practical environment constraint on such a picture: if there is a small epistemic chance that not- $p$ , one knows  $p$  only if one is in a practical environment where the difference between a small epistemic chance that not- $p$  and zero epistemic chance that not- $p$  is irrelevant to the matters at hand.

(Hawthorne, 178)

This potential solution, I believe, is more desirable. However, as it is formulated now, it seems to open the door *more fully* for the lottery problem, or at least lottery-type problems which are somehow "irrelevant" to a given agent, and we are again lead to questions concerning just how high a chance that not- $p$  has to be in order to be knowledge-harmful, a question that does not seem to have any desirable answer.

Additionally, as relevance seems to be based upon the specific concerns of the agent, we again are presented with the problems dealt with in the previous chapter; to reiterate, that

the amount of knowledge an agent has is reliant upon that agents particular interests, which may or may not be of an epistemically responsible nature. Thus in order to effectively handle lottery-type situations, this solution requires some additional tweaking: not to *cleave* to the epistemic possibility constraint, nor to reject it outright, but rather to demarcate the relevant subclasses of “chance” that are at work in so-called “mundane” cases of knowledge-worthy beliefs, as opposed to lottery-type scenarios, and apply the epistemic possibility constraint and fading principle accordingly.

### **Transient and Permanent Alternatives**

In the cases investigated thus far, relevant alternatives have been able to acquire and lose salience given evidence that either supported their relevance or went against it. These kinds of alternatives I will refer to as “transient.” The distinguishing features of a transient alternative are first, unsurprisingly, that it is able to play a significant or insignificant role in the knowledge-worthiness of an agent’s belief, possesses an activation level as discussed earlier, and is not permanently knowledge-defeating *simply by virtue of an initial instantiation*, that is to say it is reliant upon evidence for its continued viability as a relevant or irrelevant alternative<sup>45</sup>. An example of a transient alternative is that of the barn façades in the example of the same name: the knowledge-affecting power of barn façades is transient given the situation and a certain amount of evidential support/support for its defeaters. A transient alternative inherits its transience from its reliance upon evidential data: again considering the alternative of a car accident, the frequency of car accidents is itself variable, and could conceivably become so high as to be knowledge-affecting in a very significant way, or, perhaps after technological

---

<sup>45</sup> It is plausible to imagine a transient alternative that has a permanent effect, i.e. one whose salience is constantly being reinforced more than it is being denied; see note 34 for a theory that posits lottery propositions as containing this kind of an alternative.

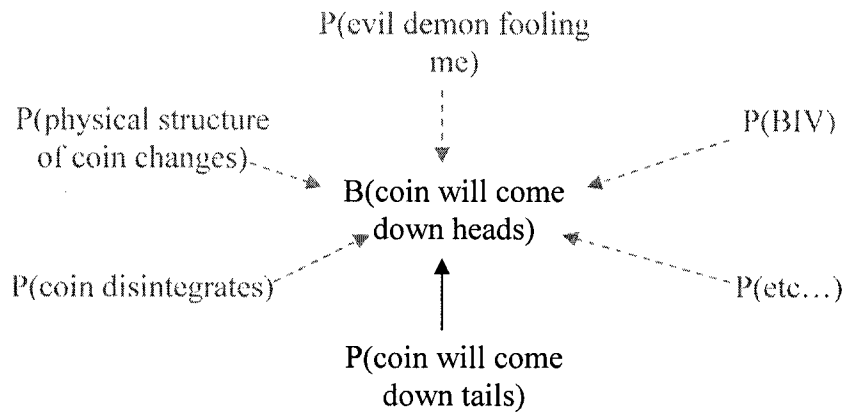
advances guaranteeing a heretofore unprecedented level of automobile safety, cease to occur at all, and likewise cease to be knowledge-affecting.

