Chapter IV

Crude Particular Evidentialism

In the previous chapter we set out the Assurance View, an alternative to our favored Particular Evidentialism. Particular Evidentialism (see section II.1) is the thesis that any particular piece of testimony can only provide a justification for believing what is told by providing evidence for what is told. The Assurance View holds, on the contrary, that a teller can give the hearer a non-evidential justification for believing what she is told by assuming responsibility for the truth of her testimony. The teller assures the hearer of the truth of her testimony, and a hearer who accepts this assurance forms a belief without relying on evidence for what is told. The remaining four chapters of the dissertation defend Particular Evidentialism against the Assurance View and the arguments that motivate it.

Our argument for Particular Evidentialism will have an offensive and a defensive component. The offensive component casts doubt on the Assurance View by arguing that it has trouble explaining the cases in which testimony fails to justify belief. By an examination of cases, we will show that testimony gives justification for belief only when it gives evidence for what is told. The teller's assurance that her testimony is true should only be accepted when this assurance does provide evidence for what is told. If the Assurance View were true, we would need an explanation of why the teller's assurance would provide a non-evidential justification for belief in exactly these cases. What is the epistemological difference between the assurances that provide both evidence and justification and the assurances that provide neither, if it is not that the former assurances

provide evidence and the latter do not? The difficulty of answering this question

provides some reason to think that the justification provided by testimony is evidential.

The offensive argument against the Assurance View will not be adequate, however,

unless the arguments that motivate the Assurance View are answered. In the previous

chapter we set out objections to Particular Evidentialism that created puzzles for any

account on which testimony is taken to provide evidence. Our discussion of intention-

dependent evidence (section III.2) raised the problem of how freely chosen testimony can

provide evidence for what is told. The Bad Faith Objection (section III.3) raised the

problem of how a teller can present her testimony as evidence while seeing herself as

responsible for its truth; the Disharmony Objection (section III.3) raised the problem of

how the hearer can take the testimony as evidence while seeing the teller as responsible

for its truth. If we cannot give an account of how freely chosen testimony can provide

evidence for what is told, and how the hearer can accept this evidence while seeing the

teller as responsible, then we must accept that testimony does provide a non-evidential

justification for belief. Our offensive argument would not prove the superiority of

Particular Evidentialism over the Assurance View if Particular Evidentialism could be

disproved on independent grounds. Accordingly we will also need to defend Particular

Evidentialism against the Bad Faith and Disharmony Objections.

The defensive component of our argument will be set forth in Chapters V through

VII. This chapter provides a preliminary sketch of the offensive component, arguing that

testimony does not provide a justification for belief unless it also provides evidence for

what is told; the sketch is preliminary because it depends on a crude conception of

evidence, which we will refine in Chapter V. The account of testimony yielded by this

conception of evidence will not allow us to defend Particular Evidentialism against the Bad Faith and Disharmony Objections; that will have to wait until after Chapter V's more sophisticated account of evidence. Indeed, the crude conception of evidence set forth in this chapter will make these objections particularly appealing.

The crude conception of evidence does, however, allow us to sketch one way in which the teller is responsible for her testimony, a way that follows from epistemology alone. On this conception, a teller who tells a falsehood will weaken all of her future testimony as evidence. Future hearers will be less justified, or in extreme cases not at all justified, in believing what she tells them. On our Gricean analysis of testimony (section I.2), testimony is an attempt to induce belief. When a hearer should not believe the teller because of past falsehoods, the teller's attempt to induce belief should be frustrated. We can think of a teller as earning credit for telling truths and discredit for telling falsehoods. Too much accumulated discredit earns disbelief, which is a penalty because it is the frustration of something that the teller is attempting to do. Having your testimony fail to induce belief in this way is what we will call the *reliability sanction*; we will call the normative structure governing the reliability sanction's exercise the *credit/discredit normative structure*.

The credit/discredit normative structure is of intrinsic interest, as is the fact that it can be derived from the epistemology of testimony without adverting to considerations of morality or personal obligation. In addition to this intrinsic interest, the structure will eventually help refute the Disharmony Objection. In Chapter VII, we will show, using the credit/discredit structure, that a hearer who takes the teller's testimony as evidence not only *can* see the teller as responsible for the truth of her testimony, but *must* treat her

as responsible for its truth. This chapter's treatment of the credit/discredit structure will not be enough to answer the Disharmony Objection, but it will lay the foundation for Chapter VII's answer.

Section 1 contains preliminary remarks concerning our entire discussion of the epistemology of particular testimony. We will be engaged in case-by-case discussions of testimony in various situations; in order to limit the discussion to a manageable number of cases, we will make some simplifying assumptions. Section 1 also lays down some general epistemological guidelines, for instance concerning our conception of justification. Section 2 then lays out the particular conception of evidence used in this chapter: the *crude enumeration conception*. It discusses how evidence can depend solely on statistical generalizations concerning cases that have already been observed, and it gives a rule for selecting the statistical generalization that serves as a basis for induction.

Section 3 then gives the offensive argument for Particular Evidentialism against the Assurance View. Given the crude enumeration conception of evidence, we can break particular pieces of testimony down into the cases in which the testimony gives evidence and the cases in which it does not. When testimony does not give evidence, it does not justify belief. In these cases the hearer specifically lacks justification for accepting the teller's assurance that the testimony is true, even though that assurance is offered. The Assurance View, that this assurance can give a non-evidential justification for belief, is thus under the burden of explaining why that non-evidential justification for belief exists only when the testimony gives evidence. Particular Evidentialism, by contrast, explains this perfectly, since it holds that there is no non-evidential justification for believing any particular testimony.

Section 4 sets out the credit/discredit normative structure on testimony. Given that testimony is a voluntary action, for which the teller can properly be held responsible, it follows from the epistemological analysis of section 3 that the teller *is* responsible for the truth of her testimony. We show this by showing that a teller who fails this responsibility can appropriately be punished with the reliability sanction, in which the hearer does not believe what she is told because of the teller's past falsehoods. Effectively, the teller stakes her future credibility on the truth of her testimony.

This discussion of the credit/discredit structure will not suffice to answer the Disharmony Objection (section III.3). To do so, we would need to show that the hearer can simultaneously see the teller as presenting evidence for what she tells and as responsible for the truth of her testimony. This will require the more sophisticated conception of evidence discussed in Chapter V. Indeed, it is part of the reason for putting forth a notion of evidence that is more sophisticated than the crude enumeration conception—besides the fact that the crude enumeration conception is too crude to be true. In this chapter, we will not yet be able to reconcile the thesis that testimonial justification is evidential with the idea that that justification depends on the teller's assumption of responsibility for her testimony. But the credit/discredit structure sketched in this chapter will lay the foundations for that reconciliation.

1. Preliminaries

Before we embark on a case-by-case analysis of the epistemology of testimony, we must lay down some ground rules for our investigation. These preliminary remarks

pertain to the epistemological analyses contained in Chapters V and VI as well as in this chapter.

First, our aim is to show that testimony provides justification for believing what is told only in the cases in which it also provides evidence for what is told. What is important here is the evidence and justification that the testimony itself provides the hearer, not any evidence or justification that the hearer may have independently of the testimony. In many cases a hearer will have independent evidence for or against what is told. If the sheriff already has strong evidence that the James Gang is planning to rob the bank, and the generally unreliable informant Bob Ford then tells him that the gang is indeed planning the robbery, the sheriff has justification and evidence for believing what Bob Ford says. The sheriff does not believe Bob Ford, though, because the sheriff's belief is based on evidence that is completely independent of Mr. Ford's testimony. The testimony makes no difference in this case. Conversely, suppose that the sheriff begins with strong evidence of the planned bank robbery, and Honest Mr. Lincoln tells him that the gang is planning to rob the stagecoach instead. The sheriff may still be justified on balance in believing that the gang will rob the bank rather than the stagecoach, though less justified than before Mr. Lincoln's testimony. Mr. Lincoln's testimony provides the sheriff with some justification for believing what he says, but this justification is outweighed by the evidence the sheriff already had.[1] In both of these questions, the sheriff's antecedent evidence has little to do with the epistemology of testimony. Any

---

[1] If Mr. Lincoln knows of the sheriff's antecedent evidence, he may not expect the sheriff to believe his testimony; he may expect that, at best, the testimony will weaken the sheriff's belief that the James gang will rob the bank. In that case, he will not intend that his testimony induce a belief in th e hearer, but only that it weaken the hearer's contrary belief. See the discussion of this sort of testimony following the definition of testimony in section I.1.

justification that a piece of testimony gives must be weighed together with everything else the believer knows; in this testimony is like any other sort of justification.

