THREE PROBLEMS OF DIRECT INFERENCE

by
Paul Darren Thorn

A Dissertation Submitted to the Faculty of the

DEPARTMENT OF PHILOSOPHY

In Partial Fulfillment of the Requirements
For the Degree of

DOCTOR OF PHILOSOPHY

In the Graduate College

THE UNIVERSITY OF ARIZONA

2 0 0 7

THE UNIVERSITY OF ARIZONA
GRADUATE COLLEGE

As members of the Dissertation Committee, we certify that we have read the dissertation

prepared by Paul Darren Thorn

entitled Three Problems of Direct Inference

and recommend that it be accepted as fulfilling the dissertation requirement for the

Degree of Doctor of Philosophy

____ Date: October 30, 2007
John Pollock

____ Date: October 30, 2007
Terence Horgan

____ Date: October 30, 2007
Shaughan Lavine

Final approval and acceptance of this dissertation is contingent upon the candidate's submission of the final copies of the dissertation to the Graduate College.

I hereby certify that I have read this dissertation prepared under my direction and recommend that it be accepted as fulfilling the dissertation requirement.

____ Date: October 30, 2007
Dissertation Director: John Pollock

STATEMENT BY AUTHOR

SIGNED: Paul Darren Thorn

ACKNOWLEDGEMENTS

There is a long list a people who made contributions to my thought on the topic of direct inference. First, I am indebted for the comments and suggestions given to me by the audiences at the University of Arizona, at the Meetings of the Society for Exact Philosophy, in 2003 and 2004, and at the Conference on Degrees of Belief, at the University of Konstanz. Particularly useful comments were made by Justin Fisher and Isaac Levi. Luc Bovens generously provided useful comments on an early draft of chapter seven, while Allen Habib and my wife, Amanda Brobbel, were available to listen to my developing thoughts regarding the content of chapters six and seven. I am thankful to Keith Lehrer who provided useful direction and professional advice as a member of my prospectus committee. I am thankful to Terry Horgan who stepped in to sit on my dissertation committee, and I am thankful to Shaughan Lavine whose persistent critiques improved the dissertation and helped shape my views on several aspects of direct inference. Most importantly, I am thankful to John Pollock, my dissertation advisor, who provided helpful feedback on every chapter of the dissertation, along with general advice as to how to set out my ideas. While John and I agree on many issues, it seems that we also disagree on many of the details of direct inference. John was very helpful in guiding me to find my view, despite our disagreements.

TABLE OF CONTENTS

TABLE OF CONTENTS – *Continued*

TABLE OF CONTENTS – *Continued*

LIST OF FIGURES

ABSTRACT

Direct inference consists in inference from a premise describing the incidence of a property among a given population to a conclusion about the likelihood of a particular element of the population having the property in question. For example, from the premise that *2% of American males are doctors* one may, in appropriate circumstances, draw the conclusion that *the probability is 0.02 that Joe, a particular American male, is a doctor*. Despite the apparent centrality of direct inference to human belief formation, the manner in which direct inference is to be justified is not well understood. Similarly, no one has succeeded (or even claimed to succeed) in articulating adequate criteria that specify the conditions under which respective instances of direct inference are correct. My dissertation addresses the three most well know problems of direct inference.

CHAPTER 1: REICHENBACH'S THEORY & THE THREE PROBLEMS OF DIRECT INFERENCE

1. Introduction

The subject matter of the present dissertation is the form of inference that is sometimes called "direct inference".[1] Direct inference consists in inference from a premise describing the incidence of a property among a population, as a whole, to a conclusion about the likelihood of a particular element of the population having the property in question. For example, from the premise that *2% of American males are doctors,* and the premise that an individual, Joe, in an American male, one may, in appropriate circumstances, draw the conclusion that *the probability is 0.02 that Joe is a doctor*.

It is apparent, to anyone who considers the topic, that direct inference is central to human belief formation and human judgment under uncertainty. In nearly all cases where we attempt to form judgments of likelihood, our reasoning appeals to frequency information, and is guided by some form of direct inference. Such inferences take place when we make

[1] The origins of the expression "direct inference" are unclear. It appears that the expression derives from the distinction between *direct* and *inverse* probabilities, which was often made in 19th Century discussions of probability and induction, though in fact the sorts of probabilities generated by direct inference do not exactly correspond to what were then regarded as direct probabilities.(Venn 1866) Direct probabilities were those probabilities that were understood to follow deductively from given probabilistic assumptions whereas inverse probabilities were those probabilities that were understood to follow non-deductively from given probabilistic assumptions.

judgments about tomorrow's weather or about who will win Sunday's game.[2] Similarly, we readily understand many of the probabilistic claims that are presented to us as based on frequency data and as grounded in direct inference. For instance, when a TV weatherman tells viewers that there is a 60% chance of rain tomorrow afternoon, we readily understand that this probabilistic claim is meant to be backed by frequency data in the manner prescribed by direct inference.

Despite the apparent centrality of direct inference to human belief formation, the manner in which direct inference is to be justified is not well understood. Similarly, no one has succeeded (or even claimed to succeed) in articulating fully adequate criteria that specify the conditions under which respective instances of direct inference are correct. While the issue of the nature of the justification for direct inference may be somewhat academic (so long as we are confident that it is in fact justified), the issue of how to correctly use direct inference is more urgent. Indeed, while the difficulties involved in using direct inference are not well known, problems abound.[3] Moreover, the very difficulty of these problems and their alleged insolubility have been cited as grounds for eschewing direct inference altogether.[4] In fairness to those who have offered such criticism, little progress has been made in addressing some of the central problems of direct inference. As we will see in

---

[2] In objecting to the present statement, one is cautioned to reflect that our reasoning about frequencies and probabilities is usually more qualitative than is typically depicted during discussions of direct inference. For example, my conclusion that it is *likely* to rain tomorrow is based on the 'frequency' judgment that it is *frequently* rains in Tucson at this time of year.

[3] Echoing the words of C.D. Broad regarding the Hume's Problem of Induction, Isaac Levi described the absence of attention to the problems of direct inference to be a scandal.(Levi 1982)

[4] See, for example, Bacchus et al., 1996.

Chapters Two through Five, large lacunae are present, even among the most sophisticated book length descriptions of the nature of direct inference.

The object of this dissertation is to address some of the problems facing direct inference. As means to articulating the central problems, I will begin by considering the theory of direct inference proposed by Hans Reichenbach, in his book *The Theory of Probability*.(1949) After sketching Reichenbach's theory, and describing the central problems of direct inference, I will spend the next four chapters examining recent leading theories of direct inference, with a particular focus on assessing the degree to which they offer adequate treatments of the central problems that arise in the course of considering Reichenbach's theory. After completing a survey of the leading contemporary theories, and the efforts of those theories to address the various problems of direct inference, I will spend the next two chapters (*Chapters Six and Seven*) weighing in on the three problems. In the final chapter, I summarize my attempt to deal with the three problems, and sketch the remaining lacunae in the work that has been done so far (including my own).

2. Reichenbach's Theory of Direct Inference

Reichenbach's account of direct inference serves as a touchstone for contemporary discussions of the topic. The influence of Reichenbach's discussion of the topic is certainly large, though it appears that the limited detail of his account, and its aptness for a quick summary, may partially account for its influence. Similarly, since all of the core problems of direct inference arise in the context of Reichenbach's account, his theory is apt for illustrating the manner in which things can go wrong.[5]

Reichenbach was an early exponent of the limiting relative frequency interpretation of probability. Reichenbach's version of this proposal includes a wrinkle or two, but the basic idea behind the *limiting relative frequency interpretation* is that probability is a function defined on a sequence of trials. According to a simplified version of the limiting frequency account, the probability of a set T relative to a set R, is the relative frequency of elements of T among the first *n* elements of a given ordering of the elements of R as *n* approaches ∞.[6]

Within his frequency interpretation of probability, Reichenbach's account of direct inference is intended as an account of the applicability of judgments about probability (a function defined on a sequence of trials) to judgments about singular propositions. An auxiliary function of Reichenbach's account of direct inference is to account for the common intuition that it makes sense to assign probabilities to singular propositions.

---

[5] Although one can find a description of direct inference in John Venn's The Logic of Chance (1866) that is virtually identical to Reichenbach's, I will stick to describing the presentation of direct inference that can be found in Reichenbach's *The Theory of Probability*, since contemporary readers are more likely to be familiar with Reichenbach's statements on the topic. It seems likely that Reichenbach's views on the topic were influenced by Venn's discussion, though no explicit reference to Venn's work on the topic is provided by Reichenbach.

[6] Reichenbach's version of the limiting frequency interpretation of probability is idiosyncratic in that he defines respective limiting frequencies relative to pairs of sequences. I will eliminate the additional generality supported by Reichenbach's official account by assuming that we only consider cases where the two sequences are identical. In that case, Reichenbach's account is equivalent to the simplified version of the limiting frequency theory.

Properly speaking, Reichenbach maintains that probability applies only to pairs of sets relative to a specified sequence of trials. And according to Reichenbach's account, the proposal that singular events may be assigned probabilities is, strictly speaking, mistaken. There is, by Reichenbach's reckoning, only one legitimate concept of probability which is a function on sets (relativized to a sequence). The values for 'probability statements' about individual events or objects are, on Reichenbach's account, constructed via a "transfer of meaning", where the value of a relative frequency *stands in* for the value for a single-case 'probability'.

Reichenbach's reluctance to recognize single-case probabilities derived largely from a commitment to the *verificationist criterion of meaningfulness*. Due to this commitment, Reichenbach was unwilling to understand statements of single-case probability as truth-valued in a manner analogous to frequency statements, since it appears that single-case probability statements are not subject to empirical verification (except perhaps in the case where the respective probability value is 0 or 1). Putative statements of single-case probability were thus understood by Reichenbach as commitments of an agent to act *as if* a respective proposition, p, is true in the case where the 'probability' of p was greater than 0.5.[7]

[7] It is likely that Reichenbach would have adopted a more sophisticated account of 'personal probabilities' defined in terms of betting quotients and dispositions to bet, had he been aware of the work of Frank Ramsey (1931), or if he had written subsequent to the publication of Leonard Savage's influential book on the topic of probability and decision theory (1954).

To emphasize his official doctrine regarding single-case probability statements, Reichenbach occasionally espoused a preference to use the term “weight” to refer to the analogue of probabilities as applied to single-case objects and events. But, for the most part, Reichenbach was amenable to the use of the term “probability” to refer to the probability-derivative measure that applied to single-case events.

From here on, I will freely use the term “probability” in describing Reichenbach’s views regarding single-case events. I will also set aside Reichenbach’s concern regarding the legitimacy of single-case probabilities. Reichenbach’s views regarding the semantics of single-case probability statements are not of major interest here. My interest lies in the broad structural features of Reichenbach’s theory, including the exact numeric ‘probability’ values that the account invites us to assign to singular propositions. It is the broad structural features of the principles that Reichenbach proposed that have served as the core of many later accounts of direct inference.—These accounts have themselves diverged from Reichenbach’s account (and each other) in their interpretation of the semantics of single-case probabilities and in their proposals for the sort of statistical statements that function as premises in correct instances of direct inference. In discussing Reichenbach’s theory, my central goal is to uncover the lacunae and the consequent problems that face the theories that share the core structural features of Reichenbach’s theory.

On Reichenbach's account, assigning a probability to the statement that a given object (or event), c, is an element of a given target class, T, consists in locating the individual within a suitable reference class, R, and then concluding the probability that c is in T is equal to r, where r is the (known) frequency of the elements of T among the elements of R.[8] Regarding this proposal, Reichenbach acknowledged that every object may be located within many different reference classes. Reichenbach christened this difficulty "the problem of the reference class". Although Reichenbach's name for the problem of choosing appropriate reference classes has caught on, some aspects of his characterization of the problem have not. For instance, in describing *the problem of the reference class*, Reichenbach stated that "different probabilities will result" when an object is incorporated in different reference classes. Reichenbach also suggested that the problem of the reference class shows that statements of single-case probability are ambiguous.

Notwithstanding his suggestion that statements of single-case probability are ambiguous, Reichenbach proposed that the choice of a reference class in deriving a single-case probability is subject to evaluation. In particular, he proposed that assigning a probability to a singular proposition is a matter of locating the object of the proposition in "the best reference class available". At the most fundamental level, Reichenbach identified the best reference class available as "the reference class that, on the basis of our present knowledge, will lead to the greatest number of successful predications". As a sort of

[8] Throughout his discussion of direct inference, Reichenbach used the term "class" where it would have been appropriate to use the word "set". Because of the influence of Reichenbach's discussion, this anachronistic usage is commonplace among discussions of direct inference. For expository reasons, I will uphold the tradition of using the word "class" in place of the word "set".

ceteris paribus rule, Reichenbach proposed that the best reference class will be the narrowest reference class, R, containing c for which one has *reliable statistics* about the incidence of elements of T among R.(1949, 374)[9]

I have described the preceding rule as a *ceteris paribus* rule due to Reichenbach's suggestion that employing the narrowest reference class available is not always the best choice. According to Reichenbach, there is often a trade-off between the accuracy of our estimate for the statistics for a given class and the desirability of selecting a narrower reference class. In articulating this point, Reichenbach described the problem of mediating such trade-offs (between statistical accuracy and narrowness) as a matter of "technical statistics", though he provided no further guidance as to how to mediate such trade-offs nor any explanation of what technical statistics consists in (nor what he meant by the expression "statistical accuracy").

As Reichenbach recognized, a further difficulty with the prescription that one prefer narrower reference classes arises in cases where one has reliable statistics for two overlapping references classes, and one does not have reliable statistics for their intersection. In the face of such examples, Reichenbach indicated that it was not his place to recommend a solution to such 'conflicts', stating: "The calculus of probability cannot help in such a case because the probabilities prob(B|A) and prob(B|C) do not determine the probability of prob(B|A∩C). The logician can only indicate a method by which our

---

[9] Narrowness, on Reichenbach's account, is judged in terms of the proper subset relation, so that a set, R′, is narrower than a set, R, just in case R′ is a proper subset of R.

knowledge may be improved. This is achieved by the rule: look for a larger number of cases in the narrowest common class at your disposal."(1949, 374)[10] Following these remarks, Reichenbach offered no clear suggestion as to what probability we should assign in the interim prior to collecting more data.

Modern commentators have typically attributed to Reichenbach the doctrine that single-case probabilities can only be assigned in cases where there is a unique narrowest reference class for which one has reliable statistics. In the course of my discussion, I will assume that this is the correct interpretation of Reichenbach's account of direct inference. The present interpretation ignores Reichenbach's suggestion that there is sometimes a trade-off between statistical accuracy and the selection of a narrower reference class. In evaluating Reichenbach's approach, I will be careful to distinguish between cases where Reichenbach's theory yields a mistaken prescription from cases where Reichenbach's theory provides no direction as to how to proceed.

3. The Orthodox Approach to Direct Inference

In the context of my investigation of the topic of direct inference, a number of Reichenbach's views on the topic are incidental, including his view on the semantics of single-case probability statements, and his view on the nature of the statistical statements

[10] In the present quote, I have altered the notation that Reichenbach used to express probabilities, and used notation that will be more familiar to contemporary readers.

that serve as the major premises for instances of direct inference. Outside of his views on these matters, Reichenbach's account of direct inference is seminal to what I will call "the orthodox approach to direct inference". The orthodox approach to direct inference is shared by John Pollock, Henry Kyburg, and myself. The core tenets of the orthodox approach are as follows:

(a) Assigning a single-case probability is a matter of locating an object in a reference class for which one has reliable statistics.

The present tenet of the orthodox view is self-explanatory, and involves the basic idea that direct inference consists in identifying the value of a respective single-case probability with the value of an appropriate statistic (most commonly the value of a frequency statement). In the case of the frequentist version of direct inference, the basic idea consists in reasoning that the probability that a given object, c, is an element of a set, T, is equal to the (accepted) frequency of elements of T among elements of a set, R, where c is an element of R.

While it may appear difficult to imagine a theory of direct inference that does not incorporate the present tenet, the *random-worlds* approach to direct inference does reject the tenet.(Bacchus, et al., 1996) On the random-worlds approach, substitutes for the orthodox rules of direct inference are recovered as theorems of a semantic account of probability statements taken in combination with a rule for assigning probabilities to an underlying space of possible worlds. Because of the way that the random-worlds approach reconstructs the rules of direct inference as theorems of an underlying system, the approach in no way incorporates a principle of direct inference as a primitive rule. Moreover, direct inference within the approach reduces to Bayesian conditionalization, subsequent to an assignment of prior probabilities to the elements of a space of possible worlds. In effect, the ampliative dimension of the random-worlds approach resides chiefly in its rule for assigning probabilities to the elements of an underlying space of possible worlds.[11]

While I have not made the point explicit in articulating the present tenet (i.e. *tenet (a)*), it is my intention that the tenet be satisfied by accounts of direct inference where the value of a given single-case probability may be determined by locating a respective object in more than one reference class. By interpreting the scope of the present tenet in this way, it applies to the theories of direct inference proposed by Kyburg and Pollock. The theories of Kyburg and Pollock recognize cases where a respective single-case probability is a function of the values of two or more statistical statements. I will follow in the steps of Kyburg and Pollock in advocating for such an approach.

[11] Pollock's account is similar to the random-worlds approach in that Pollock's principles of non-classical direct inference are derived from other principles that Pollock takes as axiomatic within his system. Nevertheless, direct inference (what Pollock calls "classical direct inference") is still grounded in a definition of single-case probabilities in terms of statistical probabilities, where each singe-case probability is identified with the statistical probability for a reference class that incorporates all of an respective agent's information about the respective object of interest. For this reason, I regard Pollock's theory as satisfying the present tenet of the orthodox approach.

(b) Direct inference does not rely on positive knowledge of the manner in which the object of inquiry is selected from the chosen reference class.

The present tenet expresses a stance on the background knowledge one is required to have before one is entitled to make a direct inference. The nature of the stance is easy to grasp by considering the contrast between Reichenbach's theory of direct inference, which accepts the tenet, and Isaac Levi's theory of direct inference which does not.

According to Reichenbach's theory, in making a direct inference, one identifies the probability that an object, c, is an element of a set, T, with the value of a frequency expression (i.e., the frequency of elements of T among elements of R), where c is an element of R. In effect, direct inference on Reichenbach's account relies on two premises:

(1) The proportion of individuals of kind R that are individuals of kind T is r.
(2) c is an individual of kind R.

In opposition to Reichenbach's account, Isaac Levi has proposed that correct instances of direct inference require that an agent accept premises of the following form:

(1) The chance of an outcome of kind *T* on a trial of kind *R* on a chance setup C is *r*.
(2) c is a trial of a kind *R* on a chance setup C.[12]

In essence, Levi proposal is that correct instances of direct inference presuppose that the object of interest, c, is presented to us as a trial of a stochastic process that generates

[12] This is a paraphrase of the conditions introduced by Levi.(Levi 1982, 193)

varying results with certain chances, and that the probability we assign to an object of interest having the respective target property be identical to the chance of a trial of the respective sort (which has the chosen reference property) having the respective target property. For example, Levi's account would not allow us to reason from the premise that *10% of American males are doctors* to the conclusion that *the probability is 0.1 that Joe, a particular American male, is a doctor*. In order to use frequency data to draw such a conclusion, Levi would require that we know that Joe was selected by some sort of stochastic process which delivers doctors 10% of the time in the case where the result of the process is an American. In effect, Levi's account requires that we know that Joe is a 'randomly' selected American as a precondition to concluding that the probability that Joe is a doctor is equal to the frequency of doctors among Americans.

Reichenbach's account, like the account of many of his successors, such as Kyburg and Pollock, does not require *positive knowledge* of the sort demanded by Levi. In the place of such positive knowledge requirements, theories of direct inference that incorporate the present tenet of the orthodox approach generally include 'negative knowledge requirements' that outline a range of cases where the possession of information of a given sort would undermine a respective instance of direct inference. In essence, one can think of Levi as requiring that an agent have knowledge of a given proposition (roughly to the effect that the object of interest is a randomly selected element of the proposed reference class) in addition to the respective statistical statement, before the agent is justified in performing direct inference. By contrast, theories that abide by *tenet (b)* propose that knowledge of the respective statistical premise is sufficient to warrant direct inference, so long as the agent does not have additional propositional knowledge which indicates that the respective statistical statement (associated with the proposed reference class) is inappropriate for direct inference (relative to the object of interest). The following tenet encapsulates the most common negative knowledge requirement.

(c) A central criterion for the selection of admissible reference classes is narrowness.[13]

Reichenbach's account epitomizes the present tenet, by proposing that, in assigning a probability to a singular proposition, we conduct direct inference by choosing the narrowest applicable reference class for which we have reliable statistics. As we will see when we survey the accounts of direct inference offered by Kyburg and Pollock, there are other ways of implementing the present tenet. On Kyburg's account, for example, frequency statements regarding a broader reference class are removed from the set of frequency statements that are considered 'relevant' to determining the value of a respective 'evidential probability' (provided that other conditions hold). Similarly, on Pollock's account, in the case of pairs of otherwise admissible instances of direct inference, an instance of direct inference based on a statistic for a broader reference class will be *defeated*, in the case where the two instances of direct inference would lead one to accept a contradictory pair of conclusions.

---

[13] In general, narrowness is determined via of the proper subset relation, so that a set, R′, is narrower than a set, R, just in case R′ is a proper subset of R.

Tenet (c) embodies a negative knowledge requirement inasmuch as the tenet specifies conditions under which an instance of direct inference based on statistics for a given reference class will be undermined. In effect, an instance of direct inference based on statistics for one reference class will be undermined if one is apprized of corresponding adequate statistics regarding a narrower reference class.

4. The Three Problems

In the context of evaluating Reichenbach's theory of direct inference, there are three problems involved in the choice of appropriate reference classes. In the early parts of this dissertation, I will survey some prominent approaches to direct inference and consider how they fare in the face of these problems. In the latter chapters, I will propose my own approach to the three problems. At this point, I will sketch the three problems. Each of the problems may be regarded as an aspect of *the problem of the reference class* (as articulated by Reichenbach) though the first problem seems to most closely track the issue that Reichenbach had in mind when he coined that phrase.

i. The Problem of Competing Statistics:

While versions of the present problem may arise for some *unorthodox* theories of direct inference, I will consider the problem as one that arises for theories of direct inference that prescribe that one prefer statistics regarding narrower reference classes in the course of performing direct inference. In this context, the problem largely concerns cases where one has reliable statistics regarding the incidence of a particular property among two overlapping references classes, and one does not have reliable statistics regarding the incidence of the property among the intersection of the two sets. Reichenbach recognized the present problem. His recommendation in such cases is that one withhold belief about the probability of the respective singular proposition.

Though Reichenbach's proposal to withhold belief in cases where there is no narrowest appropriate reference class will keep us from adopting unreasonable conclusions, it seems that there are cases where the incredulity prescribed by Reichenbach is too extreme. And, in many cases, acting in accordance with Reichenbach's proposal will prevent us from forming a reasonable substantive judgment. Moreover, it looks as though Reichenbach's proposal will leave us at a loss in a great number of cases where practical concerns call on us to make some sort of judgment as to the probability or approximate probability of a given singular proposition. Now, of course, if there is no reasonable way to form a judgment in cases where there is no narrowest applicable reference class, then the demand that we make such a judgment is irrelevant from the standpoint of epistemology. On the other hand, given the practical demands involved, we would do well to think the problem through before concluding that there is no way to form a reasonable judgment in such cases.

It seems clear that there are cases where the degree of incredulity prescribed by Reichenbach's theory of direct inference is too extreme. In recognition of this, the theories of Kyburg and Pollock both allow one to draw conclusions via direct inference in cases where one has reliable statistics for two overlapping references classes, but one does not have reliable statistics for their intersection. The sorts cases allowed by Kyburg and Pollock are ones in which the respective instances of direct inference do not lead to inconsistent conclusions. But even in cases where two instances of direct inference do lead to inconsistent conclusions, it seems that there are some cases where intuitive principles will allow one to make a reasonable judgment. For example, suppose that R_1 and R_2 are potential reference classes for drawing a conclusion about the probability that an object, c, is an element of a target set, T. Now suppose that R_1 and R_2 intersect, and one has reliable information regarding the frequency of elements of T among R_1 and reliable information regarding the frequency of elements of T among R_2, but one does not have reliable information regarding the frequency of elements of T among the intersection of R_1 and R_2. In a case so described, there are circumstances where it seems that we are able to form a reasonable judgment as to the probability that c is an element of T. Indeed, suppose that our knowledge of the statistical properties of R_1 regarding T are perfectly symmetric to our knowledge of the statistical properties of R_2 regarding $\overline{T}$. For example, suppose that the frequency of elements of T among the elements of R_1 is 0.75, and the frequency of elements of $\overline{T}$ among the elements of R_2 is 0.75 (and suppose that we have no other information that breaks the apparent symmetry in the relations between the triple c, R_1, *and* T, and the triple c, R_2, *and* $\overline{T}$). In that case, it seems that it would be reasonable to draw the conclusion that the probability that c is in T is 0.5.

In *Chapter Seven*, I will engage in an investigation of the possibility of reaching reasonable conclusions in cases where one has data that permits instances of direct inference that support mutually inconsistent conclusions. In turn, I will offer a theory which proposes the reasonableness of a wide-range of inferences in the face of such *competing* instances of direct inference. The theory that I propose treats competing instances of direct inference as generating competing reasons for belief. I propose that the *rational force* of such competing reasons may be conceived of as analogous to a vector of opposing physical forces. In turn, I propose that intuitively compelling principles may be brought to bear to determine what to believe in the face of such competing reasons.[14]

[14] I have introduced the present problem as one that arises in cases where one has reliable statistics regarding the incidence of a particular property among two overlapping references classes, and one does not have reliable statistics regarding the incidence of the property for the intersection of the two sets. However, it turns out that variants of the present problem may arise in variety of other cases *inasmuch as* distinct instances of direct inference from a consistent statistical data may give rise to a wide variety of mutually inconsistent conclusions, where a preference for narrower reference classes will not resolve the inconsistency.

ii. The Problem of Uninformative Statistics:

The proposal that statistics regarding narrower reference classes are to be preferred as a basis for direct inference leads to problems when combined with the idea that direct inference is based on frequencies. For one, consider the reference class consisting of the unit set of the object about which one wishes to draw a conclusion. If frequency data regarding narrower reference classes is to be preferred *in general*, then it seems that all interesting instances of direct inference will be defeated. Indeed, the frequency of elements of a respective target class among a unit set reference class will always be one or zero, and inference based on such reference classes would seem only to allow the conclusion that respective probabilities are one or zero.

It is possible that Reichenbach had some notion of the present problem (which would have been significant to him since he subscribed to a frequentist account of direct inference). Perhaps an awareness of the problem influenced Reichenbach's suggestion that there is a trade-off between basing our beliefs about a single-case probability upon an accurate statistic or upon a statistic regarding a narrower reference class. Aside from such remarks, and the suggestion that the issue is a problem of 'technical statistics', Reichenbach did not propose a remedy to the problem.

It is quite clear that frequency statements regarding unit set reference classes are inadequate as a basis for drawing substantive conclusions, because they are maximally uninformative (in cases where we are ignorant of their precise values). But there are many intermediate cases where a statistic is imprecise, but at the same time potentially relevant to the single-case probabilities we may assign on the basis of direct inference. For this reason, it is desirable that we have systematic rules which explain the role of all statistical statements, ranging from fully precise to fully imprecise, in conducting direct inference.

Several proposals have been offered as remedies to the present problem. All of the proposals that have been published incorporate some form of the idea that statistics for a narrower reference class may defeat an instance of direct inference based on statistics for a broader reference class *only if* the respective agent's statistics for the narrower reference class would lead to a conclusion (via direct inference) that contradicts the conclusion that the agent may be drawn (via direct inference) using the statistic for the broader reference class.[15]

Beginning in *Chapter Two*, I present variants of an example due to Mark Stone (1987) which shows that previous approaches to the Problem of Uninformative Statistics are flawed. In particular, Stone's example illustrates that in some cases statistics for a narrower reference class can defeat an instance of direct inference based on statistics for a broader reference class, even if the statistics for the narrower reference class could not lead to a conclusion (via direct inference) that contradicts a conclusion that the agent may draw

[15] Kyburg maintains that an instance of direct inference based on statistics for the broader reference class is defeated by statistics for a narrower reference class *unless* the statistics for the broader reference class are more *precise*. Kyburg's theory is a little different from the other orthodox theories that have been proposed. I will address the differences in my discussion of the Problem of Uninformative Statistics in *Chapter Two* and *Chapter Six*.

(via direct inference) using the statistic for the broader reference class (and even if the agent's statistics for the narrower reference class are less precise than agent's statistics for the broader reference class).

In *Chapter Six*, I offer a new approach to the Problem of Uninformative Statistics. The new approach employs intuitive criteria to determine when a statistic is uninformative (in a given circumstance). The approach yields the right answers in the face of examples such as the one presented by Stone, and recognizes that imprecise statistics for a narrower reference class can defeat an instance of direct inference based on statistics for a broader reference class even when instances of direct inference based on the respective statistics would not lead to mutually inconsistent conclusions.

iii. The Problem of Relevant Statistics:

Difficulties that *appear* similar to the ones that Goodman described for inductive inference arise for direct inference.(Goodman, 1955) The first illustration of these difficulties appeared in Henry Kyburg's book, *The Logical Foundations of Statistical Inference*.(Kyburg, 1974) As an aid to grasping these difficulties, consider the following situation.

We are asked to make a judgment about the likelihood that an individual, Bob, is a university graduate. As a basis for making our judgment, suppose that we are justified in accepting the following propositions:

[1] 30% of American males are university graduates.
[2] 40% of Californian males are university graduates.
[3] Bob is a Californian male.

On the supposition that the preceding three propositions (combined with general knowledge of geography) exhaust our knowledge of the factors that are relevant to assessing the probability that Bob is a graduate, it seems that we should form the belief that the probability is 0.4 that Bob is a graduate. In order to justify such a conclusion, it is typical to appeal to a principle that states that, when engaging in direct inference, one should set one's belief about the probability of an individual having a given property to be equal to the relative frequency of the property among the narrowest reference class of which the individual in question is an element. But the story does not end here, for consider the reference class, R, composed of Bob combined with the set of Californian graduates (i.e., R = Californians$\cap$Graduates)$\cup${Bob}). This reference class is narrower than the set of Californians, and Bob is, of course, an element of the set. Of concern, then, is the fact that the proportion of the elements of R that are graduates is guaranteed to be very near to one. Indeed, for simplicity's sake suppose that there are only one hundred Californians, and that the frequency of university graduates among Californians is 0.4. In that case, the frequency of graduates among R is either 40/41 or 1. The problem with which we are faced is that of explaining why we are permitted to conclude that the probability that Bob is a graduate is 0.4, and we are not permitted to conclude that probability that Bob is graduate is *one*, or in nearly *one*.

Defenders of orthodox theories of direct inference, such as Kyburg and Pollock, have tended to think that the present difficulty calls for a restriction on the set of predicates that may be used to designate reference and target classes in legitimate instances of direct inference. In *Chapters Two and Three,* I will consider Kyburg's and Pollock's proposals. The conclusion that I will reach is that their proposals for addressing the problem are quite unsatisfying, and I am sure that both Kyburg and Pollock would agree inasmuch as their proposals are extremely programmatic.

In opposition to earlier proposals, I will argue that the correct diagnosis of what goes wrong in the problematic instances of direct inference, such as in the case of Bob the Californian, is closely related to the justificatory basis of appropriate instances of direct inference. In performing an instance of direct inference, one assumes that the object about which on is reasoning, c, is as likely to be a member of the respective target class, T, as a *random element of the proposed reference class, R* (where we understand a random element of R to be an object that is described only as an element of R which was selected by a process which is equally likely to produce each element of R). In cases where direct inference is used correctly, the conclusion that c is as likely to be in T as a *random element of R* will be justifiable by appeal to the fact that c is in relevant respects indiscernible from the other elements of R. In the example of Bob the Californian, the defeasible presumption in favor of narrower reference classes is suspended, because Bob is relevantly *discernible* from Californian graduates (with regard to the property of being a graduate). The relevant difference, in this case, is illustrated by the fact that our narrowest estimate of the frequency of graduates among the unit set containing Bob is {0, 1}, but our narrowest estimate of the frequency of graduates among the set of Californian graduates is that the relative frequency is 1.

The remedy to the Problem of Relevant Statistics consists in expanding the conditions under which instances of direct inference are defeated. The condition that I propose in *Chapter Six* prescribes that an instance of direct inference is admissible *only if* the object about which one is reasoning is in relevant respects 'relevantly indiscernible' from the other elements of the reference class for the statistic which one wishes to use in performing direct inference. The application conditions for the proposed defeat condition are guided by the background knowledge possessed by a respective agent and not by a restriction on the set of predicates that may be used to designate reference and target classes in the formulation of admissible instances of direct inference. In effect, the defeat condition that is introduced serves as a *negative knowledge requirement* that is closely related to the prescription that one should prefer narrower reference classes in the course of direct inference.

CHAPTER 2: KYBURG’S THEORY OF DIRECT INFERENCE

1. Introduction

Kyburg’s theory of *evidential probability* includes a theory of direct inference as a component part. Up until the *1990s*, Kyburg’s theory of direct inference was by far the most comprehensive in print. Indeed, Kyburg intended his theory to address all of the problems associated with direct inference that were not properly addressed by Reichenbach’s theory. Before going on to consider the adequacy of Kyburg’s theory in the face of the three problems described in *Chapter One*, I will sketch the key components of Kyburg’s theory. In the end, my conclusion will be that Kyburg’s theory improves on Reichenbach’s, but that there are still a number of ways in which Kyburg’s theory could be improved itself. In the present chapter, I will describe the deficiencies of Kyburg’s theory, and in *Chapters Six and Seven*, I will make some proposals that are intended to improve upon earlier orthodox theories of direct inference such as Kyburg’s.

2. Kyburg’s Theory

Kyburg’s theory clearly falls within the family of theories of direct inference that represent developments of Reichenbach’s theory. Kyburg is quite conscious of this, and his characterization of the main idea behind his theory is closely related to the core tenets of Reichenbach's theory. According to Kyburg, "The idea behind evidential probability is a simple one. It consists in two parts: that probabilities should reflect empirical frequencies in the world, and that the probabilities that interest us – the probabilities of specific events – should be determined by everything we know about those events."(Kyburg and Teng 2001, 200) Kyburg shares Reichenbach's view that single-case probabilities are to be based on frequencies. And Kyburg's idea that probabilities regarding specific events "should be determined by everything we know about those events" is fleshed out in the manner of Reichenbach's theory, so that the demand to base our probability judgments regarding an event on *everything we know about the event* is equated with the demand that one prefer narrower reference classes as bases for direct inference. Similarly, Kyburg regards the problem of articulating an adequate theory of direct inference to be largely a problem of specifying rules for selecting correct reference classes. As we shall see, this approach is also shared by Pollock, Bacchus, and myself.[16]

Kyburg describes his theory of evidential probability as having two main parts.(2001, 201) The first part of Kyburg's theory consists of *a theory of direct inference*. The second part consists of *acceptance rules* that are intended to determine when it is correct to accept as a 'practical certainty' a proposition whose probability is less than one. My discussion will focus on Kyburg's theory of direct inference, but toward the end of the chapter, I will

[16] The approach has been criticized by the likes of Bacchus *et al* (1996). I will discuss their alternative proposal in *Chapter Five*. The approach taken by Bacchus *et al* (1996) may be seen as an alternative approach to direct inference or perhaps an alternative to direct inference, depending on one's precise understanding of the notion of *direct inference*.

briefly discuss Kyburg's acceptance rules, in connection with the idea that direct inference may serve as the foundation of statistical induction. While I am sympathetic to Kyburg's claim that direct inference may serve as a foundation for statistical induction, the theory of acceptance that Kyburg appeals to in explaining the connection between statistical induction and direct inference is deeply problematic. Despite this setback, I will argue in *Chapter Six* that acceptance rules are not needed in the grounding of statistical induction in direct inference. I will postpone presenting my reasons for this conclusion, since the explanation of why an acceptance theory is unnecessary is connected with my view regarding the sort of statements that may properly serve as the statistical premises for instances of direct inference. Moreover, I would like to consider the alternative acceptance theory presented by Pollock before discussing my proposal. The fact that my proposal for grounding statistical induction in direct inference does not appeal to an acceptance theory turns out to be a significant virtue of the proposal (both practically and as a piece of the justification for my view regarding the sort of statements that may properly serve as the statistical premises for instances of direct inference) inasmuch as acceptance theories of sort proposed by Kyburg and Pollock are questionable from the standpoint of epistemology.

3. Syntax Versus Semantics

In articulating his theory of direct inference, Kyburg proposes to focus on syntactic criteria for evaluating instances of direct inference. Kyburg professes to regard his choice to focus on syntax as a natural choice inasmuch as we are often forced by considerations of computability to focus on syntax (and proof-theoretic techniques) when evaluating arguments. However, a further substantive motivation for taking a syntactic approach derives from Kyburg's concern to deal with the problem of gerrymandered reference classes that was described in *Chapter One* (which I called "the Problem of Relevant Statistics").

By Kyburg's reckoning, the Problem of Relevant Statistics illustrates that we must limit the set of terms of our language that may appear in the frequency statements that are used in direct inference. Kyburg proposes that there are two ways by which one may attempt to limit the terms. One approach would be to articulate principles that could be used to determine which properties are "projectable".—If we could make such a determination, then we would limit the terms which appear in instances of direct inference to terms which designate projectable properties (regardless of how the properties are designated). A second approach is non-extensional, and consists in providing syntactic criteria which may be used to determine when it is permissible for a given term to appear in an instance of direct inference. According to this approach, one way of designating a given property may be permissible while another way of designating the very same property is impermissible.

Kyburg opts for the second approach, citing as justification the seeming failure of the project of specifying the set of projectable properties.

Kyburg's approach of the Problem of Relevant Statistics is best understood in the context of what is probably the received view regarding Goodman's *New Riddle of Induction*.(Goodman 1955) The received view on Goodman's problem is that difficulties arises when we mix 'normal' and 'grue-type' predicates in the course of inductive inference. The thought is then that some constraint on which predicates we allow to feature in the expression of inductive inferences will prevent the sort of problems that Goodman's examples illustrated. For many, the project of solving Goodman's riddle was conceived to be one partitioning the normal predicates from the grue-type predicates, and then exhibiting some asymmetry among the properties designated by two sorts of predicate that could be appealed to as grounds for preferring the normal predicates in the course of inductive inference.

Kyburg clearly regards the problem of gerrymandered reference classes that arises in the course of direct inference as closely connected to the problem that Goodman uncovered for induction. That Kyburg thinks that there is a connection is borne out by his assumption that the problem of delineating the set of projectable predicates with regard to direct inference is identical to the problem of delineating the set of projectable predicates with regard to induction. Moreover, since Kyburg regards the project of delineating the set of projectable properties as having failed in the case of induction, he thinks it has failed (and will continue fail) in the case of direct inference. While Kyburg does not assert that the project of delineating *the* set of projectable properties could not succeed, he does think that we must proceed in the absence of this sort of solution. Moreover, he believes that it is possible to delineate a set of predicates, such that if we restrict ourselves to those predicates, we can avoid the problem of gerrymandered reference classes.

Kyburg concedes that his approach to the problem of gerrymandered reference classes will make evidential probabilities language relative. In the face of this potential problem, Kyburg proposes that "it is important to provide an analysis of the basis on which one language is to be preferred to another as a medium for embodying knowledge."(Kyburg 2001, 202)

4. The Syntax of Kyburg's Theory

Kyburg's theory is articulated so as to apply to sentences formulated in a standard first order object language. The language may include functions and predicates of any arity, and will include separate variables for objects (or events) and for *field terms*, which are intended to be interpreted so as to designate real numbers. The central novelty of the object language for Kyburg's theory consists in the inclusion of a variable binding operator, "%", that connects formulas and field terms to form expressions of the form

'$\%_X(\tau|\rho) \in [p, q]$'.[17] In such expressions, '*X*' represents a sequence of variables, while "%" binds the elements of *X*. In turn, 'τ' and 'ρ' will be first order formulae, and 'p' and 'q' constants denoting real numbers. For example, "$\%_x$(Doctors(x)|Californians(x)) $\in$ [.05, .10]" would be interpreted as expressing the claim that between five and ten percent of Californians are doctors.

Auxiliary to a formal specification of the syntax of the theory's object language, Kyburg provides criteria for determining when a formula designates an admissible reference or target class. The admissibility conditions are quite complex, so I will not attempt to summarize them. The basic idea is to have a set of atomic formulae that may serve to designate acceptable target and reference classes, and allow other formulae to serve to designate acceptable target and reference classes, so long as the respective formulae are related to the atomic formulae in specified ways. Although Kyburg provides no explanation of the features of his criteria for admissible reference and target formulae, it is clear that one of Kyburg's goals in designing the criteria was to prevent the formation of 'disjunctive' reference formulae given a fixed stock of atomic formulae.

[17] For the purposes of presentation, I have made some minor modifications to the syntax of Kyburg's theory. Where I have used expressions of the form '$\%_X(\tau|\rho) \in [p, q]$', Kyburg uses expressions of the form '$\%_X(\tau, \rho, p, q)$'. While the presentation that I have adopted is more typical, and thereby easier to follow, Kyburg's original presentation serves to emphasize Kyburg's tenet that the statistical formulae that serve as the basis of direct inference must be interval valued.

5. Choosing the Right Reference Class(es)

The heart of Kyburg's theory of direct inference consist in its rules for determining which frequency statements have a bearing on the evidential probability of given first order statements. Unlike Reichenbach's theory, Kyburg's theory allows for the possibility that more than one statistical statement may be relevant to determining one's evidential probability for a given first order statement.

Formally, Kyburg's theory is laid out as a procedure for considering the full set of statistical statements that could *potentially* bear on the evidential probability of a given statement, and then removing elements of that set which are subsequently deemed irrelevant to our evidential probability for the statement in question. The basic idea of this aspect of Kyburg's theory was present in Reichenbach's theory of direct inference in a less developed form. In the case of Reichenbach's theory, we are told to fix our probability that an individual, c, is an element of a *target class*, T, to the frequency of elements of T among R, where R is the narrowest *reference class* containing c for which one has *reliable statistics* about the incidence of elements of T among the set.(1949, 374) One may conceptualize the algorithm proposed by Reichenbach's theory as follows: Consider each reference class for which we have *reliable statistics* about the incidence of elements of T among the set (and conceive of the corresponding statements as the set of statements that could potentially bear upon the probability of the respective statement). Now remove elements of the initial set of frequency statements that are superceded by other elements in accordance with the preference for narrower reference classes.—In the case of Reichenbach's theory, we simply remove a statement in case another element of the initial set concerns a narrower reference class. Kyburg's theory works in a similar manner, though (as we shall see in a moment) the conditions under which we remove elements from the initial set are a little more complex.