Contrast these alternatives with those that I will call “permanent” alternatives. This variety of alternative is not based upon frequency of occurrence of an event as transient alternatives are, nor do they rely upon empirical evidence for their existence. Permanent alternatives come in two additional flavors: those that are “knowledge-harmful”, and those that are “angst-inducing.” The class of angst-inducing permanent alternatives encompasses those familiar skeptical alternatives, such as brain-in-a-vat scenarios, which themselves have no evidential support for or against<sup>46</sup>, but can nevertheless cause some amount of epistemic angst<sup>47</sup>. My focus will be on the strictly knowledge-harmful variety of permanent alternatives, such as a coin coming down heads as an alternative to its coming down tails, or the possibility of winning of a lottery as an alternative to losing it. To illustrate, in a coin-tossing scenario an agent might possess the following belief structure:

---

<sup>46</sup> As, for example, evidence that purports to go *against* BIV possibilities inherently beg the question, and any evidence that purports to *support* BIV possibilities will be inherently false (i.e. if one is actually a BIV, evidence such as “seeing” a room full of brains in vats will simply be yet another image fed directly to ones brain, and thus not representative of reality).

<sup>47</sup> Recall the section “Angst and Value” in Chapter 3 for a cursory discussion of the matter, and see Pritchard, “Scepticism, Epistemic Luck, and Epistemic *Angst*” for a more thorough investigation.



**Figure 14: Potential defeaters (gray) for the belief in p: the coin will come down heads on the next toss, as modeled on previous examples presented thus far.**

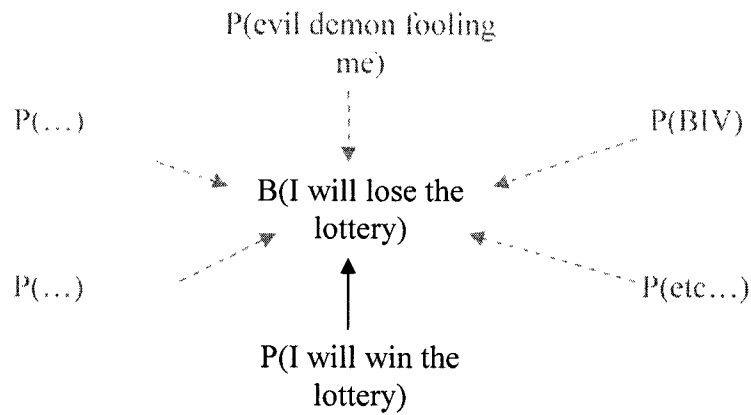
Of the above defeaters, it seems that all can be ruled out in an epistemically virtuous manner, save  $P(\text{coin will come down tails})$ . An immediate reaction might be that this is so because of the *probabilities* involved in each potentially relevant alternative: whereas the probability of  $P(\text{coin disintegrates})$  is presumably very low, the probability of  $P(\text{coin will come down tails})$  is 50%, which is very high, at least in epistemic matters<sup>48</sup>. Thus it might be assumed that it is the high probability that makes  $P(\text{coin will come down tails})$  a relevant alternative, and thus knowledge-preventing for  $B(\text{coin will come down heads})$ .

Of course, if the *sole* difference between  $P(\text{coin will come down heads})$  and  $P(\text{coin will disintegrate})$  was the degree of probability, then, in accordance with fading,  $P(\text{coin will come down heads})$  should be able to be affected by empirical evidence. After all, continuing with the theory I have been presenting, we are liable to rule out  $P(\text{coin will disintegrate})$  as a relevant alternative since it most likely has no evidence in its support<sup>49</sup>.

Consider the same scenario, although this time with a lottery proposition:

<sup>48</sup> I will assume throughout, unless otherwise specified, that the coin is fair.

<sup>49</sup> And, most likely, a great deal of implicit evidence suggesting that it is not a viable alternative.



