

Two Problems of Direct Inference

Paul D. Thorn

Abstract

The article begins by describing two longstanding problems associated with direct inference. One problem concerns the role of uninformative frequency statements in inferring probabilities by direct inference. A second problem concerns the role of frequency statements with gerrymandered reference classes. I show that past approaches to the problem associated with uninformative frequency statements yield the wrong conclusions in some cases. I propose a modification of Kyburg's approach to the problem that yields the right conclusions. Past theories of direct inference have postponed treatment of the problem associated with gerrymandered reference classes by appealing to an unexplicated notion of *projectability*. I address the lacuna in past theories by introducing criteria for being a *relevant statistic*. The prescription that only *relevant statistics* play a role in direct inference corresponds to the sort of projectability constraints envisioned by past theories.

1. Introduction

It is common to use frequency information to form probability judgments. For example, given the premise that 5% of dogs have fleas, I may (in some circumstances) justifiably infer that the probability is 0.05 that my neighbor's dog, Flint, has fleas. The inferences that we make when we draw a conclusion about the probability of a proposition on the basis of frequency information are called "direct inferences".¹ Although the expression "direct inference" is

¹ I will assume that the conclusions of direct inference are single-case probability statements whose truth conditions (or acceptability conditions) are implicitly relativized to the epistemic situation of respective agents. The proposed account of direct inference could also be formulated so that the conclusions of direct

rarely used, direct inference has many applications and is widely used in areas such as insurance pricing, weather forecasting, and medical diagnosis.

A major obstacle to defending the objectivity of direct inference is called “the Problem of the Reference Class”. This problem derives from the fact that every object may be located in many different reference classes, and from the fact that direct inference using frequency information for different reference classes will often yield mutually inconsistent conclusions. For example, in the case regarding my neighbor’s dog, the conclusion that the probability is 0.05 that Flint has fleas is based on my frequency information about the set of dogs. But Flint is a member of numerous reference classes (in addition to the set of dogs), such as the set of small-breed dogs, the set of dachshunds, the set of brown dogs, etc., and direct inference based on frequency information for the different reference classes may lead to mutually inconsistent conclusions.²

The classic response to the Problem of the Reference Class derives from Hans Reichenbach.³ When preparing to make a direct inference, Reichenbach recommended that one base one’s inference on the narrowest reference class for which one is able to make a reliable frequency judgment (Reichenbach 1949, p. 374).⁴ For example, if I am only able to make reliable judgments about the frequency of dogs having fleas and about the frequency of dachshunds having fleas, I should apply the latter frequency judgment in forming a belief about the probability that Flint, a dachshund, has fleas.

Where “ $\text{PROB}(c \in T)$ ” denotes the probability that an object, c , is an element of a set, T , and “ $\text{freq}(T|R)$ ” denotes the relative frequency of elements of a set, T , among a set, R , the essence of Reichenbach’s theory may be encapsulated by two principles. Following Pollock (1990), I call the second principle “subset defeat”, since it states the conditions under which a

inference are tantamount to defeasible prescriptions about what an agent’s degrees of belief in given propositions should be.

² The Problem of the Reference Class is often presented as a decisive objection to the objectivity of direct inference. But arguments for the claim that the problem is decisive typically go no further than describing the problem as I have here (cf. Fitelson *et al.* 2005; Rhee 2007). See also (Hájek 2007, p. 568-9), where skepticism about direct inference vis-à-vis the Problem of the Reference Class is premised on idiosyncratic features of Reichenbach’s account of direct inference. A presentation of the problem by means of an interesting example can be found in (Colyvan and Regan 2001) and (Colyvan *et al.* 2007), though the theory of direct inference presented here, and the theory presented in (Pollock 1990), have the resources to adequately address the example.

³ The key elements of Reichenbach’s account of direct inference are also present in (Venn 1866).

⁴ Almost all accounts of direct inference adopt some form of this prescription, which is closely related to the *principle of specificity* which is an element of many approaches to defeasible reasoning in the field of artificial intelligence (cf. Horty *et al.* 1990; Kraus *et al.* 1990; Geffner and Pearl 1992). Even Salmon’s proposal that direct inferences be based on the broadest homogeneous reference class is similar to Reichenbach’s proposal inasmuch as: (1) the non-homogeneity of a proposed reference class compels one to reason from statistics for a proper subset of that reference class (if a direct inference is possible), and (2) the possession of statistics for a homogeneous subset of a proposed reference class demonstrates the non-homogeneity of the proposed reference class, if the statistics for the two sets differ (Salmon 1971).

proposed direct inference, based on frequency information for a given reference class, is defeated in virtue of frequency information for a subset of that reference class:⁵

[RDI] Reichenbachian Direct Inference

If an agent, A , is justified in believing that $\text{freq}(T|R) = u$ and that $c \in R$, then A has a defeasible reason for believing that $\text{PROB}(c \in T) = u$.

[RSD] Reichenbachian Subset Defeat

A respective instance of [RDI] is defeated for A , if there is an R' such that A is justified in believing:

- (i) $c \in R'$,
- (ii) $R' \subseteq R$, and
- (iii) $\text{freq}(T|R') \neq u$.