Beyond the articulation of an object language, the analytic structure of Kyburg's theory begins with the notion of a frequency statement providing *prima facie support* for a given first order statement, and with a procedure for isolating the set of frequency statements which may bear upon the probability of a given first order statement, S. In general, Kyburg proposes that a statistical statement '$\%_X(\tau|\rho) \in [p, q]$' *provides prima facie support* for a statement 'S', for a given individual, so long as the individual accepts the statistical statement and also accepts that $\rho(c)$ and $\tau(c)\equiv S$. Where Γ is the set of statements that an individual accepts, Kyburg defines $\Delta(S, \Gamma)$ to be the set of statistical statements in Γ that may bear on a statement, S. Precisely:

$$\Delta(S, \Gamma) = \{ \sigma \mid \exists p, q, \tau, \rho, X\text{: } `\rho(c)' \in \Gamma \wedge `\tau(c)\equiv S' \in \Gamma \wedge \sigma = `\%_X(\tau|\rho) \in [p, q]' \}.$$

According to Kyburg's theory, the set $\Delta(S, \Gamma)$ describes the set of statistical statements that *could be* relevant to the probability of S, given background knowledge Γ. The set $\Delta(S, \Gamma)$ also serves as the starting point in determining the set of statistical statements that actually have a bearing on the probability of S.

Kyburg outlines three sorts of ground by which we may discount an element of $\Delta(S, \Gamma)$ as being relevant to the probability of S. In turn, the algorithmic structure of Kyburg's theory may be represented as a series of tests for removing statistical statements from $\Delta(S, \Gamma)$ in order to form a 'focused' set of statistical statements that are relevant to the probability of S. Kyburg uses the term "sharpening" to refer to the procedure of removing elements of $\Delta(S, \Gamma)$ for the purpose of forming a focused set of frequency statements relevant to the probability of S. Similarly, Kyburg uses the term "sharpens" to refer to a relation between sets of frequency statements that holds when one set is preferable to another as a basis for drawing conclusions via direct inference. Kyburg's theory incorporates three sharpening relations: *sharpens by precision*, *sharpens by specificity*, and *sharpens by richness*.

Kyburg's theory appeals to particular notion of the conditions under which a pair of intervals 'disagree' (and by extension the conditions under which a pair of interval values statistical statements disagree). Specifically, a pair of intervals, [p, q] and [r, s], are said to "disagree" *just in case* it is not the case that $[p, q] \subseteq [r, s]$, and it is not the case that $[r, s] \subseteq [p, q]$.—We write '[p, q]#[r, s]' when a pair of intervals *disagree* in the preceding sense.

The recommendation that the set of statistical statements relevant to assigning a probability to a proposition, S, be *sharpened by precision* tells us, roughly, that if the full set of statistical statements relevant to the probability of S (i.e., $\Delta(S, \Gamma)$) contains a statement $\%_X(\tau|\rho) \in [p, q]$ and a statement $\%_X(\tau|\rho') \in [p', q']$, and [p, q] is a proper subset of [p′, q′], then we should remove $\%_X(\tau|\rho') \in [p', q']$ from consideration. That is, $\%_X(\tau|\rho')$

$\in [p', q']$ will be one of the statements that is removed from $\Delta(S, \Gamma)$ in the course of forming the set of frequency statements that are ultimately relevant to assigning an evidential probability to S.

The recommendation that the set of relevant statistics be *sharpened by specificity* tells us, roughly, that if the full set of frequency statements that are potentially relevant to the probability of S contains a statement $\%_X(\tau|\rho) \in [p, q]$ and a statement $\%_X(\tau|\rho') \in [p', q']$, such that $(\forall X)(\rho' \supset \rho)$ is known, and $[p, q]\#[p', q']$, then we should remove $\%_X(\tau|\rho) \in [p, q]$ from consideration as a statistic relevant to the probability of S (unless there is an element of $\Delta(S, \Gamma)$, $\%_X(\tau|\rho'') \in [p'', q'']$, such that one knows that $(\forall X)(\rho'' \supset \rho')$, and $[p', q']\#[p'', q'']$).

Sharpening by richness is similar to sharpening by specificity, but extends the applicability of the latter condition to cover cases where the number of unbound variables appearing in the target and reference formulas for two frequency statements differ. Roughly, the condition tells us that *if* the full set of frequency statements that are potentially relevant to the probability of S contains a statement $\%_X(\tau|\rho) \in [p, q]$ and a statement $\%_{X'}(\tau'|\rho') \in [p', q']$, such that $(\forall X')(\rho' \supset \rho)$ is known, and $[p, q]\#[p', q']$, *then* we should remove $\%_X(\tau|\rho) \in [p, q]$ from consideration as a statistic relevant to the probability of S (unless there is an element of $\Delta(S, \Gamma)$, $\%_{X''}(\tau''|\rho'') \in [p'', q'']$, such that $(\forall X'')$ $(\rho'' \supset \rho')$ is known, and $[p', q']\#[p'', q'']$).

In introducing the three kinds of sharpening, I stated that they *roughly* describe the prescriptions of Kyburg's theory. The caution suggested by my use of the term "roughly" is due to the way that the different sorts of sharpening may interact in some cases. Indeed, consider a case where $\Delta(S, \Gamma)$ consists of the following statistics:

$$\begin{aligned}
&\%_x(\tau(x)|\rho_1(x)) \in [0.3, 0.6],\\
&\%_x(\tau(x)|\rho_2(x)) \in [0.2, 0.4],\\
&\%_x(\tau(x)|\rho_3(x)) \in [0.1, 0.4], \text{ and}\\
&\%_x(\tau(x)|\rho_4(x)) \in [0.3, 0.5].
\end{aligned}$$

And suppose that we known that $(\forall x)(\rho_1(x) \supset \rho_2(x))$ and that $(\forall x)(\rho_3(x) \supset \rho_4(x))$. In that case:

$\%_x(\tau(x)|\rho_1(x)) \in [0.3, 0.6]$ is more specific than $\%_x(\tau(x)|\rho_2(x)) \in [0.2, 0.4]$,
$\%_x(\tau(x)|\rho_2(x)) \in [0.2, 0.4]$ is more precise than $\%_x(\tau(x)|\rho_3(x)) \in [0.1, 0.4]$,
$\%_x(\tau(x)|\rho_3(x)) \in [0.1, 0.4]$ is more specific than $\%_x(\tau(x)|\rho_4(x)) \in [0.3, 0.6]$, and
$\%_x(\tau(x)|\rho_4(x)) \in [0.3, 0.5]$ is more precise than $\%_x(\tau(x)|\rho_1(x)) \in [0.3, 0.6]$.

If we proceed by removing sharpened elements, in the present case, then we will arrive at a different subset of $\Delta(S, \Gamma)$ depending on the order in which we apply *sharpening by specificity* and *sharpening by precision*. For example, if we first remove the elements that are sharpened by precision, then the result will be the unsharpenable set $\{\%_x(\tau(x)|\rho_2(x)) \in [0.2, 0.4], \%_x(\tau(x)|\rho_4(x)) \in [0.3, 0.5]\}$. On the other hand, if we first remove the elements that are sharpened by specificity, then the result will be the unsharpenable set $\{\%_x(\tau(x)|\rho_1(x)) \in [0.3, 0.6], \%_x(\tau(x)|\rho_3(x)) \in [0.1, 0.4]\}$.

Kyburg's solution to the present problem is straightforward. Normally, Kyburg's theory identifies the *evidential probability* of a statement, S, relative to a body of knowledge, Γ, as the narrowest interval, [p, q], that contains all of the intervals within a subset, Δ, of $\Delta(\Gamma, S)$ such that (1) Δ is not sharpened by any of its subsets, and (2) Δ sharpens each of its supersets. In cases where the order in which frequency statements are sharpened would yield different answers as to the evidential probability of given statement (i.e., there is no unique Δ satisfying (1) and (2)), Kyburg asserts that the evidential probability of the statement is *undefined*. Setting aside such cases where a respective evidential probability is undefined, let us consider a few examples of how Kyburg's theory works.

Suppose that "sixty(x)" corresponds to the property of living to be sixty, "american(x)" corresponds to the property of being an American, and "smokes(x)" corresponds to the property of being a person who smokes cigarettes. In that case, suppose that Mary is an American smoker, and "m" designates Mary. Now suppose that our knowledge base, Γ, consists of the following formulae:

$\%_x(\text{sixty}(x)|\text{american}(x)) \in [0.94, 0.96]$,
$\%_x(\text{sixty}(x)|\text{american}(x)\wedge\text{smokes}(x)) \in [0.74, 0.76]$,
american(m), and
smokes(m).

In that case, Kyburg's theory yields the result that the evidential probability that Mary will live to be sixty is equal to the interval [0.74, 0.76]. In this case, the formulae "$\%_x(\text{sixty}(x)|\text{american}(x)) \in [0.94, 0.96]$" is sharpened (by specificity) by the formulae "$\%_x(\text{sixty}(x)|\text{american}(x)\wedge\text{smokes}(x)) \in [0.74, 0.76]$".

In the preceding case, Kyburg's theory tracks the recommendation of Reichenbach's theory in prescribing that we identify the probability that an individual has a respective property with the frequency of the property among the narrowest appropriate set (known to contain the object of inquiry) for which we have reliable statistics. In other cases, the prescriptions of Kyburg's theory differ from Reichenbach's. For example, suppose that "sixty(x)" corresponds to the property of living to be sixty, "american(x)" corresponds to the property of being an American, and "miner(x)" corresponds to the property of being a person who works in a coal mine. In that case, suppose "h" designates an American named "Hank" who works in a coal mine, and suppose that our knowledge base, Γ, consists of the following formulae:

$\%_x(\text{sixty}(x) \mid \text{american}(x)) \in [0.94, 0.96]$,
$\%_x(\text{sixty}(x)|\text{miner}(x)) \in [0.52, 0.58]$,
american(h), and
miner(h).

If we now assume that we lack knowledge of the frequency with which Americans who work in coal mines live to be sixty, Reichenbach's theory prescribes that we suspend belief regarding the probability that Hank will live to be sixty. Kyburg's theory, on the other hand, yields the result that the evidential probability that Hank will live to be sixty is equal to the interval [0.52, 0.96], which is the narrowest interval covering the intervals contained in Δ(sixty(h), Γ).

We may combine the features of the two preceding examples by formulating a case where we wish to draw a conclusion about an individual, Kathleen, who is an American who smokes and works in a coal mine. In that case, suppose that “sixty(x)”, “american(x)”, “smokes(x)”, and “miner(x)” are as before, and “k” designates Kathleen. Now suppose that our knowledge base, Γ, consists of the following formulae:

$\%_x$(sixty(x)|american(x)) ∈ [0.94, 0.96],
$\%_x$(sixty(x)|american(x)∧smokes(x)) ∈ [0.74, 0.76],
$\%_x$(sixty(x)|miner(x)) ∈ [0.52, 0.58],
american(k),
smokes(k), and
minor(k).

In that case, Reichenbach’s theory prescribes that we suspend belief regarding the probability that Kathleen will live to be sixty, since there is no appropriate *narrowest* reference class for which we possess reliable statistics. Kyburg’s theory, on the other hand, yields the result that the evidential probability that Kathleen will live to be sixty is equal to the interval [0.52, 0.76], which is the narrowest interval covering the intervals contained in the unique unsharpenable subset of Δ(sixty(k), Γ). In this case, the formulae $\%_x$(sixty(x)|american(x)) ∈ [0.94, 0.96]” is sharpened by the formulae “$\%_x$(sixty(x) |american(x)∧smokes(x)) ∈ [0.74, 0.76]”, so that following formulae form the unsharpenable subset of Δ(sixty(k), Γ), upon which our probability for sixty(k) is to be determined:

$\%_x$(sixty(x)|american(x)∧smokes(x)) ∈ [0.74, 0.76], and
$\%_x$(sixty(x)|miner(x)) ∈ [0.52, 0.58].

Let us now consider the manner in which Kyburg’s theory addresses the three problems of direct inference described in *Chapter One*, and the degree to which Kyburg’s treatment of the three problems is satisfactory.

6. The Three Problems

i. The Problem of Competing Statistics:

The Problem of Competing Statistics largely concerns cases where one has reliable statistics regarding the incidence of a particular property among two overlapping references classes, and one does not have reliable statistics regarding the incidence of the property among the intersection of the two sets. In such cases, a simple theory of direct inference, such as Reichenbach's, gives one no guidance in favoring either of the two possible instances of direct inference based on the statistics for the two overlapping reference classes.

In cases where one has reliable statistics regarding the incidence of a particular property among two overlapping references classes, but no useful statistics regarding the intersection of the two sets, Reichenbach proposed that one should suspend judgment regarding the probability of the respective statement. In the face of this proposal, I have suggested that in some cases where one has divergent statistics regarding two intersecting reference sets, it is possible to use the divergent statistics to arrive at a single-case probability (or at a narrow range of possible probability values). In effect, my position is that Reichenbach would have us take a course that is too incredulous. In its favor, Kyburg's account recommends a far less incredulous approach than Reichenbach's. Even so, as I will explain later, Kyburg's account, itself, would have us take a course that is too incredulous. Another problem with Kyburg's account becomes evident largely because of the increased precision of the account. Before considering this 'new' permutation of the Problem of Competing Statistics, let us take a look at the manner in which Kyburg's account improves on Reichenbach's.

One area in which Kyburg's theory differs from Reichenbach's is in its inclusion of a *precision criterion* as well as a *specificity criterion*. In order to illustrate the manner in which Kyburg's account differs from Reichenbach's, in the face of the problem of competing reference classes, let us restrict our attention to cases where the neither the precision nor the richness criterion are in play.—In particular, let us restrict our attention to cases where no element of $\Delta(\Gamma, S)$ may be removed via sharpening by precision or sharpening by richness. In that case, there will be a unique subset of $\Delta(\Gamma, S)$ such that none of the elements of $\Delta(\Gamma, S)$ that can be sharpened by specificity are included in the set. Call this unique subset of $\Delta(\Gamma, S)$, "Γ^*". In that case, Γ^* represents the result of removing each frequency statement in $\Delta(\Gamma, S)$ such that there is another frequency statement in $\Delta(\Gamma, S)$ that is associated with a narrower reference class. In such cases, the evidential probability for S will be the narrowest interval that includes all the intervals present in Γ^*. For example, if $\Gamma^* = \{\ \%x\ (\tau(x), \rho(x), 0.1, 0.1), \%x(\tau(x), \rho'(x), 0.8, 0.8)\ \}$, then the *evidential probability* of $\tau(c)$ will be [0.1, 0.8].

Kyburg's approach is clearly more permissive than Reichenbach's in allowing one to form beliefs on the basis of direct inference, and his theory is better for this fact. Indeed, in the face of the data described in the previous example, the conclusion that the probability of $\tau(c)$ is in [0.1, 0.8] is reasonable. The conclusion makes use of the relevant data (i.e., $\Delta(\Gamma, \tau(c))$) to draw a relatively modest conclusion. And given the data it would be unreasonable to think that probability of $\tau(c)$ lies outside of [0.1, 0.8]. Despite lauding the fact that Kyburg's theory improves on Reichenbach's, I think that it is reasonable to accept a less conservative conclusion than the one prescribed by Kyburg's theory.

In evaluating the conclusion that Kyburg's theory prescribes in the present example, it is important to think of probabilities as a measure of the degree to which a given body of evidence supports (or does not support) given propositions. In this light, it seems to me that Kyburg's proposed conclusion in the present example is somewhat too conservative. In the end, the theory that I propose in *Chapter Seven* will have us accept conclusions that are less conservative than Kyburg's.[18] Ultimately, my argument for rejecting Kyburg's answer to the present example will consist in a positive argument for my theory. Because the negative argument against Kyburg's approach to the Problem of Competing Statistics turns on a positive argument for mine, I will not present that argument here. However, there is another range of cases where Kyburg's theory gives the wrong answer. These cases correspond to a new incarnation of the Problem of Competing Statistics. As I

[18] It is worth noting that in the range of cases where the results of my theory differs from Kyburg's, my theory recommends bounds that are strictly narrower than the ones that Kyburg's theory recommends.

mentioned earlier, this variation of the problem was made perspicuous by Kyburg's systematic treatment of direct inference.

To its credit, evidential probabilities, on Kyburg's theory, satisfy a principle of complementation. In particular, *the evidential probability of S given Γ* is [p, q] if and only if *the evidential probability of not S given Γ* is [1−q, 1−p]. Despite this virtue, Kyburg's theory recommends probabilistically inconsistent evidential probabilities in a wide range of cases. For example, suppose that Γ consists of the following sentences:

$$\%x(\tau(x), \rho(x)) \in [0.9, 0.9], \rho(c), \tau(c)\equiv(p\vee q),$$
$$\%x(\tau'(x), \rho'(x)) \in [0.9, 0.9], \rho'(c'), \tau'(c')\equiv\neg p, \text{ and}$$
$$\%x(\tau''(x), \rho''(x)) \in [0.9, 0.9], \rho''(c''), \text{ and } \tau''(c'')\equiv\neg q.$$

In this case, Kyburg's theory yields the result that *the evidential probability of* $p\vee q$ *given* Γ is [0.9, 0.9], *the evidential probability of* ¬p *given* Γ is [0.9, 0.9], and *the evidential probability of* ¬q *given* Γ is [0.9, 0.9]. Unfortunately, the present results are probabilistically inconsistent.

Presumably, Kyburg would recommend that one suspend belief in cases such as the preceding example. This recommendation is akin to the recommendation made by Reichenbach regarding simple cases of the Problem of Competing Statistics. As in the case of Reichenbach's proposal, I think that the recommendation to suspend belief in situations such as the preceding one is too conservative. Moreover, the prescription to suspend belief in such cases will leave us without guidance in a large number of cases where we do possess a large amount of statistical data that bears on the likely truth value of a given proposition. As I have said before, we would do well to survey the possibilities for justifying less conservative probability judgments in the cases where the theories of Kyburg and Reichenbach provide no guidance. In *Chapter Seven*, I will propose a theory that is intended to provide such guidance.

ii. The Problem of Uninformative Statistics:

The Problem of Uninformative Statistics arises particularly for theories of direct inference that propose both that statistics for narrower reference classes are to be preferred, and that direct inference is based on frequency statements. As Kyburg's theory accepts both of these proposals, his theory is potentially subject to the problem. Kyburg's means of addressing the problem grows out of a feature of his theory which marks an obvious improvement over Reichenbach's theory.

One important way in which Kyburg's theory improved on Reichenbach's was in its explicit recognition and accommodation of the fact that we often lack knowledge of the precise value of given frequencies (or limiting frequencies).(2001, 200) Kyburg also grasped and explicitly stated a fact which now seems obvious. Namely, we always have some knowledge about the possible values of a given frequency, even if our knowledge consists only in the empty knowledge that the frequency is in the interval [0, 1].

In addition to marking a comparative virtue of his theory, Kyburg's proposal that statistical statements and the conclusions of direct inference are interval valued serves as the basis of his approach to the Problem of Uninformative Statistics. In effect, Kyburg maintains that statistics regarding a broader reference class are defeated by statistics regarding a narrower reference class *only if* one's statistics regarding the narrower reference class are more precise (or differ, in the sense of #) from one's statistics regarding the broader reference class.

Kyburg's proposal appears to thwart the Problem of Uninformative Statistics by converting (non-interval valued) frequency data that is uninformative into (interval valued) frequency data that does not play a role in direct inference. For example, the statement that a respective relative frequency is in the set {0, 1} is transformed into the statement that the respective relative frequency is in the interval [0, 1].

Despite the improvement over Reichenbach's theory in recognizing that our estimates of given frequencies may be imprecise, a notable limitation of Kyburg's theory is the presumption that only interval valued statistical statements may be used in direct inference. A more significant problem deriving from Kyburg's proposal was illustrated by an example produced by Mark Stone in (Stone 1987).

Stone's example was directed to refuting Kyburg's proposal that statistics for a narrower reference class, R, may defeat an instance of direct inference based on statistics for a broader reference class, R′, only if a respective agent's statistics for R are more precise than (or differ from) the agent's statistics for R′. In opposition to Kyburg's proposal, Stone's example (especially with some minor modifications) makes a compelling case for the conclusion that statistics for a narrower reference class can defeat instances of direct inference based on statistics for a broader reference class even when the statistics for the narrower class are less precise than the statistics for the broader reference class. In order to make Stone's original example more compelling, I have taken some liberties in modifying Stone's example.

First suppose that one is presented with following information, from a source which one knows is ultra-reliable:

(1) 51% of the balls held in urns produced by the *Ace Urn Company* are red.
(2) *b* is a ball held in an urn produced by the *Ace Urn Company*.

In the case where we lack further substantive information regarding b, typical theories of direct inference, including Kyburg's prescribe that we assign probability 0.51 to the proposition that b is red.

Now suppose that one has additional information regarding the specific urn, U_b, that contains *b*. In particular, suppose that one is able to inspect the contents of U_b under conditions that allow one to count the number of balls that are in U_b, and the number of balls in U_b that are white. As a result, one determines that U_b contains exactly one hundred balls, and exactly forty nine white balls. Since one has no conclusive reason to think that U_b does not contain some balls that are neither white nor red, one is able to form the justified belief that frequency of red balls among U_b lies in the interval [0, 0.51].

Faced with the present example, Kyburg's theory prescribes that our judgment regarding the probability that *b* is red should not change, and prescribes that we draw the conclusion that the probability that *b* is red is 0.51. However, given the results of our inspection of the contents of U_b it seems probable that *b* is not red. At the very least, we should not conclude that it is likely that *b* is red.

It turns out that the preceding example also creates problems for the two other major *orthodox* theories of direct inference that have been proposed to date (i.e., the theories of Pollock and Bacchus). In *Chapter Six*, I offer a formulation of the proposal that statistics for narrower reference classes are to be preferred that gives the right answer in the face of *the Ace Urn example*. The proposal made in *Chapter Six* finds a 'middle way' between *the doctrine* that statistics for narrower reference classes are always to be preferred and *the doctrine* that statistics for a narrower reference class may defeat an instance of direct inference based on statistics for a broader reference class *only if* a respective agent's statistics for the narrower reference class are more precise than the agent's statistics for the broader reference class. The former doctrine would prescribe the defeat of virtually all instances of direct inference, while the latter doctrine leads to the wrong answer in the face of the Ace Urn example.

iii. The Problem of Relevant Statistics:

The closure conditions that Kyburg proposes as restricting the set of admissible reference and target classes provide only a template for a solution to the Problem of Relevant Statistics. Indeed, in order for these closure conditions to restrict the sort of predicates that lead to problems, we must restrict the set of predicates that we allow as atomic predicates, lest we admit the problematic predicates as atomic components. It would seem, then, that Kyburg's proposed treatment of the problem is *at best* incomplete inasmuch as Kyburg provides no guidance in the selection of the atomic predicates.

Despite my low estimation of Kyburg's proposal, I do think that Kyburg's suggestion of adopting a syntactic approach to the problem is a good one. As we will see in *Chapter Six*, the designators for reference and target classes are intensional, and it seems that we must look at individual statistical statements on a case by case basis in order to determine whether a given statistical statement may appear in an admissible instance of direct inference.

7. Induction and Acceptance Theories

The second component of Kyburg's general theory of evidential probability is his theory of statistical induction and acceptance. Though not a part of his theory of direct inference *per se*, Kyburg's theory of acceptance is closely related inasmuch as Kyburg proposes that inductive inference reduces to a species of direct inference, in the presence of his acceptance rules.

In the end, I take exception to Kyburg's theory of acceptance. Nevertheless, I think that Kyburg's view that direct inference may serve as the foundation of statistical induction is correct. For this reason, I will proceed by considering the idea that direct inference may serve as the foundation of statistical induction. After this idea has been sketched, I will consider how it is that Kyburg's acceptance theory figures in the story, and explain why his acceptance theory is problematic.

In performing statistical induction, one proceeds from a premise describing the incidence of a given property among a subset of a given population to a conclusion describing the incidence of the property among the population as a whole. For example, from the premise that "10% of the American males *that have been observed* are doctors" one may, in appropriate circumstances, draw the conclusion that "approximately 10% of American males are doctors". Since the time of Hume, the question of how, and whether, such inferences are justified has been a central concern of epistemology and the philosophy of science.

The proposal that statistical induction may be grounded in direct inference has its roots in the writings of Donald Williams (1947) and was championed by David Stove (1986).[19] The basic idea appeals to certain mathematical facts regarding the frequency with which subsets of a set will resemble the set itself, in the frequency of any given property. There are numerous ways of expressing the relevant combinatorial facts, but the following theorem, which holds for arbitrary sets A and B, is sufficient to illustrate the basic idea:

[NN] For every $\delta,\gamma > 0$, there is an n such that if B is a finite set containing at least n members, then $\mathrm{freq}(\,\mathrm{freq}(A|X) \approx_{\delta} \mathrm{freq}(A|B) \mid X \subseteq B\,) > 1-\gamma$.[20](Pollock 1990)

The basic idea here is that for sets of sufficient size, the vast majority of the subsets of the set will be very similar to the set itself in the frequency of any given property. The present fact seems ripe for use in instances of direct inference which proceed from a premise about the frequency with which observed samples (subsets) are similar to the populations from they are drawn, and leads to the conclusion that our actual observed sample is likely to be similar to the population from which it was drawn in the frequency of any given property. Let us spell out how this line of reasoning will work, in a little more detail.

[19] See (McGrew, 2001) for a recent discussion.
[20] 'freq($A|B$)' designates the frequency with which elements of B that are elements of A, and $r \approx_{\delta} s$ expresses that r differs from s by no more than δ.

Let s be our observed sample and let Ω be the population about which we wish to draw a conclusion via statistical induction. Now, let τ be the set of subsets of Ω such that s is an element of τ *just in case* the frequency of elements of another set ϕ among the set is similar (to some specified degree) to the frequency of elements of ϕ among Ω, and let ρ be the set of subsets of Ω. In that case, we may express an instance of direct inference of the following form:

$$\begin{array}{l} \mathrm{freq}(\tau|\rho) \approx 1 \\ s \in \rho \\ \hline \mathrm{PROB}(s \in \tau) \approx 1 \end{array}$$

Now the preceding form of inference allows us to conclude that, with high probability, our observed sample will resemble the population from which it is drawn in the frequency of any given property. In turn, we may use deduction to draw a conclusion about the likely incidence of elements of ϕ among the population from which our sample was drawn. That is, from the premise $\mathrm{PROB}(s \in \tau) \approx 1$ and the premise $\mathrm{freq}(\phi|s) = r$, we may conclude that $\mathrm{PROB}(\mathrm{freq}(\phi|\Omega) \approx r) \approx 1$.

The catch to the present method of grounding statistical induction is that we can do no better than assigning a high probability to a statement about the incidence of a given property (represented by, the set, ϕ in the preceding example) among the population about which we are reasoning (Ω in the preceding example). This appears to present a problem if we wish to utilize the statistical claims so generated by direct inference, since direct inference itself is generally thought to depend on ‘straight-up’ statistical statements (i.e., statements that are not qualified as merely *highly probable*). As such, it appears that the statistical claims generated in the proposed manner could not themselves be used as premises for direct inference. For instance, following the present example, it looks as though the conclusion that PROB(freq($\phi|\Omega$) ≈ r) ≈ 1 cannot be used ‘as is’ in an instance of direct inference to draw the conclusion the probability that c (an element of Ω) is an element of ϕ is r.

Kyburg’s remedy to the present problem is to introduce an acceptance theory, which tells us that when the probability of a respective proposition is sufficiently high, we are entitled to treat the proposition as a practical certainty.

While Kyburg’s acceptance theory offers a means by which we may ‘process’ the results of statistical induction for use in direct inference, there are some serious problems with acceptance theories. In effect, acceptance theories allow one to move from a justified claim that a proposition is probable to full acceptance of the proposition. While some non-deductive adjustments in one’s confidence in propositions are occasionally licensed, the adjustments proposed by acceptance theories lead directly to problems, and appear to lack any epistemic justification.

Two problems that accompany Kyburg’s acceptance theory are symptomatic of the trouble with acceptance theories in general. The obvious effect of Kyburg’s acceptance theory is that it leads to an over estimation of the likelihood of likely (yet uncertain) propositions. This is a simple restatement of exactly what Kyburg's theory does. But that the problem is transparent should not be taken as reason to tolerate the problem. Moreover, a slightly less obvious problem is the probabilistic inconsistency that ensues in the wake of such adjustments. For example, if one adjusts one's confidence in a statement whose probability is 0.99 to regard the proposition a practical certainty, the result is a mutual inconsistency in one's confidence in the proposition, and the probability that one assigns to its negation, namely 0.01.

It is possible to remedy the problem of acceptance theories yielding inconsistent probability assignments, by introducing rules which simply recalibrate one's probabilities, so that they remain consistent with the results of one's acceptance rules. Although the problem of inconsistency can be remedied, it remains that acceptance theories *unjustifiably* license an extensive recalibration of one's probabilities.

Kyburg's proposed acceptance rules play a role in grounding statistical induction, and his rules thereby play a practical role within Kyburg's larger theory of rational belief formation. However, practical considerations, in themselves, cannot be used to justify epistemic principles. More importantly, if there is some way to dispense with acceptance rules in grounding statistical induction, we should take that route in grounding statistical induction. In *Chapter Six*, I show that it is not necessary to introduce an acceptance theory as a component of a theory of rational belief that grounds statistical induction in direct inference.

8. Conclusions

Kyburg's theory of direct inference represents an improvement over Reichenbach's. The theory improves on Reichenbach's by accommodating the possibility that one has frequency data that is not point-valued, and by allowing one to form conclusions via direct inference in a wide range of cases where Reichenbach's theory recommends that one suspend judgment. Despite improvement in these areas, Kyburg's the theory yields the wrong answers in a range of cases where less precise statistics for a narrower reference class should defeat an instance of direct inference based on a broader reference class. Kyburg's theory also recommends that one suspend belief in a wide range of cases where it seems possible that we may form reasonable conclusions. In addition, Kyburg's proposal for dealing with the Problem of Relevant Statistics is too programmatic. In *Chapter Six*, I will show that Kyburg's approach to the Problem of Relevant Statistics is misguided, and offer a new approach to the problem.

CHAPTER 3: POLLOCK'S THEORY OF DIRECT INFERENCE

1. Introduction

As described in his monograph *Nomic Probability and the Foundations of Induction*, Pollock's theory of direct inference consists of one component of an extensive theory of probabilistic inference and rational belief formation. The extensive theory, of which his theory of direct inference is a part, is intended as an analysis of what Pollock calls "nomic probability", and not merely as description of the core principles rational belief formation.

Pollock's ambitious project is motivated in part by the failure of past attempts to provide a *definition* of the central notions of probability. In the place of such an approach to the analysis of probability, Pollock proposes to describe *the proper inferences licensed by statements of probability* and *the proper inferences that license acceptance of statements of probability*. By so specifying the *inferential role* of probability statements, Pollock intends to provide an informative analysis of the notion of probability.

Pollock describes his comprehensive theory as having four components, though for the purposes of describing the theory, it is best to regard the theory as having five. The five components are:

1. A theory of the inferential relations that hold between statements of nomic probability.
2. A Theory of Acceptance.

3. A Theory of the Dynamics of Defeasible Reasons.
4. A Theory of Direct Inference.
5. A Theory of Statistical Induction.

The first three components listed above form the foundations of Pollock's comprehensive theory of probabilistic inference and rational belief formation, and remaining two components are derived from the founding three.

My interest in this chapter is to evaluate Pollock's theory direct inference. However, there are many interconnections between Pollock's theory direct inference and the components of Pollock's overarching theory. In particular, Pollock's theory of direct inference is derived from the founding components, and the founding components are meant to serve as the justificatory basis of direct inference. So before considering the details of Pollock's theory of direct inference, it will be useful to survey some of the key features of the three components upon which the theory is founded.[21]

For the most part, I will not accept the founding components of Pollock's theory. But my rejection of these components is not tantamount to a rejection of Pollock's theory of direct inference, since the features of that theory can (and will be) evaluated on their own merits. Moreover, while I am strongly opposed to certain of features of the foundations of Pollock's theory, I am largely in agreement with Pollock about which features will be central to a correct theory of direct inference. Nevertheless, several of my objections to

[21] Note, however, that the justificatory relation proposed by Pollock is intended more as an explanation of the source of the justification of the principles of direct inference. So an attack on the 'justificatory basis' of the principles would be inconclusive as a refutation of Pollock's principles of direct inference.

the various features of the foundations of Pollock's theory of direct inference reflect the ways in which I think that Pollock's theory of direct inference should be modified.

2. Two Kinds of Probability

Before going on to discuss some of the specifics of Pollock's views on nomic probability and direct inference, it is useful to consider Pollock's distinction between *definite* and *indefinite probability*.

For the most part, statements of probability are understood to assert the probability that a given proposition is true. Such statements are commonly about a particular individual. Pollock uses the expression "definite probability" to describe such statements. On the other hand, some statements of probability relate (or seem to relate) sets of objects, or the respective properties that fix the membership of such sets. For example, the statement "the probability of a 50 year old male living to be 75 is 0.7." is not intended to be about a particular individual. Rather, as Pollock points out, statements such as the preceding relate the property of *being a 50 year old male* to the property of *living to be 75 years old*. Pollock uses the expression "indefinite probability" to describe such statements.

Pollock identifies *nomic probability* as the central species of *indefinite probability*, and uses the operator "prob" to express statements of nomic probability, so that the expression “prob(X|Y) = r” *states* that the nomic probability of the property X relative to the property Y is r. On the other hand, Pollock uses the operator “PROB” to express statements of *definite probability*, so that the expression “PROB(A) = r” *states* that the probability of the proposition A is r. In the end, Pollock proposes that definite probabilities are to be defined in terms of (indefinite) nomic probabilities. For the moment, let us consider Pollock’s ideas about nomic probability.

3. Nomic Probability

Although Pollock intends his comprehensive theory to serve as a conceptual role analysis of the notion of nomic probability, he begins his analysis with a sketch of the notion of *nomic probability* that will serve as the core of his broad analysis. In the course of Pollock’s analysis, the initial sketch of the notion of *nomic probability* serves to prime our intuitions for what is to come.

On Pollock’s account, statements of nomic probability are to be understood as descriptions of statistical laws of nature. For the most part, laws of nature have been understood to be universal generalizations that hold in all physically possible worlds. Statements of nomic probability are intended to generalize such ‘universal’ nomic generalizations, and permit for law-like statements describing statistical patterns in nature.

Given Pollock's proposal that nomic probabilities describe statistical laws of nature, the question of the correct analysis of nomic probabilities comes to the fore. In order to express the core content of the notion of nomic probability, Pollock uses expressions of the form '$\wp(X|Y) = r$' to symbolize statements of the form 'the proportion of the members of *Y* that are members of *X* is r'. In turn, Pollock defines prob(*F*|*G*) to be $\wp(\mathbf{F}|\mathbf{G})$, where **F** and **G** are the set of physically possible *F*s and *G*s, respectively.[22]

In the account of direct inference offered by Kyburg, frequency statements served as the major premises of instances of direct inference. Unlike such 'frequency-based' accounts of direct inference, Pollock's theory employs a proportion function, $\wp$, that is applicable to populations of infinitely large size, since for the most part the set of physically possible objects possessing a given property will be infinitely large. For fairly obvious reasons, the notion of *proportion* as applied to infinitely large sets diverges from the notion of *frequency* (though Pollock rightly requires that these measures be equivalent in the case of finite sets). Extending the notion of *proportion* to apply to infinitely large sets leads to some difficult questions regarding the character of proportions. Pollock approaches the problem by proposing a series of principles that he regards as intuitive, and which are sufficient to permit derivation of his theory of direct inference and his theory of statistical induction, in the presence of his theory of acceptance.

[22] The description of Pollock's view given here only covers those cases where there is a physically possible *G*. In order deal with cases where there is no physically possible *G*, Pollock's theory is slightly more complicated than I have described here. It seems clear to me that Pollock's analysis will work for cases where there is no physically possible *G*, if it works for the cases where there is a physically possible *G*. For this reason, it seems reasonable to focus only on the cases where there is a physically possible *G*.

My difficulty in accepting Pollock's analysis of the notion of *nomic probability* is connected with my inability to form clear intuitions about the structure of the set of physically possible worlds, or about the connection between statistical laws of nature and the size of the set of worlds that have certain sorts of property distributions. To get the feel for my concern, suppose that there is a natural law that militates against there being a large number of x-particles in any physically possible world, though there is no bound on the number of x-particles in a physically possible world. Given the present suppositions, it seems that the nomic probability of objects being x-particles should be low. Despite that intuition, I lack a clear understanding of how the structure of the set of physically possible worlds will bear on the cardinality of worlds with certain sorts of x-particle distributions. In turn, it is unclear to me that the set of statistical laws of nature regarding x-particles will entail the right nomic probabilities, if nomic probabilities are defined in the way that Pollock proposes. Put another way: I do not see that the 'natural likelihood' of certain sorts of worlds has a bearing on the 'quantity' of nomically possible worlds having respective property distributions.

Waiving the preceding concern, which may be regarded as somewhat academic, it is unclear to me that nomic probabilities (understood in the sense proposed by Pollock) are appropriate to support direct inference in a manner similar to frequencies. In particular, I have a difficulty in seeing a relevant connection between the numerousity of the set of physically possible objects possessed of a given property and the strength of one's grounds for believing that any given actual object will possess the property (unless one supposes, as I am not inclined to do, that all physically possible worlds are equally probable).

The target of the preceding objection is Pollock's claim to have provided a satisfactory characterization of the notion of a statistical law, and Pollock's claim to have described a relation between properties which could play a role in direct inference that is comparable to the role of frequencies. My conclusion is simply that I am not convinced. Even so, many of the characteristics that Pollock proposes as governing nomic probabilities are very intuitive as principles governing statistical laws. Similarly, many of the principles that he proposes as governing the relationship between frequencies and nomic probabilities go a long way in establishing the epistemic significance of the notion of *nomic probability* as it is formalized by Pollock. Let us now consider some of these principles.

4. Pollock's Principles of Nomic Probability

For the most part, the principles that Pollock proposes as governing inference between statements of nomic probability are very intuitive. Many of the intuitions that Pollock enlists are derivative of our intuitions regarding finite populations, where the notion of relative proportion may be identified with the notion relative frequency. Some of the principles that Pollock proposes are also instrumental to the attempt to establish that statements of nomic probability are fit to serve as the statistical premises in instances of direct inference. For example, the principle that Pollock calls "PFREQ" seems reasonable, and helps so that the notion of *nomic probability* may play the important epistemic role that Pollock proposes, by connecting nomic probabilities with frequencies:

> (PFREQ) If r is a nomically rigid designator of a real number and there is a physically possible world in which an object is an element of G and freq($F|G$) = r, then prob(F | G & freq($F|G$) = r) = r.(1990, 70)[23]

Basically, the present principle tells us that the nomic probability of *F*s among *G*s that exist in worlds in which the frequency of F among G is r is r. The principle is intuitively compelling. Moreover, the principle plays an extremely important role in vindicating the idea that frequency values can be used in direct inference to determine one's definite probabilities.

Despite the reasonableness of many of the principles that Pollock proposes, some of the central principles are open to debate, and I find at least one of the principles to be unacceptable. In particular, I find the principle that Pollock calls "the Principle of Agreement for Proportions" to be highly suspect.

> If B is infinite and $\wp(A|B) = p$,
> then for every $\delta>0$: $\wp(\ \wp(A|X) \approx_{\delta} p \mid X \subseteq B\ ) = 1$.(1990, 71)

[23] In Pollock discussion, properties are defined as sets of ordered pairs of object and worlds, such that a pair, $\langle c, w\rangle$, is a member of the set corresponding to a property, A, *if and only if* c is an A in w. In turn, G & freq($F|G$) = r consists of the property of being a G in a world in which freq($F|G$) = r.—The set of pairs in this case is composed of the set of pairs, $\langle c, w\rangle$, such that c is an G in w, and freq($F|G$) = r in w.

In words, the Principle of Agreement for Proportions *states* that, for any sets *A* and *B*, if *B* is a set of infinite size and the proportion of elements of *B* that are elements of *A* is *p*, then the proportion of subsets of *B* that agree with *B* (within a margin δ) on the proportion of its elements that are elements of *A* is 1. As it turns out, the present principle is central to Pollock's justification of direct inference and his justification of statistical induction. As such, the veracity of the principle is of fundamental importance for Pollock's extensive theory. In attempting to justify the Principle of Agreement for Proportions, Pollock appeals to a theorem of mathematics whose content will be familiar to statisticians:

> For every $\delta,\gamma > 0$, there is an *n* such that if *B* is a finite set containing at least *n* members, then $\text{freq}(\ \text{freq}(A|X) \approx_{\delta} \text{freq}(A|B) \mid X \subseteq B\ ) > 1-\gamma$.(1990, 71)

Citing this theorem, Pollock states that "It seems inescapable that when *B* becomes infinite, the proportion of subsets agreeing with *B* to any given degree of approximation should become 1."(1990, 71) Though Pollock admits that the Principle of Agreement for Proportions cannot be proven, he is confident that it is true.

Despite Pollock's confidence in the Principle of Agreement for Proportions, it is unclear whether the notion of *proportion* is sufficiently characterized by Pollock so that there is a fact of the matter as to whether or not the principle is true. It also seems to me that the value of the proportion of one set, A, among another set, B, should be *one* only in the case where each element of A is an element of B. More importantly, the Principle of Agreement for Proportions yields unattractive consequences given Pollock's proposal that it is statements of nomic probability (defined in terms of proportions) that are to serve as the statistical premises to instances of direct inference. The unattractive consequences of the principle are associated with the fact that statements of the form '$\wp(A|B) = 1$' do not in general imply the truth of the universal generalizations of the form '$(\forall x)(x\in B \supset x\in A)$'. One consequence of this is that there are cases where Pollock's theory of direct inference allows one to conclude that the probability that a given object, c, has a property, A, is *one* (based on the premises that c is *B* and $\wp(\mathbf{A}|\mathbf{B}) = 1$), when the agent cannot rule out the possibility that c is not-A (since it is epistemically possible that $\exists x$: $x\in B \wedge x\in A$, and epistemically possible that $c\notin A$). Faced with this kind of case, Pollock must reject certain prescriptions of decision theory associated with acting upon statements whose probability is *one*, or formulate a reason by which we may not conclude that the probability of proposition is *one*, in cases where we cannot rule out the truth of the proposition's negation. It strikes me that the appropriate response to the present dilemma is to adopt a notion of proportion so that there is a general implication between unit-valued proportion statements and corresponding universal generalizations. In that case, we must reject the Principle of Agreement for Proportions.