**Figure 15: Potential defeaters (gray) for the belief in p: I will lose the lottery, as modeled on previous examples presented thus far.**

Again we are presented with a number of potentially relevant alternatives (which I have, for the large part, not listed here), yet it is not obvious that  $P(\text{I will win the lottery})$  has any greater *probability* associated with it than any of the other alternatives. This is where the intuitions surrounding lottery cases diverge: on one hand, if the probability of losing the lottery is no greater than any of the more radical alternatives, and we are willing to ignore the latter, perhaps we should be willing to ignore the former. On the other hand, the nature of the lottery seems to be such that it is inherently not possible to definitively know that one will lose it; as Greco said, it seems that it has the concept of chance “built right in.”

The potential solution I wish to present, then, is that  $P(\text{losing a lottery})$  in a lottery case, just like  $P(\text{coin will come down tails})$  in a coin-tossing case, is a permanent alternative. Indeed both alternatives meet the same criteria: although the probability of their respective occurrences are different (most likely significantly so), that neither probability is based upon empirical evidence, or affected by previous occurrences or non-occurrences of different possible outcomes, means that it never loses its relevance as an

alternative in the belief-forming process. It is not susceptible to the fading phenomenon I have described, since, by nature, it is neither strengthened nor weakened by the amount of evidence, implicit or explicit, that is potentially shown to be in its favor or against it. Indeed it sounds strange that one could acquire evidence that is in “support” of their belief that they are going to lose a lottery; surely we would think such support (no doubt many previous instances of their losing the lottery) would not really be support at all, or would be inappropriate. The *totality* of evidence available in this situation is limited to the bare facts of ones possessing a ticket, and the odds of winning, which are not conclusive enough to result in a knowledge-worthy belief; that this is all the evidence that is available means that a knowledge-worthy belief involving a lottery proposition will never be possible. The upshot, then, is that in coin-tossing scenarios, it is not the high probability of the relevant alternative in itself that is knowledge-preventing, but rather it is the knowledge-harmful variety of permanent alternatives that is responsible. This further implies that one is unable to know the outcome of a lottery *because* it shares a common feature with coin-tossings, a situation where it is doubtful that there is any opposition to the intuition that one cannot have knowledge in said kinds of cases. Thus lottery intuitions do not *guide* the investigation, a problematic process since intuitions on the matter are not in universal agreement; rather, it is a consequence of the theory (and thus could provide a point of contention for those whose intuitions guide them to different conclusions).

The reader may be unconvinced. My reasoning so far has been as follows: one cannot know that one will lose a lottery because the relevant alternative of winning is a permanent one. Permanent alternatives are always knowledge-harmful because they



cannot be defeated through evidence, that is to say that if one forms a true belief as to the outcome of a lottery, it will be lucky, and this luck is not susceptible to fading. The alternative of winning a lottery is rightfully called permanent in virtue of the fact that what we might consider evidence that one will lose (namely, repeated losses) in no way affects the probabilities involved. Here there is a place for possible discomfort: why should lotteries, coin-tosses, etc., be kinds of phenomena which are not affected by evidence? What prevents the theory from degrading into the unfortunate consequences concerning the reasoning fallacy outlined above? In other words, I have set out to give a better explanation than the one put forth by Greco, namely that lotteries seem to have the concept of chance “built right in”, but have I merely given a different label to the same conclusion?

I believe the answer is in the negative. Consider further how the probabilities differ between that of being in a car accident and losing/winning the lottery: in the former, we typically use a relative frequency model to calculate odds, that is to say that the probability of being in a car accident is calculated on the basis of empirical evidence of the estimated number of such accidents given a certain amount of time, or some estimation of the number of safe trips taken. For lotteries, however, we use a theoretical model: in theory, there will be an expected number of winners/losers given the number of lottery players, but the probability involved is the same regardless of whether anyone wins, anyone plays, or even if the lottery exists at all. The probability involved in lotteries is strictly mathematical calculations, whereas the probability calculations such as those involved in car accidents are reliant upon empirical data. Thus any probability calculated under the theoretical model involves a permanent relevant alternative, since

these types of probabilities are unaffected by the accumulation of evidence and are, in turn, unsusceptible to fading.<sup>50</sup>