Although Reichenbach's approach to the Problem of the Reference Class has served as a touchstone for subsequent studies of direct inference, his approach is known to be limited in a number of respects. The most well known limitation of Reichenbach's theory concerns cases where one has reliable information regarding the incidence of a particular property among two overlapping reference classes, and one is unable to make a reliable judgment about the incidence of the property among the intersection of the two sets. For example, suppose I want to know how likely it is that Flint, my neighbor's dachshund, will live at least twelve years. Let us suppose that I do not have any information about the mortality rate of dachshunds, but I do know that sixty percent of small-breed dogs live at least twelve years, and I also know that forty percent of boarhounds live at least twelve years. (Suppose that I know that a dachshund is a type of small-breed boarhound, and that there are some large-breed boarhounds.) In this case, the prescription to prefer narrower reference classes is unhelpful, since neither of the two candidate reference classes is narrower than the other.

In cases where there is no narrowest relevant reference class about which one can make a reliable frequency judgment, Reichenbach prescribed that one not form a judgment regarding the probability that a given object is an element of a respective target class (Reichenbach 1949, p. 375). Reichenbach's prescription will keep us from forming unjustified beliefs. On

⁵ According to Reichenbach's official view, the statistical statements that may serve as premises for direct inference are statements of *frequency in the limit*. Roughly: *limiting frequencies* are defined relative to an infinite sequence, R , and the limiting frequency of T among R is defined as the frequency of elements of T among the first n elements of R as n approaches ∞ (provided that the frequency of T among R goes to a limit as n approaches ∞). In what follows, I will mostly ignore this detail of Reichenbach's theory, and acknowledge Reichenbach's official view only at relevant points.

the other hand, there may be cases where it is possible to draw a reasonable conclusion even if our body of data includes frequency statements that individually support mutually inconsistent conclusions. For example, in a case where one knows that $c \in R_1 \cap R_2$, that $\text{freq}(T|R_1) = 0.4$, and that $\text{freq}(T|R_2) = 0.6$, and one has no other frequency information relevant to $\text{PROB}(c \in T)$, it appears reasonable to conclude that $\text{PROB}(c \in T) \in [0.4, 0.6]$ (cf. Kyburg and Teng 2001; Thorn 2007).

Another difficulty with Reichenbach's account of direct inference concerns the role of *gerrymandered* reference classes, and is similar (at least superficially) to the problem of induction that Goodman uncovered (Goodman 1955). Henry Kyburg was the first person to notice this problem (Kyburg 1961).⁶ Typical examples of this problem, in the case of direct inference, involve a reference class that is described as the union of (1) the unit set of the object about which one wishes to draw a conclusion, and (2) a set of objects that is known to have a very high (or very low) incidence of elements of a respective target class. For example, in a case where one is trying to draw a conclusion about the probability that Flint will live twelve years, a gerrymandered reference class (that illustrates the problem) would be the set composed of Flint and all of the dachshunds whose life span is less than twelve years. The frequency of dogs that live at least twelve years among such a gerrymandered reference class is guaranteed to be nearly *zero*. While the gerrymandered reference class is narrower than the other reference classes mentioned earlier, one should not rely on frequency information for this reference class in judging how likely it is that Flint will live twelve years.

Theories of direct inference that have been proposed since Reichenbach have postponed treatment of the preceding problem.⁷ It is typical to regard the problem as analogous to the projectability problem associated with induction, and argue that a theory such as Reichenbach's must be amended to require that correct direct inferences be formulated using target and reference classes that correspond to projectable predicates (Kyburg and Teng 2001) or projectable properties (Pollock 1990).^{8,9} In a similar vein, Bacchus (1990) proposed that

⁶ Kyburg also showed that gerrymandered target classes are sufficient to lead to unreasonable direct inferences (Kyburg 1974). My approach to the 'projectability' problems associated with direct inference applies equally to problems generated by gerrymandered target classes.

⁷ Although Salmon does not discuss the problem of gerrymandered reference classes, his account of direct inference (Salmon 1971, 1977, and 1984) is adequate to address *some* of the projectability problems associated with direct inference. However, there is no obvious way to modify Salmon's restrictive paradigm, which identified reference classes with infinite sequences of temporally ordered events, in order to apply Salmon's account of direct inference to typical cases where we would like to use frequency information about a population that is not temporally ordered to make a direct inference regarding one of its members.

⁸ Hempel's closely related account of inductive-statistical explanation also appeals to an unexplicated notion of *law-like* statistical generalizations as a proxy for projectability constraints (Hempel 1968).

⁹ In (Pollock 2007) and other unpublished essays, Pollock developed an approach to direct inference that appears not to rely on projectability constraints in the same manner as his earlier approach (Pollock 1990). However, Pollock's more recent approach still relies, at the foundational level, on a variety of direct

the problem calls for a “theory of relevance” that will allow us to recognize misleading statistical statements, and thereby bar their use as premises for direct inference.