In order to isolate the evidence and justification that testimony itself gives from any other evidence or justification, we will concentrate on cases in which the hearer has no independent evidence for or against what she is told. In these cases, any justification for belief or evidence that the hearer has will have come from the testimony. The hearer will be justified *tout court* in believing what she is told when and only when the testimony itself has given her justification, and the hearer will have evidence for what she is told when and only when the testimony has given her evidence. The idea is that testimony that gives evidence and justification in these cases will give *pro tanto* evidence and justification in most cases in which the hearer has other evidence. The concept of evidence we use in this chapter is not sophisticated enough to handle cases in which the hearer has additional evidence besides the testimony, and the concept of evidence used in later chapters will allow for quite complex interactions between the testimony and other evidence that is available. As a rule of thumb, testimony that would give the teller justification for belief in the absence of other evidence will increase the teller's justification for believing what is said (or decrease her justification for believing the opposite of what is said) in the presence of other evidence. The exceptions will be unusual cases.

Second, the conceptions of evidence and justification that we will use depend only on the information that the teller's experiences make available. This is directly opposed to a reliabilist conception, on which evidence or justification for a belief could be provided by a procedure that reliably output true beliefs, even though the believer had not had the

experiences that would tell her that the procedure was reliable. We would not, for instance, claim that Aristotle could have gained evidence of a lack of oxygen in an enclosed space from the fact that a flame went out, since Aristotle had no way of knowing that flames go out when they lack oxygen. More relevantly, we will not say that whether testimony provides evidence or justification depends on how truthful and knowledgeable the teller actually is, but on whether the information available to the hearer indicates that the teller is truthful and knowledgeable. The conception of justification as dependent entirely on the information available to the teller is defended in the Appendix (which see). The conception of evidence must be entirely parallel to our conception of justification; otherwise the claim that there were non-evidential justifications would be uninteresting.

Third, the question we will focus on is whether testimony justifies belief and whether that justification is evidential or non-evidential. This need not be the same question as whether justification of a belief is in some sense *a priori* or *a posteriori*. In particular, we are not addressing Burge's contention that the status of knowledge through testimony depends on the teller's original justification for believing what she was told, so that, for instance, "[p]eople who depend on interlocution for knowledge of mathematical theorems but do not know the proofs can have apriori knowledge in [a] sense" (Burge 1993, p. 487). Obviously this knowledge is not independent of experience, because it depends on the experience of having been told that the theorem is true. If *a priori* knowledge can depend on experience, it seems possible that it can also be justified evidentially. Note in particular that, according to our arguments in favor of Burge's Acceptance Principle (section II.3), no positive evidence is necessary to justify the acceptance of testimony.

The balance of evidence may be in favor of accepting testimony even though there is no particular experience that supplies evidence in favor of the testimony. Knowledge that was gained through such a default presumption of reliability might be considered to be *a priori* in some sense of *a priori*. Since we are not concerned with *a priori* and *a posteriori* knowledge, we will not pursue the issue further.[2]

Fourth, we will speak throughout of testimony as providing evidence rather than being or serving as evidence. On the theory of justification developed in the Appendix, the believer's experiences determine what she is justified in believing; for the sake of parallelism, we should think of her experiences as also determining what evidence she has. In none of this is it necessary to say exactly what constitutes evidence. We can always think of experiences, or facts of which the believer has become aware, or testimony that the believer has heard, as providing rather than constituting evidence. This will avoid needless ontological complications.

2. The Crude Enumeration Conception of Evidence

With those preliminaries out of the way, we are now ready to set forth the simple conception of evidence that we will use as the basis for this chapter's discussion. This is the *crude enumeration conception* of evidence, which adverts exclusively to statistical generalizations. Ross remarks (as previously quoted in Section III.3):

---

[2] In order to preserve the *a priori* status of mathematical knowledge, we might invoke the "hybrid theory" proposed by Faulkner (2000): The hearer needs a justification for accepting testimony, but once she has this justification, her warrant for the belief is whatever warrant the original speaker had for accepting what was told. In terms of this theory, our concern in the text is whether the hearer's justification for believing what she is told is evidential, not with the warrant that she gets by accepting the testimony. Faulkner's hybrid theory would allow *a prioricity* to be preserved in testimonial chains, in that the hearer has an *a priori* warrant for whatever she is told if the original speaker does, though the teller's justification for accepting the testimony is not *a priori*.

> If the fact of a speaker having chosen to use certain words is to be seen as evidence of anything, it must be taken in conjunction with certain empirically established generalisations concerning the circumstances in which such a choice by such an individual is likely or unlikely—or better, a well grounded psychological theory giving us an insight into the factors which constrain or influence a speaker's choice of words (Ross 1986, p. 72).

The crude enumeration conception of evidence treats testimony as evidence in light of empirically established generalizations concerning how often teller's testimony has been true. In Chapter V, we will move to Ross's preferred conception of evidence, on which testimony's role as evidence depends on a psychological theory of why people say what they say. Before we develop that theory, the crude enumeration theory will allow us to present a sketch of why testimony sometimes fails to provide justification and how the teller must be responsible for her testimony.

The crude enumeration conception of evidence relies on inductive generalizations. To provide evidence for something, a particular object or event must be subsumed under some general type, and we must have established a correlation between that general type and the thing for which we are to have evidence.[3] Say that we are interested in whether an object or event in some reference class F also belongs to some attribute class G. (In the case of testimony, F might be the class of all pieces of testimony, and G might be the class of all true pieces of testimony, but we will sometimes need to consider other reference classes.) A crude enumerative induction is a combination of two standard types of inductive inference: *proportional induction* and *statistical syllogism*.[4] Proportional

---

[3] For simplicity of exposition, in this section I will only consider objects and events that can serve as evidence, not facts; the conception is easily enough adapted to accommodate facts. In any case, our ultimate target is testimony, which is most naturally treated as a type of event.

[4] We have already mentioned these types of inference (see section II.1). For a brief discussion, see the *Encyclopedia of Philosophy* (Black 1967).

induction is the inference from the fact that x percent of the observed[5] sample of F's are

G's, to the generalization that x percent of all F's are G's. Statistical syllogism is the

inference from the premise that x percent ($x > 50$) of F's are G's to the instance that a

particular F is a G, where the inference is stronger the higher x is. (If all F's are G's,

statistical syllogism degenerates into universal instantiation.) Crudely enumerative

induction combines these two to take us from a premise concerning a sample of F's to a

conclusion concerning a particular F outside the sample. From "x percent of the

observed sample of F's are G's" proportional induction takes us to "x percent of all F's

are G's," and from this (if $x > 50$) statistical syllogism takes us with more or less

confidence to "This F is a G."[6]

According to the crude enumeration conception of evidence, then, we gather evidence

as follows: First, we observe a sample of the reference class, and determine how many of

them belong to the attribute class. (An instance only counts as part of the sample if it is

---

[5] Note that we are concerned only with the evidence that comes from what is observed. Statements like "The symptoms of Black Plague were evidence of a bacterial infection, though no one had done observations that could have shown this" do not concern us. We are not concerned with an externalist, observation-independent notion of what constitutes real evidence that no one can know is evidence. This keeps our notion of evidence in line with our notion of justification, according to which justification depends only on the available information (see section 1 and the Appendix).