I do not regard my criticism of the principles of nomic probability proposed by Pollock as decisive. However, I think that I have shown that the Pollock's theory has intuitive liabilities. Later, in *Chapter Six*, I will defend a different proposal as to the sort of statistical statements that underlie direct inference. In the face of that alternative, any of the liabilities of Pollock's approach loom large.

5. A General Theory of Defeasible Reasoning

One may say that an agent, S, has a reason for accepting a proposition, p, if S is apprized of facts which, considered in isolation, recommend belief in p. In turn, one may say that such a reason is *defeasible,* if a wider consideration of the facts could lead to the defeat of S's grounds for believing p. Finally, in cases were one accepts a statement, q, that decisively undermines one's defeasible reason(s) for another statement, p, we say that *q is defeater for p* and that *p is defeated*.

One of Pollock's important contributions in the area of defeasible reasoning is his distinction between *rebutting* and *undercutting* defeaters. Basically, a rebutting defeater for a reason *x* for a proposition *p* is a reason for denying *p* in the face of *x*. An undercutting defeater, on the other hand, is a reason for denying that *x* is a reason for *p* (in the respective context). Pollock's characteristic example of an undercutting defeater concerns the manner in which an object appearing a certain color gives one a defeasible reason for believing that it is that color. For example, that an object appears red gives one a defeasible reason for believing that it is red. Now suppose that an object appears red to us, but we happen to know that the object is being illuminated by red-light (and objects so illuminated appear red regardless of their actual color). In that case, our normal reason for believing that the object is red is defeated. Yet in such cases we do not have a reason for denying that the object is red, and so our reason for believing the object is red is *undercut*, but not *rebutted*.

For the most part, the body of potential undercutting defeaters that may arise in any given situation will be determined by topic-specific considerations. So, for example, it appears that there is a special sort of undercutting defeater involving a preference for narrower reference classes in the case of direct inference. On the other hand, rebutting defeaters may arise in any context in which one may have reasons for mutually inconsistent propositions. There are also cases where the set of potential defeat relations among a set of statements (relative to an agent) exhibits a complex structure, such as when a reason for one statement, *p*, is a defeater for a reason for a statement, *q*, which is in turn a defeater for a reason for a statement, *r*. In cases such as the one just described, the agent should, *ceteris paribus*, accept *r* (in the case where the agent has a defeasible reason for *r*). The present example illustrates the manner in which the set of potential defeat relations among a set of statements (relative to an agent) may exhibit a complex structure. (It is possible to construct examples of far greater complexity.)

Pollock is the only scholar who has offered a theory of defeasible reasoning that is applicable to large range of cases where the set of potential defeat relations among a set of statements exhibits a structure of arbitrarily high complexity. Pollock's theory is commendable for the generality of its applicability. And an immediate virtue of his theory lies in its applicability to a range of cases that arise in the course of direct inference that Kyburg's theory failed to address.

Recall that there are cases where Kyburg's theory yields probabilistically incoherent results. Pollock's theory, on the other hand, has failsafe principles that will prevent any theory of defeasible reasoning from yielding such results. The core of this failsafe is Pollock's *Principle of Collective Defeat*:

> If we are justified in believing R and we have equally good independent prima facie reasons for each member of a minimal set of propositions deductively inconsistent with R, and none of these prima facie reasons is defeated in any other way, then none of the propositions in the set is justified on the basis of these prima facie reasons.(1990, 81)

While the Principle of Collective Defeat nicely deals with the cases that presented problems for Kyburg's theory, there are many cases in which the Principle of Collective Defeat appears to recommend a course that is too incredulous. Obvious examples, concern cases where one has perfectly symmetric defeasible reasons regarding the probability of a proposition and its negation, and where the corresponding assignment of probabilities is probabilistically inconsistent. For example, consider a case where one has a defeasible reason for believing that the probability of a proposition, *p*, is 0.8, and one has a defeasible reason for believing that the probability of a proposition, *not-p*, is 0.8.—One may imagine that these reasons were generated by competing instances of direct inference. In such cases, Pollock's theory, backed by the Principle of Collective Defeat, makes no recommendation regarding the probability that we should assign to *p* or to *not p*. Yet it seems reasonable to assign probability 0.5 to both propositions.

In effect, the Principle of Collective Defeat plays the role of 'cleaning up' any inconsistencies among the set of propositions for which one has defeasible reasons. While the principle may appear heavy handed inasmuch as it is insensitive to detectable symmetries among one's reasons, it is possible to augment one's principles of defeasible inference by auxiliary principles whose application is sensitive to such symmetries. Indeed, while the preceding example suggests that the Principle of Collective Defeat recommends a course that is too incredulous, another proposal would be that Pollock's theory is incomplete as it stands, and the theory must be augmented to include a principle of belief formation that applies in cases which are not properly handled by the theory as it is currently formulated. As such, a remedy to the incredulity prescribed by Pollock's theory does not require the rejection of the Principle Collective Defeat, but only the addition of auxiliary principles. I will propose principles of this sort in *Chapter Seven*.

6. Pollock's Acceptance Rules

We have already considered some of the features of Pollock's theory of defeasible reasoning, and seen how the theory deals with some difficulties that arose for Kyburg's theory. We have also seen that the seeming incredulity associated with the Principle of Collective Defeat could be mitigated by the addition of auxiliary rules. I now wish to consider the basic acceptance rules that Pollock proposed as a component of his general theory of rational belief formation (and theory of nomic probability).

Normally, acceptance rules describe the conditions under which the assignment of a high probability to a proposition permits one to adopt full-belief in the proposition. Pollock generalizes this conception of acceptance rules, and proposes a rule that allows one to arrive at a 'full belief' in a proposition given an assignment of a sufficiently high nomic probability regarding a pair of appropriately related properties. The most basic acceptance rule that Pollock proposes is as follows:

> (A1) If F is projectable with respect to G and $r > 0.5$, then 'Gc & prob($F|G$) $\geq r$' is a defeasible reason for believing that '*Fc*', the strength of the reason depending on the value of r.(1990, 85)

Before assessing *A1*, it is helpful to carefully reflect on what the principle expresses. First, note that the projectability constraint built into the antecedent of *A1* is intended to circumvent a variety of difficulties, including the ones that arise for direct inference (the Problem of Relevant Statistics) and for statistical induction (Goodman's New Riddle of Induction). For the moment, I will postpone discussion of this feature of *A1*. Having waived the projectability issue, we see that *A1* is a sort of rule of direct inference.[24] From a functional perspective, *A1* combines a traditional principle of direct inference with an acceptance rule of the sort proposed by Kyburg, though it is noteworthy that the threshold

[24] Pollock observes a strict distinction between direct inference, which involves definite probabilities, and statistical syllogism, which do not involve definite probabilities. For this reason, Pollock will deny that *A1* should be described as incorporating direct inference. Nevertheless, I am using the expression "direct inference" in the broad sense (typical in the writings of Levi and others) which includes statistical syllogism as a species.

for acceptance for Pollock's principle is extremely low at 0.5.[25]—Presumably, considerations of reason strength are meant to mitigate the low 'acceptance threshold'.

Like Kyburg's acceptance rule, Pollock's leads to an over estimation of the likelihood of likely (yet uncertain) propositions, since Pollock's acceptance rule effectively calls for an agent to recalibrate his or her probabilities for likely propositions. This fact is illustrated by considering, in parallel, related instances of direct inference and applications of *A1*. For example, suppose that one accepts that Gc and prob($F|G$) = 0.9. In that case, direct inference will license the conclusion that PROB(*F*c) = 0.9, while *A1* will license the conclusion that *F*c. Assuming that acceptance of *F*c licenses acceptance of the conclusion that PROB(*F*c) = 1, Pollock's acceptance rule, issues conclusions that are inconsistent with the prior judgments of the agent (given by direct inference).[26]

With the help of the Principle of Collective Defeat, I presume that Pollock's more extensive theory will call for a widespread recalibration of one's probabilities, so that they remain consistent with the results of *A1*. Although this would resolve any inconsistency, the resolution would come through a widespread recalibration of an agent's probabilities. Such a recalibration would be questionable.

[25] Viewed as a black box, which takes premises as inputs and yields conclusions as outputs, *A1* does the same work as kyburgian principles of direct inference and kyburgian acceptance rules in one step.
[26] If acceptance of *F*c does not license acceptance of the conclusion that PROB(*F*c) = 1, then it becomes difficult understand the cognitive significance of statements regarding of 'full-belief', and the role of principles such as *A1* becomes suspect.

Pollock's acceptance rules are appealed to in the derivation of his theory of direct inference and his theory of statistical induction. The derivation of direct inference from an acceptance rule is unique to Pollock, but the appeal to acceptance rules in the justification of statistical induction is not new.—Recall that Kyburg makes such an appeal. In fact, acceptance rules might be regarded as practically indispensable in providing a justification of statistical induction in the manner of Kyburg. As it turns out, however, acceptance rules are not necessary in grounding statistical induction in direct inference, as proposed by Kyburg. I will explain this point in *Chapter Six*, as it is connected with a wider discussion of the Problem of Uninformative Statistics, and with my proposal for the sort of statements that may serve as the statistical premises of direct inference. For now, I would like to register my concern for acceptance theories in general, and emphasize, again, that the perceived utility of a certain form of inference cannot be regarded as providing it with an epistemic justification.

7. Pollock's Non-Classical Theory of Direct Inference

Pollock is outspoken in crediting Kyburg's work as an influence on his ideas about direct inference. Similarly, variants of the principles of direct inference proposed by Reichenbach are set out as consequences of Pollock's theory. Pollock describes these principles as "Reichenbach's Rules", and sets himself the goal that his theory of direct inference confirm these rules. Pollock formulates Reichenbach's Rules as follows, where

'W(α)' is written to express the proposition that a respective agent is warranted in believing α:[27]

(R1) 'prob($F|G$) = r & W(Gc)' is a prima facie reason for 'PROB(Fc) = r'.

(R2) '$\forall(H \supset G)$ & W(Hc) & prob($F|H$) $\neq r$' is an undercutting defeater for (R1).

Pollock's *non-classical* theory of direct inference may rightly be regarded as a generalization of the *classical* theories of direct inference proposed by Reichenbach and Kyburg. Though the basic formal principles of direct inference proposed by Pollock are new, these principles imply limiting principles that are largely in accord with the theories defended by Reichenbach and Kyburg.

Pollock's motivation for articulating non-classical principles of direct inference begins with the idea that instances of classical direct inference from premises of the form $prob(F|G) = r$ and Gc to a conclusion of the form $PROB(Fc) = r$, presuppose for all H, such that Hc, that prob($F|G\&H$) = r. Pollock's maintains that if such an assumption were not made, then we would have to accept that it was virtually certain that prob($F|G\&H$) $\neq r$, and hence that respective instances of direct inference were *undercut* (by a principle such as (R2)). In line with the proposal that we must assume, for all such H, that prob($F|G\&H$) = r, Pollock formulates the foundational principle of direct inference as follows:

[27] Pollock's definition of warrant depends on a prior notion of *objective justification*. According to Pollock, "a person is said to be *objectively justified* in believing *P* iff he is justified in believing *P* and his justification could not be defeated 'in the long run' by any amount of reasoning proceeding entirely from other propositions he is justified in believing."(1990, 87) In turn, a person, S, is said to be *warranted* in believing *P* if and only if "*S* would become objectively justified in believing *P* through reasoning proceeding exclusively from the propositions he is objectively justified in believing."(1990, 87)

(D1) If F is projectable with respect to G then 'H is a *sub-property* of G & prob($F|G$) = r' is a prima facie reason for 'prob($F|H$) = r'.(1990, 128)[28]

In turn, the defeat condition corresponding to (R2) is as follows:

(SD) If F is projectable with respect to J then 'H is a *sub-property* of J & J is a *sub-property* of G & prob($F|J$) $\neq$ prob($F|G$)' is an undercutting defeater for (D1).(1990, 131)

Pollock's derivation of *classical direct inference* from *non-classical direct inference* depends on a definition of *definite probabilities* in terms of *nomic probabilities*. The proposed definition identifies an agent's definite probability for a singular proposition, Fc, with the indefinite probability of F conditional the property of being c under conditions **K**, where **K** is the conjunction of the set of propositions that the agent is warranted in accepting. The exact property conditioned on is the property $x = c$ & **K**:

(CDI) PROB(Fc) = prob($Fx \mid x = c$ & **K**).(1990, 133)

Pollock's non-classical theory of direct inference functions by allowing one to reason from a premise regarding the value of an indefinite probability relative to a given property to conclusions about indefinite probability relative to sub-properties.—This reasoning is encoded in (D1). (SD) plays the role is requiring that the narrowest sub-property take precedence in determining respective indefinite probabilities via direct inference.—In this way, (SD) determines that Pollock's *non-classical approach* to direct inference is nonetheless *an orthodox approach* to direct inference.

[28] In general, we say that a property H is a sub-property of a property G just in case there is a physically possible H and all physically possible Hs are Gs.

In practice, the value of the indefinite probability prob(*Fx* | *x* = *c* & **K**) (which will be used to set respective definite probabilities) will be identified with the indefinite probability, prob(*F* | *H*), where *H* corresponds to the narrowest sub-property for which we have reliable information regarding the value of prob(*F* | *H*). Pollock's approach thereby produces results that parallel classical approaches to direct inference, such as the approaches of Reichenbach and Kyburg. The main difference is in the mechanism whereby the different approaches yield 'the classical results'. Another difference is in the large number of additional inferences about the values of various indefinite probabilities supported by Pollock's theory.

In proposing that we adopt a non-classical theory of direct inference, Pollock argues that classical direct inference already presupposes the correctness of non-classical direct inference. However, it seems to me that the advocate of a classical approach could formulate his/her theory so that Pollock's argument is not compelling. In this vein, note that in Pollock's formulation of (R2), the defeat condition is satisfied *only if* a respective agent is warranted is believing that prob(*F*|*H*) ≠ *r*. If it is reasonable to substitute the narrower condition that there be an s such that the agent be warranted in believing that prob(*F*|*H*) = *s* & *s* ≠ *r*, then we need no longer accept Pollock's claim that instances of classical direct inference from premises of the form *prob(F|G) = r* and *Gc* to a conclusion of the form *PROB(Fc) = r*, presuppose for all *H*, such that *Hc*, that prob(*F*|*G&H*) = *r*. Despite this response to Pollock's argument, it may well be that the sort of inferences licensed by non-classical direct inference are correct when applied to the right sort of statistical probabilities. Having made that concession, I must now say that there is a considerable tension between the Pollock's principles of non-classical direct inference, and his proposal that it is statements of nomic probability that are to serve as the statistical premises of instances of direct inference.

As a first step in illustrating why nomic probabilities, as formulated by Pollock, are inappropriate for use in non-classical direct inference, consider the proposal to use relative frequencies in the case of non-classical direct inference. In the case of relative frequencies, it looks as thought it is usually reasonable to make an inference from premises of the form $freq(F|G) = r \ \& \ c \in G$ to a conclusion of the reach the conclusion that probability that c is an element of F is *r*. However, non-classical direct inference does not make sense for frequencies. That is, it will normally be problematic to reason from a premise of the form $freq(F|G) = r$ to a conclusion of the form $freq(F|H) = r$, where *H* is a proper subset of *G*. Such inferences are deeply problematic, as they involve reasoning in a manner that is far too credulous. The information contained in the conclusions of such inferences is far too precise in their assignment of values to relative frequencies, inasmuch as relative frequencies describe precise quantitative features of the world. For this reason, non-classical direct inference for frequencies is unreasonable.

Although nomic probabilities differ from frequencies, they are analogous to frequencies in the manner in which they describe precise quantitative features of the set of physically possible objects. And, for the same reason that non-classical direct inference is unreasonable for frequencies, it appears that non-classical direct inference for nomic probabilities is unreasonable. The feature of nomic probabilities that seems to obscure the similarity between nomic probabilities and frequencies is the tendency of nomic probabilities to apply to properties in cases where the set of nomically possible objects satisfying the property is infinitely large. However, while typical properties tend to encompass an infinitely large set of objects, not all properties have as their 'trans-world extension' an infinitely large set. When we consider that there are cases where the trans-world extension of a property is finite, we see that non-classical direct inference using nomic probabilities is problematic in precisely the same manner as non-classical direct inference using frequencies. For example, for any property that one may consider, one may designate a sub-property consisting only of the actual world objects that possess the property, and in many cases, the set of objects falling in the extension of this sub-property will be finite.[29] In those cases where the set of objects in the actual world having a property is finite, the values of respective statements of frequency and nomic probability are identical. And, in those cases, it is clear that non-classical direct inference for statements of nomic probability is unreasonable in precisely the same manner as for statements of frequency.

[29] Pollock proposes to classify such properties as non-projectable. This proposal seems to me to be *ad hoc*, so I am pursing my objection.

Having surveyed the main aspects of Pollock's theory, let us now consider the manner in which Pollock's theory addresses the three problems of direct inference, and the degree to which Pollock's treatment of the three problems is satisfactory.

8. The Three Problems of Direct Inference

i. The Problem of Uninformative Statistics:

The Problem of Uninformative Statistics arose as an issue for theories of direct inference that propose that statistics regarding narrower reference classes are to be preferred and that direct inference is based on frequency statements. Such theories face the problem of explaining how we may avoid using frequency data regarding the unit set reference class containing only the object about which one wishes to draw a conclusion. The problem with unit set reference classes is that the frequency of elements of a respective target class among such reference classes is always *one* or *zero*, and inference based on such reference classes would seem only to allow the conclusion that respective probabilities are one or zero.

In the face of the present problem, Pollock proposes that the statistical statements that may serve as the statistical premises to instances of direct inference are not frequency statements at all, and that frequency statements *cannot* serve *directly* to defeat instances of direct inference. According to Pollock's account, it is statements of nomic probability that serve as the major premises to instances of direct inference, and not statements of frequency.

Despite his position that it is statements of nomic probability that play the key role in direct inference, Pollock's theory also provides for an explanation of the manner in which frequency data is relevant to direct inference. This explanation takes the form of a foundational principle of nomic probability:

> (PFREQ) If *r* is a nomically rigid designator of a real number and there is a physically possible world in which an object is *G* and freq(*F*|*G*) = *r*, then
> prob(*F* | *G* & freq(*F*|*G*) = r) = r.(1990, 70)

In the case of relative frequencies, the Problem of Uninformative Statistics occurs because we are always entitled to conclude that the frequency with which any given object is an element of a respective target class is *one* or *zero*. On the other hand, it seems that we are not entitled to conclude, for arbitrary *F* and *c*, that prob(*Fx* | *x* = *c* & **K**) is *one* or *zero*.— Simply, prob(*Fx* | *x* = *c* & **K**) may take on values besides *one* or *zero*, since c may have different properties in different nomically possible worlds.

While it seems that Pollock's theory provides a satisfactory remedy to the Problem of Uninformative Statistics, Pollock's approach to the problem leads to the same difficulties as Kyburg's theory in the face of variants of the example presented by Mark Stone in (Stone 1987).

The Ace Urn example (from *Chapter Two*) shows that statistics for a narrower reference class *can* defeat instances of direct inference based on statistics for a broader reference class even when the conclusion of direct inference based on statistics for a broader reference class is consistent with the conclusion of direct inference based on statistics for a narrower reference class. Since Pollock's theory yields the same (incorrect) probabilities as Kyburg's in the face of the Ace Urn example, it seems that Pollock's approach to the Problem of Uninformative Statistics is not quite adequate. I will return to address the problem in *Chapter Six*.

ii. The Problem of Relevant Statistics:

Pollock characterizes the present problem as one involving the projectability and non-projectability of respective properties. And like Kyburg, Pollock likens the problem to the problem of gerrymandered properties that arises for statistical induction (i.e., Goodman's New Riddle of Induction).

Although Pollock does not provide a worked out solution to the problem, he does make a number of substantive statements about its nature. Pollock's substantive claims include the proposal that projectability is a two place relation between properties. Pollock also argues for and against various closure conditions on the set of projectable properties. In particular, Pollock maintains that the following conditions hold:

> If A and B are projectable with respect to C, then (A&B) is projectable with respect to C.
>
> If A is projectable with respect to both B and C, then A is projectable with respect to (B&C).

In turn, Pollock argues that the following conditions do not hold:

> If A and B are projectable with respect to C, then (A$\vee$B) is projectable with respect to C.
>
> If A is projectable with respect to both B and C, then A is projectable with respect to (B$\vee$C).

The most interesting feature of Pollock's account of the 'projectability' problem for direct inference is his explanation of the connection between that problem and the 'projectability' problem that arises in the case of induction. The connection, according to Pollock, derives from the existence of a projectability constraint upon acceptance rules, and from the derivation of the rules of direct inference and statistical induction from the same founding acceptance rules.

Although I have rejected Pollock's acceptance rules, and in turn his proposed justification for direct inference and statistical induction, I think that his idea about the existence of a connection between the two projectability problems is correct. But rather than thinking that the connection between the two problems derives from an underlying acceptance rule, I think that the connection traces to the derivation of statistical induction from direct inference. That is, the projectability problem that arises for statistical induction is derivative from the projectability problem that arises for direct inference. Due to this connection, I am optimistic that a remedy to the problem as it arises for direct inference will entail a remedy for the problem as it arises in the course of statistical induction. The present idea is not far from Pollock's proposal, when one considers that fact that Pollock's acceptance rule may itself be regarded as a principle of direct inference.

iii. The Problem of Competing Statistics:

Paradigm cases of the Problem of Competing Statistics arises when one has reliable statistics regarding the incidence of a particular property among two overlapping references classes, and one does not have reliable statistics regarding the incidence of the property among the intersection of the two sets. In cases where Pollock's principles of direct inference yield inconsistent conclusions, his Principle of Collective Defeat comes into play, yielding the final result that both conclusions are defeated.

The example described in section 5 was intended to illustrate one sort of case where Pollock's theory, like Kyburg's, is too conservative in the conclusions that it prescribes in cases where direct inference yields reasons for accepting mutually inconsistent conclusions. In *Chapter Seven*, I will address this issue by offering principles that allow one to reach reasonable conclusions in a wide range of cases where basic principles of direct inference, such as the ones proposed by Pollock, give one defeasible reasons for accepting mutually inconsistent conclusions.

CHAPTER 4: BACCHUS'S THEORY OF DIRECT INFERENCE

1. Introduction

In the present chapter, I discuss the theory of direct inference proposed by Fahiem Bacchus in his book *Representing and Reasoning with Probabilistic Knowledge*.(1990) The theory of direct inference presented in this book differs markedly from the approach to direct inference taken by Bacchus and his collaborators in several later essays on the topic direct inference. But for ease of presentation, I will write in the present chapter *as if* Bacchus's 1990 theory represented his final view on the topic. In *Chapter Five*, I will discuss approach advocated for by Bacchus *et al* in later essays.

Bacchus's theory of direct inference shares many features with the theory of direct inference proposed by Pollock. Clearly, Bacchus was influenced by earlier work of Pollock and others. Moreover, the work of Bacchus incorporates a number of ideas that were explored in Pollock's paper *A Theory of Direct Inference* (1983), and shares a number of central features with the theory of direct inference developed in Pollock's monograph *Nomic Probability and the Foundations of Induction*.(1990) Two similarities in the theories of Pollock and Bacchus are noteworthy. First, both theories place an emphasis on the distinction between *definite* and *indefinite* probabilities (though Bacchus does not use Pollock’s terminology). Second, both theories are *non-classical theories of direct inference*, in the sense articulated by Pollock.(Pollock 1983 & 1990)

Like Pollock’s theory, Bacchus’s theory of direct inference comes as a part of a more inclusive theory that has several distinguishable parts. In particular, Bacchus’s theory of direct inference comes as a part of a structure having the following three components:

(1) A combined semantics for statistical and propositional probability statements.

(2) A proposal for defining the values of propositional probabilities in terms of statistical probabilities (a principle of direct inference).

(3) A prescribed form of default reasoning for estimating the value of unknown statistical probabilities.

I will describe the three components of Bacchus’s theory and then go on to consider how well the theory fares in the face of the three problems of direct inference described in *Chapter One*.

2. Propositional and Statistical Probability Logic

One novelty of Bacchus’s work consists in its introduction of a thorough and precise semantics for propositional and statistical probabilities. The respective notions of propositional and statistical probability, proposed by Bacchus, fall respectively within the categories outlined by Pollock's distinction between definite and indefinite probabilities.[30]

In introducing a combined logic for propositional and statistical probabilities, Bacchus proceeded in phases, first introducing a logic for propositional probability statements, then a logic for statistical probability statements, and finally a logic that incorporated both types of probability statements. In describing Bacchus's account, I will proceed in the same order. I will omit description of some of the details of Bacchus's theory for the sake of focusing on those details that are the most relevant to evaluating his account of direct inference.

The Propositional Probability Logic:

The basic idea for dealing with propositional probabilities consists in the introduction of a probability distribution over a set of possible worlds, and the identification of the probability of an individual proposition as the sum of the values of the probabilities of the worlds in which the proposition is true. While the basic idea is pretty straight forward, there are some features of the proposal that bear mention.

[30] Note, however, that Bacchus's opinion regarding the varieties of indefinite and definite probability that are of central concern differs from Pollock's.

The syntax of the language for Bacchus's propositional probability logic is fairly standard. The object language for the logic will be a first order language, with the addition of the operator "PROB" which takes formulae as arguments, so that if 'α' is a formula, 'PROB(α)' is a numeric term.[31] The central anomaly in the syntax of the object language for Bacchus's propositional probability logic is its division of the symbols of the language into *object symbols* and *numeric symbols*. Intuitively, object symbols (variables, predicates, and function symbols) are those symbols included in the language for the purpose of describing the respective domain of interest. Numeric symbols, on the other hand, are simply those terms used for describing numeric objects and relations. In the specification of the set of well formed formulas of a propositional probability logic language, atomic formulas are restricted to what we may call "pure object formulas" and "pure numeric formulas". In particular, a string formed via the concatenation of an n-ary predicate symbol and a sequence of n terms, i.e., '$Pt_1 \ldots t_n$', is a formula of the language *if and only if* 'P' is an object predicate symbol and 't_1' through 't_n' are object terms, or 'P' is a numeric predicate symbol and 't_1' through 't_n' are numeric terms.

The semantics for the propositional probability logic rests in what Bacchus's calls "propositional probability structures". A *propositional probability structure* is an ordered quadruple $M = \langle O, S, V, \mu \rangle$, *where* O is a domain of non-numeric objects, *where* S is a set of states (or possible worlds), *where* V is a valuation function, such that, for each element, s, of S, V(s) assigns (1) a relation of the correct arity to each predicate symbol, (2) a

[31] For the sake of notational consistency, I am using the notation "PROB" where Bacchus uses the notation "prob".

function of the correct arity for each function symbol, and (3) an element of O to each constant, and *where* μ is a discrete probability function that maps the elements of S to the interval [0, 1], so that $\sum_{s \in S} \mu(s) = 1$, and where A ⊆ S, $\mu(A) = \sum_{s \in A} \mu(s)$.

Several comments about the nature of *propositional probability structures* are required to clarify the picture. First note that the domain of objects O is intended as the concrete component of the domain of discourse associated with a propositional probability structure, and that in addition to mapping non-numeric terms to the elements of O, it is intended that V will map numeric terms to the reals. A related novelty of the interpretation function for *propositional probability structures* is the stipulation that all of the symbols, save non-numeric predicate symbols, have the same interpretation in all possible worlds.

As is typical for formal semantics, a truth evaluation is defined over the formulas of the language. In the case of Bacchus's propositional probability logic, the truth value of formulas is relativized to a propositional probability structure, M, a world, s, and a variable interpretation function, υ, which assigns an interpretation to the variables of the language. In turn, one may write "(M, s, υ) ⊨ α" if the formula, α, is assigned the truth value *true* at s in M by the variable interpretation function υ. The central novelty in the truth assignments specified by Bacchus's theory is in the interpretation of terms of the form *PROB(α)*, which are assigned the value μ({ s′∈ S | (M, s′, υ) ⊨ α }).

To illustrate the character of propositional probability structures and their semantics, consider an example of relatively limited language containing only four object terms, “bob”, “carol”, “ted”, and “alice”, one two-place predicate “loves”, and a large, but finite, number of object variables “x_1”,…, “x_n”. In that case, let us define a propositional probability structure $M = \langle O, S, V, \mu \rangle$. Let O contain only four concrete objects: Bob, Carol, Ted, and Alice. Now let V be defined so that “bob”, “carol”, “ted”, and “alice” are assigned to Bob, Carol, Ted, and Alice, and so that S corresponds to the set of possible ways of assigning “loves” to pairs of concrete objects in O. In that case, S will have 2^{16} elements. Now let μ be a uniform probability function, so that for each s in S, $\mu(s) = 1/(2^{16})$.

Now that we have completed our description of M, let w be the element of S at which V assigns “loves” to all and only elements of the set of pairs { ⟨bob, carol⟩, ⟨carol, bob⟩ }. In that case, for all υ, we have $(M, w, \upsilon) \models$ (Bob)loves(Carol)$\wedge$(Carol)loves(Bob). However, it is not the case that $(M, w, \upsilon) \models$ prob((Bob)loves(Carol)$\wedge$(Carol)loves(Bob)) = 1, since prob((Bob)loves(Carol)$\wedge$(Carol)loves(Bob)) designates $n/(2^{16})$, where n is the number of elements of S, such that ⟨bob, carol⟩ and ⟨carol, bob⟩ are in the interpretation of “loves”. In fact, there are 2^{14} elements of S, where ⟨bob, carol⟩ and ⟨carol, bob⟩ are in the interpretation of “loves”, so that $(M, w, \upsilon) \models$

PROB((Bob)loves(Carol)$\wedge$(Carol)loves(Bob)) = 1/4. Moreover, in the case where a probability statement contains no free variables, its truth value is independent of the element of S in which the statement is evaluated. So, in the present case, for all υ and s, (M, s, υ) $\models$ PROB((Bob)loves(Carol)$\wedge$(Carol)loves(Bob)) = 1/4.

The Statistical Probability Logic:

Bacchus's approach to statistical probability proceeds by the introduction a probability distribution over a domain of objects. The probability and conditional probability of various sets of objects is, then, determined by the sum of the values of the probabilities of the objects falling within respective sets. In effect, Bacchus's notion of statistical probability generalizes the notion of frequency.—Frequencies correspond to statistical probabilities in the case where every object is assigned the same value, i.e., each object is assigned the value *one over the cardinality of the domain*.

As a means to representing the probability and conditional probability of respective sets, Bacchus adopts notation for denoting sets of elements of a domain, O, and notation for denoting sets of n-ary vectors of the product space, O^n, of O. Bacchus's notation uses open first order formulas to denote to the set of objects, or the set of vectors, satisfying a respective formula. Intuitively, we may think of the formula "P(x)" as designating the set of objects which satisfy "P(x)", and "R(x,y)" as the set of ordered pairs of objects which satisfy "R(x,y)".

The syntax of the language for Bacchus's statistical probability logic is similar to the syntax of the propositional probability logic. The syntax of statistical probability logic differs in that it does not include the operator "PROB", but instead includes a variable binding statistical probability operator. Square brackets "[" and "]" surrounding a formula along with subscripted variables indicates the application of the operator, so that the expression "$[P(x)]_x$" denotes the real number corresponding to the statistical probability of the set of objects satisfying the open formula "P(x)". In order to formulate more complex expressions designating statistical probabilities, Bacchus's uses the "$\vec{x}$" to represent a vector of variables $\langle x_1, \ldots, x_n \rangle$. In turn, terms of the form '$[\alpha]_{\vec{x}}$' are used to designate the statistical probability of the set of instantiations of $\vec{x}$ that satisfy 'α'. Conditional statistical probabilities are, in turn, introduced by a definitional extension:

Definition: $[\alpha|\beta]_{\vec{x}} \times [\beta]_{\vec{x}} = [\alpha \wedge \beta]_{\vec{x}}$, if $[\beta]_{\vec{x}} \neq 0$, else $[\alpha|\beta]_{\vec{x}} = 0$.

The semantics for the statistical probability logic rests in what Bacchus's calls "statistical probability structures". A *statistical probability structure* is an ordered triple M = $\langle O, V, \mu \rangle$, *where* O is a domain of objects, *where* V is a valuation function that assigns (1) a relation of the correct arity to each predicate symbol, (2) a function of the correct arity for each function symbol, and (3) an element of O to each constant, and *where* μ is a discrete probability function that maps the elements of O to the interval [0, 1], so that $\sum_{o \in O} \mu(o) = 1$, and where $A \subseteq O$, $\mu(A) = \sum_{o \in A} \mu(o)$.

In the case of Bacchus's statistical probability logic, the truth value of a proposition is relativized to statistical probability structure, M, and a variable interpretation function, υ. In turn, one may write '$(M, \upsilon) \models \alpha$' if the formula, "$\alpha$", is assigned the truth value *true* in M by the variable interpretation function υ. A central novelty of the semantics of the statistical probability logic is the interpretation of formulas of the form '$[\alpha]_{\vec{x}}$', which are assigned the value $\mu^n(\{\vec{\alpha} \mid (M, \upsilon[\vec{x}/\vec{a}]) \models \alpha\})$, where μ^n represents the n-fold product measure formed from μ, and where $\upsilon[\vec{x}/\vec{a}]$ is the variable assignment function that differs from υ only in that $\upsilon(x_i) = a_i$ (for $i = 1, \ldots, n$).

To get the feel for the character of statistical probability structures and their semantics, consider a variant of the example of Bob and Carol, and company. Once again, suppose that our language contains only four object terms, "bob", "carol", "ted", and "alice", one one-place predicate "is-a-person", one two-place predicate "loves", and a large, but finite, number of object variables "x_1",…, "x_n". Now let $M = \langle O, V, \mu \rangle$, *where* O contains only four concretes objects, Bob, Carol, Ted, and Alice, and *where* V assigns "bob", "carol", "ted", and "alice" to Bob, Carol, Ted, and Alice, assigns "loves" to all and only elements of the set of pairs { ⟨bob, carol⟩, ⟨carol, bob⟩, ⟨bob, bob⟩ }, and assigns "is-a-person" to all concrete elements of O. In this case, let us suppose that μ(bob) = 1/2, while μ(carol) = μ(ted) = μ(alice) = 1/6. In that case, $(M, w, \upsilon) \models$ [(x)love(bob) | (x)is-a-person $]_x$ =2/3, since Carol and Bob both stand in the *loves* relation to Bob, and Carol and Bob collectively count for two-thirds of the domain according to the weighting prescribed by μ.

The Combined Propositional & Statistical Probability Logic:

Bacchus's combined statistical and propositional logic is generated by augmenting the syntax of the propositional probability logic with the introduction of the variable binding statistical probability operator characteristic of the statistical probability logic, and by augmenting the four components of characteristic of propositional probability structures so as to include both a discrete probability function, μ_O, defined on the elements of the object domain, O, as well as a discrete probability function, μ_S, defined on the elements of the set of possible worlds, S. The combined logic then allows for the application of propositional probability operator to statements formulated in terms of the statistical probability operator, to form statements such as "PROB($[\alpha]_{\vec{x}} = 0.75$) > 0.5". The combined probability logic also permits for the introduction of an *expectation operator* through a definitional extension of the language. The expectation operator is a monadic function that applies to any term, t, to form a new term E(t). The value of the term E(t) is equal a weighted average of the measure of the interpretation of t across the elements of S. That is, $E(t)^{(M, s, \upsilon)} = \sum_{s' \in S} \mu_S(s') \times t^{(M,s,\upsilon)}$. The expectation operator is of greatest interest when its object is a statistical probability term or a conditional statistical probability term. The expectation of such terms feature largely in Bacchus's account of direct inference.

3. Bacchus's Direct Inference Principle

Bacchus's theory of direct inference is applicable to draw conclusions about the values of respective propositional probabilities from a body of accepted statements expressed in the language of Bacchus's Statistical Probability Logic. For purposes of articulating Bacchus's theory of direct inference, call a formula "objective" just in cases it is a formula of Bacchus's Statistical Probability Logic. In that case, the core of Bacchus's theory of direct inference is captured by the following principle:

[Direct Inference Principle] *If* α is an objective formula and KB is the conjunction of the complete set of objective formulae that an agent believes, *then* the agent's degree of belief in α should be determined by the equality $\mathrm{PROB}(\alpha) = E([\alpha^V|KB^V]_{\vec{v}})$.[32]

The principle is presented as a definition within Bacchus's work, and though it is possible for the stated equality to fail in the context of Bacchus's combined probability logic, it is demonstrable that the function set out by the principle is a probability function, i.e., for all KB and $\vec{v}$, $E([-^V|KB^V]_{\vec{v}})$ is a probability function.

[32] Where $\langle c_1, \ldots, c_n\rangle$ are the *n* distinct constants that appear in $\alpha\wedge KB$, $\langle v_1, \ldots, v_n\rangle$ are *n* distinct variables, α^V and KB^V are the result of uniformly substituting v_i for c_i in α and KB, for all i.

On its own, Bacchus's Direct Inference Principle will allow one to draw useful conclusions in many cases. This fact is facilitated by a number of theorems that permit us to deduce the value of $E([\alpha^V|KB^V]_{\vec{v}})$ from the value of the expectation of α^V conditional on consequences of KB^V. For example, the following theorem is very useful in applying Bacchus's Direct Inference Principle in a wide range of cases:

> If no element of $\vec{x}$ which appears in $\alpha \wedge \beta$ is also free in λ, then
> $\models [\beta \wedge \lambda]_{\vec{x}} \neq 0 \supset [\alpha|\beta \wedge \lambda]_{\vec{x}} = [\alpha|\beta]_{\vec{x}}$.(1990, 107)

Due to the preceding theorem, Bacchus's Principle of Direct Inference is applicable to a wide range of simple cases. For instance, suppose that one's knowledge base, KB, is Californian(Joe) $\wedge$ [doctor(x)|Californian(x)]$_x$ = 0.1. In that case, we may reach the correct conclusion as to the likelihood that Joe is a doctor by reasoning as follows:

1. PROB(doctor(Joe)) = E([doctor(v)| Californian(v) $\wedge$ [doctor(x)|Californian(x)]$_x$ = 0.1]$_v$) [Principle of Direct Inference]

2. PROB([doctor(v)| Californian(v) $\wedge$ [doctor(x)|Californian(x)]$_x$ = 0.1]$_v$ = [doctor(v)|Californian(v)]$_v$) =1 [the Theorem]

3. E([doctor(v)| Californian(v) $\wedge$ [doctor(x)|Californian(x)]$_x$ = 0.1]$_v$) = E([doctor(v)|Californian(v)]$_v$) [Definition of Expectation]

4. PROB(doctor(Joe)) = E([doctor(v)|Californian(v)]$_v$) [Transitivity of Identity]

5. E([doctor(v)|Californian(v)]$_v$) = E(0.1) = 0.1 [Definition of Expectation]

6. PROB(doctor(Joe)) = 0.1 [Transitivity of Identity]

Despite the usefulness of the Direct Inference Principle on its own (in the context of the combined *the Combined Propositional & Statistical Probability Logic)*, the principle by itself will not allow the range on inferences permitted by the orthodox theories of direct inference that have already been considered. To extend the range of inferences permitted by the Direct Inference Principle, Bacchus's proposes a default rule that may be used to infer the value of expectations, and a method for reaching a consistent set of conclusions in cases where the default rule yields inconsistent conclusions. The default rule, which is very similar to Pollock's principle of non-classical direct inference, is as follows:

Let D_0 be the set of instances of direct inference, relative to a given knowledge base, KB. That is, for each objective formula, α, D_0 contains $\text{PROB}(\alpha) = E([\alpha^V|KB^V]_{\vec{\nu}})$. Now let T_0 be the closure of $D_0 \cup KB$ under deductive consequences. Now generate a set of non-monotonic assumptions, NA, in accordance with the following conditional:

If $\text{PROB}((\forall\vec{\nu})(KB^V \supset \beta)) = 1 \in T_0$, *then* $E([\alpha|KB^V]_{\vec{\nu}}) = E([\alpha|\beta]_{\vec{\nu}}) \in NA$.

In general, the set of default identities, NA, generated determined by the preceding conditional will not form a consistent set. As such, Bacchus proposed a procedure for selecting an acceptable (consistent) subset of NA. Bacchus first proposes an enumeration, $\{ D_1, \ldots D_n, \ldots\}$ of the set of subsets of NA, and then defines the notion of a default theory, as follows:

> Let T_i be the closure under logical consequences of $T_0 \cup D_i$, where D_i is the i-th finite set of default assumptions. That is, $T_i = \{ \alpha \mid T_0 \cup D_i \models \alpha \}$.

Many of the respective T_i will be inconsistent. But, intuitively, a permissible set of conclusions must be consistent. To this end, Bacchus develops criteria for determining whether a theory is 'viable'. The criteria articulated by Bacchus rule that inconsistent theories are unviable. The criteria also incorporate a rule which favors default assumptions based on more precise information.—This rule corresponds to the preference for narrower reference classes espoused by the orthodox approach to direct inference. Bacchus's criteria for selecting viable default theories are outlined through the following series of definitions.

Definition - Contradicted Theories:

A theory T_i contradicts a theory T_j if there exists a formula α such that $\alpha \in T_i$ and $\neg\alpha \in T_j$.

Definition - Preferred Assumptions:

The assumption $E([\alpha|KB^V]_{\vec{\nu}}) = E([\alpha|\beta_1]_{\vec{\nu}})$ is preferred to the assumption $E([\alpha|KB^V]_{\vec{\nu}}) = E([\alpha|\beta_2]_{\vec{\nu}})$, if $PROB((\forall\vec{\nu})(\beta_1 \supset \beta_2)) = 1 \in T_0$.

Definition - Preferred Theories:

The theory T_i is preferred the theory T_j if for every default assumption $d \in D_i$, there exists a default assumption $d' \in D_j$ such that d is preferred to d′.