## Final Thoughts: Revisiting the Incompatibility Thesis

At the beginning of the investigation I briefly mentioned what has been called the incompatibility thesis, which states that luck, in one form or another, is incompatible with knowledge. As Pritchard was right to point out, it is only a certain kind of luck, one that has to do directly with the truth of the beliefs involved that is knowledge-harmful, and I have attempted to show that even this kind of luck is either not as harmful as has been previously assumed (through the fading effect), or is still able to be knowledge-supportive in perhaps some unexpected ways (as in Judy and Trudy cases). All of these kinds of luck involved transient alternatives; with the introduction of permanent alternatives, however, there does seem to be a variety of luck that is always incompatible with knowledge, namely that where one's belief is true despite the presence of a permanent knowledge-harmful alternative.

The investigation is, of course, not without its gaps. Questions which I have not been able to tackle include: Bayesian treatments of probability and its role in the conflict between the epistemic possibility constraint and the objective chance principle, subjective theories of objective chance as a potential solution to Bayesian problems, the possible correlation between statistical probability and activation levels, and a more thorough treatment of the nature of “randomness” that is involved in situations like coin tosses and

---

<sup>50</sup> It is due to these factors that I largely reject Bayesian epistemology, as it seems that Bayesians are prone to equivocating between relative frequency probability and theoretical probability when assigning said probabilities to beliefs. Such a discussion is, unfortunately, outside the scope of the current investigation.

lotteries. These questions have largely been left untouched due to the constraints of a relatively short project, but I believe comprise a list of the most obvious “next steps”<sup>51</sup>.

---

<sup>51</sup> Issues of epistemic closure and the compatibility of stakes theory and the theory presented in this investigation, which I have alluded to more than once in the course of this investigation, is also an area of study that deserves further investigation.

### Bibliography

DeRose, Keith. "Solving the Skeptical Problem." *Philosophical Review*. 104(1): 1-52.

---. "Contextualism and Knowledge Attributions." *Philosophy and Phenomenological Research*. 52(4): 913-929.

Greco, John. "Knowledge as Credit for True Belief," in *Intellectual Virtue: Perspectives from Ethics and Epistemology*, Michael DePaul and Linda Zagzebski, eds. Oxford: Oxford University Press, 2004.

Hawthorne, John. *Knowledge and Lotteries*. Oxford: Clarendon Press, 2004.

Kornblith, Hilary. *Knowledge and its Place in Nature*. Oxford: Clarendon Press, 2002.

Latus, Andrew. "Moral and Epistemic Luck." *Journal of Philosophical Research*. 25: 149-172.

Lewis, David. "Elusive Knowledge." *Australasian Journal of Philosophy*. 74: 549-567.

Pritchard, Duncan. *Epistemic Luck*. Oxford: Clarendon Press, 2005a.

---. "Scepticism, Epistemic Luck, and Epistemic *Angst*." *Australasian Journal of Philosophy*. 83.2(2005b): 185-205.

Rescher, Nicholas. *Luck*. New York: Farrar Straus Giroux, 1995.

Riggs, Wayne D. "Reliability and the Value of Knowledge." *Philosophy and Phenomenological Research*. 64(2002): 1-24.

---. "What are the 'Chances' of Being Justified?" *Monist*. 81.3(1998): 452-473.

---. "Why Epistemologists are so Down on Their Luck." *Synthese*. Forthcoming.

Sosa, Ernest. "Skepticism and Contextualism." *Philosophical Issues*. 10(2000): 1-18.

Stanley, Jason. *Knowledge and Practical Interests*. Oxford: Clarendon Press, 2005.

Steup, Matthias and Ernest Sosa. *Contemporary Debates in Epistemology*. Oxford: Blackwell Publishing, 2005.

Williamson, Timothy. *Knowledge and its Limits*. Oxford: Oxford University Press, 2000.