[6] "Is" is tenseless throughout. In the case of testimony, which usually vanishes quickly, the F's will all be past, and our observations will concern whether past testimony *was* true when it was spoken. For instance, if the speaker said on January 1 "It will snow within the next week," and we are evaluating the utterance on February 1, we are concerned with whether it was true on January 1 that it would snow within the next week, which is to say whether it snowed between January 1 and January 8.

Belnap (2001) discusses how to evaluate such sentences using double time references, to the time of speech act and the time of evaluation. John MacFarlane (2003) gives an elegant solution to the problem of evaluating future-tensed sentences whose truth value has not been settled at the time of utterance. The context of assessment of such a sentence need not be the same as the context of its original utterance. Suppose that on February 1 we wish to determine whether the January 1 utterance "It will snow within the next week" has been settled true yet. We must evaluate the indexical "next week" with respect to the context of utterance (January 1) rather than the context of assessment (February 1), but the possibilities relevant to the utterance's settled truth are those that are still open in the context of assessment (February 1) rather than all those open in the context of utterance (January 1). MacFarlane's treatment thus allows us to evaluate utterances whose truth becomes determinate between the time of utterance and the time of assessment. The details of his solution do not matter for our purposes; we will not deal with any hard cases.

either determined to belong to the attribute class or determined not to belong to it.) When we are confronted with a new instance of the reference class, the strength of our evidence that it belongs to the attribute class corresponds exactly to the proportion of the sample that belonged to the attribute class. Suppose that the reference class is all coin flips and the attribute class is coin flips that come up heads. If exactly half of the coin flips that we observe have come up heads, then we have no evidence that the next coin flip will come up heads and no evidence that it will not.

The crude enumeration conception of evidence is not only crude, but also faces several problems of application. In the rest of this section, we will address some of these problems and bracket others, in order to yield sharp criteria for our preliminary sketch of when testimony provides evidence.

The first problem is the problem of selecting a reference class. Any given particular belongs to many different general categories, and each of these general categories may have a different proportion of observed instances that belong to the attribute class. Each of these different general categories could be used as the reference class for a crude enumeration induction, but the resulting inductions would yield conflicting evidence about the same particular. Suppose that exactly half of the coin flips that Alice has observed have come up heads, but 99% of the flips of one particular coin that Alice has observed have come up heads. She now wonders whether this particular coin will come up heads the next time she flips it. The particular in question is the next flip of this coin, which can be subsumed under the reference class "All coin flips" or the reference class "All flips of this coin." A crude enumeration induction using "All coin flips" as a reference class would yield that Alice has no evidence that the coin will come up heads,

while a crude enumeration induction using "All flips of this coin" would yield that Alice has strong evidence that the coin will come up heads. These two inferences are actively in conflict. (Note that it is not that one inference fails to give evidence which the other supplies; they cannot both be good inferences.) Which should Alice believe?

   The problem of selecting a reference class is analogous to the problem, as described by Wesley Salmon (Salmon 1971), of applying the frequency theory of probability to single cases. The frequency theory defines probability as "the limit of the relative frequency of an attribute in an infinite sequence of events" (Salmon 1971, pp. 38-9), but in application we need to determine the probability of a single event. Assigning an event to a reference class will yield a (potentially) infinite sequence of events within which we can determine the relative frequency of the attribute, yet the event will belong to many different potential reference classes and many different infinite sequences. On the frequency theory, then, the probability of a particular event may depend on which of the many possible reference classes we assign it to. Consider the coin-flip example, supposing that Alice's observations could be extended to infinite sequences. Assigning the next coin flip to the sequence of all coin flips, there is a probability of 0.5 that it will come up heads; assigning it to the sequence of all flips of *this* coin, there is a probability of 0.99 that the same coin flip will come up heads.

   In this example, the correct answer is obvious: Alice has strong evidence that this coin will come up heads the next time it is flipped.[7] The statistics she should rely on are the more specific ones about this coin as opposed to the less specific ones about all coins.

---

[7] In fact, she has strong evidence that the coin is unevenly weighted, but this evidence goes beyond what the crude enumeration conception can capture. On the crude enumeration conception, such evidence could only be obtained by examining many coins that come up heads 99 percent of the time, and discovering that most of them were unevenly weighted.

The induction that uses all coin flips as a reference class ignores relevant information

about this particular coin. In general, the induction should use the reference class that

embodies all the relevant factors and none of the irrelevant factors concerning the thing

we are testing. This means using the equivalent of Carnap's requirement of total

evidence for applying inductive logic: "The probability value to be chosen for practical

applications is that given by the degree of confirmation statement involving the total

available (relevant) evidence" (Salmon 1971, p. 37; see Carnap 1950, p. 211).

Incorporating all available relevant evidence amounts to making the reference class, in

Salmon's terms, *epistemically homogeneous*.[8] The definition of epistemic homogeneity

is roughly as follows:

> A reference class R is epistemically homogeneous with respect to an
> instance f and an attribute class A iff:
> f is known to belong to R (this is necessary for R to be a reference class at
> all), and
> R has no uncontrived subclass S such that both
> (i) f is known to belong to S
> and (ii) the proportion of R's that are A is known to differ from the
> proportion of S's that are A. [9]

---

[8] Salmon himself is concerned with determining probability according to the frequency definition, for which he uses homogeneity *simpliciter* rather than epistemic homogeneity. Since probability is meant to be an objective notion, it must take into account even those reference classes that we do not know about; so homogeneity *simpliciter* essentially requires that there exist no subclass of the reference class in which the proportion of the attribute class differs from the proportion of the attribute class within the reference class, whether or not we know of that subclass (Salmon 1971, p. 43). Our notion of evidence only considers the information available to the subject (see note 5 above), so it need not take into account reference classes for which information has not been gathered; thus epistemic homogeneity is the appropriate criterion for selecting the reference class when evaluating evidence.

[9] See (Salmon 1971), p. 44. Salmon does not include the requirement that f be known to belong to S, but the requirement is in the spirit of Salmon's suggestion that we exclude subclasses such that we cannot know whether an instance belongs to the subclass without also knowing whether it belongs to the attribute class (Salmon 1971, p. 50; see the discussion of "the color at the opposite end of the spectrum from violet"). Note that this definition makes epistemic homogeneity relative to the instance as well as to the attribute class. Epistemic homogeneity will also be relative to the reasoner, because it is affected by the knowledge the reasoner has concerning the various classes and instances.

This definition does not exclude classes that contain subclasses S that fulfill condition (ii) and such that f is known *not* to belong to S, because the fact that f belongs to S need not be relevant evidence. Suppose that the proportion of A in S differs significantly from the proportion of A in R, but S is so small compared to R that the proportion of A in (R but not S) does not differ significantly from the proportion of A in R.

An induction that used an epistemically inhomogeneous reference class would fail to take into account relevant evidence. By the definition of epistemic homogeneity, we know that *f* belongs to *S*, and we know that the proportion of *A* in *S* is different from the proportion of *A* in *R*. Then *f*'s membership in *S* is relevant to the question of whether *f* is *A*, but the induction using *R* as a reference class does not take into account *f*'s membership in *S*.

The more relevant characteristics of *f* we include in our definition of the reference class, the narrower the reference class will be. For instance, if our reference class is all flips of a particular coin, we will have observed fewer instances than if the reference class were all flips of any coin. This introduces a problem mentioned by Reichenbach and Salmon: Narrowing the reference class can reduce the reliability of the statistics we have gathered about that class (Reichenbach 1949, p. 375; Salmon 1971, p. 41). Less reliable statistics weaken the first step of the crude enumerative induction, the proportional induction to a generalization about the reference class. If too few data have been collected we will not be able to perform a reliable proportional induction. For instance, observing that a particular coin has come up heads on six out of eight flips is not adequate evidence that 75% of all flips of this coin will come up heads, because the sample is too small.[10] Thus the narrowest possible reference class will not always yield a suitable crude enumerative induction.