Definition - Viable Theories:

T_i is viable if it is not contradicted by any preferred theory.

In essence, Bacchus's system regards a theory, T_i, as viable just in case there is no theory, T_j, based on more specific (or equally specific) information that yields different conclusions.—Note that Bacchus's definitions entail that no inconsistent theories are viable, since all theories are preferred to themselves. In cases where there are multiple distinct viable theories, Bacchus proposes that one may accept any one of the viable theories.

In addition to leading to 'credulity problems' (which I discuss below), Bacchus proposal that one may accept any viable theory is a little perplexing, since the theory T_i for the case where $D_i = \emptyset$ is always a viable theory.—This means that it is always permissible to accept a theory which incorporates no default assumptions. It would seem, then, that Bacchus's theory gives little in the form of positive direction as to what conclusions one ought to draw using direct inference. That said, in his presentation of examples to illustrate the virtues of his approach, Bacchus tends to assume that one would like to accept a viable theory that incorporates the greatest number of non-monotonic assumptions, so long as the theory is not inconsistent with another viable theory.—This is something to bear in mind as one considers the following examples.

Example 1 – Preference for Narrower Reference Classes:

Suppose that KB is $\beta(c)\wedge\chi(c)\wedge[\alpha(x)|\beta(x)]_x=0.6\wedge[\alpha(x)|\beta(x)\wedge\chi(c)]_x=0.3$, and we wish to assign a value to $\mathrm{PROB}(\alpha(c))$.

In this case, D_0 contains PROB(α) = E([α(x)|KB(x)]$_x$). and NA contains both E([α(x)|KB(x)]$_x$) = E([α(x)|β(x)$\wedge\chi$(c)]$_x$) and E([α(x)|KB(x)]$_x$) = E([α(x)|β(x)]$_x$).

In this case, there will be viable theories corresponding to KB that contain E([α(x)|KB(x)]$_x$) = E([α(x)|β(x)$\wedge\chi$(c)]$_x$), and no viable theory will contain E([α(x)|KB(x)]$_x$) = E([α(x)|β(x)]$_x$), since the two expressions contradict one another (in the presence of KB) and the former equality is *preferred* to the latter. Thus, Bacchus's theory of direct inference permits one to accept that PROB(α) = 0.3, but it does not permit one to accept PROB(α) = 0.6.

Example 2 – Ambivalence in the Case of Competing Statistics:

Suppose that KB is β(c)$\wedge\chi$(c)$\wedge\delta$(c)$\wedge$[α(x)|β(x)$\wedge\delta$(c)]$_x$=0.6$\wedge$[α(x)|β(x)$\wedge\chi$(c)]$_x$=0.3, and we wish to assign a value to PROB(α(c)).

In this case, D_0 contains PROB(α) = E([α(x)|KB(x)]$_x$). and NA contains both E([α(x)|KB(x)]$_x$) = E([α(x)|β(x)$\wedge\chi$(c)]$_x$) and E([α(x)|KB(x)]$_x$) = E([α(x)|β(x)$\wedge\delta$(c)]$_x$).

In this case, some viable theories corresponding to KB will contain E([α(x)|KB(x)]$_x$) = E([α(x)|β(x)$\wedge\chi$(c)]$_x$) and some will contain E([α(x)|KB(x)]$_x$) = E([α(x)|β(x)$\wedge\delta$(c)]$_x$) (and no viable theory will contain both of the equalities). Thus, Bacchus's theory of direct inference permits that one accept PROB(α) = 0.3 and permits that one accept PROB(α) = 0.6 (while it does not permit that one accept both equalities).

Example 3 – Broader Reference Classes with Precise Statistics:

Suppose that KB is $\beta(c)\wedge\chi(c)\wedge[\alpha(x)|\beta(x)]_x=0.51\wedge[\alpha(x)|\beta(x)\wedge\chi(c)]_x\in[0,0.51]$, and we wish to assign a value to PROB(α(c)).

In this case, D_0 contains PROB(α) = $E([\alpha(x)|KB(x)]_x)$. and NA contains both $E([\alpha(x)|KB(x)]_x) = E([\alpha(x)|\beta(x)\wedge\chi(c)]_x)$ and $E([\alpha(x)|KB(x)]_x) = E([\alpha(x)|\beta(x)]_x)$.

In this case, some viable theories corresponding to KB will contain $E([\alpha(x)|KB(x)]_x) = E([\alpha(x)|\beta(x)\wedge\chi(c)]_x)$, others will contain $E([\alpha(x)|KB(x)]_x) = E([\alpha(x)|\beta(x)]_x)$, and some will contain both of the equalities. Thus, Bacchus's theory of direct inference permits that one accept PROB(α) $\in$ [0,0.51] and permits that one accept PROB(α) = 0.51.

One feature of Bacchus's approach that is novel in comparison to the theories that we have considered so far is that it does not propose a way of selecting a unique set of conclusions. Instead, Bacchus proposes that it is acceptable to accept any viable theory. This approach corresponds to what has been called a "credulous" approach to belief formation, in the context of default reasoning. In the context of the Problem of Competing Statistics, Bacchus's approach is tantamount to a recommendation that one accept the logical consequences of any collection of instances of direct inference, so long as doing so does not lead to inconsistency. Pollock has argued that this approach to the Problem of Competing Statistics is unacceptable because it allows agents to arbitrarily choose one set of conclusions of over other possible sets of conclusions.[33] Like Pollock, I think that these sorts of choices should not be arbitrary. While I think that Pollock's reason for rejecting approaches such as Bacchus's is compelling, additional compelling objections can also be raised. I will describe these objections below.

4. Treatment of the Three Problems

i. The Problem of Competing Statistics:

Recall that the present problem arises largely in cases where one has reliable statistics regarding the incidence of a particular property among two overlapping references classes, and one does not have reliable statistics regarding the incidence of the property among the intersection of the two sets. The most common approach to this problem, championed by Reichenbach and Pollock, is to recommend that we suspend belief in such cases. This sort of approach has been characterized as the 'skeptical policy' inasmuch as recommends

[33] See (Pollock 1995, 92-94).

suspension of belief in cases where it seems that one has defeasible grounds for two (or more) incompatible conclusions and one has no grounds for favoring either of the conclusions. On Bacchus's approach, there are many cases where each of the two incompatible conclusions will be members of competing viable theories.—This is illustrated by Example 2, above. In that case, Bacchus proposes that is permissible to accept either, but not both, of the viable theories.

I agree with Pollock's reasons for rejecting approaches such as Bacchus's. Epistemology is not arbitrary, and cases where there are no grounds for preferring one of two incompatible conclusions, it is unreasonable to accept either conclusion. Sound thinking also indicates that if there are defeasible grounds for mutually inconsistent conclusions, and one has no grounds for favoring either of the conclusions, then the truth must lie somewhere in between. Probability has long offered the possibility of setting our credence to a middle ground. It was for this reason that I rejected Pollock's skeptical approach, and for similar reasons I now reject Bacchus's approach.

My approach to the Problem of Competing Statistics (that is developed in *Chapter Seven*) attempts to moderate the extremes of skepticism and credulity. Moreover, the same example that I pressed against Pollock's approach can be enlisted to illustrate the problem with Bacchus's approach. Once again, suppose that one has reliable statistics regarding the incidence of a particular property among two overlapping references classes, and one does not have reliable statistics regarding the incidence of the property among the intersection of the two sets. Furthermore, suppose that the instance of direct inference supported by the first statistic supports the conclusion that the probability is 0.8 that the object of inquiry has the property of interest and that the instance of direct inference supported by the second statistic supports the conclusion that the probability is 0.8 that the object of inquiry does not have the property of interest. Ceteris paribus, it seems that one should adopt the belief that the probability is 0.5 that the object of inquiry has the property of interest. Yet Bacchus's proposes that it is reasonable to accept either of the two conflicting conclusions regarding the likelihood that the object of inquiry has the property of interest. Clearly, in a case such as the one just described, neither of these conclusions is reasonable inasmuch as both contradict the conclusion that is reasonable.

ii. The Problem of Uninformative Statistics:

Bacchus proposes to address the Problem of Uninformative Statistics through the doctrine that it is statements of expected frequency (or expected statistical probability) that serve as the statistical premises to instances of direct inference. The basic idea behind Bacchus's proposal is that we will rarely be able to assign a precise value to the expectation for a respective unit set reference class (or to other reference classes for which we lack information). Along with the present idea, Bacchus's theory implements the tenet that an statistic for a narrower reference class may defeat an instance of direct inference based on a statistic for a broader reference class *only if* the instances of direct inference supported by the two reference classes yield mutually inconsistent conclusions.—This tenet is implemented through Bacchus's definitions of *preferred assumptions* and *preferred theories*.

I accept Bacchus's proposal that it is statements of expected frequency that serve as the statistical premises to instances of direct inference, and I will defend the proposal in *Chapter Six*. Despite my acceptance of Bacchus's proposal that it is statements of expected frequency that serve as the major premises to instances of direct inference, the maneuver does not yield an adequate remedy to the Problem of Uninformative Statistics.

Once again, the Ace Urn example (described in *Chapter Two*) illustrates the problem with a theory of direct inference that incorporates the idea that statistics for a narrower reference class may defeat an instance of direct inference based on statistics for a broader reference class *only if* the instances of direct inference supported by the two reference classes yield inconsistent conclusions.—The problem arises for Bacchus's approach, since his approach rules that a theory associated with a default for a broader reference class is *unviable* only in cases where one's statistics for the narrower reference class are incompatible with one's statistics for the broader reference class. This point was illustrated in Example 3 (above), where the theory associated with the default, $E([\alpha(x)|KB(x)]_x) = E([\alpha(x)|\beta(x)]_x)$, is not determined to be *unviable* via the *preferred default* $E([\alpha(x)|KB(x)]_x) = E([\alpha(x)|\beta(x)\wedge\chi(c)]_x)$, since the values for $[\alpha(x)|\beta(x)]_x$ and

$[\alpha(x)|\beta(x)\wedge\chi(c)]_x$ that are recorded in KB are compatible (in the sense that they do no lead to inconsistent conclusions via direct inference).

The statistical information represented in KB in Example 3 is formally identical to the statistical information available in the Ace Urn example, and we see that Bacchus's theory gives the exact same answer as Kyburg's in the face of the example. As such, we must conclude that Bacchus's treatment of the Problem of Uninformative Statistics is unsuccessful.

iii. The Problem of Relevant Statistics:

Bacchus acknowledges this difficulty, but does not provide a concrete proposal for how to deal with it. He suggests two possible routes by which the problem could be remedied within the context of his theory. One approach is in line with the suggestions of Kyburg and Pollock, and would place syntactic criteria on the set of legitimate non-monotonic assumptions. Bacchus's other suggestion is to limit the set of non-monotonic assumptions (corresponding a respective body of accepted statements, KB) so that we allow a given disjunctive reference formula only if the disjunct appears in KB.

Bacchus professes to regard both of the approaches just suggested as *ad hoc*.(1990, 216) In turn, Bacchus proposes that it would be ideal to articulate a *theory of relevance* that would permit us to recognize misleading reference classes as irrelevant in the formation of degrees of belief. Although Bacchus provides no concrete guidance as to how one might formulate a theory of relevance, I think that his proposal is a good one. In *Chapter Six*, I propose *a theory of relevant statistics* that is in keeping with Bacchus's proposal.

CHAPTER 5: UNORTHODOX THEORIES OF DIRECT INFERENCE

1. Introduction

So far, I have surveyed the theories of direct inference offered by Reichenbach, Kyburg, Pollock, and Bacchus. These theories fit the criteria for what I have described as *orthodox theories of direct inference*. Reichenbach's account of direct inference is the touchstone of the orthodox approach, while the theories of Kyburg, Pollock, and Bacchus represent three detailed attempts to develop the approach.

In the present chapter, I will briefly survey three *unorthodox* approaches of direct inference. In considering each of the unorthodox approaches, I will try to give reasons for preferring the orthodox approach. This course is not without its difficulties, since I have already outlined some substantive problems for the orthodox approach in the preceding chapters, and the reader will have to wait until *Chapters Six* and *Seven* to see how I propose to remedy those problems. Even so, I think that we may dispense with the various unorthodox approaches, prior to actually seeing how the problems facing the orthodox approach are to be addressed. Indeed, in the case of each of the unorthodox theories that I consider, I will show either (1) that the approach is beset with difficulties that are just as troubling as the difficulties facing the orthodox theories, or (2) that the approach is extremely limited in its applicability to forming non-arbitrary probability judgments in real world cases.

My main purpose in surveying the prominent unorthodox approaches is to give readers a sense of the alternatives that are out there, and reveal the attractiveness of the orthodox approach by favorable contrast. Before going on to consider the three unorthodox approaches to direct inference, I summarize the features that are central to the orthodox approach. The orthodox approach to direct inference is characterized by the following three tenets:

(a) Assigning a single-case probability is a matter of locating an object within a reference class for which one has reliable statistics.

The present tenet involves the basic idea that direct inference consists in identifying the value of the single-case probability with the value of *an appropriate statistic*. In the case of the frequentist version of direct inference, the basic idea is to reason that the probability that a given object, c, is an element of a set, T, is equal to the (accepted) frequency of elements of T among elements of a set, R, where c is an element of R.

In the case of the *non-classical* approaches to direct inference espoused by Pollock and Bacchus, the present tenet is embodied in the identification of the values of single-case probabilities with the value of a statistic for a reference class that incorporates everything we know about the object of interest. In particular, where c is the object of interest and φ is the property of interest, PROB(φ(c)) is to be identified with the statistical probability of the property φ conditioned on the property ψ, where ψ incorporates everything we know about c.

Of the three unorthodox approaches that I consider, only the *random-worlds approach* rejects the present tenet.

(b) Direct inference does not rely on positive knowledge of the manner in which the object of inquiry is selected from the chosen reference class.

The present tenet embodies a particular stance on the background knowledge one is required to have before one is entitled to make a direct inference. According to opponents of this tenet such as Isaac Levi, correct instances of direct inference presuppose that the object of interest, c, is presented to us is as a trial of a stochastic process that generates varying results with certain chances, and presuppose that the probability we assign to the object of interest having the respective target property should be identical to the chance of a trial of the respective sort (i.e., a trial having the chosen reference property) having the respective target property.

Accounts of direct inference that accept *tenet (b)* do not accept that direct inference presupposes *positive knowledge* of the sort demanded by Levi. In the place of such a requirement, orthodox accounts of direct inference generally include a *negative knowledge requirement* that outlines the range of cases where the possession of information of a given sort would undermine a respective instance of direct inference.

The characteristic feature of Levi's theory of direct inference is the rejection of *tenet (b)*. While Salmon accepts *tenet (b)*, the tenet is not applicable to the random-worlds approach, due to its rejection of the first tenet of the orthodox approach.

(c) A central rule for the selection of admissible reference classes is narrowness.

This tenet can be implemented in various ways, depending on whether one is an advocate of a *classical* or *non-classical* approach to direct inference. In either case, the tenet prescribes that, by default, one should prefer statistics regarding narrower reference classes, especially in cases where the instances of direct inference based on different reference classes would yield mutually inconsistent conclusions

Of the three unorthodox approaches to direct inference, only Levi's approach accepts this tenet. Let us now consider the prominent alternatives to the orthodox approach.

2. Levi's Theory

Levi's theory sets an extremely high threshold to surpass before one is permitted to make a direct inference. In this respect, Levi is outspoken in rejecting the second tenet of the orthodox approach.

In order for an instance of direct inference to be justified, Levi requires that one have *positive knowledge* regarding the manner in which the object of interest, c, came to be presented to us, and *positive knowledge* regarding the chance that an object, so presented, is a member of the respective target class, T. In particular, Levi argues that legitimate instances of direct inference require that an agent accept premises of the following form:

(1) The chance of an outcome of kind *T* on a trial of kind *R* on a chance setup C is *r*.

(2) c is a trial of a kind *R* on a chance setup C.

In essence, Levi proposes that an instance of direct inference is legitimate *only if* one knows that a respective chance statement, such as (1), is true, and not a merely that a respective frequency statement (or corresponding statement of statistical probability) is true.

I share the view of other advocates of the orthodox approach, and I believe that Levi sets the bar too high regarding the sort of knowledge one must have before one is entitled to make a direct inference. Moreover, Levi's 'restrictive' theory of direct inference is a lot less useful than the style of direct inference permitted by orthodox approaches.

Unlike Levi, I am optimistic of the possibility of articulating the justificatory foundations of a robust orthodox theory of direct inference. While I do not deny that the form of inference described by Levi is reasonable, I think that there are more ambitious forms of inference that also reasonable. The burden is to illustrate that an orthodox approach can surmount the problems described in *Chapter One*.

3. Salmon's Theory

Salmon's theory of probability is similar to Reichenbach's in several respects. Like Reichenbach, Salmon regards *probability* as measure that is applicable to infinite sequences, and regards *the probability of a property* in a given sequence to be identical to *the relative frequency of the property among the elements of the sequence* (as the number of members of the sequence considered approaches infinity). Similar to Reichenbach, Salmon believes that we may assign a probability (or weight) to the proposition that a given individual has a given property (or is a member of a given set) by locating the individual in a suitable reference class (in Salmon's mind, an infinite sequence), and then 'transferring' the probability of the property of interest within the reference class to the corresponding proposition. Despite these similarities, Salmon's theory departs significantly from Reichenbach's. The specific difference in the two theories lies in their prescriptions regarding the selection of suitable reference classes.

In articulating his rule for the selection of suitable reference classes, Salmon appeals to von Mises's notion of a *place selection*. As Salmon employs the notion, a *place selection* is simply a description which affects a partition of a sequence into two cells, where the description in question does not make reference to the target property which is of interest. So, for example, if the relevant reference class is the sequence of draws from a given earn, and the property of concern is *redness*, then “was drawn on a Sunday” is a valid place selection and “is a red ball” is not.

According to Salmon’s account, a place selection, P, is statistically relevant, relative to a target class, T, and reference class, R, *just in case* there is a difference between the relative frequency of elements of T among P and among R–P. Using the notion of a *place selection*, Salmon calls a reference class, R, *homogeneous* for a target class, T, just in case there is no *statistically relevant place section* relative to T and R. Salmon proposes the following rule of reference class selection: In performing direct inference, one should “choose the broadest homogeneous reference class” of which the object of inquiry is an element.(Salmon 1970, 43)

Before going on to describe the caveats that Salmon appends to the preceding ‘basic’ proposal, it is worth considering why he thinks that his proposal to choose the broadest homogeneous reference class improves on Reichenbach’s proposal to choose the narrowest class for which we have reliable statistics. By Salmon’s reckoning, the ‘probabilities’ associated with broader reference classes are more reliable, with the basic assumption being that larger samples serve as better grounds for inductive inference.—This assumption leads to the proposal to favor broader reference classes.(Salmon 1970, 41-42) On the other hand, Salmon recognizes that there are situations in which the right reference class should not include certain objects. As a way of implementing this idea, Salmon proposes that a suitable reference class must be *homogenous* (in the technical sense defined above), with the underlying idea being that *no description should affect a partition of a suitable reference class into two cells such that the incidence of the respective target property differs among the cells* (save in the case where the partition is affected by reference to the respective target property, itself).

Subsequent to offering his basic rule for reference class selection, Salmon offers two *more restrictive* notions of homogeneity that are intended to guide our choice of reference classes, given our practical limitations as finite beings. According to the first 'practical' notion, we say that a reference class is "epistemically homogeneous" for a given target class, just in case we are *unable to formulate a description* that would affect a statistically relevant place selection (regardless of whether a statistically relevant place selection exists).(Salmon 1970, 44) According to the second 'practical' notion, we say that a reference class is "practically homogeneous" for a given target class, just in case there is no description of a statistically relevant place selection that we could practically deploy in classifying the elements of a given reference class. For example, suppose that we are considering a sequence of flips of a coin. In that case, we may reasonably accept that some description of the initial conditions of the coin at a time preceding flips of the coin would affect a statistically relevant place selection. Salmon proposes that we may ignore such place selections and regard the sequence of coin flips as practically homogeneous, since we cannot practically deploy a place selection based on the initial conditions of the coin in classifying the elements of the coin-flip sequence.

In relation to the official 'reference class rule' that prescribes that one select the broadest homogeneous reference class, the two more restrictive notions of homogeneity are intended to guide our behavior in the course of 'real life'. In this context, it seems that Salmon would have us regard the official rule as an ideal that we may approach for practical purposes by employing the criteria outlined by the two weaker notions.(Salmon 1970, 44)

Salmon's proposals for the selection of suitable statistics are very interesting, and we will see in the following chapter that some of my ideas about direct inference are similar to Salmon's.—In particular, the notion of *indiscernbility* that I appeal to plays a similar role to Salmon's notion of *homogeneity*. Despite its interesting features, Salmon's account is flawed.

In criticizing Salmon's theory, I will focus exclusively on his rule of reference class selection. Although I regard Salmon's understanding of probability as a property of infinite sequences to be unacceptable, Salmon's rule of reference class selection may be evaluated independently, and could be appropriated by a theory of direct inference that eschews interest in limiting frequencies. In assessing Salmon's rule of reference class selection, I will focus on: (1) the claim that there is a default assumption favoring broader reference classes, and (2) the proposal that there is a homogeneity condition for suitable statistics.

4. Assessment of Salmon's Reference Class Rule

i. Is broader better?

The most striking feature of Salmon's reference class rule is the default to favor broader reference classes. Now, as it turns out, Salmon's rule seems to emulate the preference for narrower reference classes prescribed by orthodox theories, due to Salmon's additional prescription that suitable reference classes be homogeneous. So, for example, if we consider a case where we are trying to predict whether a given Californian has a university degree, and we only have statistics about the incidence of graduates among Californians and among Americans (and those statistics differ), then Salmon would recommend that we use the statistics corresponding to the narrower reference class as opposed to the broader reference class, as homogeneity fails for the broader one.

It is difficult to see the point of Salmon's prescription to prefer broader reference classes, for, if Salmon's theory worked as he intended, it appears that the existence of statistics for a reference class, R′, will prove the unsuitability of a reference class, R, so long as R′ is narrower than R, and R′ is determined by a suitable place selection. That said, it seems that Salmon's idea that broader reference classes are *prima facie* superior is itself based on a confused understanding of the nature of direct inference. Indeed, Salmon maintains that broader reference classes are superior, on the grounds that the number of instances of examination in a reference class constitutes the evidence required to conduct an inductive inference. This line of reasoning conflates inductive inference and direct inference. While it may be that inductive inference is somehow grounded in direct inference, direct inference is not grounded in induction. Once one has a proper understanding of how direct inference functions, there is no temptation to think that the size of a reference class has a bearing on credibility of the conclusions generated by statistics for the reference class.

In performing an instance of direct inference, one assumes that the object about which on is reasoning, c, is as likely to be a member of the respective target class, T, as a *random element of the proposed reference class, R* (where we understand a random element of R to be an object that is minimally described as an element of R which was selected by a process which is equally likely to produce each element of R). Ultimately, the reasonableness of drawing such a conclusion depends on the degree to which the object of interest is indiscernible from a *random element* of the reference class, relative to the property of being an element of the target class. The problem of determining when it is reasonable to accept that the object of inquiry is relevantly indiscernible is problem which requires attention, but breadth is not a factor in attaining such indiscernibility. Such indiscernibility is premised on an agent's relative lack of information about the object of inquiry *vis a vis* its membership in the target class.

In proposing that broadness is a desideratum for reference classes, it is likely that Salmon is failing to distinguish between the sample by which we may infer the statistics for a given reference class and the statistics for the reference class itself. It is obvious that we wish the sample that we use in inferring the statistics for a given reference class to be relatively large, for if the sample is relatively large, then *ceteris paribus* we will have good reason to think that our estimate of the statistics for the reference class is accurate. Despite the value of having a relatively large sample as a basis for one's estimate for the statistics for a reference class, there is no virtue associated with statistics for large reference classes when it comes to direct inference.

ii. Must suitable reference classes be homogeneous?

I find Salmon's proposal that suitable reference classes must be homogeneous very interesting. That said, the condition (as defined by Salmon) is not quite correct. The inadequacy of the proposal is demonstrable, since there are cases where it is reasonable to conduct direct inference using a non-*homogeneous* reference classes. Worse still, there are technical problems with the notion of a *place selection* (associated with the proposal to limit place selections to descriptions that do not make reference to the target property of interest).[34]

We can see by example that homogeneity, as defined by Salmon, is not required of suitable reference classes. Indeed, suppose that we are considering a reference class, R, consisting of a sequence of dice having various numbers of sides. Suppose that we know that the relative frequency of six-sided dice among the sequence is 0.4. However, let us also suppose *that* each member of the sequence is either green or blue, *that* the relative frequency of six-sided dice among the green dice is 0.3, and *that* the relative frequency of six-sided dice among the blue dice is 0.6. In this case, the place selection “is green” and the place selection “is blue” are each statistically relevant, relative to the sequence of dice and the property of being *six-sided*. Yet, in the case where one’s object of inquiry, c, is an element of R, and one does not have any evidence regarding the likelihood that c is blue or the likelihood that c is green, it is reasonable to use our statistics regarding the reference class, R, to conclude that the probability that c is six-sided is 0.4 (in spite of the fact that R is not homogeneous).

In the face of the preceding example, one might propose that the Salmon’s rule is saved by the fact that R (in the example) is *practically homogeneous* inasmuch as we are not able to determine the color of at least one of its elements (namely the object of inquiry, c). While

[34] For one, there are many cases where we should be able to perform direct inference using statistics for a given reference class, but we are nonetheless able to articulate a statistically relevant place selection. Similarly, the constraint that place selections not make reference to respective target properties is insufficiently restrictive. For example, the constraint proposed by Salmon does not restrict place selections that make reference to a property that is deductively related to a respective target property.

it is unclear whether Salmon's notion of *practical homogeneity* is meant to apply in the manner proposed by this response, it is trivial to formulate other statistically relevant place selections (such as "x is c or x is green") that show that R is not practically homogeneous.[35, 36]

It could be argued that the preceding counter response is only made possible through the use of gerrymandered place selections. Unfortunately, Salmon provides no account of how we may restrict such selections. Regardless, reflection of the original version of the example of the sequence of dice should help us to see that it is not homogeneity (as spelled out by Salmon) that is necessary for suitable reference classes. What the example shows is that '*heterogeneity'* need not be a problem, so long as we are *not* able to locate the object of inquiry within one of the place selections that proves the heterogeneity of the reference class.—So in the example, it is acceptable that there is a place selection, "x is green", that is statistically relevant, because we cannot locate the objection of inquiry among the elements of the set affected by the place selection.

The problem associated with place selections, or reference classes, affected by descriptions such is "x is c or x is green" is one which Salmon's theory is not equipped to deal.

[35] The proposed response has the unwelcome consequence that we may regard any reference class as practically homogeneous so long as c is an element of the class and it is not practically determinable whether c has the respective target property.

[36] In order for the proposed place selection, *x is c or x is green*, to be practically homogeneous, we must suppose that we are able to determine for each element of R whether it is c or green.—There is no difficulty in formulating the present example so that this is the case.

Salmon's theory shares this problem with all of the orthodox theories that have been proposed up until now. I will endeavor to address the problem in *Chapter Six*.

5. The Random-World Approach to Direct Inference

The *random-worlds approach* to direct inference has been presented in a number of papers and books that have been authored and co-authored by an array of researchers working in the field of artificial intelligence. In my discussion of the approach, I will focus on the presentation that appeared in the paper "From Statistical Knowledge Basis to Degrees of Belief", authored by Bacchus, Grove, Halpern, and Koller.(1996)

The syntax and the semantics of object language proposed for the random-worlds approach are nearly identical to the syntax and the semantics of the object language used in (Bacchus 1990). Moreover, the differences between the object languages for the two approaches are not ones that are of concern to me, so I will proceed as if the two approaches employed the same object language (and I will adjust my presentation to accommodate this identification).

The language of the random-worlds approach is simply a first order language augmented by the addition of two sorts of probability operators (which correspond to the two sorts of probabilities marked by Pollock's distinction between *indefinite* and *definite* probabilities).

The approach allows for statistical probability statements of the form "[doctor(x)|Californian(x)]$_x$ = 0.4", which would be understood to express that 40% of Californians are doctors. The approach also allows for propositional probability statements of the form "PROB(doctor(Bob)) = 0.4", which would be understood to express that the probability that Bob is a doctor is 0.4.

I will not go into detail in describing the semantics of the approach, as those details are included in the preceding chapter. The basic idea is that statistical probability statements are evaluated at individual worlds (or models) *whereas* propositional probability statements are evaluated by consideration of the full set of worlds (or models). In the case of propositional probability statements, it is understood that there is an underlying assignment of probability to the space of possible worlds, so that a statement "PROB(doctor(Bob)) = 0.4" expresses that the sum of the probabilities of the possible worlds in which Bob is a doctor is 0.4.

As in the case of Bacchus, the space of possible worlds is identified with the set of first-order models for a given language for a domain a given size. In the case of the random-worlds approach, each possible world/model is assumed to be equally likely (and in general it is assumed that the domain is arbitrarily large).[37] As such, the random-worlds

[37] In the detailed presentation of the random-worlds approach, probability is defined relatively to a domain of certain size. In order to focus on the critical aspects of the approach, I will suppress mention of domain sizes, and in all cases, state the results of the random-worlds approach for the limiting case where the domain size approaches ∞.—There are additional complications that also lie beyond my immediate interest. Those who are interested in further details should consult Bacchus, *et al.* (1996).

approach differs from the approach to direct inference proposed in Bacchus (1990) in that it prescribes a way of assigning probabilities to individual worlds.

The ampliative dimension of the random-worlds approach is encapsulated in the assumption that each possible world/model is equally likely. And with a uniform probability distribution in place, it is supposed that belief formation will be conducted *via* conditionalization, so that the probability one assigns to a given proposition is identical to the prior probability of the proposition conditioned on one's background assumptions. In turn, the random-worlds approach does not include a rule of direct inference. Rather patterns of implication similar to the ones prescribed by orthodox theories of direct inference are entailed by the probability distribution stipulated by the approach.

A central virtue of the random-worlds approach lies in the fact that it validates many intuitively correct patterns of direct inference. In addition to this virtue, the approach suffers from none of the so called "three problems of direct inference". In evaluating the random-worlds approach, I will proceed as follows: First, I will summarize some of the key theorems which illustrate that the random-worlds approach validates numerous desirable patterns of direct inference. Next, I will briefly explain why the approach is not affected by any of the three problems of direct inference. Finally, I will explain why the random-worlds approach is deeply flawed despite its desirable features. In particular, the random-worlds approach is tainted with the problem of language dependence, which is typical of theories that appeal to indifference reasoning to justify the use of uniform probability assignments.

i. Virtuous Patterns

In the case of orthodox approaches to direct inference, it is typical that the prescriptions of a theory take the form of a rule that says that in cases where one is warranted in accepting a relevant statistical premise, and premise stating that the object of interest is a member of the corresponding reference class, one may defeasibly accept a conclusion about the probability that the object of inference is a member of the corresponding target class. The random-worlds approach 'validates' many patterns of direct inference that have generally been regarded as correct. This desideratum is achieved by showing that it is a theorem that when the relevant premises are assumed, the related conclusions generally thought to be licensed by direct inference follow deductively. For example, where KB is the set of propositions that one accepts as fact, KB has the form $\{ \psi(c) \wedge KB' \}$, and KB logically implies $[\varphi(x)|\psi(x)]_x \in [r, s]$, it follows that $PROB(\varphi(c)|KB) \in [r, s]$, provided that c does not appear in '$\varphi(x)$', '$\psi(x)$', or KB.

In addition to validating patterns of direct inference such as the preceding (where it seems that only one statistical statement relevant to $PROB(\varphi(c)|KB)$ is a consequence of KB), the random worlds approach validates the idea that there is a default preference for narrower reference classes. In particular, where KB implies $[\varphi(x)|\psi_0(x)]_x \in [r, s]$, it follows that PROB($\varphi$(c)|KB) $\in$ [r, s], so long as the following conditions obtain:

(1) KB implies $\psi_0(c)$,
(2) for any expression of the form '$[\varphi(x)|\psi(x)]_x \in [t, u]$' in KB, it is the case that KB implies $(\forall x)(\psi_0(c) \supset \psi(x))$ or implies $(\forall x)(\psi_0(c) \supset \neg\psi(x))$,
(3) the symbols that appear in '$\varphi(x)$' appear in KB only on the left side of the proportion expression described in (2), and
(4) c does not appear in '$\varphi(x)$'.

ii. The Three Problems

All of the three problems introduced in *Chapter One* are problems that arise through the procedure of identifying single case probabilities with statistical probabilities for appropriate references classes. The Problem of Uninformative Statistics concerns the possibility of setting a single-case probability to the value of a frequency for an overly narrow reference class, while the Problem of Relevant Statistics is connected the possibility of identifying a single-case probability with the value of a statistic for a gerrymandered reference class. The Problem of Competing Statistics arises in cases where one is aware of divergent statistics for two references classes (and neither reference class is preferable as a basis for direct inference). Because the random-worlds approach does not attempt to identify single case probabilities with the values of statistical statements, none of the three problems arise for the approach.

iii. The Problem of Language Dependence

The problem of 'language dependence' characteristic of the random-worlds approach is most easily grasped by considering how the probabilities prescribed by the approach vary according to one's language, and, in particular, according to one's choice of a set of atomic predicates. For example, suppose that one's language includes a constant, c, and we regard, $\varphi(c)$, to be an atomic formula. In that case, the random-worlds approach yields the conclusion that $PROB(\varphi(c)) = 1/2$. Now suppose that we do not regard $\varphi(c)$ as an atomic formula, but regard $\varphi(x)$ as logically equivalent to a disjunction of a pair of atomic formulae, $\varphi_1(c)$ and $\varphi_2(c)$, so that $\varphi(c)$ is logically equivalent to $\varphi_1(c) \vee \varphi_2(c)$. In that case, the random-worlds approach yields the conclusion that $PROB(\varphi(c)) = 3/4$.

The fact that the probability of a proposition could vary according to one's choice of a set of atomic predicates is deeply problematic, as the decision to regard a predicate as atomic as opposed to as a disjunction of atomic predicates is one that may be made by fiat. That is, given a vocabulary, Φ, of atomic formulae, there is nothing to prevent one from taking any formula, $\varphi(x)$, in Φ and introducing a new pair of formulae, $\varphi_1(x)$ and $\varphi_2(x)$, so that $\varphi(x)$ is logically equivalent to $\varphi_1(x) \vee \varphi_2(x)$, and $\varphi_1(x)$ and $\varphi_2(x)$ may be regarded as atomic in place of $\varphi(x)$. This flexibility in our attitude toward the atomicity of predicates derives from the relative lack of restrictions on our choice to regard a predicate as atomic. The following constitutes the only objective constraint: A set of formulae, Φ, may be regarded as one's set of atomic formulae *only if* the extension of any element of Φ is not deductively constrained by the extension of the other elements of Φ.

Bas van Fraassen has proposed that the problem with language dependent systems is that they allow differences in prescription to result from differences in situation that should not make prescriptive differences.(van Fraassen, 1989) I agree with this diagnosis. In the case of the random-worlds approach, the prescribed probabilities vary according to what set of predicates one takes as atomic. But what probabilities are prescribed should not depend on the choice of atomic predicates, since that choice is largely arbitrary (since, for any choice of atomic predicates, one may augment one's language to change it in a way that affects which predicates are regarded as atomic).

There have been some attempts to defend language dependent systems of inference, in the face of the preceding sort of objection, by claiming that the choice of one's language is affected by epistemic considerations, and similarly that the fact that one has chosen a given language reflects one's epistemic state.[38] Despite arguments of the preceding sorts, it is clear that the choice of atomic vocabulary is not ordinarily an epistemic concern, and, more importantly, a given choice of vocabulary cannot ordinarily be deemed epistemically incorrect. That said, it is clear that if one is determined to employ a language dependent system of inference (so that choosing a given language entails choosing certain beliefs), then one's choice of an atomic vocabulary would be of epistemic importance. But in a

[38] See Williamson (2003), and Halpern and Koller (1995), for example.

case where one is determined to employ a particular language dependent system, the problem of rational belief formation reduces to the problem of language choice. And in such a situation, it is unclear how one would go about choosing an atomic vocabulary unless one will simply 'reverse engineer' a choice of vocabulary that will determine the probabilities that one antecedently regards as reasonable (assuming that there is some choice of atomic vocabulary which will enable the given language dependent system to deliver the correct results). If the choice of a correct language to use in employing a language dependent system is *not* simply a matter of reverse engineering what one antecedently accepts to be correct, then the means to correct language choice is a mystery. As such, I think we would do well to set the random-worlds approach to the side, and work to surmount the difficulties that currently face orthodox approaches to direct inference. The primary difficulty with the orthodox approach to direct inference is associated with the problem of choosing appropriate reference classes. While it has proven difficult to solve the problem, it appears that the problem of choosing appropriate reference classes is one that is not arbitrary. Moreover, I will argue, in the following chapter, that the choice of appropriate reference classes is one that is guided by epistemic considerations that are directly tied to the justificatory bases of direct inference.

6. Conclusions

None of the unorthodox approaches to direct inference would be worthy competitors to an orthodox approach that performed well in the face of the three problems outlined in *Chapter One*. Salmon's approach does little to improve on the theory proposed by Reichenbach, and the one factor that places the approach outside of the domain of orthodox approaches (a default for broader reference classes) appears to rest on confusion. The random-worlds approach surmounts the three problems of direct inference, but leaves us with prescriptions that rely on the choice of an atomic vocabulary. At worst, the selection of an atomic vocabulary will be arbitrary. At best, our selection of an atomic vocabulary will be driven by antecedent intuitions about what conclusions we would like to reach. For this reason, the random-worlds approach does not represent progress. The principle of inference articulated by Levi is reasonable. However, the range of conclusions permitted by Levi's principle is severely limited. The limitations built into Levi's theory reflect his belief that the problems associated with orthodox theories are insurmountable. In the following two chapters, I will try to make progress in showing that the problems facing orthodox theories can be surmounted.

CHAPTER 6: TOWARD A THEORY OF DIRECT INFERENCE

1. Introduction

It is common to use frequency information to form judgments of likelihood, e.g., it's usually sunny in Tucson, so it will probably be sunny there tomorrow. But the correct rules for making such inferences are difficult to spell out. Indeed, problems quickly surface when we reflect carefully on everyday examples.

Say I want to know whether my dachshund, Flint, will live at least twelve years, but no information about the mortality rates for dachshunds is available. However, I discover that forty percent of dogs live at least twelve years, and that seventy percent of small breed dogs live at least twelve years. I decide it is better to rely on the frequency information for the narrower set of small breed dogs as opposed to the wider set of all dogs, and conclude that there is a seventy percent chance that Flint will make it to his twelfth birthday.

When I made my judgment regarding how long Flint would live, I relied on a rule that told me to prefer frequency information for narrower sets when I have relevant frequency information for two or more sets. But I cannot always rely on this rule. For example, consider the gerrymandered set formed from Flint along with all of the small breed dogs who will not live twelve years. While this set is narrower than the set of small breed dogs, clearly, I should not rely on it to help me decide whether Flint will live twelve years.

Another problem with always preferring narrower sets is that some sets are too narrow to be informative. For example, the narrowest set that includes Flint is simply the set containing Flint and no other objects. All I that I know about this set is that either none or all of its members will live twelve years. Clearly, that information is unhelpful: my frequency information for the set is uninformative.

Even if I disregard uninformative and gerrymandered sets, the prescription to prefer narrower sets provides no direction in cases where I have frequency information for two sets and neither set is narrower than the other. For example, in addition to the information already described, suppose I know that dachshunds are one of a number of breeds (including several large breeds) that are classified as boar hounds, and that only thirty percent of boar hounds live at least twelve years. In that case, I have competing frequency information for two relevant sets, where neither set is narrower than the other. Moreover, inference using the two sets leads to mutually contradictory conclusions. My information about the set of small breed dogs supports the conclusion that there is a seventy percent chance that Flint will make it to his twelfth birthday, and my information about the set of boar hounds supports the conclusion that there is a thirty percent chance that Flint will make it to his twelfth birthday.

The sorts of inferences that we make when we form probabilistic beliefs about singular propositions on the basis of frequency information are called "direct inferences." In many cases where we make judgments of likelihood, we appeal to frequency information and make such inferences, e.g., when we make judgments about tomorrow's weather or about who will win Sunday's game. Similarly, we readily understand many of the probabilistic claims that are presented to us as based on frequency data and as grounded in direct inference. For instance, when a physician tells me that there is a twenty percent chance that my case of influenza will result in pneumonia, I readily understand that this probabilistic claim is meant to be backed by frequency data in the manner prescribed by direct inference.

Hans Reichenbach was one of the first individuals to articulate an explicit theory of direct inference.[39] Reichenbach's theory recommended that one set one's probability that an individual, c, is an element of a *target class*, T, to the frequency of elements of T among R, where R is the narrowest *reference class* containing c for which one has *reliable statistics* about the incidence of elements of T among the set.(Reichenbach 1949, 374)

Where "PROB($c \in T$)" denotes the probability that an object, c, is an element of a set, T, and "freq(T|R)" denotes the relative frequency of elements of a set, T, among a set, R, the core of Reichenbach's theory (with a minimal amount of anachronism) may be encapsulated by two principles:

[39] Many of the elements of Reichenbach's theory are also present in John Venn's book *the Logic of Chance* (1866).

[RDI] Reichenbachian Direct Inference: *A*'s being justified in believing that freq(T|R) = u and that $c \in R$ gives *A* a defeasible reason for believing that PROB($c\in T$) = u.

[RSD] Reichenbachian Subset Defeat: A respective instance of [RDI] is defeated for *A*, if $\exists R'$,v: *A* is justified in believing:
(i) $c \in R'$,
(ii) $R' \subseteq R$, and
(iii) freq(T|R′) = v (and $v \neq u$).

A significant limitation of Reichenbach's theory concerns cases where one has reliable information about the frequency of a given property among two overlapping references classes, and one does not have reliable information about the frequency of the property among the intersection of the two sets. In cases where there is no narrowest reference class, R, about which one has adequate statistics, Reichenbach recommended that one not form a judgment regarding the likelihood that c is in T.(Reichenbach 1949, 375)[40]

In the present chapter, I will set aside the issue of what to believe in cases where one's frequency information supports conclusions that contradict one another. Reichenbach's proposal to suspend belief in such cases will keep us from adopting unreasonable conclusions. For this reason, I will give priority to addressing the other two problems that cropped-up during my reflections on Flint's lifespan. In the following chapter, I will propose a theory that addresses the question of what to believe in cases where one's frequency information supports conclusions that contradict one another.