Then *f*'s failure to belong to *S* does not give any relevant evidence to the question of whether *f* is an *A*. Say there are one trillion *R*'s, of which one billion are *A*'s, and one million *S*'s, all of which are *A*'s; then 0.1% of all *R*'s are *A*'s, 100% of all *S*'s are *A*'s, and 0.099% of all *R*'s that are not *S*'s are *A*'s; so we do not learn much more about whether *f* is a *A* when we learn that *f* is not an *S*.

Of course, if the proportion of *A* in (*R*'s but not *S*'s) does differ significantly from the proportion of *A* in *R*, then (*R* but not *S*) will be a subclass, such that *f*'s membership in it will be relevant evidence. *R* will fail to be epistemically homogeneous because (*R* but not *S*) will fulfill conditions (i) and (ii).

[10] In surveys, a larger sample size reduces the expected margin of error. If we were to take many random samples of equal size from a population, the proportion of the attribute class within the samples

The definition of epistemic homogeneity, however, does not force us to look at the narrowest reference class concerning which we have any data; rather it forces us to look at the narrowest reference class such that we know the proportion of the attribute class within the reference class. If we do not have enough data concerning a potential reference class S, then we do not really know the proportion of the attribute class A within S, owing to the lack of reliable statistics.[11] Then a wider class R may be epistemically homogeneous, even though the observed proportion of A in R may differ from the proportion of A in S observed in our inadequate sample. Given the available data, f's membership in S does not provide relevant evidence concerning whether f is A. When a coin has come up heads six of eight times, we do not yet have evidence that it will come up heads more or less often than any other coin, so an induction that uses all coin flips as a reference class will not ignore any relevant evidence. Incorporating all relevant evidence only requires, as Reichenbach proposes, that we take "the narrowest reference class for which reliable statistics can be compiled" (Reichenbach 1949, p. 374; quoted in Salmon 1971, p. 41). We need not consider narrower classes, for which reliable statistics cannot be compiled.[12]

So our rule for selecting reference classes is: Choose the most specific uncontrived reference class such that the class is known to contain the instance in question and there are enough data to perform a reliable proportional induction as to the frequency of the

---

would be normally distributed around the actual proportion of the attribute class within the entire population. The larger the samples, the smaller the standard deviation; 95% of all samples will be within two standard deviations of the actual proportion. If the standard deviation is too high because the sample size is too low, one survey will not provide meaningful information about the population.

[11] Of course, reliability of statistics is a matter of degree, so there will be cases in which it is not clear-cut which reference class we should use.

[12] Note that the issue is statistics that can be compiled from already known data, not statistics that could be gathered on further investigation.

attribute class within the reference class.[13] This reference class will incorporate all the

evidence about this particular instance that is known to be relevant. If we have ample

data that show that a particular coin comes up heads 99% of the time, we should take as

the reference class flips of this coin rather than all flips of all coins; this will be the most

specific class for which we have reliable data, and it will incorporate the relevant data

that we have concerning this coin.

Note that this rule does not always determine a unique reference class. Sometimes

there may be no unique narrowest class for which we have reliable statistics; we have

observed enough to derive reliable statistics for two separate classes, but we have not

observed enough of their intersection to derive reliable statistics for the intersection.

Suppose that we have observed that a particular coin comes up heads 99% of the time,

and we have also observed that 99% of the time that we have seen a particular person flip

any coin it has come up tails. Yet we have never seen this particular person flip this

particular coin. The "narrowest class for which we have reliable statistics" rule will not

determine whether we should take as the reference class flips of this coin or flips by this

person. This sort of situation is simply beyond the power of the crude enumeration

approach. In order to determine what evidence we do have, we would have to advert to

the most likely explanation of our data (as in Chapter V). For instance, we might

conclude both that this coin is unevenly weighted and that this person knows how to

control coins (and has always made them come up tails); we would then have to try to

---

[13] This solution conflicts with Salmon's recommendation that we take the *broadest* homogeneous reference class (Salmon 1971, p. 43). These two approaches, however, can never yield conflicting inferences. For a reference class to be homogeneous, it must have the same proportion of the attribute class as every non-contrived subclass we have reliable statistics for, and the narrowest homogeneous reference class will be one of these subclasses. The narrowest homogeneous reference class and the broadest homogeneous reference class will thus contain the same proportion of the attribute class, and inductions that take either class as the reference class will yield the same result concerning how much evidence we have that the instance belongs to the attribute class.

determine whether she can make this unevenly weighted coin come up tails. Crude enumeration alone will not indicate the answer.

The rule for selecting reference classes raises another problem for applying the crude enumeration conception: What is it for a reference class to be uncontrived? This is Goodman's new problem of induction (Goodman 1954). The problem is that we can select reference and attribute classes so as to make crude enumeration inductions yield counterintuitive and contradictory results. Suppose for instance that, as of Jan. 1, 2001, we have flipped a particular loaded coin 100 times, and it has come up heads every time. We want to know whether we have evidence that it will come up heads the next time we flip it. Clearly we do; a crude enumerative induction, using "flips of this coin" as the reference class and "heads" as the attribute class, yields strong evidence that this next flip will come up heads. Unfortunately, we can also perform a crude enumerative induction using "hails" as the attribute class, where a coin flip is *hails* if it takes place before 2001 and comes up heads, or if it takes place in 2001 or later and comes up tails. So far every flip of this coin has come up hails (heads before 2001), so this crude enumeration would strong evidence that the next flip will come up hails, which is to say tails (since it is now 2001). We could also choose the reference class so that it consisted of all coin flips that had ever come up tails, plus this one coming up; everything we have ever observed in this reference class has come up tails, so this crude enumeration also yields strong evidence that the next coin flip will come up tails. (Note that we do know that this upcoming coin flip belongs to this reference class, as required by the reference class selection rule.) With appropriately contrived reference classes and attribute classes (reminiscent of Goodman's "emeroses" and "grue"), we could construct crude enumerations to give us

strong evidence that this next coin flip will be a baby elephant. The new problem of induction is how to exclude these inductions while keeping the good ones.

We will not even gesture at a solution of the new problem of induction. The difference between a contrived class and an uncontrived class is a matter of common sense: It will be more or less clear whether a class is suitable for use in an induction, or whether an induction that uses the class cannot provide real evidence. What is important for us is that the attribute and reference classes used in our preliminary analysis of testimony should be suitable, and it should be uncontroversial that they are. The attribute of testimony that we will be concerned with is truth; the reference classes we use will be the sum of all testimony and the sum of testimony told by the particular teller. (For simplicity's sake, we will be neglecting another suitable reference class, that of testimony of a particular teller on a particular topic. This class will effectively be considered in our more refined analysis of testimony as evidence [Chapter VI].) At the level of approximation aspired to by the crude enumeration conception of evidence, the resulting inductions should yield commonsense results about when testimony provides evidence for the truth of what is told.

3. Testimony as Evidence (on the Crude View)

In this section, we will employ the crude enumeration conception of evidence in a preliminary sketch of the offensive component of our argument for Particular Evidentialism. At the level of detail captured by the crude enumeration conception, any testimony that is not evidence for what is told will not give the hearer justification for believing what she is told. Thus a hearer who only believed a piece of testimony when it

gave her evidence, as Particular Evidentialism counsels, would never miss out on any testimonially justified beliefs. This is the offensive component of our argument, because it shows that Particular Evidentialism can explain which beliefs are testimonially justified, while the Assurance View has further explanation to do; the Assurance View's idea that the teller's assurance gives the hearer non-evidential justification for belief does not explain why certain assurances do not provide this justification.