[40] We could incorporate the present recommendation into [RSD] by rewriting condition (ii) as: "(ii) $R \not\subset R'$".

Two urgent difficulties facing Reichenbach's account of direct inference were papered-over by the proposal that one consider only 'reliable statistics' in formulating instances of direct inference. One problem concerns the role of uninformative frequency information in licensing and defeating direct inferences. I call this problem "the Problem of Uninformative Statistics". A further difficulty that Reichenbach did not address involves the use of gerrymandered reference and target classes, and is similar (at least cosmetically) to the problem of induction that Goodman uncovered. I call this problem "the Problem of Relevant Statistics".[41]

In attempting to remedy the two *urgent* problems of direct inference, I will be guided by a common sense understanding of the justificatory basis of direct inference. This common sense understanding derives from reflection on the nature of the 'inferential jumps' prescribed by principles of direct inference, and on the background presuppositions that underlie such jumps. In particular, when one makes a direct inference about an object, c, one assumes that c is as likely to be a member of a given target class, T, as a random element of the proposed reference class, R (where we understand a random element of R to be an object that is minimally described as an element of R that was selected by a process which is equally likely to produce each element of R). In cases where a given instance of direct inference is permissible, the conclusion that c is as likely to be in T as a *random element of R* will be justified *only if* c is in relevant respects *indiscernible* from the other

[41] Here I am self-consciously alluding to the concluding sections of Bacchus (1990), where Bacchus calls for a "Theory of Relevance" that will allow us to recognize and, consequently, bar misleading statistical statements from use as premises in direct inference.

elements of R (from the standpoint of the person who is making the direct inference). Similarly, if an agent is aware of features of an object of interest, c, so that c is in relevant respects *discernible* among the other elements of a proposed reference class, R (with respect to being a member of a target class, T), then a statistical statement that features R as its reference class should not be used by the agent to draw a conclusion about the probability that c is an element of T.[42]

The long held tenet that there is a default preference for narrower reference classes may be regarded as a partial implementation of an indiscernibility condition. Indeed, in cases where a direct inference is defeated by appeal to frequency information for a narrower reference class the fact that the inference is defeated may be attributed to a failure of indiscernibility. For example, as a basis for making a judgment about Flint, I initially know:

[1] 40% of dogs live at least twelve years.
[2] 70% of small-breed dogs live at least twelve years.
[3] Flint is a dachshund (and thus a small breed dog).

In this case, I am not entitled to suppose that Flint is *indiscernible* among the set of dogs, because Flint's membership in the set of small breed dogs relevantly distinguishes Flint from the members of the set of dogs that are not also members of the set of small breed dogs.

[42] I will not attempt to provide an analysis of the conditions under which the present sort of indiscernibilty obtains. Rather, in the course of the chapter, I will spell out some *necessary* conditions for *relevant indiscernibility* that may be used to address some of the classic problems of direct inference.

As it turns out, it is possible to remedy the two urgent problems of direct inference by implementing the prescription that respective objects of interest be relevantly indiscernible among elements of proposed reference classes. The particular measures that I propose as the means to addressing the two urgent problems take the form of modifications to Reichenbach's principles [RDI] and [RSD].

An additional amendment that I will make to Reichenbach's theory (as expressed by [RDI] and [RSD]) is in regard to the type of statistical statements that are appropriate to serve as the statistical premises for direct inference. I will adopt Faheim Bacchus's proposal that it is statements of expected frequency that play the key role.[43] As proposed by Bacchus, the prescription that it is statements of expected frequency that serve as the statistical premises for direct inference is intended as a remedy to the Problem of Uninformative Statistics. After I have described the reasons why Bacchus's proposal does not remedy the Problem of Uninformative Statistics (and presented my approach to the Problem of Uninformative Statistics), I will present alternate reasons for accepting Bacchus's proposal.

Before going on to discuss the two urgent problems of direct inference, it is essential to stress a simple, yet non-trivial, feature of direct inference, namely: the manner in which given reference and target classes are described makes a difference regarding whether corresponding direct inferences are permissible. We will see more about this a little later.

[43] Bacchus's proposal is a little more general than this. I will ignore the details of Bacchus's system that allow for greater flexibility, as those details will not extricate Bacchus from any of the criticisms presented here.

For now, it is enough to see that the manner in which a set is described may determine (by logical implication) whether it is reasonable to believe that a given object *is* or *is not* an element of a proposed reference class, and, in turn, whether a proposed direct inference that employs a given statistical statement is permissible (*inasmuch as* an object which is the subject of a proposed direct inference must be an element of the proposed reference class).[44] As a very simple solution to this issue, I will assume that an agent's thoughts regarding respective reference classes, target classes, and objects of interest are always mediated by associated descriptions, and that the admissibility of any particular direct inference is determined relative to the descriptions employed by an agent in formulating the inference.[45]

2. The Problem of Uninformative Statistics

The case of the unit set reference class is a paradigmatic example of the Problem of Uninformative Statistics. A solution to the difficulty must accomplish two things. First, a solution must explain why frequency information about such reference classes may not

---

[44] For example, where c is an element of R, the names "R" and "R∪{c}" are coextensive. However, while we may suppose that any rational agent should believe that *c is an element of R∪{c}*, it may be perfectly reasonable to believe that *c is not an element of R*, in situation where c is an element of R (so that R = R∪{c}).

[45] In supposing here that an agent's thoughts regarding respective reference classes, target classes, and objects of interest are always mediated by associated descriptions, it is not my intention to propose a substantive thesis about the nature of thought or language. But it seems uncontroversial that in the course of ratiocination an agent's thoughts about respective sets are *mediated* by some sort of mental representations. The choice to suppose that thoughts about sets are mediated by descriptions is made for convenience sake, and as a placeholder for a more sophisticated account of mental representation.

generally serve as a basis for direct inference. Second, a solution must explain why frequency information about such reference classes does not ordinarily undermine (via a principle such as *Reichenbachian Subset Defeat*) direct inferences that are based on informative frequency information for broader reference classes.

Kyburg's proposed remedy to the Problem of Uninformative Statistics was the first to appear in the literature, and has two parts.[46] First, Kyburg proposes that the only statements that may serve as the statistical premises for direct inferences are statements that describe a relevant frequency as residing within an interval. Second, Kyburg maintains that direct inference based on statistics for a broader reference class is defeated by statistics for a narrower reference class *unless* the statistics for the broader reference class are more *precise* (i.e., unless the interval associated with the frequency for the broader reference class is a proper superset of the interval for the frequency for the narrower reference class).

Kyburg's approach promises to thwart the Problem of Uninformative Statistics by converting frequency data that is uninformative into frequency data that does not play a role in direct inference. For example, the statement that a given relative frequency is in the set {0, 1} is transformed into the statement that the relative frequency is in the interval [0, 1]. Kyburg's approach dissolves the problem associated with unit set reference classes,

[46] See Kyburg 1974, and Kyburg and Teng 2001.

since the interval [0, 1] will be less precise than any interval associated with a statistical statement that we would like to use as a premise for direct inference.

A variety of proposal that is similar to Kyburg's maintains that the statistical statements that may serve as the statistical premises for direct inferences are not frequency statements at all, and that frequency statements *cannot* serve *directly* to defeat direct inferences. To achieve the desired effect, the present proposal is combined with the proposal that statistics regarding a narrower reference class may defeat statistics regarding a broader reference class *only if* the instances of direct inference supported by statistics for the two reference classes yield inconsistent conclusions. Both Pollock (1990) and Bacchus (1990) have made proposals of the present sort. According to Pollock, it is statements of *nomic probability* that serve as the statistical premises to instances of direct inference. Bacchus maintains that it is statements of expected frequency that serve as the relevant statistical premises. In establishing the viability of their proposals, both Pollock and Bacchus show (1) that statistical statements of the preferred sort (statements of nomic probability and expected frequency, respectively) may take on values other than *one* and *zero* in the case of unit set reference classes, and (2) that the use of known frequencies in the course of direct inference is usually permitted, since point-valued frequency statements usually entail a respective preferred statistical statement whose value is identical to the value of the respective frequency statement.[47]

[47] Section 4 of the present chapter includes a discussion of the nature of expected frequencies and some the deductive relations that hold between frequencies and expected frequencies.

A problem with past approaches to the Problem of Uninformative Statistics can be illustrated by a simple example.[48] Indeed, suppose that one is presented with following information, from a source which one knows is ultra-reliable:

[1] 51% of the balls held in urns produced by the *Ace Urn Company* are red.
[2] *b* is a ball held in an urn produced by the *Ace Urn Company*.

In the case where we lack further substantive information regarding *b*, theories of direct inference prescribe that we assign the probability 0.51 to the proposition that *b* is red (and the conclusion that the probability is 0.51 that *b* is red is reasonable).

But now suppose that one has additional information regarding the specific urn, U_b, that contains *b*. In particular, suppose that one is able to inspect the contents of U_b under conditions that allow one to count the number of balls that are in U_b, and the number of balls in U_b that are white. As a result, one determines that U_b contains exactly one hundred balls, and exactly forty nine white balls. Since one has no reason to think that U_b does not contain some balls that are neither white nor red, one is able to form the justified belief that the frequency of red balls among U_b lies in the interval [0, 0.51].

Faced with the present example, the theories of Kyburg, Pollock, and Bacchus all agree that one's judgment regarding the probability that *b* is red should not change, and each theory allows one to draw the conclusion that the probability that *b* is red is 0.51.

[48] The example presented here is adapted from an example presented by Mark Stone.(1987)

However, given the results of our inspection of the contents of U_b, it seems probable that *b* is not red. At the very least, one should *not* conclude that it is likely that *b* is red.

We are faced with a dilemma when we consider past approaches to the Problem of Uninformative Statistics. On the one hand, we cannot maintain that statistical information for a narrower reference class defeats a direct inference based on statistical information for a broader reference in all cases where our best estimates of the values of the respective statistical statements *merely* differ. If we require *identity* in our estimates of the values of the respective statistical statements, then all useful instances of direct inference will be defeated by statistical information for narrower reference classes. Indeed, in any case where direct inference allows us to draw a useful conclusion, the range of epistemically possible statistical values for proposed reference class will differ from the range of epistemically possible statistical values for the unit set containing only the given object of interest. On the other hand, we must require more than that our frequency information for the two reference classes yield *consistent* conclusions.[49] If we require only *compatibility* of this modest sort, then we will form an incorrect conclusion in the Ace Urn example, and in similar cases.

We are faced with the problem of determining the sort of *compatibility* that must obtain between a proposed reference class, R, and the subsets of R that are known to contain a given object of interest. Since the *compatibility conditions* that have appeared in the

[49] Similarly, we must require more than that our frequency information for the broader reference class be more precise.

literature so far are each either too restrictive (e.g., Reichenbach) or too permissive (e.g., Kyburg, Pollock, and Bacchus), it seems that we must formulate a new condition whose degree of restrictiveness is intermediate to the ones that have appeared so far. In proposing an intermediate condition, I will take a new approach that does not require comparisons of our statistics for proposed reference classes and its subsets.

The idea for my approach originates from one view of the problem of requiring identity in the judged frequency for a proposed reference class and its subsets. For example, suppose that we know that the frequency of elements of a given target class, T, among a proposed reference class, R, is 0.9. In that case, the requirement of identity *cannot* be satisfied, since identity requires that the cardinality of any relevant subset of R be divisible by ten. Simply: we should not expect identity in the values of the judged frequency of a set and its subsets, since the very fact that the sets are of different sizes will assure that identity does not obtain.

Faced with the preceding fact, I propose that when attempting to determine whether frequency information for a narrower reference class, R′, defeats an instance of direct inference, for a target class, T, based on a broader reference class, R, we must compare our justified beliefs regarding the possible values of the frequency of the elements of T among R′ to our justified beliefs regarding the possible values of the frequency of T among a set, R*, where R* is described only as a subset of R whose size is identical to the size of R′. We then conclude (subject to some qualifications) that an instance of direct inference based on the reference class, R, is defeated *if* the narrowest set of values that we are justified in accepting for the frequency of elements of T among R′ *differs from* the narrowest set of values that we are justified in accepting for the frequency of elements of T among R*.[50,51]

As we will see, the preceding test protects correct instances of direct inference from being defeated by *genuinely* uninformative statistics for narrower reference classes, and also yields the right answer in the Ace Urn example, and in similar examples. Before I explain how the proposal deals with the two sorts of problem, let us take a moment to consider how one may reason about the statistical properties of a subset, R*, of a set, R, when one knows R* only as a subset of R of a given size. For example, if one is justified in accepting that R* ⊆ R, freq(T|R) = 0.6, R has ten members, and R* has five members, then *ceteris paribus* the narrowest set of values that one will be justified in accepting for freq(T|R*) is that the value lies in { 0.2, 0.4, 0.6, 0.8, 1 }. Similarly, if one is justified in accepting that R* ⊆ R, freq(T|R) = 0.6, and R has ten members, and the most precise judgment that one can make about the size of R* is that it has either four or five members, then *ceteris paribus* the narrowest set of values that one will be justified in accepting for freq(T|R*) is that the value lies in { 0, 0.2, 0.25, 0.4, 0.5, 0.6, 0.75, 0.8, 1 }.

---

[50] The reason why the condition's satisfaction is not sufficient for the defeat of a respective instance of direct inference will become evident when we consider the Problem of Relevant Statistics.

[51] The present condition is a simplified version of a more complex condition. More generally, we must require, for all U ⊆ [0, 1] and V ⊆ [0, 1], that A is justified in accepting PROB(freq(T|R′)∈U) ∈V *if and only if* A is justified in accepting PROB(freq(T|R*)∈U) ∈V.

My proposal for dealing with the Problem of Uninformative Statistics (without at the same time failing in the face of the Ace Urn example) is to require of each subset, R′, of a proposed reference class, R, where R′ is known to contain the object of interest, c, that the narrowest set of values that one is justified in accepting for the frequency of elements of T among R′ *not differ* from the narrowest set of values that one is justified in accepting for the frequency of elements of T among R*, where R* is known only as a subset of R whose size is identical to the size of R′. Note that the present condition holds for a reference class, R, and a respective subset of R, R′, *if and only if* one's judgments regarding value of freq(T|R′) reflect no information that is not *deductively implied* by one's prior judgments regarding the value of freq(T|R), and the respective sizes of R and R′. For this reason, it appears that the proposed condition precisely delineates the set of cases where one's judgments about the value of freq(T|R′) are genuinely uninformative to one's judgments regarding the value of PROB(c∈T) in the context of one's judgments about the value of freq(T|R). Moreover, (subject to some qualifications) cases where the narrowest set of values that an agent is justified in accepting for the frequency of elements of T among R′ *differ* from *narrowest set of values* that the agent is justified in accepting for the frequency of elements of T among R* (where R* is known to the agent only as a subset of R whose size is identical to the size of R′) are cases where an object of interest, c, is relevantly discernible among R (assuming that c is an element of R′).—In cases where our judgments regarding R′ and R* differ, we are not entitled to suppose that c is *indiscernible* among the elements R, because c's membership in R′ distinguishes c from the elements of R that are not also elements of R′.

The proposed condition also yields the right answer for the Ace Urn example. In the Ace Urn example, the lowest upper bound, 0.51, corresponding to our estimate of the frequency of red balls among balls that are in U_b differs from the lowest upper bound that we are justified in accepting for the frequency of red balls among the set, R*, where R* is known only as a subset of *the set of balls held in urns produced by the Ace Urn Company*, where the size of R* is identical to the size of the set of balls that are in U_b. In particular, for all values for the size of the set of balls held in urns produced by the Ace Urn Company, we may not deduce that freq(*red balls*|R*) ≤ 0.51 (save where the size of *the set of balls held in urns produced by the Ace Urn Company* is identical to the size of *the set of balls the are in* U_b).

In the Ace Urn example, the divergence in our estimate of the frequency of red balls among balls that are in U_b and our estimate of the frequency of red balls among R* reveals that our judgment regarding the possible values of freq(*red balls*|*balls that are in* U_b) is more informative than we would expect given only our judgments regarding the size of the set of balls held in urns produced by the Ace Urn Company, our judgments regarding the size of the set of balls that are in U_b, and our judgments regarding the frequency of red balls among the set of balls held in urns produced by the Ace Urn Company. Moreover, our knowledge that *b* is an element of U_b relevantly distinguishes *b* from the members of the set of balls held in urns produced by the Ace Urn Company that are not also elements of U_b.[52]

3. Overly Informative Statistics

A problem related to the one illustrated by the Ace Urn example arises in cases where our frequency information for a given subset of a proposed reference class is overly informative in a way that permits us to draw incorrect conclusions through the application of traditional principles of direct inference (and through the principles proposed by Reichenbach (1949), Kyburg (1974), Pollock (1990), Bacchus (1990), and Kyburg & Teng

[52] One possible concern regarding my proposal has do to with the potential difference between the judged statistics for respective sets R′ and R* that may result from one's possession of minimally informative information regarding the frequency of elements of a target class, T, among R′. For example, suppose that one knows that Bob is a California resident, and one is interested in assigning a probability to the proposition that Bob is a university graduate. As a basis for forming one's judgment, suppose that one knows that 50% of Californian residents have university degrees. In addition, suppose that one possesses the somewhat uninformative information that Bob lives in Oakhurst, and that at least one resident of Oakhurst has a university degree. In that case, one may deduce that the frequency of university graduates among Oakhurst residents is greater than *zero*. However, where R* is described only as a subset of the set of California residents whose size is equal to the size of the set of Oakhurst residents, one may not deduce that the frequency of university graduates among R* is greater than *zero* (assuming that the number of California residents is at least an order of magnitude greater than the number of Oakhurst residents). So it seems, in the present case, that my proposal yields the *unfortunate* result that we should not conclude that the probability is 0.5 that Bob has a university degree on the basis of our statistics regarding California residents.
The present case highlights the fact that relevant discernibility is a matter of degree. Indeed, I think we are correct to conclude that Bob is relevantly distinguished from other Californians by his membership in a set that is known to contain at least one university graduate (given that the set is not gerrymandered). However, in the present example, the degree to which Bob is distinguished is minimal, so that it would be reasonable to ignore the information that distinguishes Bob from other Californians and conclude that the probability is 0.5 that Bob has a university degree. On the other hand, such a maneuver seems inappropriate when we are dealing with smaller reference classes. For example, suppose that a token, t, is the element of a set, R, of twenty tokens, where we know that 50% of the elements of R are red. Now suppose that we also know that t is an element of a two membered subset of R that contains at least one red token (and the subset is not gerrymandered). In such a situation, it seems that it would be unreasonable to ignore our information regarding t's membership in the given two membered subset, and conclude that the probability is 0.5 that t is red.

(2001)). The examples that exhibit this problem are due to Pollock.[53, 54] According to one example,

> [because] the vast majority of birds can fly and because there are many more birds than giant sea tortoises, it follows that most things that are either birds or giant sea tortoises can fly. If Herman is a giant sea tortoise, [statistical syllogism] would give us a reason for thinking that Herman can fly, but notice that this is based simply on the fact that most birds can fly.(Pollock 1990, 84)

The present example causes problems for Pollock's theory, and other theories of direct inference, in cases where our frequency information regarding *creatures that fly* among *the set of sea tortoises* is less precise that than our frequency information regarding *creatures that fly* among *the set of creatures that are birds or giant sea tortoises* (so that instances of direct inference based on the respective reference classes yield *consistent* conclusions).[55]

In the case of Herman, we know that the frequency of *flying creatures* among *creatures that are birds or sea tortoises* is high. One potential means to defeat a direct inference that appeals to such information would have us appeal to frequency information for a reference class that contains Herman and is narrower than *the set of creatures that are birds or sea*

[53] A second example that I will not discuss concerns the inference to the conclusion that a bird with a broken wing is likely to be able to swim the English Channel, by appeal to the statistic that *most birds can fly or swim the English Channel* (which is true in virtue of the fact that most birds can fly). The condition that remedies the case of Herman applies equally to the case of the bird with the broken wing.

[54] Although of the orthodox theories of direct inference that I considered earlier in the dissertation yield the wrong answers in the face of Pollock's example, I did not present the example earlier as exhibiting a flaw of those theories. Indeed, the present example was presented by Pollock as an example of a projectability problem. Since the framers of each of the major orthodox theories of direct inference admit that their theories are inadequate to addressing such problems, it seemed inappropriate to push the example as an objection to the respective theories (in addition to making the point that the theories are not capable of dealing with direct inference based on statistics for gerrymandered reference classes).

[55] If the instance of direct inference based on our statistics for the incidence of *creatures able to fly* among *the set of sea tortoises* allowed us to draw a conclusion that contradicted the conclusion based on our statistics for the incidence of *creatures able to fly* among *the set of creature that are bird or sea tortoises*, then (according to Pollock's theory) the statistics for the former inference would defeat the latter.

tortoises. The reference class that is the obvious candidate is *the set of creatures that are sea tortoises*. But according to the features of Pollock's example (on which the difficulty is premised), our information regarding the frequency of *flying creatures* among *the set of sea tortoises* is very imprecise, so that direct inference based the two reference classes does not yield mutually inconsistent conclusions. However, in addressing the Ace Urn example, we saw that statistics for a narrower reference class may defeat statistics regarding a broader reference class *even if* direct inference using the two reference classes does not yield mutually inconsistent conclusions. So it seems that a remedy to the present problem may come via the articulation of a condition that entails that our information regarding the frequency of *flying creatures* among *sea tortoises* defeats the direct inference based on our information regarding the frequency of *flying creatures* among *the set of creatures that are birds or sea tortoises*.

The condition proposed as a remedy to the Ace Urn example is not adequate to deal with Pollock's example. In the case of the Ace Urn example, our frequency information regarding a narrower reference class is very imprecise, and yet the imprecise data is actually more informative than we would expect. Indeed, in the Ace Urn example, the lowest upper bound for the judged frequency for the narrower reference class is lower than it should be, assuming that we have no special information about the narrower reference class *vis a vis* the incidence of elements of the respective target class. The situation is different in Pollock's example concerning Herman. In the example concerning Herman, we have no, or very little, information regarding the frequency of *flying creatures* among *the set of sea tortoises*.

In the Ace Urn example, we found that an instance of direct inference may be defeated by the statistics for a narrower reference class in cases where the statistics for the narrower reference class are more informative than they should be. On the other hand, in Pollock's example, we find that our frequency information for a narrower reference class (the set of sea tortoises) is not in the least bit informative, while our frequency information for the complement of the narrower reference class (the set of birds) is more informative than it should be. In light of the preceding observation, I propose that an instance of direct inference based on a reference class, R, (for an agent, A, relative to a target class, T, and an object of interest, c) may be defeated by statistics for a subset, R′, of R *if* the narrowest set of values that A is justified in accepting for the frequency of elements of T among R–R′ *differs from* the narrowest set of values that A is justified in accepting for the frequency of elements of T among R*, where R* is described only as a subset of R whose size is identical to the size of R–R′.[56]

The preceding condition (like the condition used to respond to the Ace Urn example) is defensible by appeal to the idea that justified instances of direct inference presuppose that the given object of interest is relevantly indiscernible among a proposed reference class.

[56] As with the condition introduced as a remedy to the Ace Urn example, the present condition is a simplified version of a more complex condition. More generally, we must require, for all $U \subseteq [0, 1]$ and $V \subseteq [0, 1]$, that A is justified in accepting $PROB(freq(T|(R–R′)) \in U) \in V$ *if and only if* A is justified in accepting $PROB(freq(T|R^*) \in U) \in V$.

In the case at hand, our knowledge that Herman is an element of *the set of sea tortoises* relevantly distinguishes Herman from members of *the set of creatures that are birds or sea tortoises* that are not members of *the set of sea tortoises* (relative to the target class, *creatures able to fly*).—This seems pretty obvious. The official *informativeness condition* also applies in this case, since Herman (Herman = c) is a member of *the set of sea tortoises* (*the set of sea tortoises* = R′) and our judgments regarding the possible values of the frequency of *creatures able to fly* (*creatures able to fly* = T) among *the set of birds* (*the set of birds* = R–R′) are more precise than we would expect, given our judgments regarding the size of *the set of creatures that are birds or sea tortoises*, our judgments regarding the size of *the set of birds*, and our judgments regarding the frequency of *creatures able to fly* among *the set of creatures that are birds or sea tortoises*.[57]

Let A be an agent, c be an object of interest, T be a target class, R be a proposed reference class, and R′ be a subset of R that contains c. In that case, the condition appealed to in

[57] In order to demonstrate the applicability of the official *informativeness condition* to the case of Herman, we need to fill in the details regarding our judgments of the size of *the set of birds*, the size of *the set of creatures that are birds or sea tortoises*, the frequency of *creatures capable of flight* among *the set of birds*, and the frequency of *creatures capable of flight* among *the set of creatures that are birds or sea tortoises*. In cases where our judgments regarding the values of the four preceding quantities are such that a statement of the frequency of *creatures capable of flight* among *the set of creatures that are birds or sea tortoises* is inappropriate as a basis for drawing a conclusion regarding the probability that Herman is a *creature capable of flight*, we find that the greatest lower bound that we are justified in accepting for *the frequency of creatures capable of flight* among *the set of birds* is higher than it should be.—In particular, the greatest lower bound that we are justified in accepting for the frequency of *creatures capable of flight* among *the set of birds* is greater than the greatest lower bound that we are justified in accepting for the frequency of *creatures capable of flight* among R*, where R* is known only as a subset of *the set of creatures that are birds or sea tortoises* whose size is identical to the size of *the set of birds*. For example, suppose that we know that there are at least ninety birds, between one and ten tortoises, and the frequency of *creatures capable of flight* among *the set of birds* is at least 0.9. Now if we suppose that we have no information regarding the frequency *creatures capable of flight* among *the set of tortoises*, then the greatest lower bound that we are justified in accepting for the frequency of *creatures capable of flight* among *the set of creatures that are birds or sea tortoises* is 0.81. And, given the present assumptions, the greatest lower bound that we are justified in accepting for the frequency of *creatures capable of flight* among R* is 71/90.

order to address the Ace Urn example requires that the narrowest set of values that A is justified in accepting for freq(T|R′) *not differ* from the narrowest set of values that A is justified in accepting for freq(T|R*), where R* is known to A only as a subset of R whose size is identical to the size of R′. The condition appealed to in order to address Pollock's example requires that the narrowest set of values that A is justified in accepting for freq(T|R–R′) *not differ* from the narrowest set of values that A is justified in accepting for freq(T|R*), where R* is known to A only as a subset of R whose size is identical to the size of R–R′. When both of the preceding conditions hold for a reference class, R, and one of its subsets, R′, I will say that R and R′ are *i-compatible* (or *informativeness-compatible*) for the agent, A, and the target class, T.

Described as a modification to Reichenbach's theory as expressed by [RDI] and [RSD], my proposed remedy to the Problem of Uninformative Statistics is to replace (iii) of [RSD] (that currently requires that the one's statistics for R and R′ be identical), so that we require only that R and R′ be *i-compatible* relative to a respective agent and target class.

4. Expected Frequencies

Bacchus's proposal that it is statements of expected frequency that serve as the statistical premises to direct inferences does not lead to an adequate remedy to the Problem of

Uninformative Statistics. Nevertheless, Bacchus's proposal has considerable merit. Since Bacchus did not detail the merits of his proposal, I will do so here.

In probability theory, a random variable is identified with a range of possible numeric values corresponding to the possible outcomes of a trial. In turn, random variables may be assigned an expectation (or expected value). The expected value of a random variable is simply the sum of the possible values of the random variable multiplied by the probabilities of the respective values. In general, the probability of a proposition may be identified with the expectation of the proposition's truth-value, where being *true* is identified with the value *one*, and being *false* is identified with the value *zero*. Similarly, one may speak of the expected value of a relative frequency. Here the expectation is identified with a weighted average of the set of possible values of the relative frequency. As a special case, probability statements regarding singular propositions are equivalent to statements of expected relative frequency regarding unit set reference classes.

Before explaining the virtues of the idea that it is statements of expected frequency that serve as the preferred statistical premises to direct inferences, let me explain why the use of frequencies in direct inference consists of a special case of the use of expected frequencies. By doing so, I will discharge the demand to explain the manner in which frequency statements are relevant to direct inference.

In general, if one knows only a set of possible values for a relative frequency, then one's best estimate of the expectation of the relative frequency will be that the expectation lies within the narrowest interval that covers the range of possible relative frequencies.[58] Moreover, in circumstances where a set of possible values is assigned to a given relative frequency, upper and lower bounds on the possible values of the expectation of the relative frequency may be easily calculated, by appeal to the following theorem. (I will use the notation 'E[freq(T|R)]' to denote the expectation of the relative frequency of T among R.)

Theorem:

$$\forall T, R, S, U\text{: PROB}(\text{freq}(T|R) \in S) = 1 \;\&\; U \text{ is the smallest interval such that } S \subseteq U \Rightarrow E[\text{freq}(T|R)] \in U.$$

The present theorem illustrates the relevance of frequency information to direct inference inasmuch as it describes an important deductive relationship between frequencies and expected frequencies, and thereby accounts for the use of point-valued and interval-valued frequency statements in the course of direct inference. For example, note the implication between PROB(freq(T|R) = r) = 1 and E[freq(T|R)] = r.

On the proposal that it is statements of expected frequency that serve as the basis for direct inference, direct inference may be conceived of as the 'transfer' of the expected value of the relative frequency of a property among a set to an expected value of the relative frequency of the property among an element of the set. Such 'transfers' of expectation are reasonable so long as one does not possess better information upon which to base one's

[58] For this reason, Kyburg's proposal to restrict the frequency statements that may serve as the statistical premises for direct inference to statements that describe the value of a frequency as within an interval makes some sense.

estimates of the likely properties of the given object. Moreover, since statements of probability regarding singular propositions are equivalent to statements of expected relative frequency regarding a unit set reference class, it appears reasonable to regard statements of expected relative frequency as the basis for direct inference, since the values that are transferred are values of exactly the same function. Even so, the main reason for treating expected frequencies as the basis of direct inference is connected to the intuitive justification of direct inference.

In making a direct inference, one assumes that the object about which one is reasoning, c, is as likely to be a member of the respective target class, T, as a *random element* of the proposed reference class, R. On the assumption that c is as likely to be in T as a *random element of R*, one is obliged to conclude that the probability that c is in T is equal to the frequency of elements of T among R, in cases where one is aware of the value of this frequency. For parallel reasons, one is obliged to conclude that the probability that c is in T is equal to the expectation of the frequency of elements of T among R.

The expected value of a relative frequency derives from an assignment of probabilities to the relative frequency's possible values, and the application of expected frequencies in direct inference trades upon the manner in which expected frequencies encode information about the probability of various respective frequencies. In the case where one is aware of an assignment of probabilities to the values of a given relative frequency, one should calculate the probability of a random element of R being in T by considering the chances that freq(T|R) takes on respective values. The best way to see this point is to imagine the situation as a two-tiered lottery, where, first, the frequency of elements of T among R is selected and, next, an element of R is selected at random. Since the expected value of a relative frequency simply encodes a weighting of respective relative frequencies according to probability, expected frequencies may be used in place of a corresponding assignment of probabilities to the values of a given relative frequency. Moreover, on the assumption that each element of a set, R, is equally likely to be selected, the likelihood that the selected element of R is an element of T is equal to the expected frequency of elements of T among R.

In accordance with the preceding observations, I will now assume that proper instances of direct inference proceed from premises of the form ‘E[freq(T|R)] $\in$ U’ and ‘c $\in$ R’ to conclusions of the form ‘PROB(c$\in$T) $\in$ U’. But given the deductive relations between statements of frequency and statements of expected frequency, it is usually permissible to formulate instances of direct inference using frequency statement, and I will occasionally do so.

5. Direct Inference and the Foundations of Induction

In *Chapter Two*, I described the proposal that statistical induction may be grounded in direct inference. While I expressed my approval for the proposal, I also described an apparent limitation of the proposal, and described the attempt of Kyburg to deal with the limitation by the introduction of an *acceptance theory*. At that time, I expressed my aversion to Kyburg's acceptance theory, and my aversion to acceptance theories, in general. I also promised, at that time, that I would explain why it was not necessary to appeal to an acceptance theory in order to address the limitation. I am now prepared to fulfill that promise.

The proposal that direct inference may be used to ground statistical induction appeals to certain mathematical facts regarding the frequency with which subsets of a set will resemble the set itself, in the frequency of any given property. The following theorem, which holds for any pair of sets A and B, is sufficient to illustrate the basic idea:

[NN] For every $\delta,\gamma > 0$, there is an *n* such that *if B* is a finite set containing at least *n* members, *then* $\mathrm{freq}(\, \mathrm{freq}(A|X) \approx_{\delta} \mathrm{freq}(A|B) \mid X \subseteq B\,) > 1-\gamma$.(Pollock 1990)[59]

The basic idea expressed by [NN] is that for any set of sufficient size, the vast majority of the subsets of the set will be very similar to the set itself in the frequency of any given property.

According to the sort of proposal advocated for by Kyburg, the statistical consequences of [NN] are ripe for use in instances of direct inference which proceed from a premise about the frequency with which observed samples (subsets) are similar to the populations from

[59] "$r \approx_{\delta} s$" means that *r* differs from *s* by no more than δ.

which they are drawn, and leads to the conclusion that it is likely that our actual observed sample is likely to be similar to the population from which it was drawn in the frequency of any given property. Let us review the details of this line of reasoning.

Let Ω be a sufficiently large population about which we wish to draw a conclusion via statistical induction, and let *s* be our observed sample. Now, let τ be the set of sets such that a set is an element of τ *just in case* the frequency of elements of ϕ among the set is similar (to some specified degree) to the frequency of elements of ϕ among Ω.

In the present case, [NN] tells us that $\text{freq}(\tau|\rho) \approx 1$, and so it seems that we may express an instance of direct inference of the following form:

$$\begin{array}{l} \text{freq}(\tau|\rho) \approx 1 \\ s \in \rho \\ \hline \text{PROB}(s \in \tau) \approx 1 \end{array}$$

Now the preceding form of inference allows us to conclude that, with high probability, our observed sample will resemble the population from which it is drawn in the frequency of objects that are elements of ϕ. In turn, we may use deduction to draw a conclusion about the likely incidence of ϕ among the population from which our sample was drawn. That is, from the premise that $\text{PROB}(s \in \tau) \approx 1$ and the premise $\text{freq}(\phi|s) = r$, we may deduce that $\text{PROB}(\text{freq}(\phi|\Omega) \approx r) \approx 1$.

The limitation to the preceding method of grounding statistical induction is that we can do no better than assigning a high probability to a statement about the incidence of a given property (represented by, the set, ϕ in the preceding example) among the population about which we are reasoning (Ω in the preceding example). This *appears* to present a problem if we wish to use the statistical statements generated by statistical induction as the statistical premises of further instances of direct inference, since direct inference itself appears to depend on the use of unqualified statistical statements (i.e., statements that are not qualified as merely *highly probable*). As such, it appears that the statistical claims generated in the proposed manner could not themselves be used as premises in direct inference. For instance, following the present example, it looks as though the conclusion $\mathrm{PROB}(\mathrm{freq}(\phi|\Omega) \approx r) \approx 1$ cannot be used *as is* in direct inference to draw the conclusion that $\mathrm{PROB}(c \in \phi) \approx r$ (where c is an element of Ω).

Kyburg's remedy to the preceding problem is to introduce an acceptance theory, which simply tells us that when the probability of a proposition is sufficiently high, we are entitled to treat the proposition as a 'practical certainty'. For reasons that I described in *Chapter Two*, I think that acceptance theories are deeply problematic, and that Kyburg's proposed remedy is unacceptable.

My alternative to Kyburg's proposal derives from the idea that it is expected frequencies and not frequencies that play the central role in direct inference. In that case, it is not a problem that we cannot use statistical induction (grounded in direct inference) to derive unqualified frequency statements. Indeed, what is required for direct inference is only that it is possible to derive unqualified statements of expected frequency, and that is possible. In particular, the same sort of reasoning that underlies [NN] may be used to establish that the expected frequency of a given property lies within a small interval. That this is so follows immediately from the definition of expected frequency, and is best understood by focusing one's attention on the fact that expected frequencies consist in an average of the various possible values of a given frequency weighted by the probabilities of the respective values. In that case, we see that elementary arithmetic can be used to reason from the statement that a given frequency lies within a given interval with a high probability, r, to the conclusion that the expected value of the given frequency lies in a similar interval. For example, if we know that a given frequency lies in the interval [0.49, 0.51] with probability, 0.95, we may deduce that the expected frequency lies in the interval [0.4655, 0.5345]. The lower bound of the interval [0.4655, 0.5345] may be calculated by assuming that the probability is 0.95 that the frequency is 0.49, and the probability is 0.05 that the frequency is 0. Similarly, the upper bound of the interval [0.4655, 0.5345] is calculated by assuming that the probability is 0.95 that the frequency is 0.51, and the probability is 0.05 that the frequency is 1. It is obvious how we may generalize the preceding inferences, though in most cases we will have sufficient data to establish that the expected frequency lies within a much narrower interval than in the present example.

6. The Problem of Relevant Statistics

Recall the case of Flint, where I am justified in accepting the following propositions:

> [1] 40% of dogs live at least twelve years.
> [2] 70% of small-breed dogs live at least twelve years.
> [3] Flint is a dachshund (and thus a small-breed dog).

On the supposition that the preceding three propositions encapsulate my knowledge of the factors that are relevant to judging the probability that Flint will live at least twelve years, it seems that I should form the belief that the probability is 0.7 that Flint will live at least twelve years. In order to justify such a conclusion, it is typical to appeal to a principle that tells me to prefer frequency information for narrower sets when I have relevant frequency information for two or more sets. But recall that the story does not end there, for it is possible to formulate a gerrymandered reference class R_G (where R_G is formed from Flint along with all of the small breed dogs who will not live twelve years), so that R_G is narrower than the set of small breed dogs, and Flint is an element of R_G. The problem, then, is that the proportion of the elements of R_G that will not live twelve years is guaranteed to be very near to one. Indeed, if we suppose that there are only one hundred dogs, then freq(creatures-that-will-live-twelve-years|R_G) $\in$ {0, 1/31}. The problem with which we are faced is that of explaining why I am permitted to conclude that the probability that Flint is a small-breed dog is 0.7, and I am not permitted to conclude that the probability that Flint will live twelve years is in the interval [0, 1/31].[60]

---

[60] The present example is similar to one that is described in Bacchus et al (1996), but Henry Kyburg (1974) was the first to publish examples that illustrated difficulties of the present sort.

I call the present problem “the Problem of Relevant Statistics” with the idea being that certain statistical statements are not relevant to direct inference. My solution to the present difficulty will take the form of criteria for determining when a statistical statement is irrelevant to the determination of the probability of a singular proposition.

7. Criteria for when Statistics are Relevant

Past proposals regarding the Problem of Relevant Statistics (e.g., the proposals of Pollock (1990), and Kyburg & Teng (2001)) have been extremely programmatic, and have simply proposed that the rules of direct inference must include projectability constraints (akin to the ones that are thought to be needed in the case of statistical induction). I offer an alternative approach to the problem that derives from a natural understanding of the justificatory basis of direct inference.

The correct explanation of what goes wrong in the case of Flint and the gerrymandered reference class, R_G, flows from the assumptions that underlie justified instances of direct inference. Recall that in making a direct inference, one assumes that the object about which on is reasoning, c, is as likely to be a member of the respective target class, T, as a *random element* of the proposed reference class, R. In cases where direct inference is used correctly, the conclusion that c is as likely to be in T as a random element of R will be justifiable by appeal to the fact that c is in relevant respects indiscernible from the other elements of R.

In the case of Flint, the inference to the conclusion that Flint is as likely to live twelve years as a randomly selected small breed dog is not defeated by the statistical fact that a very high proportion of things that are elements of the gerrymandered reference class, R_G, will not live twelve years. Similarly, we are not permitted to make a direct inference using frequency information for R_G to draw the conclusion that it is probable that Flint will not live twelve years. In the case of Flint and R_G, the defeasible presumption in favor of narrower reference classes is superceded, because Flint is relevantly discernible among R_G (relative to the property of being a creature that will live at least twelve years). A relevant difference, in this case, is demonstrable from the fact that our narrowest estimate of freq(creatures-that-will-live-twelve-years |{Flint}) is that the relative frequency is in {0, 1}, but our narrowest estimate of freq(creatures-that-will-live-twelve-years |R_G −{Flint}) is that the relative frequency is 0. In other words, we are aware of a relevant difference between Flint and the other elements of R_G.

There is one feature that is characteristic of all the examples that have appeared in the literature to illustrate the 'projectability' problems associated direct inference. In each example, the reference or target class for the key statistical premise is formulated using a description that is known to pick out a proper subset of the proposed reference class that is known to contain the object of interest. Through the use of such a description the value of the key statistic is computed via a reference class that is *gerrymandered* relative to the given target class. In particular, the value of the statistical statement is computed by appeal to the sizes and statistical values for its subsets, where one of the subsets is known to contain the object of interest. For example, in the case of Flint, the range of possible values for the frequency of creatures that will live twelve years among the gerrymandered set, R_G, is computed by the possible values for the frequency of creatures that will live twelve years among R_G−{Flint}, and by the possible values for frequency of creatures that will live twelve years among the unit set of Flint. Where "L12" stands for the set of creatures that will live twelve years, and "SB" stands for the set of small breed dogs (so that $R_G = (SB \cap \sim L12) \cup \{Flint\}$), the computation proceeds by cases:

> If Flint $\notin$ L12, then
> $freq(L12|(SB\cap\sim L12)-\{Flint\}) = 0/29$,
> $|(SB\cap\sim L12)-\{Flint\}| = 29$,
> $freq(L12|\{Flint\}) = 0/1$, and
> $freq(L12|(SB\cap\sim L12)\cup\{Flint\}) = 0/30$.
>
> If Flint $\in$ L12, then
> $freq(L12|(SB\cap\sim L12)-\{Flint\}) = 0/30$,
> $|(SB\cap\sim L12)-\{Flint\}| = 30$,
> $freq(L12|\{Flint\}) = 1/1$, and
> $freq(L12|(SB\cap\sim L12)\cup\{Flint\}) = 1/31$.
>
> Therefore, $freq(L12|(SB\cap\sim L12)\cup\{Flint\}) \in \{ 0/30, 1/31 \}$.

It is easy to see why the preceding sort of gerrymandering violates the indiscernibility condition that is tacitly assumed when we use direct inference. In the computation just described, Flint is treated separately from other elements of the proposed reference class, so that Flint is literally *discerned* from other elements of the reference class in the chain of reasoning that leads to our judgment regarding the possible frequency values for the proposed reference class.

To remedy the problem associated with gerrymandered statistics, we must restrict the use of descriptions that are known to pick out a proper subset of the proposed reference class that is known to contain the object about which we are reasoning. Now, in 'real life', an agent may allow all sorts of extraneous descriptions to appear in her computation of the value of a given statistical statement. Since we do not wish the results of a theory of direct inference to depend on accidental features of an agent's computation of a given statistical statement, we should not restrict the use of any particular description in the course of reasoning. Indeed, rather than concern ourselves with the actual descriptions that an agent employs in the computation of a given statistic, we require only that the agent could have justified her conclusion through a chain of reasoning that does not rely on a problematic description.

In keeping with my characterization of the *problem in view* as "the Problem of Relevant Statistics", I will describe cases where a statistic is gerrymandered (in problematic way) as cases where the statistic is *irrelevant* to the probability that a given object is a member of a respective target class:

[IR] $E[freq(T|R)] \in V$ *is irrelevant* to the value of $PROB(c \in T)$ for an agent, A, *just in case* for all chains of inference, C, sufficient for justifying A's belief that $E[freq(T|R)] \in V$, $\exists R'$:

(i) R' is describable using only vocabulary used in the course of C,[61]
(i) A is justified in believing that $R' \subseteq R$,
(ii) A is justified in believing that $c \in R'$,
(iii) R and R' are *not i-compatible,* relative to A and T.

The present notion of *irrelevance* is obviously quite narrow, and does not encompass all of the cases where a statistic may rightly be called "irrelevant" to the value of a respective probability. Rather the present notion of *irrelevance* is technical and applies to a narrow range of cases where a statistic for a given reference class is computed through the values of statistics for its subsets, and where the object of interest is known to be a member of one of the subsets whose statistics are relied upon in the computation.

The prescription that *irrelevant statistics* not be used in the course of direct inference (or to defeat instances of direct inference) makes sense in light of the precept that direct inference relies on the relevant indiscernibility of the object of interest from the other elements of the reference class associated with a respective statistical statement. Indeed, in cases where a statistic is deemed to be *irrelevant*, we know that the object of interest is discernable from other elements of the proposed reference class, R, in the course of the respective agent's reasoning about the value of statistics regarding R.

---

[61] The sort of logical combinations permitted in the description of R' would depend on the richness of the vocabulary used in the course of respective chains of inference. Assuming that chains of inference resemble proofs in a first order language, it is intended that R' be describable by a first order formula, φ, with a single free variable x. We may then regard R' as the set of objects that satisfy $\varphi(x)$.

8. Conclusion

With the notion of an *irrelevant statistic* and the notion of *i-compatibility* in place, I now propose some fairly traditional looking principles of direct inference. The principles incorporate three amendments to Reichenbach's theory of direct inference:

[DI] Direct Inference: A's being justified in believing that $c \in R$ and that $E[freq(T|R)] \in V$ gives A a defeasible reason to believe that $PROB(c \in T) \in V$, so long as $E[freq(T|R)] \in V$ is not *irrelevant* to value of $PROB(c \in T)$ for A.

[SD] Subset Defeat: A respective instance of [DI] is defeated for an agent, A, if $\exists R'$:
(i) A is justified in believing that $R' \subseteq R$,
(ii) A is justified in believing that $c \in R'$,
(iii) R and R′ are *not i-compatible*, relative to A and T, and
(iv) $E[freq(T|R')] \in U$ is not *irrelevant* to value of $PROB(c \in T)$ for A, where U is the narrowest set of values that A is justified in accepting for $E[freq(T|R')]$.

I do not maintain that [DI] and [SD] constitute of a complete theory of direct inference. One particular area which calls for further work concerns cases where one's frequency information supports conclusions that contradict one another. I will address this problem in the following chapter.

CHAPTER 7: THE KINETIC THEORY OF COMPETING REASONS

PART I: The Kinetic Analogy

1. Introduction

To say that one has a *reason* for believing a proposition, p, is to say that one is apprised of facts which, considered in isolation, recommend belief in p. We say that such a reason is *defeasible* if a wider consideration of the facts could lead to the defeat of one's grounds for believing p, otherwise we say that the reason is *indefeasible*. A *theory of competing reasons* is simply a theory (or a set of principles) which tells one what conclusion(s) to adopt in cases where one has defeasible reasons for accepting mutually inconsistent conclusions.

Paradigm examples of competing reasons occur in the course of direct inference when an individual has statistical data regarding multiple overlapping references classes, but no significant data regarding their intersection. For example, suppose that one is informed that c is an element of a large urn, U, which contains a large number of coin-shaped tokens. Now say that one learns that each of the tokens in U is (1) either blue or red, (2) either plastic or glass, and (3) either 10-edged or 12-edged. In addition, suppose that one learns that 85% of the plastic tokens are red, 75% of the 12-edged tokens are not red, but one does not possess informative statistical data regarding the frequency of red tokens among tokens that are plastic and 12-edged.[62] Now, suppose that one discovers that c is plastic and 12-edged. The fact that c is plastic gives one a defeasible reason to believe that *c is probably red*, and the fact that c is 12-edged gives one a defeasible reason to believe that *c is probably not red*. The present case is one where different components of one's body of evidence support mutually inconsistent conclusions, and no answer about what to believe is dictated by traditional principles of direct inference. This is a case where a broad examination of the evidence leaves us with a view of 'opposing evidence' and competing reasons for belief. This is the sort of case to which a theory of competing reasons could be applied. Such a theory would provide relief to the Problem of Competing Statistics that was introduced in *Chapter One*.

A kinetic theory of competing reasons is a theory of competing reasons upon which the features of individual reasons are in some way conceived according to an analogy with the physical characteristics, and competing reasons are conceived of as physical objects possessed of opposing physical characteristics. The applicability of a kinetic theory of competing reasons will depend on the source and character of respective competing reasons. I maintain that a kinetic theory is applicable to a wide range of cases where competing reasons arise in the course of direct inference, such as the one described in the previous paragraph.

[62] In some cases, an agent may have sufficient data to draw a defeasible conclusion about the expected frequency of red tokens among the tokens that are both plastic and 12-edged. While it may be that direct inference, itself, will license such inferences in some cases, it appears that there are many cases where such an inference is not possible (save for uninformative conclusions which place the expected frequency which interests us within a very large interval).

Part I of the present chapter prepares the way for the kinetic theory of competing reasons that is proposed in *Part II*. The preparations in *Part I* take the form of a kinetic analogy along with two conditions whose satisfaction is intended to ensure the applicability of the analogy. With the analogy and the two conditions in hand, it is possible to delineate a range of cases that may arise in the course of direct inference, and argue that the two conditions are satisfied in such cases. After delineating a range of cases to which the analogy is applicable, I articulate four further principles that describe the dynamics of defeasible reasons in cases where the Kinetic Analogy applies. The four conditions are applicable in corroborating the theory that is proposed in *Part II*, since the Kinetic Theory satisfies the four conditions.

Throughout the course of the chapter, I will have need to refer to the proposition (or propositions) for which a given reason (or set of reasons) provide reason to accept. I will refer to such propositions as *the content(s)* of the given reasons, so that when one has a reason to belief p, p is the content of one's reason.

2. The Analogy

In proposing a kinetic theory of competing reasons, I will employ an analogy upon which reasons for propositions are treated as analogous to voluminous objects. In turn, the demand for consistency among the propositions that one believes is represented by a constraint that the voluminous objects (representing one's reasons) be made to occupy a space of fixed size. According to the analogy, each object/reason possesses a usual volume and a *relatively permeable membrane*. The membrane of each object/reason is *relatively permeable* in the sense that some of the objects/reasons may co-locate as facilitated by the reciprocal permeability their membranes and others may not (in accordance with whether the respective objects correspond to reasons for consistent propositions). According to the analogy, each object/reason has a natural volume, and is relatively resistant to compression. Yet an object/reason may be subject to compression in the case where it is forced to occupy a space with other objects/reasons with which it cannot share location (due to inconsistency).

Once I have outlined some cases where the analogy is applicable, the analogy will be used to guide our intuitions regarding the problem of arbitrating between competing reasons. The operative mechanism for the arbitration of competing reasons will be that of compression. According to the analogy, compression of an object below its natural volume corresponds to the revision of the contents of one's reasons.

The Kinetic Analogy represents an attempt to moderate the received view concerning what one should believe in cases where one has evidence that provides defeasible reasons for conclusions that are mutually inconsistent. The received view is that one should adopt a 'skeptical policy' in dealing with defeasible reasons for contradictory conclusions. The *skeptical policy* states that, in cases where one has defeasible reasons for accepting each element of a minimal inconsistent set of propositions (i.e., an inconsistent set that has no inconsistent proper subsets), one should refrain from accepting any element of the set, unless one has undercutting evidence that defeats one's reason(s) for an element of the set.[63] The Kinetic Analogy suggests a more credulous response to competing reasons.

3. Sufficient Conditions for Application of the Analogy

Both the Kinetic Analogy and the theory that I propose in *Part II* are intended to apply to a wide range of cases that may arise in the course of direct inference. In service of making this application of the analogy plausible, I now propose that the analogy apply to a collection of reasons in cases where the collection satisfies two conditions. I call the conditions "volumetric representability" and "deductive assessability". I take the satisfaction of the two conditions to be sufficient, but perhaps not necessary, for determining the applicability of the Kinetic Analogy. The two conditions are intended to connect the features of the analogy to features that may be possessed by the elements of a set of reasons. Roughly, the condition of *volumetric representability* is intended to ensure that the elements of a collection of reasons can be 'pictured' via the analogy. On the other hand, *deductive assessability* is intended to ensure that the array of features 'picturable' via

[63] The skeptical policy corresponds to the behavior of *skeptical reasoners*, as discussed by Touretzky, Horty, and Thomason (1987). A version of the skeptical policy is applied in the theory of direct inference articulated by Kyburg (1974), and Kyburg & Teng (2001). An explicit formulation and endorsement of the skeptical policy, for cases where one's defeasible reasons are un-weighted, can be found in Pollock (1990) and Pollock (1995, 62-3).

the analogy exhaust the features that may be relevant to resolving conflicts among a given set of reasons.

I will say that a set of reasons, R, is amenable to *volumetric representation* just in case (1) there is a parameter common to the contents of the elements of R which is appropriately analogous to volume, (2) the values of these parameters determine which elements of R stand in relations of conflict, and (3) moderate downward revision of the values of these parameters (compression) makes sense as a response to competing defeasible reasons, and is a means to resolving conflicts among the elements of R.

For the purposes of grasping my intention in articulating the present condition, it is useful to consider the sort of case where the contents of the elements of R are lower probability bounds on propositions, i.e., the objects of the elements of R are propositions of the form '$\mathrm{PROB}(\alpha) \geq s$', where s is an element of [0, 1]. In such cases, the value of a respective bound is an obvious of analogue of a measure of volume, where a space of possible worlds is an analogue of the fixed space in which the voluminous objects are to be placed. Permeability, in this case, corresponds to deductive relations between the propositions that are the objects of the probability bounds, where equivalent propositions are fully permeable, and inconsistent propositions are fully impermeable. In the case where one has defeasible reasons for inconsistent probability bounds, *compression* makes sense as an analogue of the operation of moderate downward revision of the bounds, and is a means to reaching consistency.

I will say that a set of defeasible reasons, R, is amenable to *deductive assessment* just in case the relations relevant to arbitrating conflict among the elements of R are limited to deductive relations between the contents of the elements of R.

The component conditions of *volumetric representability* are tailored to the Kinetic Analogy, and it is transparent that volumetric representability is *necessary* for the applicability of the Kinetic Analogy. The connection between *deductive assessability* and the Kinetic Analogy is less transparent. In fact, I do not think that deductive assessability is *necessary* for the applicability of the analogy. However, it seems that, in the presence of volumetric representability, deductive assessability is *sufficient* for the applicability of the analogy. Indeed, the Kinetic Analogy invites us to consider a collection of reasons under the guise of a collection of voluminous objects which (in the case of inconsistency) must be compressed in order to fit into a fixed space. The features of a collection of defeasible reasons that the analogy would have us 'picture' are simply the volumes associated with the respective reasons (which correspond to parameters of the contents of the elements of R) and the reciprocal permeability of various objects according to the deductive relations between the contents of the defeasible reasons. That said, in the case where deductive assessability holds, we know that the range of features 'picturable' via the Kinetic Analogy is adequate to represent all of the information relevant to resolving inconsistencies among the contents of a set of competing defeasible reasons.

I will now articulate some conditions under which *deductive assessability* and *volumetric representability* hold for a collection of defeasible reasons, and explain why it is normal for those conditions to be satisfied in the case of defeasible reasons that are generated by direct inference.

4. Conditions for Volumetric Representation

It is normally possible to encapsulate a set of reasons that have been generated by direct inference so that the first two components of volumetric representability are satisfied for the set. Indeed, defeasible reasons that have been generated by direct inference will usually have as their object either a point probability on a proposition or the specification of an interval on the range of possible values for the probability of a proposition.[64] In either case, the content of a respective reason may be decomposed into a pair of reasons for corresponding lower probability bounds, so that a defeasible reason for a proposition of the form $PROB(\alpha) \in [a, b]$ may be represented by a pair of reasons for the propositions $PROB(\alpha) \geq a$ and $PROB(\neg\alpha) \geq b{-}1$.

The *third component* of *volumetric representability* requires that compression of the volumetric parameter associated with the content of a defeasible reason makes sense as a

[64] In general, the basis for an instance of direct inference is an expected frequency (which is itself based on a frequency statement). So long as this expected frequency statement is interval valued (including point valued), the conclusions based on the expected frequency statement will also be interval valued.

response to inconsistency. We can satisfy ourselves that this condition is fulfilled in the case of defeasible reasons for lower probability bounds by considering examples. The clearest examples occur when we are able to apply ‘indifference reasoning’.

According to indifference principles, we should attach equal degrees of ‘credence’ to the elements of a collection of propositions in the case where the evidence for each of the propositions is equivalent (i.e., one should be indifferent in the absence of relevant difference). The application of this principle in the context of the Kinetic Analogy provides us with guidance in choosing a moderate revision of the objects of one’s reasons in a wide range of cases. For example, imagine a situation in which we are limited in our statistical data and in the propositions for which we have reasons. In particular, suppose that by symmetrical instances of direct inference we are limited in drawing the defeasible conclusion that the probability is at least 0.75 that a given object *is red* and the defeasible conclusion that the probability is at least 0.75 that the object *is not red*.[65] In this case, our evidence for the conflicting propositions (*o is red* and *o is not red*) is equivalent, and so indifference reasoning compels us to reach equivalent conclusions regarding the two propositions. One possibility, consistent with the prescriptions of indifference, would be to withhold judgment regarding each of the two propositions, and accept no more for each of the two propositions than that their probability is *greater than or equal to zero.* An alternate possibility is to accept, for each proposition, that its probability is at least 0.5. The latter possibility is the one suggested by the Kinetic Analogy and backed by the third

[65] For the purposes of the example, I am supposing that one has no other evidence that bears on the two propositions.

component of volumetric representability (as setting the respective lower probability bounds to 0.5 consists of a minimal revision of the objects of our original defeasible reasons).[66]

In looking beyond cases where it is possible to apply indifference reasoning, it appears that there may be difficulties in eliciting precise prescriptions of the sort that would be in keeping with the mechanism of compression associated with the Kinetic Analogy. Although the analogy suggests that one should not compress/revise more than is necessary to achieve consistency among the contents of one's individual reasons, the analogy rarely suggests specific prescriptions as to how to make such revisions, and it seems that, in many cases, any procedure for making such a revision will be arbitrary.

In proposing a theory that is guided by the Kinetic Analogy, I will honor the intuition that there are many cases where there is no non-arbitrary way to select a minimal revision of the contents of one's reasons. However, despite limitations in the guidance that the analogy may provide, the analogy will often rule out a range of revisions that are not in keeping with the mechanism of compression suggested by the Kinetic Analogy.

In the case where we are considering defeasible reasons for mutually inconsistent lower probability bounds, the Kinetic Analogy would have us imagine a process of compression/revision to a point where consistency is reached. And, in accordance with the

[66] The proposed revision is minimal in the sense that it is not possible to adjust either of the proposed bounds upward while maintaining consistency.

proposed analogy, it seems that the resistance of the objects/bounds to compression/revision would require that compression/revision be minimal in the sense of not going beyond what is necessary to achieve consistency. To get the idea, consider an example where direct inference gives us a defeasible reason for believing that the probability is at least 0.85 that *the token is red*, and a defeasible reason for believing that the probability is at least 0.75 that *the token is not red*. Now where 0.85 and 0.75 represent the volumes associated with the propositions *the token is red* and *the token is not red* (relative to the space of possibilities represented via the Kinetic Analogy), we assume that we must compress the respective volumes to the point were the sum of their values is less than or equal to 1.0. Moreover, the mechanism suggested by the Kinetic Analogy leads one to expect that the volume associated with the two propositions will occupy the entire space (and the sum of revised values will be 1.0). Despite this seeming consequence of the analogy, the analogy itself does not suggest which values to accept, nor does there seem to be any non-arbitrary principle that will. Yet this lacuna in the prescriptions suggested by the analogy need not cause us despair, since the analogy still provides guidance. For example, in the case described above, the analogy prescribes that the greatest lower probability bound for the proposition that *the token is red* not be lower than 0.25, which is the greatest lower bound that is consistent the original bound, 0.75, on the proposition that *the token is not red*. Similarly, the analogy prescribes that the greatest lower probability bound for the proposition that *the token is not red* not be lower than 0.15, which is the greatest lower bound that is consistent the original bound, 0.85, on the proposition that *the token is red*.

The preceding conclusions are suggestive of some of the non-arbitrary prescriptions suggested by the Kinetic Analogy. There are additional principles that can be used to determine further non-arbitrary prescriptions. Many of these prescriptions proceed from varieties of indifference reasoning. I will explore these prescriptions in detail in *Part II* of the chapter.

5. Conditions for Deductive Assessability

It is normally possible to isolate sets of defeasible reasons that have been generated by direct inference so that deductive assessability is satisfied for the set. Explaining why this is the case will require a detailed explanation. As a means to this explanation, I will appeal to a variant of John Pollock's distinction between *rebutting* and *undercutting* defeaters.(Pollock 1995) According to my variant of Pollock's distinction, there are two ways by which a defeasible reason may be defeated. The most well know cases of defeat occur when the object of a given defeasible reason is inconsistent with other propositions for which an agent has reasons. In cases where inconsistency among the object of one's reasons leads to the defeat of a defeasible reason, we say that the affected reason is subject to *rebutting defeat*. A less well known means to defeat occurs when an agent (1) possesses information that attacks the connection between the agent's evidence for a conclusion and the conclusion itself, or (2) possesses information that attacks the considerations upon which a given conclusion is based. In the case where a defeasible reason is defeated in the preceding manner, one says that the agent's reason for a respective conclusion is "undercut".—The characteristic feature of undercutting defeaters is that such defeaters 'remove' a respective reason for a proposition without providing a reason for the proposition's negation.

My explanation of why of sets defeasible reasons generated by direct inference will often satisfy deductive assessability requires the notion of one set of reasons encapsulating another set of reasons:

A set of reasons, R*, *encapsulates* the complete set of reasons of an agent, A, *just in case* every element of R that is not an element of R* is redundant to the determination of what A should believe.

Obvious cases where a reason for belief, r, is redundant relative to a set of reasons, R, possessed by an agent A, occur when the object of r is a logical consequence of the object of an element, r′, of R, where the fact that A has r is due to the fact that A has r′.

Finally, my explanation requires a notion of one set being deductively independent of another set:

A is *deductively independent* of B *just in case*, for all a_1, a_2, and b, if a_1 and a_2 are consistent subsets of A and b is a consistent subset of B, then
(a) $a_1 \cup b \models \wedge a_2 \Rightarrow a_1 \models \wedge a_2$, and
(b) $a_1 \cup b \models \neg\wedge a_2 \Rightarrow a_1 \models \neg\wedge a_2$.[67]

[67] For all sets, s, $\wedge$s is simply the conjunction of the s's elements.

With these notions in hand, I am prepared to illustrate that the following is a sufficient condition for a set of reasons being deductively assessable:

If (i) R is a subset of a set of reasons, R*, (ii) R* *encapsulates* A's full set of reasons, (iii) the set of contents of R is *deductively independent* of the set of contents of R*–R, and (iv) there is no undercutting defeater for any element of R among the elements of R*, *then* deductive assessability holds for R for A.

That the preceding condition holds can be seen by considering one way of arraying the facts that are sufficient for determining what any agent should believe:

(1) facts as to which propositions the agent has indefeasible reasons (combined with the status of the propositions as such),
(2) facts as to which propositions the agent has defeasible reasons (combined with the status of the propositions as such), and
(3) facts as to which of the agent's defeasible reasons are defeated.[68]

Now in the case where we are considering a set of defeasible reasons, R, such that conditions (i) through (iv) hold, it is reasonable to conclude that the elements of R*–R are irrelevant to determining A's beliefs regarding the contents of the elements of R, since the contents of the element R*–R bear no significant deductive relations to the elements of R,

[68] One possible objection to my assessment to the range of factors that may be relevant to resolving inconsistency among the objects of an agent's defeasible reasons proposes that individual defeasible reasons may have varied associated strengths in addition to their objects and statuses as indefeasible and defeasible reasons. In that case, these variable strengths would consist of factors that transcend deductive assessability.

I am inclined to think that such factors exist, the issues is one that need to be considered more carefully than is possible at present. That said, both the kinetic analogy and the theory of competing reasons proposed in *Part II* could easily be adapted to account for variations in the strength of individual reasons.

At the moment it is impractical to explain in detail how it is that an assignment of strengths to individual reasons would be incorporated into the kinetic analogy and the theory that will be proposed. However, incorporating such factors is not difficult due to the way in which the prescriptions of both the analogy and the theory to come in *Part II* rely on species of dominance reasoning. As such, incorporating strength assignments to individual defeasible reasons requires only that we take account of an additional type of positive factor in making comparative judgments as to the balance of considerations bearing on respective propositions.

and there are no undercutting defeaters for elements of R among the elements of R* (and R* *encapsulates* A's full set of reasons).

Where (i) through (iv) hold, the factors relevant to the objects of R are limited to facts about which propositions are the contents of the elements of R (combined with their associated statuses as the contents of defeasible and indefeasible reasons, as the case may be), and facts about which contents of the elements of R form inconsistent sets (and may thereby be subject to rebutting defeat). Moreover, in cases where (i) through (iv) hold for a set of defeasible reasons, R, of an agent, A, we see that deductive assessability holds for R for A. What remains to be seen is that it is normally possible to isolate sets of defeasible reasons that have been generated by direct inference so that (i) through (iv) are satisfied for such sets.

In assessing the claim that it is normally possible to isolate sets of defeasible reasons that have been generated by direct inference so that (i) through (iv) are satisfied for the sets, we see that conditions (i) and (ii) simply describe the preconditions of the sort of case that I want to consider, so that we are asked to imagine cases where R is a set of reasons that has been generated by direct inference by an agent, A, where R is a subset of R* and R* encapsulates A's full set of reasons for belief. To see that (iii) will frequently be satisfied, where R is a sufficiently extensive subset of the set of reasons that have been generated by direct inference by an agent A, we need only consider the usual deductive relations that obtain between an agent's background information and the conclusions that the agent reaches by direct inference. When we do so, we see that it is normal that the background information possessed by an agent is consistent with and is not deductively relevant to any of the contents of the elements of R (and not deductively relevant in a way which will amount to a failure of R's *deductive independence* of the background information). So it seems that (iii) is normally satisfied where R is a suitably large subset of the set of reasons that may be generated by direct inference by a particular agent. Finally, it is normally possible to isolate sets of defeasible reasons that have been generated by direct inference so that (iv) will be satisfied in addition to (i), (ii), and (iii). The ability to array a set of reasons that have been generated by direct inference in a way that satisfies condition (iv) is facilitated by the fact that any defeasible reason that is subject to undercutting defeat has no bearing on what an agent should believe. As such, in cases where an agent has a defeasible reason that is undercut, it is admissible to remove the defeasible reason from consideration in the course of determining what the agent should believe. In the normal case, removal of reasons that are undercut from a set of reasons that are generated by direct inference will leave us with a set of reasons, R, that satisfies conditions (i) through (iv). And so, in the typical case, it is possible to represent the outputs of direct inference so that *deductive assessability* holds for the set.

In light of the discussion in section 4, we also know that it is normally possible to represent the outputs of direct inference as defeasible reasons for accepting respective lower probability bounds. For this reason, it is typical for *volumetric representatity* to hold for a set of defeasible reasons that have been generated by direct inference. In light of the immediately preceding discussion, we see that it is also typical for *deductive assessability* to hold for a suitably extensive representation of the set of defeasible reasons that may generated by an agent by direct inference.

6. Principles that Hold When the Analogy Applies

In cases where the Kinetic Analogy is applicable (inasmuch as volumetic representability and deductive assessability apply), there are a number of principles that describe the dynamics of reasons.

i. Acceptance for Consistent Propostions

In cases where deductive assessability holds for a collection of reasons, R, it is reasonable to accept the objects p, of any element of R, so long as p is not a contradiction and p is consistent with every consistent subset of the objects of elements of R.

As we are considering only cases where deductive assessability holds for the elements of R, the fact that p is not a contradiction and is consistent with other all consistent subsets of the objects of R indicates that the p is in no way suspect. In such cases, it is reasonable to accept p.

ii. Aggregativity

Clearly, it is not always reasonable to accept all of the deductive consequences of the propositions for which one has defeasible reasons. The preceding fact is illustrated by cases where an agent has defeasible reasons for each of the elements of an inconsistent set. However, in cases where the set of propositions for which one has reasons is consistent (and there are no undercutting defeaters for any of the respective reasons), it reasonable to accept all of the deductive consequences of those propositions. Similarly, in the case where deductive assessability holds for a set of reasons, R, it is reasonable to accept conjunctions of the objects of the elements of R, so long as the conjunction is not a contradiction and the conjunction is consistent with the object of every consistent subset of the objects of the elements of R.

As we are only considering cases where deductive assessability holds for a set of defeasible reasons, R, the fact that a conjunction of the objects of the elements of R is not a contradiction and is consistent with all other consistent subsets of the objects of the elements of R indicates that the conjunction in no way suspect. In that case, it is reasonable to accept the conjunction.

iii. Non-skepticism

In cases where deductive assessability *does not* hold for a set of defeasible reasons, it is generally conceded that inconsistency among the objects of the elements of the set does not always lead to mutual defeat of the respective reasons. But in cases where deductive assessability does hold, it is often thought that inconsistency among the propositions for which one has defeasible reasons implies the mutual defeat of the respective reasons. Despite common impressions, it follows immediately from the third component of volumetric assessability that inconsistency among a set of propositions for which one has defeasible reasons does not imply mutual defeat (in the case where deductive assessability holds for the reasons). Moreover, that mutual defeat does not obtain in such cases is illustrated by the examples described in section 4.

iv. Limited Deductive Closure for Defeasible Reasons

A defeasible reason, r, for a proposition, α, will said to be *independent* of the elements of a set of reasons, R, just in case it is reasonable to accept α in the case where α is not a contradiction and α is consistent with every consistent subset of the objects of the elements of R.

In cases where *deductive assessability* and *volumetric representability* hold for a set of reasons, R, I maintain that the set of *independent* defeasible reasons associated with R satisfies a modest closure condition. My proposal is that in cases where deductive assessability holds for a set of reasons, R, the elements of R generate auxiliary independent defeasible reasons which may go undefeated in the presence of inconsistency even in cases where the initial element of R that ‘generated’ the auxiliary reason is itself defeated. The specific closure condition that I recommend is as follows:

> In the case where deductive assessability holds for R and the content of each element of R is a lower probability bound, the set of objects of independent reasons, I_R, associated with R is such that for all α, if PROB(α) $\geq$ t $\in I_R$ and t $\geq$ s, then PROB(α) $\geq$ s $\in I_R$.

The present closure condition is meant to capture the idea (implicit in the Kinetic Analogy and the third component of volumetric representability) that it makes sense to respond to inconsistency among a set of lower probability bounds (in the case where one has a defeasible reason for accepting each of the bounds) by the sort of modest downward revision of the bounds as illustrated by the second example presented in section 4. In that example, we supposed that we had a defeasible reason for believing that PROB(the token is red) $\geq$ 0.85, and a defeasible reason for believing that PROB(the token is not red) $\geq$ 0.75. Assuming that the present pair of reasons is deductively assessable, and each of the reasons is *independent*, it follows (by the closure condition) that we have *independent* reasons for each of the propositions of the form *PROB(the token is red) ≥ r*, for every r less than or equal to 0.85, and independent reasons for each of the propositions of the form *PROB(the token is not red) ≥ r*, for every r less than or equal to 0.75. And given our assumptions, it follows from the definition of what it is to be an *independent reason* that it is reasonable for us to accept the position that PROB(the token is red) ≥ 0.25, and the position that PROB(the token is not red) ≥ 0.15 (since neither proposition is a contradiction and each proposition is consistent with every consistent subset of { PROB(the token is red) ≥ 0.85, PROB(the token is not red) ≥ 0.75}).

7. Summary of Part I

In *Part I* of the present chapter, I set out of the foundational intuitions for a theory that will lead an agent to reasonable conclusions in cases where the agent's data supports instances of direct inference that yield mutually inconsistent conclusions. As an approach to the problem of competing reasons, I proposed a 'kinetic analogy' as a heuristic to shape our intuitions as to what an agent should believe in cases where the agent has defeasible reasons for mutually inconsistent propositions. To facilitate use of the analogy, I proposed two conditions whose satisfaction is sufficient to indicate the applicability of the analogy in guiding an agent to reasonable beliefs. After explaining why it is that collections of defeasible reasons generated by direct inference *routinely* satisfy the two conditions, I proposed four further principles that characterize a set of reasons in cases where the two conditions are satisfied. In *Part II* of the chapter, I will articulate a theory of competing reasons that is intended to apply in a wide range of cases where multiple instances of direct inference yield defeasible reasons for mutually inconsistent conclusions. The theory satisfies the four principles that characterize cases where the Kinetic Analogy applies.

Part II: The Kinetic Theory

8. Preliminaries

Recall that a *theory of competing reasons* is a theory (or a set of principles) which tells an agent what conclusion(s) to adopt in cases where the agent has defeasible reasons for accepting mutually inconsistent conclusions. It is intended that the theory of competing reasons articulated in this chapter, hereafter *the Kinetic Theory*, be applicable to any collection of defeasible reasons, R, where *volumetric representation* and *deductive assessability* hold for R, and where the objects of the elements of R are statements of lower probability bounds. In *Part I* of the chapter, I argued that it is normally possible to represent the outputs of direct inference as defeasible reasons for accepting respective lower probability bounds, and that *volumetric representation* and *deductive assessability* normally hold for suitably extensive representations of the set defeasible reasons that are generated by a body of data via direct inference. It is intended that Kinetic Theory be applicable to resolving conflicts between sets of defeasible reasons that have been generated by direct inference, and to others cases where conflicting defeasible reasons arise and satisfy the preconditions for the applicability of the theory, if such cases exist.

The Kinetic Theory will take as input a set of defeasible reasons for belief in the elements of a set of lower probability bounds. For ease of presentation, I will identify reasons for belief with their contents, so that a collection of defeasible reasons for belief, Σ, will be represented by a finite set of propositions of the form { $PROB(\alpha_1) \geq s_1, \ldots, PROB(\alpha_n) \geq s_n$ }, where *PROB* is a finitely additive probability function, each α_i is a well formed formula of a countable language, and each s_i is in [0, 1].—For convenience, it is also stipulated that none of the non-probabilistic propositions that are the objects of the elements of Σ (i.e., the respective α_i are repeated), though there is no restriction forbidding instances of logically equivalent non-probabilistic propositions. Rather than represent the full set of defeasible reasons that an agent may have in a given situation, it is intended that Kinetic Theory be applicable to sets of defeasible reasons, Σ, where deductive assessability holds for the set.—In the case were *deductive assessability* holds for a set of defeasible reasons, the only relations relevant to arbitrating conflict among the elements of the set (for the purposes of determining what to believe) are deductive relations between its elements.

Although the theory that I propose is intended to be applicable to defeasible reasons whose objects are statements of lower probability bounds, my approach will focus on evaluating the *aggregate degree of support* for the non-probabilistic propositions that are the objects of the lower probability bounds that are the official inputs to the theory. In particular, my approach will be to apply comparative assessments of the aggregate degree of support for the non-probabilistic propositions as a tool for reaching reasonable conclusions about what lower probability bounds an agent should accept for the respective non-probabilistic propositions.

The foundational principle, to which I will appeal for the sake of justifying the prescriptions of the Kinetic Theory, is called the "principle of non-arbitrariness". The *principle of non-arbitrariness* is expressed as follows: When making judgments about the likely truth value of a proposition, the greatest lower probability bound that one accepts for the proposition should *correspond* to the aggregate degree of support for the proposition. Though the present principle is justified upon any reasonable interpretation, the requirement of 'correspondence' demanded by the principle is, without a doubt, somewhat equivocal. I will draw on *three* consequences of the principle as a means to a partial analysis of the relevant notion of *correspondence*. In the course of justifying the Kinetic Theory, I will appeal only *indirectly* to the principle of non-arbitrariness. The three derivative principles to which I will appeal *directly* are: *the principle of indifference, the principle of difference, and the principle of least amendment*.

The principle of indifference that I endorse is a variation of Keynes's principle of indifference.[69] Keynes's principle states that if one has equal evidence regarding the likelihood of each element of an exclusive and jointly inclusive set of possibilities, then one should assign equal probability to each of the possibilities. My principle states that *if* the aggregate structure of one's reasons bearing on one proposition is equivalent to the

[69] For a historical discussion of *the principle of indifference*, see Hacking (1975).

aggregate structure of one's reasons bearing on another proposition (so that the aggregate degree of support for the two propositions is the same), then the greatest lower probability bound that one accepts for the two propositions should be the same.

There is a widespread belief that, in many cases, applying Keynes's principle of indifference leads to contradictory conclusions. And this has been taken by some to be symptomatic of the fact that the principle itself is inconsistent.[70] Setting aside the possible defects and limitations of Keynes's principle, it is possible to demonstrate that the principle of indifference proposed in the present chapter *cannot* lead to contradictory results. Furthermore, unlike other variants of the principle of indifference, the principle that I propose does not depend for its application on the selection of an underlying language or on the selection of an underlying partition of the space of possible worlds.[71]

In addition to the principle of indifference, I will appeal to what I call "the principle of difference". This principle states that *if* the structure of one's aggregate reasons bearing on one proposition diverges from equivalence in comparison to the structure of one's aggregate reasons bearing on another proposition, and the character of the divergence can

[70] By Howson (1976), and by Pollock & Cruz (1999), for example.
[71] The theories of logical probability proposed by Carnap (1952) are *representation dependent*, as is the currently popular *maximum entropy* inference rule. For a critical discussion of representation dependent inference procedures, see van Fraassen (1989). For a concise summary and critical discussion of Carnap (1952), see Kyburg & Teng (2001). For an attempt to defend the maximum entropy inference rule in the face of the difficulties associated with representation dependence, see Paris & Vencovská (1990) and (1997). For attempts to defend representation dependent inference procedures generally, see Jaynes (1973), Halpern & Koller (1995), and Williamson (2003). Representation dependent theories that address the problems of direct inference have been proposed by Paris & Vencovská (1992) and by Bacchus *at al* (1996). I think that representation dependent inference procedures are, in significant respects, deeply flawed. The Kinetic Theory represents an attempt to articulate a non-skeptical theory which is not representation dependent.

only favor the aggregate degree of support for the first proposition, then the greatest lower probability bound that one accepts for the first proposition should be *at least as great as* the greatest lower probability bound that one accepts for the second.

In general, the principles of indifference and difference place ordering constraints on the greatest lower probability bounds that one should adopt for respective propositions.

The third consequence of the principle of non-arbitrariness that I will appeal to is called “the principle of least amendment”. This principle states that one should accept any proposition for which one has a defeasible reason, so long as one has no *credible reason* for rejecting it.[72] In the context of the Kinetic Theory, which is intended to apply to defeasible reasons for lower probability bounds, I take the principle of least amendment to be closely connected with the third component of the condition of volumetric representability, as described in *Part I* of the chapter. According to that component, reasons for mutually inconsistent lower probability bounds are to be reconciled via a moderate downward revision of the respective bounds. The principle of least amendment is proposed here as a vehicle for determining the degree of such amendments.

In the following two sections of the chapter, I propose criteria for judging whether the aggregate structure of an agent’s reasons bearing on one proposition is equivalent to the

[72] The notion of a *credible reason* is non-technical, and my intention is that our understanding of the notion be partially determined by a commitment to maintain the *principle of least amendment* as an *a priori* truth. In turn, specific applications of the principle of least amendment will require that one make the case that a proposed reason (or type of reason) is non-credible.

aggregate structure of the agent's reasons bearing on another proposition. In section 11, I will propose criteria for judging whether the structure of an agent's aggregate reasons bearing on one proposition diverges from equivalence from the structure of the agent's aggregate reasons bearing on another proposition, and the divergence can only favor the aggregate degree of support for the first proposition. In section 12, I will explain my general proposal for how we may appeal to conclusions about comparative degrees of support as a basis for determining the greatest lower probability bounds an agent should accept for respective propositions.

9. A Canonical Representation of Reason Characteristics

A *format* for representing sets of defeasible reasons provides an adequate basis for applying *the principle of indifference* if *equivalence* in the represented features of one's *reasons* (relative to the format) *implies identity* in one's aggregate degree of support for the respective propositions. We have this implication from *equivalence of represented features* to *identity in aggregate degree of support*, if a *representation format* encodes all of the factors that are determinative of the strength of one's aggregate degree of support for respective propositions.[73]

[73] In fact, the only difficulty is to find a format that encodes the right features in such a way that equivalences among reasons can be discerned.

Recall that my approach will focus on evaluating the *aggregate degree of support* for the non-probabilistic propositions that are the objects of the lower probability bounds that are the official inputs to the theory. Due to this focus, I will frequently want to refer to an ordered version of the set of non-probabilistic propositions corresponding to a set of defeasible reasons.—So where Σ is $\{ \text{PROB}(\alpha_1) \geq b_1, \ldots , \text{PROB}(\alpha_n) \geq b_n \}$, and $\langle \text{PROB}(\alpha_1) \geq b_1, \ldots , \text{PROB}(\alpha_n) \geq b_n \rangle$ is a canonical ordering of the elements of Σ, $[\Sigma]$ will be the corresponding ordering of the non-probabilistic propositions associated with Σ, i.e., $[\Sigma] = \langle \alpha_1, \ldots , \alpha_n \rangle$.

In the case where an agent's defeasible reasons are for lower probability bounds on propositions and *deductive assessability* holds for the given reasons, the values of these bounds, in combination with the set of possible rows for the truth table corresponding to the respective non-probabilistic propositions encodes all of the factors that are determinative of the aggregate degree of support for the respective non-probabilistic propositions. For example, if the set $\{ \text{PROB}(p \vee q) \geq .8, \text{PROB}(p \vee \neg q) \geq .8, \text{PROB}(\neg p) \geq .7 \}$ represents a set of defeasible reasons for which *deductive assessability* holds, then the outlined portion of *figure 1* represents all of the factors that are determinative of one's aggregate degree of support for the propositions $p \vee q$, $p \vee \neg q$, and $\neg p$ (with the first row representing the defeasible lower probability bounds on the corresponding non-probabilistic propositions, and the subsequent rows representing the truth table for the non-probabilistic propositions).

$p \vee q$	$p \vee \neg q$	$\neg p$
0.8	0.8	0.7
1	1	0
1	0	1
0	1	1

Figure 1

In general, I will describe the information represented in a matrix such as the one in *figure 1* as the *reason matrix* for the corresponding set of reasons.

Definition 1: The *reason matrix* for a set, Σ, of defeasible reasons is the triple consisting of (1) [Σ] (i.e., the canonical ordering of the set of non-probabilistic propositions corresponding to the elements of Σ), (2) a corresponding ordering of the set of *lower probability bound values* for the given non-probabilistic propositions, and (3) the set of possible rows of a truth table associated with the respective non-probabilistic propositions.[74]

Given the reason matrix for a set of defeasible reasons, Σ, one may treat the elements of [Σ] as *mere* labels. In turn, the values that a reason matrix *associates* with these labels represent *all* of the features that are relevant to assessing the degree of support for each of the elements of [Σ] (assuming that *deductive assessability* holds for Σ). In addition to

[74] A more precise definition can be found in the Appendix.

representing the lower probability bounds associated with the respective elements of [Σ], all of the inferential relationships that hold between the elements of [Σ] are represented. Finally, it is demonstrable, for all propositions, α, that the *reason matrix* for Σ encodes precisely the information necessary for determining the greatest lower probability bound provable for α from Σ.

10. Support Strength Equivalence

We may conclude that the *aggregate* degree of support for two propositions is identical, in cases where the character of our reasons bearing on the two propositions exhibits a specific sort of symmetry. As an illustration of the sort of cases where such symmetry is present, suppose that our set of defeasible reasons is Σ, where Σ = { PROB(¬p) ≥ .6, PROB(p∧¬q) ≥ .8, PROB(q∧r) ≥ .6 }. In this case, the reason matrix for Σ is given by *figure 2*. And in this case, we see that we can obtain a matrix that is structurally equivalent to the one presented by *figure 2*, by switching the columns corresponding to ¬p and q∧r, and then switching the rows labeled "w_4" and "w_5". This fact is illustrated by *figure 3*.