This section's argument is not enough to defend Particular Evidentialism against the objections that motivate the Assurance View (chapter III). The crude enumeration conception of evidence will not be sophisticated enough to show how the teller or hearer can take testimony as evidence while seeing it as a freely chosen action for which the teller assumes responsibility. This will have to wait until we have a more refined conception of evidence (Chapters V and VI). In subsequent sections of this chapter, however, we will use this section's analysis of testimony as evidence to derive a normative structure on which the teller must *be* responsible for the truth of her testimony. This structure, which will be preserved and refined after we move to Chapter V's more sophisticated conception of evidence, will eventually be used in our defense of Particular Evidentialism against these objections (chapter VII).

On the crude enumeration conception of evidence, whether a particular piece of testimony provides evidence for what is told will depend on the reference class to which that testimony is assigned and on the percentage of testimony within that reference class that is known to be true. We will consider two possible reference classes for a particular piece of testimony: the class of all testimony and the class of all testimony by that particular speaker. Both of these classes are clearly uncontrived. Other uncontrived

reference classes are possible, such as the class of all testimony by this particular speaker on this particular subject. We will bracket these for simplicity's sake; considering them would raise issues that we will address when we deal with a speaker's authority on a particular subject (section V.4).

By the rule for selecting a reference class (section 2), we should select the narrowest reference class for which reliable data are available. So the class of all testimony should be selected when the hearer lacks data concerning the speaker's past testimony, and the class of the speaker's testimony should be considered when the hearer has such data. We will examine each of these cases in turn to show that testimony only gives justification for belief when it provides evidence for what is told.

First let us consider the case in which the hearer does not have data concerning the teller's testimony. Suppose that Alice tells Sarah that it is raining, that Sarah has no other evidence concerning whether it is raining, and that Sarah knows nothing about Alice's past testimony.[14] According to the Acceptance Principle for Testimony (section II.3), Sarah is justified in believing what she is told in this case. In the absence of positive evidence against the testimony, the hearer is justified in believing whatever she is told.

On the crude enumeration conception, the testimony also gives Sarah evidence that it is raining. We have argued (section II.4) that General Evidentialism is false, which is to say that (even without gathering evidence) a hearer is justified in believing the claim that most testimony is true. In order to determine whether Alice's testimony is true, Sarah

---

[14] Excluding other evidence about whether it is raining is in accord with our focus on cases in which testimony is the only possible source of evidence; see section 1. Note also that most cases in which Sarah has other evidence will be beyond the power of the crude enumeration view. On the crude enumeration conception, the other evidence must take the form of a reference class under which the current situation could be subsumed; unless Sarah has reliable statistics concerning cases that fall into the intersection of that class and situations in which she has been told that it is raining, it will be impossible to select a unique narrowest reference class for which Sarah has reliable statistics. See section 2.

should carry out a statistical syllogism using the class of all testimony as the reference

class. Since she has no data concerning Alice's past testimony, the class of all testimony

is the narrowest class for which she has data. The appropriate statistical syllogism is:

> (I1) This is a piece of testimony that it is raining.
> (I2) Most testimony is true (by the falsity of General Evidentialism).
> (I3) By (I2), there is good evidence that this is true testimony, which is to
> say by (I1) that there is good evidence that it is raining.

Thus any piece of testimony will yield evidence when the hearer does not have data about

the teller's past testimony.

Note that this is not quite the paradigmatic process of evidence-gathering described in

the previous section, because (I2) is known non-evidentially rather than by observation of

a sample of testimony. Nevertheless, the hearer's experience of being told something

may be seen as evidence that has been gathered. At the very least, justification that

comes from an argument such as (I1)-(I3) can be seen as continuous with justification

that comes from an induction in which all premises are supported by evidence. It is

radically different from the sort of non-evidential justification that supports (I2) itself (see

the end of section II.4). This justification also differs from the sort of non-evidential

justification envisioned by the Assurance View, which comes from the teller's acceptance

of responsibility for the testimony. The argument for the Acceptance Principle and

against General Evidentialism showed that we need not gather evidence to be justified in

believing that most testimony was true. This does not show that the testimony itself

yields non-evidential justification, but rather, as Moran puts it, that "we may have some

non-empirical right to treat this phenomenon [sc. testimony] as evidence" (Moran 1999,

p. 7).

Let us now consider the case in which Sarah does have enough data about Alice's past testimony to compile reliable statistics. Here there are two subcases, depending on whether or not Alice's past testimony (as Sarah knows of it) has mostly been true. When Alice's past testimony has mostly been true, her current testimony gives Sarah evidence that it is raining. When the past testimony has not mostly been true, the current testimony does not provide evidence that it is raining, but neither does it justify Sarah in believing that it is raining. In either subcase, testimony only provides justification when it provides evidence.

Consider first the subcase in which Sarah knows that most of Alice's past testimony has been true. She can compile reliable statistics about Alice's testimony, which is a narrower class than testimony in general. So by our rule for reference class selection, the suitable reference class for induction is the class of Alice's testimony rather than the class of all testimony. This means that Sarah should not apply the inference (I1)-(I3), but should instead apply the following induction:

> (I1a) This is testimony of Alice's that it is raining.
> (I2a) Most testimony of Alice's has been true.
> (I3a) By (I2a), there is good evidence that this is true testimony, which is
> to say by (I1a) that there is good evidence that it is raining.

As in (I1)-(I3), Alice's current testimony gives Sarah evidence that it is raining. This evidence depends on (I2a), the data Sarah has concerning Alice's past testimony. This is a paradigmatic case of evidential justification; Sarah first must gather instances of Alice's past true testimony, and then apply what she has learned to Alice's current testimony.[15]

---

[15] As remarked in our discussion of General Evidentialism (section II.2), in general Sarah will not be able to gather this evidence without somehow relying on someone's testimony; for instance, using someone else's testimony to check a piece of Alice's past testimony.

In the other subcase, Sarah knows that much of Alice's past testimony has not been true. Again Sarah should take as a reference class the class of Alice's testimony rather than the class of all testimony. So the inference (I1)-(I3) is blocked and replaced by the following inference concerning Alice's testimony:

(I1a) This is testimony of Alice's that it is raining.
(I2n) Testimony of Alice's is true about as often as it is false.
(I3n) By (I2n), there is no more evidence that this is true testimony than that this is false testimony, which is to say by (I1a) that there no evidence (from the testimony) whether it is raining or not.

Of course, Alice must have told many falsehoods to make (I2n) true. If Alice's testimony has mostly been false, her testimony will give positive evidence that it is not raining. In general there is no hard and fast line between tellers whose testimony will provide evidence and tellers whose testimony will not; the less of Alice's past testimony that Sarah knows to be true, the weaker the evidence that her current testimony provides Sarah. I claim that the same holds for justification: The less of Alice's past testimony that Sarah knows to be true, the weaker the justification for belief that her current testimony provides Sarah.

According to the Assurance View (chapter III), testimony gives the hearer a non-evidential justification for believing what she is told that is based in the teller's assumption of responsibility for the her testimony's truth. The teller offers an assurance, and the hearer gains justification when she accepts it. In our example, no matter how much of her past testimony has been false, Alice is still offering an assurance and can be seen as responsible for her testimony. (When we explain the credit/discredit normative structure [section 4], we will see how in a certain way she must be responsible for her testimony.) The question, however, is whether Alice's responsibility for her testimony gives Sarah a justification for believing what she has been told. When Alice has told so

many falsehoods, Sarah has no epistemic reason to accept her assurance. To believe an informant with so bad a track record of inaccuracy would not be a good way to maximize true beliefs and minimize false beliefs given the available information. That means, on our conception of justification (see Appendix), that believing such an informant would not be justified.

The fact that Sarah is unjustified in believing Alice does not remove Alice's responsibility for the truth of her testimony. In some situations, for instance, Alice's testimony may make her liable to reimburse Sarah for any bad consequences that befall Sarah if she acts on the false testimony.[16] This right of reimbursement may even make it practical for Sarah to act on what Alice tells her. If, for instance, Alice guarantees that her trucks will transport Sarah's avocadoes without spoilage, and this testimony makes her liable if the avocadoes do spoil, then Sarah can safely ship her avocadoes in Alice's trucks without regard to Alice's past record of testimony. If Alice's testimony is true, then the avocadoes will arrive safely, and if Alice's testimony is false, then she will have to pay Sarah for the avocadoes; so Sarah comes to no harm either way. Such a situation, however, would not mean that Sarah was justified (even in some practical sense) in believing what Alice told her, that the trucks would transport the avocadoes without spoilage. The justified belief that Sarah acts on would be the belief that she will be able to recover damages from Alice if Alice's testimony turns out to be false. If this belief were not justified, then as a practical matter Sarah should only ship her avocadoes in Alice's trucks if she is justified in believing that they will not spoil. The fact that Alice is liable for her testimony will not do Sarah any practical good if Alice will not pay up.

---

[16] This is the sort of responsibility that Thomson (1990) thinks we incur when we give someone our word. The responsibility Thomson discusses is distinct from the credit/discredit normative structure, discussed in section 4.

This shows that, even if justification for belief were defined pragmatically, in terms of justification for acting on belief, the teller's responsibility for her testimony would not justify the hearer in believing what she was told if the teller had been inaccurate in the past. At the most, it would change the circumstances of action so that the hearer was justified in acting on the expectation that she would be reimbursed when the testimony turned out to be false. This change will only occur when the hearer has reasonable expectation of full compensation; otherwise it would be practically unwise to act except when the teller had proved herself reliable.

Part of the appeal of the Assurance View is the idea that, when someone believes false testimony, the teller is responsible for the hearer's false belief. As Moran puts it, "When someone gives me his assurance that it's cold out he explicitly assumes a certain responsibility for what I believe" (Moran 1999, p. 12). If the teller had *sole* responsibility for the hearer's belief, then the teller's assurance would indeed always justify a hearer in believing what she was told, even when it failed to provide evidence. Insofar as the concept of blame applies to an epistemological evaluation of beliefs at all, it would not be appropriate to blame the hearer for believing what she was told if she were not responsible for that belief. It would be inaccurate, however, to assign the teller sole responsibility for the hearer's belief; the hearer also has the epistemological responsibility to avoid being taken in by an unreliable teller. As Pamela Hieronymi puts it, expounding the Assurance View, "[i]t seems that the responsibility for the truth of that belief [gained by trusting testimony] is somehow shared [between hearer and teller]" (Hieronymi 2000, p. 23). We can criticize a hearer on epistemological grounds if she

believes a teller whom she should realize to be untrustworthy, even as we criticize the

teller for defaulting on her responsibility to tell the truth.

Note that this shared responsibility for belief is not unique to the case of testimony. If

Alice plants fake evidence that convinces Sarah of something false, then Alice bears at

least partial responsibility for Sarah's false belief. If, furthermore, Sarah is unjustified in

taking the "evidence" at face value—she should realize that it is planted—then Sarah

would be open to epistemological criticism for forming an unjustified belief. In this case

both Alice and Sarah bear responsibility for Sarah's false belief. Alice's responsibility is

practical, because her actions intentionally led to Sarah's false belief, while Sarah's

responsibility is epistemological, because her false belief is not epistemically justified.

Testimony does not essentially differ from other sorts of evidence in this structure of

shared responsibility. The difference is a matter of frequency. It is fairly common for an

unreliable speaker and a gullible hearer to share responsibility for the false belief induced

in the hearer by the speaker's testimony, whereas it is probably rare for a devious planter

of evidence and a credulous observer to share responsibility for the false belief induced in

the observer by the planted evidence.

The teller's responsibility for her testimony, then, is not enough to justify the hearer

in accepting her testimony when the teller has a record of false testimony. The teller

remains responsible for her testimony, but the hearer is not off the hook with regard to

whether she should believe what she is told; she may still be held epistemically

responsible for believing the teller. Only when the teller has not proved untrustworthy

will the hearer be justified in accepting what she is told. These cases will be exactly

those cases in which the testimony provides evidence, on the level of detail captured by the crude enumeration conception of evidence.

This section has shown that a hearer who always treats testimony as evidence need never miss out on a belief that would be justified, nor believe something that would be unjustified. Particular Evidentialism thus yields the right results about which beliefs are justified. If the Assurance View were correct, by contrast, we would need to explain why those assurances that the hearer is justified in accepting are exactly those that provide evidence of what is told. We have not yet shown that testimony never provides a non-evidential justification for belief, as Particular Evidentialism requires. It could be that testimony that provides evidence also provides a non-evidential justification for belief. The Bad Faith and Disharmony Objections to Particular Evidentialism (section III.3) argue that this must be so, for an evidential view of testimony could never account for the way in which a teller assumes responsibility for her testimony and a hearer accepts this assurance. We have not yet answered these objections.

Indeed, the crude enumeration conception of evidence is not fit to answer these objections. This section's account of how testimony provides evidence has treated testimony simply as a kind of phenomenon that is correlated with the truth, like any natural sign. The teller's responsibility for her testimony plays no role in this account of how testimony serves as evidence. In order to show how taking testimony as evidence is compatible with treating the teller as responsible for her testimony, we will need the more refined conception of evidence put forth in Chapter V.

We can, however, make a start on showing how evidence and responsibility can be reconciled. We have not yet shown how taking someone's testimony as evidence is

compatible with seeing it as a free action, the sort of thing for which she may be held responsible. In the next section, however, we will show that, if taking someone's testimony as evidence is compatible with treating her as responsible for the testimony, then a hearer who takes the testimony as evidence *must* hold her responsible for the truth of the testimony. This is the credit/discredit normative structure on testimony, which will hold even when we analyze testimony with a more sophisticated conception of evidence (chapters VI and VII).

4. The Reliability Sanction and the Credit/Discredit Structure

This section puts forth the credit/discredit normative structure on testimony. The credit/discredit structure analyzes how the teller must be responsible for the truth of her testimony, on the assumption that testimony is a voluntary action for which the teller can be responsible. Given this, it will follow from our analysis of testimony as evidence that the teller must be responsible for the truth of her testimony. We will derive from the epistemology of testimony that the teller stakes her credibility on the truth of her testimony. Should she tell falsehoods, what we call the *reliability sanction* becomes appropriate: Her future hearers should not believe what she tells them.

Throughout this section, we will assume that testimony is a voluntary action for which the teller can properly be held responsible. We have not yet shown how a hearer can take the teller's testimony simultaneously as evidence and as a voluntary action for which she is responsible. To show this will require a more sophisticated conception of evidence than the crude enumeration conception; this more sophisticated conception will be set forth in Chapter V and applied in Chapters VI and VII. This section, accordingly,

will not be able to answer the Disharmony Objection (section III.3) that taking testimony as evidence prevents us from seeing the teller as responsible. It will, however, lay down the outlines of how the teller must in fact be responsible; eventually (section VII.1) we will be able to show that a hearer who takes testimony as evidence does hold the teller responsible for her testimony in this way. The more sophisticated conception of evidence that we use later will change the details of the credit/discredit normative structure, but the main ideas will remain the same.

We must, then, derive the teller's responsibility for her testimony from the epistemology of testimony. Effectively, anyone who follows a policy of treating a teller's testimony as evidence is holding the teller responsible for the truth of her testimony, and this policy is the epistemologically correct policy. The hearer can hold the teller responsible for telling a falsehood by not believing her future testimony, which disbelief we call the *reliability sanction*. We must explain, however, why applying this reliability sanction counts as holding someone responsible for the truth of her testimony.

Our analysis of responsibility in terms of sanctions comes from A. R. Anderson's logic of "ought." Anderson analyzes Ought(p) as $\sim p \Rightarrow S$, where S, for "sanction," is a sentential constant that can be interpreted as "wrongdoing has occurred" (Anderson 1956).[17] To transfer this analysis to the domain of responsibility, when an agent is responsible for doing something she ought to do it; when she does not do it, some sanction applies.[18] In our case, we will argue that the teller is responsible for telling the

---

[17] The implication must be strict implication, or something else that is stronger than material implication.