	$\neg p$	$p \wedge \neg q$	$q \wedge r$
	0.6	0.8	0.6
w_1	1	0	1
w_2	0	1	0
w_3	0	0	0
w_4	1	0	0
w_5	1	0	1

Figure 2

	$q \wedge r$	$p \wedge \neg q$	$\neg p$
	0.6	0.8	0.6
w_1	1	0	1
w_2	0	1	0
w_3	0	0	0
w_5	1	0	0
w_4	1	0	1

Figure 3

Because of the symmetry that holds between the *features* of our reasons relative to ¬p and q∧r (and because we are assuming that deductive assessability holds in the present case), we may conclude that the *aggregate* degree of support for ¬p and for q∧r is identical. The following is a general definition of the relation that holds between ¬p and q∧r (relative to Σ) in the present example:

Definition 2: Relative to Σ, the structure of one's aggregate reasons bearing on α and β is equivalent (written "$\alpha \approx_{\Sigma} \beta$") *just in case* it is possible to permute the rows and columns of the *reason matrix* for Σ so that (1) the column corresponding to α takes the position formerly held by the column corresponding to β, (2) the arrangement of 1s and 0s composing the permuted matrix is identical to that of the original matrix, and (3) the column corresponding to the *i*th element of [Σ] takes the position previously held by the column corresponding to the *k*th element of [Σ] *only if* the bound attached to the proposition in *i*th position is identical to the bound attached to the proposition in the *k*th position.[75]

The idea underlying the relation, $\approx$, is that the aggregate reasons bearing on a pair of positions are equivalent, if the 'position' of the respective propositions is equivalent within a sufficiently rich representation of the qualitative and quantitative features of a set of defeasible reasons. While the sort of equivalence which is presently at issue is obviously not a species of logical equivalence, we note that $\approx$ is an equivalence relation.

Theorem 1: $\forall\Sigma$: $\approx_\Sigma$ is reflexive, symmetric and transitive.[76]

Moreover, it is demonstrable that *if* the character of one's reasons bearing on a pair of propositions, α and β, is *equivalent* in the present sense, then the greatest lower probability bound that one can prove for α from Σ and for β from Σ is identical. That is:

Theorem 2:
$\forall\alpha,\beta,\Sigma$: $\alpha \approx_\Sigma \beta \;\Rightarrow\; \max\{\, r \mid \Sigma \models \mathrm{PROB}(\alpha) \geq r \,\} = \max\{\, r \mid \Sigma \models \mathrm{PROB}(\beta) \geq r \,\}$.

The relation, $\approx$, provides us with a criterion for regarding our *aggregate* degree of support for a pair of propositions to be equivalent. In turn, the relation supplies a basis for applying the *principle of indifference*. Indeed, if the character of one's aggregate reasons bearing on α and on β is equivalent, then the greatest lower probability bound that one

[75] A precise definition of the conditions under which $\alpha \approx_\Sigma \beta$ holds for a respective triple is given in the Appendix.
[76] A sketch of the proof is provided in the Appendix.

accepts for α should be the same as the greatest lower probability bound that one accepts for β. Given theorem 2, we know that this application of the principle of indifference is a generalization of a valid rule of inference.[77]

11. Dominance Among Reasons

It is apparent that there are many situations where the aggregate degree of support for a proposition, α, is at least as great as the aggregate degree of support for another proposition, β. Moreover, there are clear cases where the aggregate structure of our reasons bearing on respective propositions diverges from *equivalence* in a way that can only favor the conclusion that aggregate degree of support for one proposition is at least as great as the aggregate degree of support for another. For example: Suppose that deductive assessability holds for the following set defeasible reasons: $\{ \text{PROB}(p) \geq .8, \text{PROB}(\neg p) \geq .4, \text{PROB}(q) \geq .8, \text{PROB}(\neg q) \geq .6 \}$. In this case, the structure of one's reasons bearing on p and q would be *equivalent* (so that $p \approx_{\Sigma} q$), if it were not for a divergence in the bounds attached to $\neg p$ and $\neg q$. Here the divergence favors the judgment that the aggregate degree of support for p is at least as great as the aggregate degree of support for q (since the only

[77] Note that, in the present context, it would not have been acceptable to define the relation $\approx_{\Sigma}$ so that it holds between a pair of formulae α and β *just in case* $\max\{ r \mid \Sigma \models \text{PROB}(\alpha) \geq r \} = \max\{ r \mid \Sigma \models \text{PROB}(\beta) \geq r \}$. Indeed, our present aim is to articulate a principle that will be applicable to a set of defeasible reasons even when the objects of those reasons form an inconsistent set, and we do not wish to conclude that the strength of our aggregate reasons for each element of $[\Sigma]$ is identical in each case where Σ is inconsistent.

difference in the structure of reasons bearing on the two propositions is in the magnitude of a lower probability bound on an 'opposing' proposition).

In order to identify cases where the character of our aggregate reasons diverge from *equivalence* in a way that favors the aggregate degree of support for one proposition over another, I will articulate some conditions under which a divergence in the lower probability bound on one proposition may positively or negatively affect the aggregate degree of support for another proposition. Ultimately, the objective will be to identity which sorts of divergence from a *base-line*, represented by $\approx$, may make a positive or negative differences regarding the aggregate degree of support for a proposition. In pursuing this objective, my intention is to err on the side of caution. In particular, I will attempt to identify types of divergence from equivalence which can only have a positive effect on the degree of support for a proposition, and types of divergence which can only have a negative effect on the degree of support for a proposition.[78]

In cases where a proposition, β, implies a proposition, α, it appears that a defeasible reason for a lower probability bound on β has the *potential* to promote the aggregate degree of support for α. The following definition is intended as a generalization of the present judgment:

[78] With regard to the present strategy, it is important to notice that in the case of a finite set of lower probability bounds, Σ, the highest lower probability bound, r, that one may prove on a proposition, α, is a monotone increasing function of the bounds on the elements of $[\Sigma]$. In other words, if $\mathrm{PROB}(\alpha) \geq r$ is derivable from a set, Σ, then $\mathrm{PROB}(\alpha) \geq r$ is derivable from all supersets of Σ.

Definition 3: Relative to Σ, (a lower probability bound on) β *has the potential to directly promote* (the aggregate degree of support for) α *just in case* it is possible to increase the greatest lower bound provable for α by increasing the greatest lower bound on β (relative to an assignment of lower probability bounds to the elements of $[\Sigma]$).

To get the idea, suppose that $\Sigma = \{$ PROB(p) ≥ 1, PROB(p$\supset$q) ≥ 1, PROB(q) $\geq 1\}$, and that $[\Sigma] = \langle$ p, p$\supset$q, q $\rangle$. In that case, p *has the potential to directly promote* q (relative to Σ), since there exists an assignment, A, of bounds to the elements of $[\Sigma]$ so that by increasing the bound on p, it is possible to increase the greatest lower bound provable for q. For example, let A = { PROB(p) ≥ 0, PROB(p$\supset$q) ≥ 1, PROB(q) ≥ 0 }). In that case, the greatest lower bound on q provable from A is PROB(q) ≥ 0. But suppose we increase the greatest bound on p present in A, so that result is A′ = { PROB(p) ≥ 1, PROB(p$\supset$q) ≥ 1, PROB(q) ≥ 0 }). The greatest lower bound on q provable from A′ is PROB(q) ≥ 1.

Correlative to the notion of *potential direct promotion*, we have a notion of *potential direct inhibition*:

Definition 4: Relative to Σ, (a lower probability bound on) β *has the potential to directly inhibit* (the aggregate degree of support for) α *just in case* it is possible to increase the greatest lower bound provable for $\neg\alpha$ by increasing the greatest lower bound on β (relative to an assignment of lower probability bounds to the elements of $[\Sigma]$).

Although the two preceding definitions would be complete as a characterization of the conditions under which a defeasible reason for lower probability bound on one proposition may have a positive or negative impact on the aggregate degree of support for another proposition relative to a set of defeasible reasons, Σ, in the case where Σ *is consistent* (and deductive assessability holds), it seems that more subtlety is required in the case where Σ is inconsistent. To illustrate this fact, suppose that we possess the following set of defeasible reasons: { PROB(p) $\geq$ .9, PROB(p$\supset\neg$q) $\geq$ .9, PROB(q) $\geq$ .9, PROB(q$\supset$r) $\geq$ .9, PROB(r) $\geq$ .9 }. In this case, both p and p$\supset\neg$q have the potential to directly inhibit q, and both q and q$\supset$r have the potential to directly promote r, and although it seems that p and p$\supset\neg$q are negatively relevant to r (since p and p$\supset\neg$q have the potential to directly inhibit a proposition that has the potential to directly promote r), neither p nor p$\supset\neg$q have the potential to directly inhibit r. We must generalize the definitions of direct promotion and inhibition, in order to capture the fact that a lower probability bound on a proposition, β, may have the potential to promote or inhibit the aggregate degree of support for another proposition, α, even when a lower probability bound on β does not have the potential to *directly* promote or inhibit α.

In the case of defeasible reasons for mutually inconsistent propositions, relations of promotion and inhibition between reasons may be *indirect*. For example, *if* a lower bound on χ has the potential to *directly inhibit* a proposition, β, and a lower bound on β has the potential to *directly inhibit* a proposition, α, *then* a lower bound on χ has the potential to (indirectly) increase the aggregate degree of support for α (*figure 4*). And, similarly, if a lower bound on χ has the potential to *directly inhibit* a proposition, β, and a lower bound on β has the potential to *directly promote* a proposition, α, then a lower bound on χ has the potential to indirectly decrease the aggregate degree of support for α (*figure 5*). On the other hand, if a lower bound on χ has the potential to *directly* promote a proposition, β, and a lower bound on β has the potential to *directly* promote a proposition, α, then a lower bound on χ has the potential to indirectly increase the aggregate degree of support for α (*figure 6*). Finally, if a lower bound on χ has the potential to *directly* promote a proposition, β, and a lower bound on β has the potential to *directly inhibit* a proposition, α, then a lower bound on χ has the potential to indirectly decrease our aggregate degree of support for α (*figure 7*).

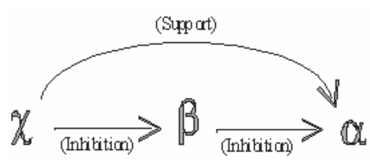

Figure 4

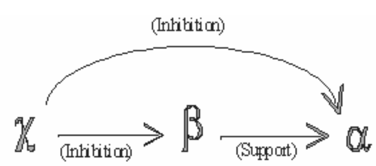

Figure 5

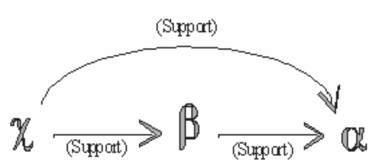

Figure 6

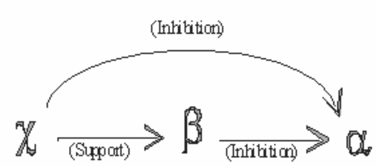

Figure 7

The following conditions codify (and generalize) the ideas behind the preceding examples, and provide a means of identifying cases where a defeasible reason for a lower probability bound on one proposition may *directly* or *indirectly* influence the aggregate degree of support for another proposition:

(1) If α has the potential to *directly inhibit* β, relative to Σ, then α has the potential to *inhibit* β, relative to Σ (written α inh$_\Sigma$ β).

(2) If α has the potential to *directly promote* β, relative to Σ, then α has the potential to *promote* β, relative to Σ (written α prom_Σ β).

(3) If (α inh_Σ β & β inh_Σ χ), then α prom_Σ χ.

(4) If (α inh_Σ β & β prom_Σ χ), then α inh_Σ χ.

(5) If (α prom_Σ β & β prom_Σ χ), then α prom_Σ χ.

(6) If (α prom_Σ β & β inh_Σ χ), then α inh_Σ χ.

(7) If it does not follow from (1) – (6) that α prom_Σ β or α inh_Σ β, then not(α prom_Σ β) and not(α inh_Σ β).

We now note that a single proposition, β, may have both the *potential to promote* and the *potential to inhibit* the degree of support for a proposition, α, relative to a set, Σ. For example, relative to { $\text{PROB}(p) \geq .9$, $\text{PROB}(p \supset q) \geq .9$, $\text{PROB}(p \supset \neg q) \geq .9$, $\text{PROB}(q) \geq .9$ }, p has the *potential to promote* and the *potential to inhibit* q. In light of the fact that a proposition can have both the *potential to promote* and the *potential to inhibit* the degree of support for another proposition, I propose the following taxonomy of relations in which propositions may stand, relative to a set of defeasible reasons:

Definition 5: α is *a friend* of β, relative to Σ, *just in case* α pro_Σ β and not(α inh_Σ β).

Definition 6: α is *an enemy* of β, relative to Σ, *just in case* not(α pro_Σ β) and α inh_Σ β.

Definition 7: α is *ambivalent to* β, relative to Σ, *just in case* α pro_Σ β and α inh_Σ β.

Using the present vocabulary, we may assert that increasing the lower bound for an *enemy* of a proposition, α, is *guaranteed* not to increase the aggregate degree of support for α, and increasing the lower bound for a *friend* of a proposition, α, is *guaranteed* not decrease the aggregate degree of support for α. On the other hand, we have no guarantee about the effect of varying the bound on a proposition that is ambivalent to α. The preceding observations suggest how we may apply dominance reasoning to recognize cases where the aggregate degree of support for one proposition is *at least as great as* the aggregate degree of support for another proposition. To see how this works consider the set of defeasible reasons Σ, where Σ = { PROB(p) $\geq$ .8, PROB($\neg$p) $\geq$ .4, PROB(q) $\geq$ .8, PROB($\neg$q) $\geq$ .6 }. Now it is not the case that p $\approx_\Sigma$ q, due to the divergence in the bounds attached to $\neg$p and $\neg$q. Figures 8 and 9 illustrates the problem, when we try to map the columns corresponding p and $\neg$p to columns corresponding q and $\neg$q.

	p	$\neg p$	q	$\neg q$
	0.8	0.4	0.8	0.6
w_1	1	0	1	0
w_2	0	1	1	0
w_3	1	0	0	1
w_4	0	1	0	1

Figure 8

	q	$\neg q$	p	$\neg p$
	0.8	0.6	0.8	0.4
w_1	1	0	1	0
w_2	0	1	1	0
w_3	1	0	0	1
w_4	0	1	0	1

Figure 9

But note that (relative to Σ) $\neg$p is an enemy of p, so that it would be the case that p $\approx_\Sigma$ q if it were not for the fact that the corresponding enemy p has a lower than the corresponding enemy of q, so the failure of p $\approx_\Sigma$ q is due to a difference that could only speak in favor of p in comparison to q. The following notion of 'reason dominance' generalizes the present observation to take into consideration divergences from $\alpha \approx_\Sigma \beta$ that can only speak in favor of α in comparison to β.

Definition 8: Relative to Σ, the structure of one's aggregate reasons bearing on α and β is equivalent or deviate from equivalence in a manner that can only favor the aggregate degree of support for α compared with the aggregate degree of support for β (written "$\alpha \succ_{\Sigma} \beta$") *just in case* it is possible to permute the rows and columns of the *reason matrix* for Σ so that (1) the column corresponding to α takes the position formerly held by the column corresponding to β, (2) the arrangement of 1s and 0s composing the permuted matrix is identical to that of the original matrix, and (3) the column corresponding to the *i*th element of [Σ] takes the position previously held by the column corresponding to the *k*th element of [Σ] *only if* (a) the proposition in the ith position *is ambivalent to* α and the bound attached to the proposition in *i*th position is identical to the bound attached to the proposition in the *k*th position, (b) the proposition in the *i*th position *is a friend of* α and the bound attached to the proposition in *i*th position is greater than or equal to the bound attached to the proposition in the *k*th position, or (c) the proposition in the ith position *is an enemy of* α and the bound attached to the proposition in *i*th position is less than or equal to the bound attached to the proposition in the *k*th position.

The plausibility of the preceding definition (given its intended application) derives from the idea that the aggregate degree of support for two propositions is *identical* just in case the 'position' of the two propositions, within an adequate representation of the structure of one's defeasible reasons, is equivalent. Note, then, that conditions (1) and (2) of the present definition are identical to conditions (1) and (2) of the definition of 'reason equivalence' given in the previous section. In turn, the present definition codifies the idea that the aggregate degree of support for a proposition, α, is at least as great as the aggregate degree of support for a position, β, *if* the character of one's aggregate reasons bearing on α and β deviate from *equivalence* in a manner that can only favor the aggregate degree of support for α compared with the aggregate degree of support for β. In particular, the aggregate degree of support for α is at least as great as the aggregate degree of support for β, *if* the relevant structure of one's reasons bearing on two propositions diverge from *equivalence* in a manner which results in *friends* of α having at least as high a bound as *corresponding friends* of β, and *enemies* of α having at least as a high a bound as *corresponding enemies* of β.

The conditions for the applicability of the relation, $\succ$, reflect our ability to apply dominance reasoning as a means to recognizing cases where the aggregate degree of support for one proposition is *at least as great as* the aggregate degree of support for another. As we would expect of a weak dominance relation, the following condition holds:

Theorem 4: $\forall\Sigma$: $\succ_\Sigma$ is reflexive and transitive.

The notion of 'reason domination' codified in the definition of $\succ$ provides us with a sufficient condition for applying *the principle of difference*. The following theorem demonstrates that the proposed application of the principle of difference is a generalization of a valid rule of inference:

Theorem 5:
$\forall\alpha,\beta,\Sigma$: $\alpha \succ_\Sigma \beta \;\Rightarrow\;$ max{ r | $\Sigma \models$ PROB(α) $\geq$ r} $\geq$ max{ r | $\Sigma \models$ PROB(β) $\geq$ r}.

For the remainder of the dissertation, the expression "$\succ_\Sigma$" will be used to designate the set of pairs, $\langle\alpha, \beta\rangle$, that stand in the relation $\succ_\Sigma$ (i.e., $\succ_\Sigma = \{ \langle\alpha, \beta\rangle \mid \alpha \succ_\Sigma \beta \}$).

The following definition specifies the general conditions under which a set of lower probability bounds will be said to *conform to* the 'constraint set', $\succ_\Sigma$:

Definition 9: Γ *conforms to* $\succ_\Sigma$ just in case (i) Γ is consistent, and
(ii) $\forall\alpha,\beta$: $\langle\alpha, \beta\rangle \in \succ_\Sigma \;\Rightarrow\;$ max{ s | PROB(α) $\geq$ s $\in \Gamma$ } $\geq$ max{ s | PROB(β) $\geq$ s $\in \Gamma$ }.

The fact that mention of "$\approx$" is omitted from the discussion that follows is due to fact that $\approx$ is subsumed by $\succ$, as a basis for ordering the elements of a respective set, [Σ], in accordance with the aggregate degree of support for its elements. In particular:

Theorem 6: $\forall\alpha,\beta,\Sigma$: $\alpha \approx_\Sigma \beta \Rightarrow (\alpha \succ_\Sigma \beta \;\&\; \beta \succ_\Sigma \alpha)$.

12. Application of the Kinetic Analogy

The Kinetic Theory is intended to apply to defeasible reasons for lower probability bounds, where volumetric representability and deductive assessability hold for the set. In the accordance with the Kinetic Analogy, the bounds to which the Kinetic Theory applies are thought of as representing the size of voluminous objects, where the correct response to ‘incompatible’ bounds/volumes is moderate downward revision of the values of the respective bounds (compression). In part, the demand for moderate revision of the lower probability bounds (called for by the Kinetic Analogy) is implemented by the Kinetic Theory through a closure condition on independent defeasible reasons.

Recall that a defeasible reason, r, for a proposition, φ, is *independent* of the elements of a set of defeasible reasons, Σ, just in case it is reasonable to accept φ (on the basis of r) in the case where φ is not a contradiction and φ is consistent with every consistent subset of the objects of the elements of Σ. In the case where deductive assessability holds for Σ and the objects of each element of Σ is a lower probability bound, it reasonable to accept that the set of independent reasons, I_Σ, associated with Σ is such that for all α, if $PROB(\alpha) \geq t \in I_\Sigma$ and $t \geq s$, then $PROB(\alpha) \geq s \in I_\Sigma$.

As the Kinetic Theory is intended to apply only reasons for lower probability bounds in cases where deductive assessability holds, I will assume that the set of independent reasons corresponding to inputs to the Kinetic Theory satisfy the present closure condition. Moreover, I will require that the proposed theory recommend belief in each of the independent reasons corresponding to a set of defeasible reasons, Σ, in the case where the content of a respective independent reason is not a contradiction and is consistent with every consistent subset of the objects of the elements of Σ. To facilitate the present requirement, I will use the expression "EXT(Σ)" to denote the set of probability bounds PROB(α) ≥ s, such that PROB(α) ≥ t is an element of the set, Σ, of defeasible reasons that is under consideration, and t ≥ s. More precisely:

Definition 10: The *extended set* of defeasible reasons corresponding to a *set*, Σ, of defeasible reasons is the set $\text{EXT}(\Sigma) = \{ \text{PROB}(\alpha) \geq s \mid \text{PROB}(\alpha) \geq t \in \Sigma \;\&\; t \geq s \}$.

In accord with the preceding discussion, it is hereafter assumed that one should accept any element of EXT(Σ) in the absence of conflicting defeasible reasons (in the case where the set of defeasible reasons originally under consideration is Σ). The present assumption is central to the Kinetic Theory. My other proposal for applying the Kinetic Analogy involves the appeal to symmetries and asymmetries in the structure of a set of defeasible reasons, as grounds for accepting a proposition for which one has a defeasible reason, even if one also has a defeasible reason for rejecting it.

Relative to a set, Σ, of defeasible reasons, ≻ specifies a partial preordering of the elements of [Σ], in accordance with the aggregate degree of support for the respective propositions. Clearly, ≻ places ordering constraints on the greatest lower probability bounds one should accept for respective elements of [Σ]. Importantly, ≻ may also be used to place constraints on what counts as a 'credible defeater'. The present idea can be illustrated by contrast to the skeptical policy. According to the *skeptical policy*, any proposition, φ, that is derivable from a consistent set of propositions, Γ, where one has a defeasible reason for each element of Γ, will be a defeater for any proposition that it contradicts. My alternative proposal is that a proposition, φ, is a credible defeater for an element of EXT(Σ) *only if* φ is derivable from a set of propositions, Γ, where one has a defeasible reason for each element of Γ, and Γ *conforms to* $\succ_\Sigma$. The precise idea is as follows:

> Definition 11: Relative to a set, Σ, of defeasible reasons, PROB(β) $\geq$ t constitutes a *credible defeater* for PROB(α) $\geq$ s, *just in case* { PROB(β) $\geq$ t, PROB(α) $\geq$ s } is inconsistent, and $\exists\Gamma$: (i) $\Gamma \subseteq$ EXT(Σ), (ii) $\Gamma \models$ PROB(β) $\geq$ t, and (iii) Γ *conforms to* $\succ_\Sigma$.[79]

I now propose that one should accept any element of EXT(Σ), as long as there is no *credible defeater* for the proposition. Assuming that we identify *credible defeaters* with *credible reasons for rejecting a proposition*, the present proposal is justified by appeal to *the principle of least amendment*. This identification is justified inasmuch as the 'non-credible defeaters' recognized by the *skeptical policy* tacitly rely on the assumption, for some χ and δ, that the degree of support for χ is greater than the degree of support for δ, despite the fact that the degree of support for δ is at least as great as the degree of support for χ. The fact that non-credible defeaters rely on such an assumption is illustrated by the fact that it is impossible to 'ground' a non-credible defeater in a 'basis', Γ, which *conforms to* the constraints present in $\succ_\Sigma$. For example, suppose that the set of defeasible reasons under consideration is Σ, where Σ = { PROB(p) $\geq$ 0.8, PROB($\neg$p) $\geq$ 0.7 }. In that case, p $\succ_\Sigma$ $\neg$p. Now suppose that one is trying to decide whether or not to accept PROB(p) $\geq$ 0.5. According to the skeptical policy one has a defeasible reason for PROB($\neg$p) $\geq$ 0.7, and so one should not accept that PROB(p) $\geq$ .5. According to my proposal, PROB($\neg$p) $\geq$ 0.7 is

[79] By this definition, no self-contradiction may be a credible defeater. Similarly, any formula, PROB(α) $\geq$ r, that satisfies (i), (ii) and (iii) is a credible defeater for any self-contradiction.

not a *credible defeater*. Indeed, according to Definition 10, $PROB(\neg p) \geq 0.7$ would be a credible defeater *only if* $PROB(\neg p) \geq 0.7$ could be derived from a subset of $EXT(\Sigma)$ that *conforms to* $\succ_\Sigma$. However, there is no such set, since the only consistent subsets of $EXT(\Sigma)$ that entail $PROB(\neg p) \geq 0.7$ are subsets of $EXT(\Sigma)$ whose greatest lower probability bound on $\neg p$ is greater than the greatest bound on p.—And for this reason I say that proposal that $PROB(\neg p) \geq 0.7$ is a defeater tacitly relies on the assumption that the degree of support for $\neg p$ is greater than the degree of support for p.

In what follows, it will be useful to have a compact way of designating the set of propositions that are *potential* credible defeaters. (These are propositions that will defeat any proposition that they contradict.) I will use the expression "$D(\Sigma)$" for this purpose.

> Definition 12: The set of *potential credible defeaters* corresponding to a set, Σ, of defeasible reasons is $D(\Sigma) = \{ PROB(\beta) \geq t \mid \exists\Gamma\colon \Gamma \subseteq EXT(\Sigma)\ \&\ \Gamma \models PROB(\beta) \geq t\ \&\ \Gamma$ *conforms to* $\succ_\Sigma \}$.

Having generated the set, $D(\Sigma)$, of potential credible defeaters, corresponding to a set, Σ, of defeasible reasons, one may extract the elements of $EXT(\Sigma)$ that are inconsistent with elements of $D(\Sigma)$ in order to form a 'revised' set of defeasible reasons. I use the expression "modified reason set" to refer to the result of this extraction, and use the expression "$MRS(\Sigma)$" to denote that set. More precisely:

Definition 13: The *modified reason set* corresponding to a set, Σ, of defeasible reasons is $MRS(\Sigma) = \{ \varphi \mid \varphi \in EXT(\Sigma)\ \&\ \forall\psi \in D(\Sigma)\colon \{\psi, \varphi\}$ is consistent $\}$.

As I will discuss later, it is significant that *modified reason sets* satisfy the following equivalence:

Theorem 7: $\forall\Sigma$: MRS(Σ) = $\cap\{\ \Gamma \mid \ \Gamma$ *conforms to* $\succ_\Sigma$ & $\Gamma \subseteq$ EXT(Σ) & ($\forall\Gamma'$: $\Gamma \subset \Gamma' \Rightarrow$ ~(Γ' *conforms to* $\succ_\Sigma$ & $\Gamma' \subseteq$ EXT(Σ))) $\}$.

It follows immediately from theorem 7 that *modified reason sets* are consistent:

Corollary: $\forall\Sigma$: MRS(Σ) is consistent.

I propose that the set of lower probability bounds entailed by MRS(Σ) is an adequate representation of the set of rationally acceptable probability bounds on the elements of [Σ]. In other words:

Definition 14: The set of rational consequences of a set of defeasible reasons, Σ, is R(Σ) = $\{\ \varphi \mid$ MRS(Σ) $\models \varphi\ \}$.

In the case where Σ is consistent, R(Σ) will simply be the set of logical consequences of Σ. Indeed, if Σ is consistent, then MRS(Σ) will be EXT(Σ) (since the consistency of Σ will guarantee that no element of D(Σ) is inconsistent with any element of EXT(Σ)). Moreover:

Theorem 8: $\forall\Sigma$: Σ is consistent $\Rightarrow$ R(Σ) = $\{\ \varphi \mid \Sigma \models \varphi\ \}$.

The identification of R(Σ) as the set of rational consequences of Σ satisfies the demands placed on rational belief formation that derive from the principles of indifference and difference. Indeed:

Theorem 9: $\forall\Sigma$: R(Σ) *conforms to* $\succ_\Sigma$.

I must admit that the present definition of R(Σ) is provisional, for the specification of $\succ_\Sigma$ that is used in the generation of R(Σ) may be incomplete as an ordering of the set of rational consequences of Σ that may be established by appeal to the principle of difference. In that case, R(Σ) may be a proper subset of the 'true' set of rational consequences of Σ.[80]

13. Behavior of the Kinetic Theory

One means to appreciating the nature of the Kinetic Theory, is to consider its relationship to an exposition of the *skeptical policy* that was proposed by Pollock.[81] Pollock's proposal, for the range of cases that are under consideration, is that the set of rational conclusions corresponding to a set of defeasible reasons, Σ, is the set of deductive consequences of the intersection of the set of maximal consistent subsets of Σ. To be precise, the skeptical policy enjoins us to accept the set of propositions Skp(Σ), where

[80] Note that Theorem 7 holds independently of the way that $\succ$ is defined, so that any definition of $\succ$ that increases the number of formulae that stand in the relation will increase the set of 'rational consequences' one may draw from a given set defeasible reasons, Σ, since MRS(Σ) is equivalent to the *intersection* of the set of maximal subsets of EXT(Σ) that satisfy $\succ_\Sigma$ (assuming that the set of such maximal subsets is non-empty).

[81] Pollock's proposal appears as a part of a criticism of a certain approach to *default logics* (Reiter (1980)). What I have expressed here is a simplified version of Pollock's proposal. Pollock's general proposal is intended to apply to a larger range of cases than the one's to which the theory that I have proposed is intended to be applicable. See Pollock (1990, 239-41) and (1995, 106-8).

$\mathrm{Skp}(\Sigma) = \{ \varphi \mid \cap\text{M-Con}(\Sigma) \models \varphi \}$, and where $\text{M-Con}(\Sigma) = \{ \Gamma \mid (\Gamma \text{ is consistent } \& \ \Gamma \subseteq \Sigma) \ \& \ (\forall\Gamma'\colon \Gamma \subset \Gamma' \Rightarrow \sim(\Gamma' \text{ is consistent } \& \ \Gamma' \subseteq \Sigma)) \}$.

As an approach to belief formation in the face of conflicting reasons, the skeptical policy is quite limited. However, a significant virtue of the skeptical policy is that it is conservative, and it never recommends belief in a proposition when there is credible conflicting reason. I claim the same virtue on behalf of the Kinetic Theory. At the same time, the prescriptions of the Kinetic Theory is less limited than the prescriptions of the skeptical policy. In fact, the Kinetic Theory differs from the skeptical policy in exactly two respects, each of which contributes to increases to the ‘power’ of the Kinetic Theory as compared to the skeptical policy. Indeed, as theorem 7 illustrates, the Kinetic Theory is equivalent to the proposal that it is rational to accept any proposition that follows from the intersection of the set of maximal subsets of EXT(Σ) that *conform to* $\succ_\Sigma$. To be more precise: $\forall\Sigma\colon \mathrm{MRS}(\Sigma) = \cap\{ \Gamma \mid (\Gamma \textit{ conforms to } \succ_\Sigma \ \& \ \Gamma \subseteq \mathrm{EXT}(\Sigma)) \ \& \ (\forall\Gamma'\colon \Gamma \subset \Gamma' \Rightarrow \sim(\Gamma' \textit{ conforms to } \succ_\Sigma \ \& \ \Gamma' \subseteq \mathrm{EXT}(\Sigma))) \}$. Similarly, inasmuch as the recommendations of the Kinetic Theory are reasonable, we see that the proposal represents a strict improvement over the skeptical policy. Indeed:

Theorem 11: $\forall\Sigma\colon \mathrm{Skp}(\Sigma) \subseteq \mathrm{R}(\Sigma)$.

I will now briefly describe the manner in which the Kinetic Theory responds to given sets of defeasible reasons, by describing the greatest lower probability bounds that the theory deems acceptable in respective circumstances.

> Definition 15: The greatest acceptable lower probability bound for a proposition, α, given a set, Σ, of defeasible reasons is $R(\Sigma)(\alpha) = \max\{ s \mid PROB(\alpha) \geq s \in R(\Sigma) \}$.[82]

If applied to a set of defeasible reasons, Σ, where (1) the structure of one's aggregate reasons bearing on each element of $[\Sigma]$ is equivalent, (2) the elements of Σ are individually consistent but pair-wise inconsistent, and (3) Σ contains *n* elements, then the Kinetic Theory prescribes that greatest acceptable lower probability bound for each of the elements of $[\Sigma]$ is *1/n*. For example, suppose that $\Sigma = \{ PROB(p \wedge q) \geq .6, PROB(\neg p \wedge q) \geq .6, PROB(\neg q) \geq .6 \}$. In that case, we have $\succ_\Sigma = \{ \langle p \wedge q, \neg p \wedge q \rangle, \langle \neg p \wedge q, p \wedge q \rangle, \langle p \wedge q, \neg q \rangle, \langle \neg q, p \wedge q \rangle, \langle \neg p \wedge q, \neg q \rangle, \langle \neg q, \neg p \wedge q \rangle \}$. Cleary, the intersection of the set of maximal consistent subsets of $EXT(\Sigma)$ which *conform to* $\succ_\Sigma$, will contain a maximum upper probability bound of 1/3 on $p \wedge q$, $\neg p \wedge q$, and $\neg q$. In other words, $R(\Sigma)(p \wedge q) = R(\Sigma)(\neg p \wedge q) = R(\Sigma)(\neg q) = 1/3$.

If applied to a set of defeasible reasons, Σ, where (1) the structure of one's aggregate reasons bearing on each element of $[\Sigma]$ is equivalent, (2) Σ is a minimal inconsistent set, and (3) Σ contains *n* elements, *then* the Kinetic Theory counsels that greatest acceptable lower probability bound for each of the elements of $[\Sigma]$ is *(n-1)/n*. For example, if $\Sigma = \{$

[82] Note that the bound $PROB(\alpha) \geq R(\Sigma)(\alpha)$ may only be regarded as the greatest lower bound on α consequent to the elements of Σ. In the case where an individual is possessed of considerations outside of Σ relevant to α, it may be reasonable for the agent to accept a greater bound. That said, in the case where α is an element of $[\Sigma]$, we may include that $PROB(\alpha) \geq R(\Sigma)(\alpha)$ is the greatest lower bound on α that the agent should accept.—This conclusion is closely connected to the condition of *deductive assessability*, which is a precondition for the applicability of the kinetic theory.

PROB(p) $\geq$ .75, PROB($\neg$p) $\geq$ .75 }, then R(Σ)(p) = .5, and R(Σ)($\neg$p) = .5. Similarly, suppose that Σ = { PROB(p$\vee$q) $\geq$ .7, PROB($\neg$p$\vee$q) $\geq$ .7, PROB($\neg$q) $\geq$ .7 }. In that case, we have $\succ_\Sigma$ = { $\langle$p$\vee$q, $\neg$p$\vee$q$\rangle$, $\langle\neg$p$\vee$q, p$\vee$q$\rangle$, $\langle$p$\wedge$q, $\neg$q$\rangle$, $\langle\neg$q, p$\vee$q$\rangle$, $\langle\neg$p$\vee$q, $\neg$q$\rangle$, $\langle\neg$q, $\neg$p$\vee$q$\rangle$ }. The intersection of the set of maximal consistent subsets of EXT(Σ) which *conform to* $\succ_\Sigma$, will contain a maximum upper probability bound of 2/3 on p$\vee$q, $\neg$p$\vee$q, and $\neg$q, so that R(Σ)(p$\vee$q) = R(Σ)($\neg$p$\vee$q) = R(Σ)($\neg$q) = 2/3.

In cases where Σ is an *inconsistent* pair, {α_1, α_2}, and one's aggregate reasons favor one element of [Σ], say α_1, over the other, the Kinetic Theory counsels that the greatest acceptable lower probability bound for α_1 is the *greater* of .5 and 1 *minus* the original lower bound on α_2. Conversely, the Kinetic Theory counsels that the greatest acceptable lower probability bound for α_2 is 1 *minus* the original lower bound on α_1. For example, suppose that Σ = { PROB(p) $\geq$ .85, PROB($\neg$p) $\geq$ .75 }. In that case, we have $\succ_\Sigma$ = { $\langle$p, $\neg$p$\rangle$ }. Cleary, the intersection of the set of maximal consistent subsets of EXT(Σ) which *conform to* $\succ_\Sigma$ will contain a maximum upper probability bound of .5 on p and .15 on $\neg$p. In other words, R(Σ)(p) = 0.5, and R(Σ)($\neg$p) = 0.15. On the other hand, if Σ = { PROB(p) $\geq$ .85, PROB($\neg$p) $\geq$ .3 }, then R(Σ)(p) = 0.7, and R(Σ)($\neg$p) = 0.15.

Given a set of defeasible reasons, Σ, the following examples illustrate some simple cases of how the Kinetic Theory responds when one has non-equivalent aggregate reasons bearing on the elements of [Σ]. Suppose that Σ = { PROB(p$\wedge$q) $\geq$ .6, PROB($\neg$p$\wedge$q) $\geq$ .6, PROB(¬q) ≥ .5 }. In that case, we have $\succ_\Sigma$ = { ⟨p∧q, ¬p∧q⟩, ⟨¬p∧q, p∧q⟩, ⟨p∧q, ¬q⟩, ⟨¬p∧q, ¬q⟩ }. And in that case, the intersection of the maximal consistent subsets of EXT(Σ) that *conform to* $\succ_\Sigma$, will contain a maximum upper probability bound of 1/3 on p∧q and ¬p∧q, and maximum upper probability bound of 0 on ¬q. In other words, R(Σ)(p∧q) = R(Σ)(¬p∧q) = 1/3, and R(Σ)(¬q) = 0. Similarly, if Σ = { PROB(p∨q) ≥ .7, PROB(¬p∨q) ≥ .7, PROB(¬q) ≥ .5 }, then R(Σ)(p∨q) = 2/3, R(Σ)(¬p∨q) = 2/3, and R(Σ)(¬q) = .3.

Let us now turn to some cases that illustrate the limitations of the Kinetic Theory. In the face of these limitations, I maintain that the theory recommends acceptance of a probability bound only when the bound is reasonable, even if the theory fails to recommend every bound that it is reasonable to accept.

One possible concern for the adequacy of the present theory has to do with its seeming oversensitivity to small changes in inputs. For example, suppose that Σ = { PROB(p) ≥ .8, PROB(¬p) ≥ .8 }. In that case, R(Σ)(p) = .5 and R(Σ)(¬p) = .5. However, consider Σ′ = { PROB(p) ≥ .8, PROB(¬p) ≥ .799999 }. In that case, R(Σ)(p) = .5 and R(Σ)(¬p) = .2. It would seem, then, that a minimal change in the character of our aggregate reason yields a disproportionate change in the proposed greatest lower probability bound on ¬p.

Rather than illustrate any deep flaw in the proposed theory, the present example merely illustrates the paucity of our bases for forming rational beliefs in the face of conflicting defeasible reasons. In the case of Σ and Σ', the difference in the greatest acceptable lower probability bound for $\neg p$ results from our ability to apply a dominance principle in the first case (since $\neg p \succ_{\Sigma} p$) and not in the second (since not($\neg p \succ_{\Sigma'} p$)). In other words, the dramatic quantitative difference in the two cases results from a significant qualitative difference (dominance no longer obtains). That a significant qualitative difference may result from a small quantitative difference is familiar from game theory. On the other hand, the present theory awaits amendment by someone who is able to find some principled way to assign a greater greatest acceptable lower probability bound to $\neg p$, in the second case. If such an amendment is possible, then, assuming that the prescriptions of the Kinetic Theory are correct, the amended theory will entail the results of the Kinetic Theory (in the sense of recommending acceptance of each of the elements of $R(\Sigma)$, for any given Σ).

A further concern for the adequacy of the Kinetic Theory concerns its treatment of sets of defeasible reasons, Σ, where $[\Sigma]$ contains numerically distinct instances of logically equivalent formulae. The Kinetic Theory seems not to provide an adequate treatment of such sets, in some cases. For example, if $\Sigma = \langle\ PROB(p) \geq .8, PROB(\neg p) \geq .8\ \rangle$, then we have $R(\Sigma)(p) = .5$ and $R(\Sigma)(\neg p) = .5$. But suppose we have $\Sigma' = \langle\ PROB(p) \geq .8,$ $PROB(\neg p) \geq .8, PROB(p \vee p) \geq .8\ \rangle$, and suppose that we take this to mean that we have two ‘distinct’ reasons for accepting a probability bound on p. In its present state, the proposed theory yields the output $R(\Sigma')(p) = .2$ and $R(\Sigma')(\neg p) = .2$. The problem in the present case is that the addition of $PROB(p \vee p) \geq .8$ to Σ breaks the symmetry in the character of our reasons bearing on p and $\neg$p, in such a way that $\succ_{\Sigma'}$ can no longer ‘detect’ the fact that the degree of support for p is at least as great as the degree of support $\neg$p. In turn, the Kinetic Theory fails to recommend the acceptance of a sufficiently high lower probability bound on p.

A remedy to the present difficulty may consist in fabricating the sort of structural symmetry to which the proposed dominance notion (i.e., $\succ$) is applicable. To illustrate the idea, suppose that we were to augment Σ' so as to form the set $\Sigma'' = \{ PROB(p) \geq .8, PROB(\neg p) \geq .8, PROB(p \vee p) \geq .8, PROB(\neg p \vee \neg p) \geq .0 \}$. In that case, we would have the result that $p \succ_{\Sigma''} \neg p$, and the result that $R(\Sigma'')(p) = .5$ and $R(\Sigma'')(\neg p) = .2$. The upshot of the present observation is that it may be possible to amend the Kinetic Theory so as to give the right answers in cases where one is presented with a set such as Σ'. But the problems would not end there, for we can imagine cases where a person has numerous numerically distinct reasons for accepting one propositions and fewer distinct reasons for accepting another contradictory proposition. For example, suppose that $\Sigma = \{ PROB(p) \geq .8, PROB(p \vee p) \geq .8, PROB(p \vee p \vee p) \geq .8, PROB(\neg p) \geq .8 \}$, and suppose that we take this to mean that we have three distinct defeasible reasons for accepting a probability bound on p. In the present case, we may manufacture the structure requisite for reaching the conclusion that $p \succ_{\Sigma'} \neg p$ (for a relevant Σ'), as we did in the earlier example by adding to Σ vacuous defeasible reasons for accepting a probability bound on $\neg p$. In that case, we could justify an assignment of the value .5 as the maximum lower probability bound on p. Nevertheless, there is no obvious way by which the present theory could be amended to reflect the intuitive idea that one's magnitude of belief in a proposition should increase in proportion to the number of distinct reasons that one has for accepting it.

The present difficulty is not all that surprising inasmuch as the 'power' of the Kinetic Theory is grounded entirely in the application of dominance reasoning. Indeed, a theory of the present sort will be (in many cases) at a loss when it comes to recommending differences in belief on the basis of differences that fail to result in the applicability of a new dominance condition. In fact, the supposed failing of the theory strikes me as an appropriate failing inasmuch as one cannot, in many cases, discern a principled way to derive quantitative consequences about what lower probability bounds one should accept on the basis of the number of distinct reasons one has for a proposition. Yet as with other cases where the Kinetic Theory demonstrates insensitivity toward what seem to be significant differences in the structure of sets of defeasible reasons, I am not averse to accepting proposed amendments to the theory, provided that they can be defended on principled grounds. Moreover, I maintain that if a reasonable amendment to the theory is possible, then the amended theory will entail the prescriptions of the present theory (as the present theory entails the prescriptions of the *skeptical policy*).