[18] There is another sense of "responsible for," in which something that someone is responsible for may not be an action that they are obliged to do, but more akin to an action that they have already done, or an effect of such an action. For instance, we may say that someone is responsible for the door's being open; this may be the effect of an omission rather than an action, but the responsibility will be assigned

truth; when she does not tell the truth, some sanction will apply.[19] The sanction will be

that, if the teller has told enough falsehoods, then future hearers should not believe her

testimony.

To show that this reliability sanction truly is a sanction for failing a responsibility, we

must say more about the sanction analysis of responsibility. In particular, we will have to

distinguish between sanctions that are deserved for failing a responsibility and unpleasant

consequences that happen when something foolish is done. Someone who fails to take

care of her health may fall ill, and this may be expressed by a conditional of the form $\sim p$

$\Rightarrow S$, "If you abuse your health sufficiently, you will fall ill." Falling ill, however, is not

a deserved sanction for failing to take care of your health. There is nothing normative in

the operation of your immune system; no one falls ill just because they deserve to. If we

are to show that the teller has a responsibility to tell the truth, we must show that the

reliability sanction for falsehoods is a sanction for falsehoods rather than an unpleasant

consequence.

We should first clarify that, for many kinds of responsibility, matched with the

appropriate sanction, the $\sim p \Rightarrow S$ analysis is trivially true. This is suggested by

Anderson's gloss on S, "wrongdoing has occurred." To say that wrongdoing has

occurred is simply to say that someone has failed to fulfill a responsibility. If someone

has failed a responsibility, it follows that she deserves to be viewed as having failed a

---

retrospectively rather than prospectively. Similarly, we may say that someone is responsible for the
introduction of modern sanitary techniques into hospitals after they have introduced those techniques. In
the text, we are only concerned with the sense of "responsible for" that is akin to obligation.

[19] Note that the argument is that the teller is responsible for telling the truth, not that she is responsible
for guaranteeing that her testimony is true. This leaves open the possibility that a teller can, without
violating her responsibility, make a prediction whose truth is not determined at the time of her utterance. If
the prediction turns out to be true, she will not have violated her responsibility. By contrast, if we required
that the teller guarantee that her testimony is true, she would have to either restrict herself to statements
whose truth was already determined or bring about the truth of her prediction herself.

responsibility (since it is true that she has failed), but it need not follow that any other more concrete sanction is appropriate. We are not offering a reductive account of responsibility. For instance, where someone has failed to fulfill a moral responsibility, it is true that moral wrongdoing has occurred, and it is true that she deserves to be seen as having committed moral wrongdoing.[20] It need not be true that any other punishment or sanction is appropriate for the moral wrongdoing. Certainly we do not mean to assert that there is a kind of punishment that is always appropriate whenever moral wrongdoing has occurred.

There are, however, cases in which there is a concrete sanction that is made appropriate by the failure to fulfill a responsibility. Legal responsibilities are the most obvious example. If someone breaks a law, then a specific punishment becomes appropriate; the lawbreaker may be prosecuted or sued. The penalties laid down in the law, such as jail time, fines, or restitution, are sanctions that go beyond the fact that lawbreaking has occurred, or that others may rightly see the lawbreaker as having failed a legal responsibility.

Legal sanctions are sanctions rather than mere unpleasant consequences of lawbreaking because their application is normative rather than natural. Breaking the law does not lead by any purely natural process to confinement in jail or deprivation of property; the legal authorities must intervene. And it is possible for the authorities to punish someone who has broken no law or to fail to punish someone who has broken the law. Punishment of the innocent is not merited, however, and the guilty who go free

---

[20] Though only those who know of her moral wrongdoing should see her as having committed moral wrongdoing. A similar point will apply to the reliability sanction; see the discussion below of the difference between absolute and personal reliability scores, below.

nevertheless deserve punishment.[21] When lawbreaking leads to legal sanctions, the

sanctions paradigmatically come about because the lawbreaking has made them

appropriate and the relevant authorities have applied those appropriate sanctions. When

we analyze "A is legally responsible that p" as $\sim p \Rightarrow S$, the appropriate gloss on S is

"Legal sanctions against A will be appropriate." Since S is still normative, this preserves

the normative character of legal responsibility and avoids reducing it to unpleasant

consequences that *will* follow.

Our account of the teller's responsibility to tell the truth will be analogous to this

account of legal responsibility. We will show that, when enough of a teller's testimony is

false, future tellers should not believe what she tells them. This loss of credibility, which

we will call the *reliability sanction*, is an unpleasant consequence for the teller.

Furthermore, its operation is normative rather than natural: It is not that hearers *will* not

believe a teller who has told past falsehoods, but that according to epistemic norms they

should not believe her. Thus we have a situation of the form $\sim p \Rightarrow S$ with S stating that a

sanction is appropriate: If someone does not tell the truth, she should not be believed in

the future. This is enough to establish that the teller has some sort of responsibility to tell

the truth. (All this is conditional on the idea that testimony is the sort of thing for which

the teller can be held responsible. In this chapter we are not yet in a position to solve the

---

[21] Note that it need not be a legal responsibility of the authorities that they see that all and only
lawbreakers are punished. When we say that punishment is or is not merited, we are not saying that there is
a law that requires or forbids the legal authorities to mete out punishment. Not only may the authorities be
justifiably ignorant of the facts of the case (see below), the law itself may not require them to act. It is
notorious that under U.S. law an inmate may continue to be punished even in the face of evidence of his
innocence; at a certain stage of legal proceedings, it is necessary to prove previous procedural irregularity
in order to force a new trial. Nevertheless, the punishment of an innocent is not justified by the laws that
mandate punishment for the crime that he did not commit. It need not be that anyone has violated a legal
responsibility for a miscarriage of justice to take place.

problem of how a hearer can treat the teller's testimony as evidence while seeing it as a voluntary action for which the teller can be held responsible.)

That it is unpleasant not to be believed follows from our definition of telling. Telling is meant to provide a reason for belief based on the teller's say-so alone (section I.1); on our Gricean analysis, testimony is an attempt to induce belief by means of the hearer's recognition of the testimony (section I.2). So when a teller is not believed, her attempt to induce belief is frustrated. On the $\sim p \Rightarrow S$ analysis of responsibility, a teller is responsible to tell the truth if failing to tell the truth should lead to an unpleasant consequence through some normatively governed process. The reliability sanction qualifies as an unpleasant consequence because it is the frustration of the teller's attempt. It remains to show that the sanction follows from false testimony by a normatively governed process.

Thus we must show why a teller who tells too many falsehoods should not be believed. According to the analysis of the previous section, when a hearer considers testimony as evidence, she should take into account whatever she knows about the teller's past testimony. If the teller's past testimony has mostly turned out to be true, the current testimony provides her with evidence for what she has been told, as in inference (I1a)-(I2a)-(I3a). If the teller's past testimony has been false as often as it has been true, the current testimony provides her with no evidence for what she has been told, as in inference (I1a)-(I2n)-(I3n). The greater the proportion of falsehoods that the teller has told, the weaker the evidence that the current testimony provides, if any. In the previous section we argued that testimony only provides justification when it provides evidence (even if, as the Assurance View has it, testimony that provides evidence can also provide

a non-evidential justification for belief). So past falsehoods weaken the extent to which the believer is epistemically justified in believing what she has been told. When the hearer does not believe what she has been told because the teller's past falsehoods have weakened her testimony as evidence, that disbelief is what we will call the *reliability sanction*.

Epistemic justification is the norm that makes the reliability sanction a true sanction rather than a mere natural consequence of telling falsehoods. It is not that someone who tells too many falsehoods will never be believed. Even a hearer who knows of her past false testimony may foolishly believe her. Similarly, a teller who has told the truth in the past may fail to convince her hearers. An epistemologically sound hearer, if she knows the teller to have been unreliable, should not believe what she is told; that is a normative connection rather than a natural one. The teller stakes her credibility on the truth of her testimony, but credibility is a normative notion, determining whether her hearers *should* believe what she says.

A teller of falsehoods may avoid the reliability sanction not only if hearers fail to apply their knowledge of her unreliability, believing her without justification, but also if the hearers do not know of her past false testimony. On our conception of evidence and justification (see Appendix), the hearer's evidence depends only on the information that is available to her. Whether current testimony gives a hearer evidence can depend only on the past testimony that the hearer already knows about. So if Alice has told many falsehoods, but Sarah does not know about the falsehoods that Alice has told, then Alice's testimony does give Sarah evidence and justification for believing what she is told. (The evidence may be misleading, but it is evidence nevertheless.) In this case it

would seem that the reliability sanction does not apply: Alice has told falsehoods, but Sarah would not be justified in disbelieving Alice's current testimony. If Sarah did not believe what she was told, she would violate the epistemic norm of believing what you are justified in believing. This might seem to violate the $\sim p \Rightarrow S$ analysis of the teller's responsibility for the truth of her testimony; the teller has told falsehoods, but it is not appropriate for the hearer to apply the reliability sanction by refusing to believe her current testimony.

In this respect the reliability sanction is no different from legal sanctions. For the machinery of the law to be brought to bear on an offender, it is not enough that a law be violated; the officers of the law must know about the violation. It would be inappropriate for the authorities to prosecute someone whom they were not justified in believing to have broken the law, even if that person had in fact broken the law. We can distinguish between the fact that the lawbreaker is objectively liable to legal sanctions and the fact that sanctions may not be appropriately imposed by specific authorities who do not have the appropriate information. Similarly, it will be inappropriate for a hearer to exercise the reliability sanction against a teller whom she is not justified in believing to have told falsehoods, even if that teller has in fact told falsehoods.[22]

---

[22] There is a disanalogy between the undetected criminal and the undetected false teller, stemming from the different role of information in the two cases. The criminal legally deserves punishment whenever the authorities would be justified in punishing her if they had full information. We cannot similarly use the idea of full information to explain the cases in which someone deserves the reliability sanction, even though her hearers are not justified in exercising it. If the hearers had full information, they would already know whether what she told them was true, and it would be a moot question whether her testimony served as evidence. We can repair the analogy by saying that someone legally deserves punishment for conduct whenever the authorities would be justified in punishing her if they had full information *about that conduct*. Then we could analogously say that a teller deserves the reliability sanction for her past testimony whenever a hearer would by justified in enforcing it if she had full information *about the teller's past testimony*.

We can think of an individual's reliability in terms of scorekeeping: She accumulates credit for every truth she tells and discredit for every falsehood she tells. We will call the accumulated credit and discredit the teller's *reliability score*.[23] Because individual hearers will not know about all of a teller's testimony, each teller will have two kinds of reliability score. The teller's *absolute* reliability score will reflect every piece of true or false testimony that the teller has given. Even if she paints something on a sign that is swept out to sea before anyone sees it, it will count toward her absolute reliability score. The teller's reliability score *relative* to an individual hearer, or *to* the hearer, will reflect only the information that the hearer has concerning the teller's testimony. Alice's reliability score to Sarah will reflect only the testimony that Sarah is justified in believing Alice to have given, and how much of that Sarah is justified in believing true. Whether Sarah is justified in believing what Alice tells her, on our conception of justification, will depend on Alice's reliability score to Sarah rather than on Alice's absolute reliability score.[24]

---

[23] This notion of keeping score of a speaker's normative status has considerable precedent in the literature. Brandom (1994) bases an account of the philosophy of language on the idea that hearers keep score of the commitments and entitlements of speakers. Brandom draws on the account of scorekeeping in Lewis (1979a), who in turn draws on the account of presupposition in Stalnaker (1973). These accounts do not see the speaker's score as a record of the truth and falsity of her utterances, as our account of the score does, but rather as a record of the propositions to which she is committed, or that she presupposes, or that she must make true.

Our account of scorekeeping is closer to that of the account of assertion in Belnap, Perloff, and Xu (2001). On this provisional account of assertion on which "'α asserts 'A'' is true just in case, if A is true then α is vindicated, and if A is false then α is impugned" (Belnap, Perloff, and Xu 2001, p. 174). (This account is based on Belnap and Green (1994); see also Belnap (2001), p. 19.) The project of this essay can be thought of as specifying how this vindication and impugnment work for testimony.

[24] It might be thought that, on a reliabilist conception of justification, the hearer's justification for believing what she was told *would* depend on the teller's absolute reliability score. On the reliabilist conception, the hearer would be justified in believing what the teller said if believing what that teller said were a reliable method of obtaining true beliefs; the absolute reliability score might be thought to determine how reliable this method was. This, however, is not obvious. The teller's absolute reliability score only reflects the truth and falsity of her past testimony. The reliability of believing the teller's testimony might better be calculated by taking into account all of her testimony, past, present, and future, or by considering internal factors that affect her propensity to tell the truth. To determine which of these is correct would require a worked-out theory of reliabilism. If the absolute reliability score does determine the reliability of

This difference in reliability scores might seem to be a problem for the idea that the reliability sanction is a sanction made appropriate by false testimony. Not every past piece of false testimony makes it more appropriate for the current hearer to apply the reliability sanction. Only those pieces of false testimony that the hearer knows about make the reliability sanction more appropriate. This makes it seem as though the teller's responsibility, rather than to tell the truth, is to avoid getting caught.

In response, the teller's absolute reliability score is not completely disconnected from her reliability score to particular hearers. Accounts of your past testimony eventually are likely to reach the ears of your future hearers. As they do, your reliability score to those hearers will approximate your absolute reliability score, so that it will be appropriate for them to apply the reliability sanction if your absolute reliability score is low enough. Indeed a teller's reliability score to a hearer is simply the absolute reliability score that the hearer is justified in believing the teller to have. The reliability sanction itself is only appropriate when the hearer has a more or less accurate picture of the teller's past false testimony; only tellers who are caught telling falsehoods should have the reliability sanction applied to them. But this does not prove that the teller's responsibility is not to be caught, any more than our legal responsibility is not to be caught, even though only people who are caught should be punished.

We thus have a normative structure for testimony, which we call the *credit/discredit* normative structure. In this structure, the teller is responsible for the truth of her testimony (again, assuming that testimony is the sort of thing for which she is responsible). When she tells a truth, her absolute reliability score is credited, and when

---

believing the teller's testimony, then the credit/discredit normative structure will hold on a reliabilist conception of justification, as well as on the conception of justification used in the text.

she tells a falsehood, her absolute reliability score is debited. In addition, her reliability scores to particular hearers are credited or debited as they learn of her true or false testimony. The sanction for failing her responsibility to tell the truth is the reliability sanction, that a hearer should disbelieve what she tells her. If the teller accumulates too much discredit on her reliability score to a particular hearer, then the hearer will not be justified in believing what the teller tells her, and so should apply the reliability sanction. This "should" is a matter of following epistemic norms on belief.

The applicability of the reliability sanction depends on Particular Evidentialism; if testimony could give the hearer a non-evidential justification for belief, the teller's weakening of her testimony as evidence need not weaken the hearer's justification for believing what she is told. (The Assurance View may be able to accommodate an account on which the hearer's justification for belief does diminish if the teller has told falsehoods, but such an account would need to explain why a non-evidential justification should mirror the testimony's evidential status.) The credit/discredit structure thus lays the foundations for an account of how Particular Evidentialism forces the teller to assume responsibility for her testimony. To give this account, however, we will have to answer the Disharmony Objection and show that it is possible to take testimony as evidence while seeing the teller as responsible for the testimony. This will require the more refined conception of evidence described in the next chapter.