14. Summary

$R(\Sigma)$ describes a set of lower probability bounds that one should accept for the elements of $[\Sigma]$ in the face of a set, Σ, of defeasible reasons, in the case where deductive assessability holds for Σ. It is possible that the complete set of lower probability bounds that one should accept in the face of a set, Σ, is a *proper superset* of the one's prescribed by $R(\Sigma)$. This possibility remains because we have no guarantee that $\succ_{\Sigma}$ represents all of the ordering constraints that may be justified by appeal to the principle of difference. Nevertheless the policy of accepting each element of $R(\Sigma)$ represents an improvement over the skeptical policy. Moreover, there are strong arguments in favor of the claim that it is rational to accept each element of $R(\Sigma)$, and irrational to reject any element of $R(\Sigma)$. Let us take a moment to review these arguments.

Given a consistent set, Σ, of defeasible reasons, $R(\Sigma)$ is simply the set of propositions that follow validly from Σ. In the case where Σ is inconsistent, $R(\Sigma)$ is the set of propositions that is generated in accordance with a general principle (the Principle of Non-Arbitrariness) that states that the greatest lower probability bound that one accepts on a proposition should correspond to the aggregate degree of support for the proposition. In selecting $R(\Sigma)$, the preceding principle was applied at two key points. First, *non-arbitrariness* was invoked as a basis for requiring that one's rational magnitudes of belief *conform to* the ordering constraints encoded in $\succ_\Sigma$. I then argued that one should accept any proposition for which one has a defeasible reason, so long as there is no *credible defeater* for the proposition. I maintain that a lower probability bound is a *credible defeater* only if the probability bound does not 'violate' *the principle of difference*. A lower probability bound is *credible* in this sense, if the bound can be deduced from a set of probability bounds, Γ, where one has a defeasible reason for each element of Γ, and Γ *conform to* the set of constraints that are encoded in $\succ_\Sigma$. Since R(Σ) corresponds to a set of propositions for which one has defeasible reasons and no *credible* contradictory reasons, rationality requires that one accept each element of R(Σ).

CHAPTER 8: EPILOGUE - STATUS OF THE THREE PROBLEMS

1. The Problem of Uninformative Statistics

The Problem of Uninformative Statistics first presented itself as a difficulty for theories of direct inference that implement the tenet that it is the frequency statements serve as the statistical premises for direct inference, and the tenet that there is a default preference to use a frequency statement regarding a narrower reference class as one's statistical premise when engaging in direct inference. The paradigmatic example of the Problem of Uninformative Statistics concerns the role of statistics regarding a reference class containing only of the object about which one would like to draw a conclusion.

Previous approaches to the Problem of Uninformative Statistics have proposed alternative accounts of the sort of statistical statements that are fit to serve as the statistical premises for direct inference, and alternative *compatibility conditions* regarding the range of cases where statistics for a narrower reference class will defeat an instance of direct inference based on statistics for a broader reference class. One sort of compatibility condition maintains that statistics for a narrower reference class may defeat direct inference based on statistics for a broader reference class *only if* direct inference based on statistics for the narrower reference class would yield a conclusion that contradicts the conclusion that one could draw using statistics for the broader class. Another compatibility condition maintains that statistics for a narrower reference class may defeat direct inference based on statistics for a broader reference class *unless* the statistics for the broader reference class are more precise.

While the past approaches to the Problem of Uninformative Statistics have been successful in addressing the difficulties associated with unit set reference classes, the compatibility conditions for such approaches have proven to be insufficiently restrictive. In particular, a variant of example due to Mark Stone (the Ace Urn example) shows that statistics for a narrower reference class may defeat an instance of direct inference based on statistics for a broader reference class *even if* instances of direct inference based our statistics for the broader and narrower reference class yield *consistent* conclusions (and our statistics for the broader reference class be more precise).

In response to the Problem of Uninformative Statistics, I proposed a new compatibility condition that is derived from one way of thinking about the conditions under which statistics for a narrower reference class are genuinely uninformative in comparison to statistics for a broader reference class. The new compatibility condition yields the right conclusions in the face of unit set reference classes, and in the face of the Ace Urn example. Moreover, the new compatibility condition is defensible by appeal to the idea that direct inference depends on the assumption that a respective object of interest is relevantly indiscernible among the elements of a proposed reference class.

2. The Problem of Relevant Statistics

The Problem of Relevant Statistics arises in a range of cases where the use of *gerrymandered* reference or target classes allows one to draw incorrect conclusions using simple principles of direct inference, such as the ones proposed by Reichenbach.

My approach to the Problem of Relevant Statistics appeals to the same common sense understanding of the justificatory basis of direct inference involved in my approach to the Problem of Uninformative Statistics. As a remedy to the Problem of Relevant Statistics, I require, roughly, that an agent's judgment of the value (or possible values) of a respective statistical statement, S, not depend for its justification on reflection on the values of statistics for a proper subset of the reference class for S, in the case where the object of interest is known to be a member of the subset. In the case where such a condition fails, the object of interest is relevantly differentiated from other elements of the proposed reference class.

On the assumption that statistical induction is to be grounded in direct inference, one might expect the proposed remedy to Problem of Relevant Statistics to ramify and provide a remedy to Goodman's New Riddle of Induction. Unfortunately, while my proposed remedy to the Problem of Relevant Statistics handles all of the cases of gerrymandered statistics that have been presented in the literature on direct inference, the remedy, as currently formulated, appears not to cover Goodman cases. However, it seems likely that the two rules that I propose at the end of *Chapter Six* could be amended so that the resulting theory did yield a remedy to New Riddle of Induction. Working out the details of such an amendment is left as problem for future study.

3. The Problem of Competing Statistics

The Problem of Competing Statistics arises from the fact that multiple instances of direct inference may yield defeasible reasons for mutually inconsistent conclusions, where none of the respective inferences takes priority over any of the others. Past commentaries on the problem have proposed a *skeptical policy* which states that one should suspend judgment regarding respective outputs of direct inference in the cases where the outputs are inconsistent. While the recommendation to suspend judgment in cases where inconsistencies arise will keep us from accepting unjustified conclusions, it seems that suspending judgment in all such cases is too conservative, since there appear to be cases where we may reach reasonable substantive conclusions in cases where we have reasons for accepting each element of an inconsistent set propositions.

In response to Problem of Competing Statistics, I proposed an analogy that treats competing reasons as the analogues of physical objects possessed of competing dispositions. The chief function of the analogy was to guide our thinking in motivating an approach to the Problem of Competing Statistics that is less conservative than the skeptical policy. In light of the Kinetic Analogy, I proposed the Kinetic Theory of Competing Reasons. Through the application of dominance reasoning applied to the qualitative and quantitative features of competing defeasible reasons, the Kinetic Theory improves upon the skeptical policy in a wide range of cases. That said, the Kinetic Theory may itself be subject to improvements as it is likely that the belief formation policies recommended by the Kinetic Theory are themselves too conservative.

APPENDIX A: DEFINITIONS & PROOFS

A more precise definition of a *reason matrix* applies the notion of a *possibility of co-satisfiability* corresponding to a set of propositions.

Definition 0: w is a *possibility of co-satisfiability* with respect to a set of propositions, S, *just in case* w is consistent & ($\forall \alpha_k \in$S: $\alpha_k \in$w or $\neg\alpha_k \in$w) & ($\forall$w′$\subset$w: $\exists \alpha_k \in$S: $\alpha_k \notin$w′ & $\neg\alpha_k \notin$w′).

We also need conventions for denoting *ordered* versions of various sorts of important sets. Given a set, Σ, of defeasible reasons, "⟨Σ⟩" will denote the *canonical ordering* of the elements of Σ. So, ⟨Σ⟩ is an ordered set of the form ⟨ PROB(α_1) ≥ b_1, … , PROB(α_n) ≥ b_n ⟩. Given an ordered set ⟨Σ⟩, "[Σ]" will be used to denote the corresponding ordering of the object language propositions that are bounded by the elements of ⟨Σ⟩, and "B(Σ)" will be used to denote the corresponding ordering of the values of the probability *bounds* on the elements of ⟨Σ⟩. So, in the case where ⟨Σ⟩ = ⟨ PROB(α_1) ≥ b_1, … , PROB(α_n) ≥ b_n ⟩, we have [Σ] = ⟨ α_1, … , α_n ⟩, and B(Σ) = ⟨ b_1, … , b_n ⟩.

Given an ordered set of propositions, S (imagine that S is [Σ] corresponding to a set, Σ, of defeasible reasons), *define* W(S) = ⟨ w_1, … , w_m ⟩ to be the canonical ordering of the set of *possibilities of co-satisfiability* for the elements of S. Given an ordered set of propositions, S, where S = ⟨ α_1, … , α_n ⟩, *define* r(w_i) to be the ordered mulitset ⟨ v(α_1, w_i), … v(α_n, w_i) ⟩, where v(α_k, w_i) = 1, if $\alpha_k \in$ w_i, and v(α_k, w_i) = 0, otherwise. Now *define* ρ(S) to be ⟨ r(w_1), …, r(w_m) ⟩, and note that ρ(S), represents the set of possible *rows* of a truth table for the ordered set of propositions, S, where r(w_1) represents the first *row* of the table, r(w_2) the second, and so on. *Figure A* illustrates the manner in which ρ(S) represents the possible rows of a truth table for the set of propositions, S.

$$
\begin{array}{ll}
 & \begin{array}{ccccccc} S & = & \alpha_1 & \alpha_2 & \alpha_3 & \cdots & \alpha_n \end{array} \\
\rho(S) & \left\{ \begin{array}{ccccccc}
r(w_1) & = & v(\alpha_1, w_1) & v(\alpha_2, w_1) & v(\alpha_3, w_1) & \cdots & v(\alpha_n, w_1) \\
r(w_2) & = & v(\alpha_1, w_2) & v(\alpha_2, w_2) & v(\alpha_3, w_2) & \cdots & v(\alpha_n, w_2) \\
r(w_3) & = & v(\alpha_1, w_3) & v(\alpha_2, w_3) & v(\alpha_3, w_3) & \cdots & v(\alpha_n, w_3) \\
\vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\
r(w_m) & = & v(\alpha_1, w_m) & v(\alpha_2, w_m) & v(\alpha_3, w_m) & \cdots & v(\alpha_n, w_m)
\end{array} \right.
\end{array}
$$

Figure A

It is now possible to give a precise definition of the notion of a reason matrix:

Definition 1: The *reason matrix* for a set, Σ, of defeasible reasons is an ordered triple RM(Σ) = ⟨ [Σ], B(Σ), ρ([Σ]) ⟩.

Definition 2: $\alpha \approx_\Sigma \beta$ *if and only if* $\exists f,g$:
(i) f is a bijection from $[\Sigma]$ to $[\Sigma]$,
(ii) g is a bijection from $W([\Sigma])$ to $W([\Sigma])$,
(iii) $f(\alpha) = \beta$,
(iv) $\forall \alpha_i \in [\Sigma]$: $\forall w_k \in W([\Sigma])$: $v(\alpha_i, w_k) = v(f(\alpha_i), g(w_k))$, and
(v) $\forall \alpha_i \in [\Sigma]$: $f(\alpha_i) = \alpha_k \Rightarrow b_i = b_k$.

Theorem 1: $\forall \Sigma$: $\approx_\Sigma$ is (i) reflexive, (ii) symmetric and (iii) transitive.

Proof: (i) Assume that f and g are identity functions on $[\Sigma]$ and $W([\Sigma])$, respectively. Then f and g validate $\alpha \approx_\Sigma \alpha$, for all α in $[\Sigma]$. (ii) Assume that f and g validate $\alpha \approx_\Sigma \beta$. In that case, f^{-1} and g^{-1} validate $\beta \approx_\Sigma \alpha$. (iii) Assume that f and g validate $\alpha \approx_\Sigma \beta$, and that f' and g' validate $\beta \approx_\Sigma \chi$. In that case, $f \circ f'$ and $g \circ g'$ validate $\alpha \approx_\Sigma \chi$. ■

Theorem 2:
$\forall \alpha,\beta,\Sigma$: $\alpha \approx_\Sigma \beta \;\Rightarrow\; \max\{ r \mid \Sigma \models PROB(\alpha) \geq r \} = \max\{ r \mid \Sigma \models PROB(\beta) \geq r \}$.

Proof: Since $\approx_\Sigma$ is symmetric, its is sufficient to assume that $\alpha \approx_\Sigma \beta$, and prove that $\forall r$: $\Sigma \models PROB(\alpha) \geq r \Rightarrow \Sigma \models PROB(\beta) \geq r$. Assume that $\Sigma \models PROB(\alpha) \geq r$. Then $\exists S \subseteq \Sigma$: $S \models PROB(\alpha) \geq r$, where $S = \{ PROB(\alpha_1) \geq r_1, \ldots, PROB(\alpha_n) \geq r_n \}$. Assume that f and g validate $\alpha \approx_\Sigma \beta$. Then $\exists T \subseteq \Sigma$: $T = \{ PROB(\beta_1) \geq r_1, \ldots, PROB(\beta_n) \geq r_n \}$, and
$\forall i$: $f(\alpha_i) = \beta_i$. And, since $S \models PROB(\alpha) \geq r$, $T \models PROB(\beta) \geq r$. ■

Definition 3: Relative to Σ, a defeasible reason for a lower probability bound on β *has the potential to directly promote* the aggregate degree of support for α *just in case* $\exists \Gamma, s$:
(i) $\beta \in [\Sigma]$ & $\alpha \in [\Sigma]$,
(ii) Γ is a set of lower probability bounds, such that $[\Gamma] = [\Sigma]$,
(iii) $\Gamma \cup \{ PROB(\beta) \geq s \}$ is consistent, and
(iv) $\max\{ r \mid \Gamma \cup \{ PROB(\beta) \geq s \} \models PROB(\alpha) \geq r \} > \max\{ r \mid \Gamma \models PROB(\alpha) \geq r \}$.

Definition 4: Relative to Σ, a defeasible reason for a lower probability bound on β *has the potential to directly inhibit* the aggregate degree of support for α, relative to Σ, just in case $\exists \Gamma, s$:
(i) $\beta \in [\Sigma]$ & $\alpha \in [\Sigma]$,
(ii) Γ is set of lower probability bounds, such that $[\Gamma] = [\Sigma]$,
(iii) $\Gamma \cup \{ PROB(\beta) \geq s \}$ is consistent, and
(iv) $\max\{ r \mid \Gamma \cup \{ PROB(\beta) \geq s \} \models PROB(\neg\alpha) \geq r \} > \max\{ r \mid \Gamma \models PROB(\neg\alpha) \geq r \}$.

Definition 8: $\alpha \succ_\Sigma \beta$ *just in case* $\exists f,g$:
(i) f is a bijection from $[\Sigma]$ to $[\Sigma]$,
(ii) g is a bijection from $W([\Sigma])$ to $W([\Sigma])$,

(iii) $f(\alpha) = \beta$,
(iv) $\forall\alpha_i\in[\Sigma]:\forall w_k\in W([\Sigma])$: $v(\alpha_k, w_i) = v(f(\alpha_k), g(w_i))$,
(v) $\forall\alpha_i\in[\Sigma]$: ($\alpha_i$ is a friend of α (relative to Σ) & $f(\alpha_i) = \alpha_k) \Rightarrow b_i \geq b_k$,
(vi) $\forall\alpha_i\in[\Sigma]$: ($\alpha_i$ is an enemy of α (relative to Σ) & $f(\alpha_i) = \alpha_k) \Rightarrow b_i \leq b_k$, and
(vii) $\forall\alpha_i\in[\Sigma]$: ($\alpha_i$ is ambivalent to α (relative to Σ) & $f(\alpha_i) = \alpha_k) \Rightarrow b_i = b_k$.

Lemma 1: $\forall f,g$: (f and g satisfy conditions (i) thru (iv) (of the definition of $\succ$) relative to α, β and Σ) $\Rightarrow$ ($\forall$i,j: (α_i is a friend of/enemy of/ambivalent to α_j (relative to Σ)) $\Leftrightarrow$ ($f(\alpha_i)$ is a friend of/enemy of/ambivalent to $f(\alpha_j)$ (relative to Σ))).

Proof sketch: Note that application conditions for the relations 'α is a friend of β', etc. depend only the inferential relations between the elements of $[\Sigma]$. But the inferential relationship between α_i and α_j is identical to the inferential relationship between $f(\alpha_i)$ and $f(\alpha_j)$ (relative to $[\Sigma]$), when conditions (i) thru (iv) hold. ■

Lemma 2: $\forall\alpha,\beta,\Sigma,f,g$: if f and g validate $\alpha \succ_\Sigma \beta$, then
(i) $\forall\beta_i\in[\Sigma]$: β_i is a friend of $\beta \Rightarrow \exists\alpha_j\in[\Sigma]$: $f(\alpha_j) = \beta_i$ & α_j is a friend of α & $\alpha_j \succ_\Sigma \beta_i$,
(ii) $\forall\beta_i\in[\Sigma]$: β_i is an enemy of $\beta \Rightarrow \exists\alpha_j\in[\Sigma]$: $f(\alpha_j) = \beta_i$ & α_j is an enemy of α & $\beta_i \succ_\Sigma \alpha_j$, and
(iii) $\forall\beta_i\in[\Sigma]$: β_i is ambivalent to $\beta \Rightarrow$
$\exists\alpha_j\in[\Sigma]$: $f(\alpha_j) = \beta_i$ & α_j is ambivalent to α & $\alpha_j \succ_\Sigma \beta_i$ & $\beta_i \succ_\Sigma \alpha_j$.

Proof sketch for (i): We assume for arbitrary α, β, Σ that f and g validate $\alpha \succ_\Sigma \beta$, and assume for an arbitrary β_i in $[\Sigma]$ that β_i is a friend of β. Then $\exists\alpha_i$: $f(\alpha_i) = \beta_i$ and α_i is a friend of α.[Lemma 1] Now $\forall\chi\in[\Sigma]$: (χ is a friend of $\alpha_i \Rightarrow \chi$ is a friend of α) & (χ is an enemy of $\alpha_i \Rightarrow \chi$ is an enemy of α) & (χ is ambivalent to $\alpha_i \Rightarrow \chi$ is ambivalent to α). So, f and g prove that $\alpha_i \succ_\Sigma \beta_i$. Parts (ii) and (iii) follow by similar reasoning. ■

Theorem 4: $\forall\Sigma$: $\succ_\Sigma$ is (i) reflexive and (ii) transitive.

Proof: Let Σ be arbitrary. (i) If f and g are identity functions on $[\Sigma]$ and $W([\Sigma])$, then f and g validate $\alpha \succ_\Sigma \alpha$, for all α in $[\Sigma]$. (ii) Assume that f and g validate $\alpha \succ_\Sigma \beta$, and that f' and g' validate $\beta \succ_\Sigma \chi$. In that case, $\forall\alpha_i\in[\Sigma]$: ($\alpha_i$ is a friend of α (relative to Σ) & $f{\circ}f'(\alpha_i) = \alpha_k) \Rightarrow b_i \geq b_k$ (since $\forall\alpha_i\in[\Sigma]$: ($\alpha_i$ is a friend of α (relative to Σ) & $f(\alpha_i) = \alpha_j) \Rightarrow b_i \geq b_j$, and $\forall\alpha_j\in[\Sigma]$: ($\alpha_j$ is a friend of β (relative to Σ) & $f'(\alpha_j) = \alpha_k) \Rightarrow b_j \geq b_k$) [By Lemma 1]. Similarly, $\forall\alpha_i\in[\Sigma]$: ($\alpha_i$ is an enemy of α (relative to Σ) & $f{\circ}f'(\alpha_i) = \alpha_k) \Rightarrow b_i \leq b_k$, and $\forall\alpha_i\in[\Sigma]$: ($\alpha_i$ is ambivalent to α (relative to Σ) & $f{\circ}f'(\alpha_i) = \alpha_k) \Rightarrow b_i = b_k$. ■

Theorem 5:
$\forall\alpha,\beta,\Sigma$: $\alpha \succ_\Sigma \beta \Rightarrow \max\{ r \mid \Sigma \models \text{PROB}(\alpha) \geq r \} \geq \max\{ r \mid \Sigma \models \text{PROB}(\beta) \geq r \}$.

Proof: Assume that $\alpha \succ_\Sigma \beta$, and $\Sigma \models PROB(\beta) \geq r$. Then $\exists S \subseteq \Sigma$: $S \models PROB(\beta) \geq r$, and $S = \{ PROB(\beta_1) \geq r_1, \ldots, PROB(\beta_n) \geq r_n \}$, where each element of $\{ \beta_1, \ldots, \beta_n \}$ is a friend of β or ambivalent to β. Assume that f and g validate $\alpha \succ_\Sigma \beta$. Then $\exists T \subseteq \Sigma$:
$T = \{ PROB(\alpha_1) \geq s_1, \ldots, PROB(\alpha_n) \geq s_n \}$, where $\forall i$: $f(\alpha_i) = \beta_i$ & $s_i \geq r_i$. And, since $S \models PROB(\beta) \geq r$, $T \models PROB(\alpha) \geq r$.■

Lemma 3: $\forall \Gamma,\Sigma,\alpha,\beta$: ($\{ \alpha \mid \exists r: PROB(\alpha) \geq r \in \Gamma \} = \{ \alpha \mid \exists r: PROB(\alpha) \geq r \in \Sigma \}$ & Γ *conforms to* $\succ_\Sigma$ & $\alpha \succ_\Sigma \beta$ & $\Gamma' = \{ PROB(\chi) \geq s \mid s = \max\{ r \mid PROB(\chi) \geq r \in \Gamma \} \}$) $\Rightarrow$
$\alpha \succ_{\Gamma'} \beta$.

Proof: Suppose that f and g validate $\alpha \succ_\Sigma \beta$. Then f and g validate $\alpha \succ_{\Gamma'} \beta$. Indeed, conditions (i) through (iv) (of the definition of $\succ$) are satisfied, since $[\Gamma'] = [\Sigma]$. Now note that $\forall \alpha_j \in [\Sigma]$: ($\alpha_j$ is a friend of α & $f(\alpha_j) = \beta_i$) $\Rightarrow \alpha_j \succ_\Sigma \beta_i$.[Lemma 1 & the definition of $\succ$] And so, since Γ *conforms to* $\succ_\Sigma$, we have $\forall \alpha_j \in [\Gamma']$: ($\alpha_j$ is a friend of α (relative to Γ') & $f(\alpha_j) = \beta_i$) $\Rightarrow \max\{ s \mid PROB(\alpha_j) \geq s \in \Gamma\} \geq$
$\max\{ s \mid PROB(\beta_i) \geq s \in \Gamma \}$.[By Theorem 5] Thus, condition (v) is satisfied. Conditions (vi) and (vii) are satisfied, for reasons similar to condition (v).■

Lemma 4: $\forall \Gamma,\Sigma,\alpha,\beta$: ($\{ \alpha \mid \exists r: PROB(\alpha) \geq r \in \Gamma \} = \{ \alpha \mid \exists r: PROB(\alpha) \geq r \in \Sigma \}$ &
Γ *conforms to* $\succ_\Sigma$ & $\alpha \succ_\Sigma \beta$) $\Rightarrow \max\{ r \mid \Gamma \models PROB(\neg\beta) \geq s \} \geq$
$\max\{ r \mid \Gamma \models PROB(\neg\alpha) \geq s \}$.

Proof: Let $\Gamma' = \{ PROB(\chi) \geq s \mid s = \max\{ r \mid PROB(\chi) \geq r \in \Gamma \} \}$. Let f and g be a pair of functions that validate $\alpha \succ_{\Gamma'} \beta$. By Lemma 3, f and g exist. Now assume
$\Sigma \models PROB(\neg\alpha) \geq r$. Then $\exists S \subseteq \Sigma$: $S \models PROB(\neg\alpha) \geq r$, and $S = \{ PROB(\chi_1) \geq r_1, \ldots, PROB(\chi_n) \geq r_n \}$, where each element of $\{ \chi_1, \ldots, \chi_n \}$ is a enemy of β or ambivalent to β. In that case, $\exists T \subseteq \Sigma$: $T = \{ PROB(\delta_1) \geq s_1, \ldots, PROB(\delta_n) \geq s_n \}$, where $\forall i$: $f(\chi_i) = \delta_i$ & $s_i \geq r_i$. And since $S \models PROB(\neg\alpha) \geq r$, $T \models PROB(\neg\beta) \geq r$.■

Lemma 5:
$\forall \Gamma$: (Γ *conforms to* $\succ_\Sigma$ & $\Gamma \subseteq EXT(\Sigma)$, and $\forall \Gamma'$: $\Gamma \subset \Gamma' \Rightarrow \sim(\Gamma'$ *conforms to* $\succ_\Sigma$ &
$\Gamma' \subseteq EXT(\Sigma)$)) $\Rightarrow \forall \Gamma'$: $\Gamma \subset \Gamma'$: $\sim(\Gamma'$ is consistent & $\Gamma' \subseteq EXT(\Sigma)$).

Proof: Assume not. Then $\exists \Gamma$: Γ *conforms to* $\succ_\Sigma$ & $\Gamma \subseteq EXT(\Sigma)$ & ($\forall \Gamma'$: $\Gamma \subset \Gamma' \Rightarrow \sim(\Gamma'$ *conforms to* $\succ_\Sigma$ & $\Gamma' \subseteq EXT(\Sigma)$) & $\exists \Gamma'$: $\Gamma \subset \Gamma'$: Γ' is consistent & $\Gamma' \subseteq EXT(\Sigma)$. In that case, $\exists \alpha \in [\Sigma]$: $PROB(\alpha) \geq r \in EXT(\Sigma)$ & $PROB(\alpha) \geq r \notin \Gamma$ & $\Gamma \cup \{ PROB(\alpha) \geq r\}$ is consistent. In that case, $\exists \beta \in [\Sigma]$: $\beta \succ_\Sigma \alpha$ & $PROB(\beta) \geq r \notin \Gamma$ (or else $\Gamma \cup \{ PROB(\alpha) \geq r\}$ *conforms to* $\succ_\Sigma$). Let χ_1 be such that $\chi_1 \in [\Sigma]$ & $\chi_1 \succ_\Sigma \alpha$ & $PROB(\chi_1) \geq r \notin \Gamma$. Then $PROB(\chi_1) \geq r \in EXT(\Sigma)$ (since $PROB(\alpha) \geq r \in EXT(\Sigma)$ & $\chi_1 \succ_\Sigma \alpha$), and $\Gamma \cup \{ PROB(\chi_1) \geq r\}$ is consistent [by Lemma 4]. In that case, $\exists \beta \in [\Sigma]$: $\beta \succ_\Sigma \chi_1$ & $PROB(\beta) \geq r \notin \Gamma$ (or else

$\Gamma\cup\{$ PROB$(\chi_1) \geq r\}$ *conforms to* $\succ_\Sigma$). Let χ_2 be such that $\chi_2\in[\Sigma]$ & $\chi_2 \succ_\Sigma \chi_1$ & PROB$(\chi_2) \geq r \notin \Gamma$. (Moreover, PROB$(\chi_2) \geq r \in$ EXT(Σ), and $\Gamma\cup\{$ PROB$(\chi_2) \geq r\}$ is consistent.) In that case, $\exists\beta\in[\Sigma]$: $\beta \succ_\Sigma \chi_2$ & PROB$(\beta) \geq r \notin \Gamma$. However, the present chain of existent χ_i cannot continue infinitely, since Σ is finite. So, in the end, we have a final case, $\exists\beta\in[\Sigma]$: $\beta \succ_\Sigma \chi_n$ & PROB$(\beta) \geq r \notin \Gamma$. Let χ_{n+1} be such that $\chi_{n+1}\in[\Sigma]$ & $\chi_{n+1} \succ_\Sigma \chi_n$ & PROB$(\chi_{n+1}) \geq r \notin \Gamma$. But, in that case, PROB$(\chi_{n+1}) \geq r \in$ EXT(Σ), and $\Gamma\cup\{$ PROB$(\chi_{n+1}) \geq r\}$ is consistent. But it is *not* the case that $\exists\beta\in[\Sigma]$: $\beta \succ_\Sigma \chi_{n+1}$ & PROB$(\beta) \geq r \notin \Gamma$. Then, contrary to our hypothesis, $\exists\Gamma'$: $\Gamma\subset \Gamma'$ & Γ' *conforms to* $\succ_\Sigma$ & $\Gamma' \subseteq$ EXT(Σ).■

Theorem 7: $\forall\Sigma$: MRS$(\Sigma) = \cap\{\, \Gamma \mid \Gamma$ *conforms to* $\succ_\Sigma$ & $\Gamma \subseteq$ EXT(Σ) &
($\forall\Gamma'$: $\Gamma\subset \Gamma' \Rightarrow \sim(\Gamma'$ *conforms to* $\succ_\Sigma$ & $\Gamma' \subseteq$ EXT$(\Sigma)))\,\}$.

Proof: $\forall\varphi,\Sigma$: $\varphi\in$MRS$(\Sigma) \Leftrightarrow (\varphi\in$EXT$(\Sigma)$ & $\forall\psi\in$D(Σ): $\{\varphi, \psi\}$ is consistent). But ($\varphi\in$EXT(Σ) & $\forall\psi\in$D(Σ): $\{\varphi, \psi\}$ is consistent) $\Leftrightarrow$ ($\varphi\in$EXT(Σ) & ($\forall\Gamma$: ($\Gamma\subseteq$EXT(Σ) & Γ *conforms to* $\succ_\Sigma$) $\Rightarrow \Gamma\cup\{\varphi\}$ is consistent)). But ($\varphi\in$EXT(Σ) & ($\forall\Gamma$: ($\Gamma\subseteq$EXT(Σ) & Γ *conforms to* $\succ_\Sigma$) $\Rightarrow \Gamma\cup\{\varphi\}$ is consistent)) $\Leftrightarrow$ ($\varphi\in$EXT(Σ) & ($\forall\Gamma$:(Γ *conforms to* $\succ_\Sigma$ & $\Gamma \subseteq$ EXT(Σ), and $\forall\Gamma'$: $\Gamma\subset \Gamma' \Rightarrow \sim(\Gamma'$ *conforms to* $\succ_\Sigma$ & $\Gamma' \subseteq$ EXT$(\Sigma))) \Rightarrow \Gamma\cup\{\varphi\}$ is consistent)). But ($\varphi\in$EXT(Σ) & ($\forall\Gamma$:(Γ *conforms to* $\succ_\Sigma$ & $\Gamma \subseteq$ EXT(Σ), and $\forall\Gamma'$: $\Gamma\subset \Gamma' \Rightarrow \sim(\Gamma'$ *conforms to* $\succ_\Sigma$ & $\Gamma' \subseteq$ EXT$(\Sigma))) \Rightarrow \Gamma\cup\{\varphi\}$ is consistent)) $\Leftrightarrow$ ($\forall\Gamma$:(Γ *conforms to* $\succ_\Sigma$ & $\Gamma \subseteq$ EXT(Σ), and $\forall\Gamma'$: $\Gamma\subset\Gamma' \Rightarrow \sim(\Gamma'$ *conforms to* $\succ_\Sigma$ & $\Gamma'\subseteq$EXT$(\Sigma))) \Rightarrow \varphi\in \Gamma$).[Lemma 5]■

Lemma 6: $\forall\alpha,\beta,\Sigma$: $\alpha \succ_\Sigma \beta \Rightarrow \sup\{\, s \mid$ PROB$(\neg\beta)\geq s \in$ D$(\Sigma)\,\} \geq \sup\{\, s \mid$ PROB$(\neg\alpha)\geq s \in$ D$(\Sigma)\,\}$.

Proof: We consider arbitrary α, β, and Σ, such that $\alpha \succ_\Sigma \beta$. Let f and g be a pair of functions that validate that $\alpha \succ_\Sigma \beta$. Then $\forall\beta_i\in[\Sigma]$: β_i is an enemy of (or is ambivalent to) $\beta \Rightarrow \exists\alpha_i$: $f(\alpha_i) = \beta_i$ & α_i is an enemy of (or is ambivalent to) α & $\beta_i \succ_\Sigma \alpha_i$.[By Lemma 2] (Note that if χ has the potential to directly inhibit δ, then χ is an enemy of δ or χ is ambivalent to δ.) So, $\forall\Gamma\subseteq$EXT(Σ): (Γ satisfies $\succ_\Sigma$ & $\Gamma \models$PROB$(\neg\alpha)\geq s) \Rightarrow \exists\Gamma'\subseteq$EXT$(\Sigma)$: ($\Gamma'$*conforms to* $\succ_\Sigma$ & $\Gamma' \models$ PROB$(\neg\beta)\geq s$).■

Lemma 7:
$\forall\alpha,\beta,\Sigma$: $\alpha \succ_\Sigma \beta \Rightarrow \max\{\, s \mid$ PROB$(\alpha) \geq s \in$ MRS$(\Sigma)\,\} \geq \max\{\, s \mid$ PROB$(\beta) \geq s \in$ MRS$(\Sigma)\,\}$.

Proof: For arbitrary α, β, and Σ, we assume $\alpha \succ_\Sigma \beta$, and show that $\max\{\, s \mid$ PROB$(\alpha) \geq s \in$ MRS$(\Sigma)\,\} \geq \max\{\, s \mid$ PROB$(\beta) \geq s \in$ MRS$(\Sigma)\,\}$. We note that $\max\{\, s \mid$ PROB$(\chi) \geq s \in$ MRS$(\Sigma)\,\} = \min\{\, \max\{\, s \mid$ PROB$(\chi) \geq s \in$ EXT$(\Sigma)\,\}, 1 - \sup\{\, s \mid$ PROB$(\neg\chi) \geq s \in$ D$(\Sigma)\,\}\,\}$. But if $\alpha \succ_\Sigma \beta$, then $\sup\{\, s \mid$ PROB$(\neg\beta) \geq s \in$ D$(\Sigma)\,\} \geq \sup\{\, s \mid$ PROB$(\neg\alpha) \geq s \in$ D(Σ)

}.[By Lemma 6] But clearly, if $\alpha \succ_\Sigma \beta$, then max{ s | PROB($\alpha$) $\geq$ s $\in$ EXT(Σ) } $\geq$ max{ s | PROB(β) $\geq$ s $\in$ EXT(Σ) }.■

Theorem 9: $\forall\Sigma$: R(Σ) *conforms to* $\succ_\Sigma$.

Proof: We prove, for arbitrary α, β, and Σ, such that $\alpha \succ_\Sigma \beta$, that max{ s | PROB($\alpha$)$\geq$s $\in$ R(Σ) } $\geq$ max{ s | PROB(β)$\geq$s $\in$ R(Σ) }. Let *f* and *g* be a pair of functions that validate that $\alpha \succ_\Sigma \beta$. Then $\forall\beta_i\in[\Sigma]$: β_i is a friend of (or ambivalent to) $\beta \Rightarrow \exists\alpha_i$: ($f(\alpha_i) = \beta_i$ & α_j is a friend of (or ambivalent to) α & $\alpha_i \succ_\Sigma \beta_i$).[By Lemma 2] (Note that if χ has the potential to directly support δ, then χ is a friend of δ or χ is ambivalent to δ.) Moreover, $\forall\chi,\delta,\Sigma$: $\chi \succ_\Sigma \delta \Rightarrow$ max{ s | PROB(χ) $\geq$ s $\in$ MRS(Σ) } $\geq$ max{ s | PROB(δ) $\geq$ s $\in$ MRS(Σ) }.[By Lemma 7] So, $\forall\Gamma\subseteq$MRS(Σ): $\Gamma \models$ PROB(β) $\geq$ s $\Rightarrow \exists\Gamma'\subseteq$MRS($\Sigma$): $\Gamma' \models$ PROB(α) $\geq$ s.■

Theorem 11: $\forall\Sigma$: Skp(Σ) $\subseteq$ R(Σ).

Proof: Let MRS$^-$(Σ) = $\cap${ Γ | Γ is consistent & $\Gamma \subseteq$ EXT(Σ), and $\forall\Gamma'$: $\Gamma\subset\Gamma' \Rightarrow$ ~(Γ' is consistent & $\Gamma' \subseteq$ EXT(Σ)) }. In that case, $\forall\Sigma$: $\cap$M-con(Σ) $\subseteq$ MRS$^-$(Σ), since, for all Σ, MRS$^-$(Σ) is just the intersection of a set of supersets of elements of M-con(Σ). But MRS$^-$(Σ) $\subseteq$ MRS(Σ), since, MRS(Σ) is the intersection of a set of sets, S, and MRS$^-$(Σ) is the intersection of a set of sets, T, where S $\subseteq$ T.■

REFERENCES

Bacchus, Fahiem: 1990, *Representing and Reasoning with Probabilistic Knowledge*, The MIT Press, Cambridge, Massachusetts.

Bacchus, F., Grove, A.J., Halpern, J., and Koller, D.: 1996, ‘From statistical knowledge bases to degrees of belief’, *Artificial Intelligence* 87:1-2, 75-143.

Carnap, Rudolph: 1952, *The Continuum of Inductive Methods*, University of Chicago Press, Chicago.

Goodman, N.: 1955, *Fact, Fiction, and Forecast*, Harvard University Press, Cambridge.

Hacking, I.: 1975, *The emergence of probability: a philosophical study of early ideas about probability, induction, and statistical inference*, Cambridge University Press, New York.

Halpern, J. and Koller, D.: 1995, ‘Representation Dependence and Probabilistic Inference’ in *Proceedings of the 14th International Joint Conference on Artificial Intelligence (IJCAI-95)*, 1853-1860.

Howson, C.: 1976, ‘The Development of Logical Probability’, in Cohen, R.S. et al. (eds.), *Essays in Memory of Imre Lakatos*, D. Reidel Publishing Company, Dordrecht-Holland.

Jackson, F.: 1975, ‘Grue’, *Journal of Philosophy*, v. 72, 113-131.

Jaynes, E.T.: 1973, 'The Well-Posed Problem', Foundations of Physics, 3, 477-493.

Keynes, John Maynard: 1921, *A Treatise on Probability*, Macmillan & Co, London.

Jeffrey, Richard: 1966, ‘Goodman’s Query’, *Journal of Philosophy*, v. 60, 452-462.

Kyburg, H.E.: 1956, ‘The Justification of Induction’, *Journal of Philosophy*, 53, 394-400.

Kyburg, H.E.: 1974, *The Logical Foundations of Statistical Inference*, Reidel Publishing Company, Dordrecht-Holland.

Kyburg, H.E., and Teng, C.M: 2001, *Uncertain Inference*, Cambridge University Press, Cambridge.

Levi, I.: 1982, ‘Direct Inference’, *Journal of Philosophy*, v. 74, 5-29.

McGrew, T.: 2001, ‘Direct Inference and the Problem of Induction’, *The Monist*, Volume 84, Number 2.

Paris, J.B. and Vencovská, A.: 1990, ‘A Note on the Inevitability of Maximum Entropy’, *International Journal of Approximate Reasoning* 4, 183-223.

Paris, J.B. and Vencovská, A.: 1992, ‘A Method for Updating that Justifies Minimum Cross Entropy’, *International Journal of Approximate Reasoning* 7, 1-18.

Paris, J.B. and Vencovská, A.: 1997, ‘In defence of the Maximum Entropy Inference Process’, *International Journal of Approximate Reasoning* 17, 77-103.

Pollock, John L.: 1983, ‘A Theory of Direct Inference’, Theory and Decision, Volume 15, Number 1, 29-95.

Pollock, John L.: 1990, *Nomic Probability and the Foundations of Induction*, Oxford University Press, New York.

Pollock, John: 1994, 'The Projectability Constraint', in D. Stalker (ed.), *Grue!: the New Riddle of Induction*, Open Court Press, Chicago, Illinois.

Pollock, John L.: 1995, *Cognitive Carpentry: A Blueprint for How to Build a Person*, The MIT Press, Cambridge, Massachusetts.

Pollock, John L. and Cruz, Joseph: 1999, *Contemporary Theories of Knowledge*, Rowman &Littlefield Publishers, Inc., Lanham, Maryland.

Ramsey, F. P.: 1931, ‘Truth and probability’, in *The Foundations of Mathematics: Collected Papers of Frank P. Ramsey*, Routledge and Kegan, London.

Reichenbach, Hans: 1949, *A Theory of Probability*, Berkeley University Press, Berkeley.

Reiter, R.: 1980, 'A logic for default reasoning', *Artificial Intelligence* 13, 181-132.

Salmon, Wesley: 1966, *The Foundations of Scientific inference*, University of Pittsburgh Press, Pittsburgh.

Salmon, Wesley: 1970, 'Statistical Explanation', in W. Salmon (ed.), *Statistical Explanation and Statistical Relevance*, University of Pittsburgh Press, Pittsburgh.

Savage, L.J.: 1954, *The Foundations of Statistics*, Wiley, New York, New York.

Shafer, Glen: 1976, *A Mathematical Theory of Evidence*, Princeton University Press, Princeton.

Skyrms, Brian: 1986, *Choice and Chance: an Introduction to Inductive Logic* (third edition), Wadsworth Publishing.

Stone, Mark: 1987, 'Kyburg, Levi, and Petersen' *Philosophy of Science*, Vol. 54, No. 2., 244-255.

Stove, David: 1986, *The Rationality of Induction*, Oxford: Clarendon.

Touretzky, D., Horty, J. and Thomason, R.: 1987, 'A clash of intuitions: the current state of monotonic multiple inheritance systems', *Proceedings of the Tenth international Joint Conference on Artificial Intelligence (Proceedings IJCAI-87)*, 476-482.

Van Fraassen, Bas: 1989, *Laws and Symmetry*. Oxford University Press, Oxford.

Venn, John: 1866, *The Logic of Chance*. Chelsea Publishing Company, New York, New York.

von Mises, Richard: 1939, *Probability, Statistics, and Truth*. Translated by J. Neyman, D. Sholl, and E. Rabinowitsch, William Hodge & Co. ,Ltd., Edinburgh.

Williams, Donald: 1947, *The Ground of Induction*. Harvard University Press, Cambridge, Massachusetts.

Williamson, Jon: 2003, 'Bayesianism and Language Change', *Journal of Language, Logic and Information* 12, 53-97.