A final inadequacy of Reichenbach’s theory concerns the role of uninformative and less informative frequency statements. This problem was first described by Kyburg, and is easily grasped when one reflects carefully on Reichenbach’s proposal that frequency data regarding narrower reference classes is to be preferred as a basis for direct inference. For one, consider the reference class consisting of the unit set containing 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 we should always prefer frequency data about unit set reference classes, and in that case, all interesting instances of direct inference would be defeated. Indeed, the frequency of elements of a given target class among a unit set reference class will always be one or zero, and direct inference based on such reference classes would seem only to allow the conclusion that a respective probability is one or zero.¹⁰

The focus of the discussion that follows will be on outlining a theory of direct inference that remedies the problem associated with uninformative frequency statements. I will also briefly address the problem associated with gerrymandered reference classes. Although I regard the problem involving gerrymandered reference classes as distinct from the one involving uninformative frequency statements, it is difficult to address the latter problem without touching upon the former, since the former problem is omnipresent, in the following sense: Every case is a case where it is possible to introduce a frequency statement with a gerrymandered reference class that will lead to an unreasonable conclusion if it is used as a premise for a direct inference. I will not provide a detailed treatment of the problem of making a direct inference in cases where there is no narrowest relevant reference class about which one can make a reliable frequency judgment, although I will comment briefly on this problem in the closing section of the article.

inference (what Pollock calls “the statistical syllogism”) that invokes an unexplicated notion of *projectability*. Pollock’s recent work also bears similarities to the *random-worlds* approach to inferring single-case probabilities, as proposed in (Bacchus *et al.* 1996) and (Halpern 2003), inasmuch as the outputs of Pollock’s approach are sensitive to the manner in which inputs to the theory are represented. The indifference principles presented in (Pollock *unpublished*) are particularly suggestive of this problem.

¹⁰ To generate a corresponding problem for Reichenbach’s official view (which identifies reference classes with infinite sequences), we need only consider the limiting frequency for a given target property among a reference class defined by an exhaustive listing of the known properties of the individual about which we wish to make a direct inference. For such reference classes, we are rarely in a position to make an informative judgment regarding the value of the relevant limiting frequency. See (Fetzer 1977), for a discussion of the present problem in connection with a close reading of Reichenbach’s official view.

2. Past Approaches to 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 direct inferences based on uninformative frequency information generally yield conclusions that are consistent with the intuitively correct conclusions based on informative frequency statements for broader reference classes (or, alternatively, explain why direct inferences based on uninformative frequency information for narrow reference classes are generally defeated). Second, a solution must explain why uninformative frequency information for narrow reference classes does not ordinarily undermine (via a principle such as [RSD] Reichenbachian Subset Defeat) direct inferences that are based on informative frequency statements for broader reference classes.

Kyburg's proposed remedy to the Problem of Uninformative Statistics is the most well known, and was the first to appear in the literature. It is also representative of other proposals that have appeared since. Kyburg's approach has two parts. First, Kyburg proposed that the only statements that may serve as statistical premises for direct inference are statements that describe a relevant frequency as residing within an interval. In other words, a statistical premise for direct inference is always of the form: $\text{freq}(T|R) \in [r, s]$. Second, Kyburg maintained (modulo projectability considerations) that a frequency statement for a relevant narrower reference class will defeat a direct inference based on a broader reference class if and only if the range of values judged to be possible for the broader class is not a subset of the range of values judged to be possible for the narrower class.¹¹ For example, if one knows that the frequency of university degree holders among Californians is 0.25, and $[0.1, 0.4]$ is the narrowest interval within which one may locate the frequency of university degree holders among southern Californians, then Kyburg's theory deems it permissible to use one's frequency information about Californians to draw a conclusion about the likelihood that a particular southern Californian has a university degree.

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 thereby provides a possible means to dissolving the problem associated with unit set reference classes, since the interval $[0, 1]$ will be less precise than any interval associated with a frequency statement that we would like to use as a premise for direct inference.

¹¹ It is understood that the range of values for the broader class and for the narrower class are interval-valued.

Two Problems of Direct Inference

Other approaches to the Problem of Uninformative Statistics (including my own) are similar to Kyburg's. Like Kyburg's approach, the other approaches combine two parts: (1) a thesis about the sort of statistical statements that may serve as premises for direct inference, and (2) a thesis about the conditions under which a statistical statement about a narrower reference class will defeat an instance of direct inference based on statistics for a broader class. For each approach, the first thesis is intended to ensure that uninformative statistical statements for narrower reference classes do not yield conclusions (via direct inference) that will contradict the correct conclusions based on informative statistical statements for broader reference classes. The second thesis is intended to ensure that uninformative statistical statements for narrow reference classes do not undercut the correct direct inferences (via subset defeat). I will refer to the conditions under which a statistical statement about a narrower reference class results in the subset defeat of a direct inference based on statistics for a broader class as the "incompatibility conditions" for proposed reference classes and their subsets. The idea is that statistics for a narrower reference class, R' , will defeat a direct inference based on statistics for a broader class, R , *only if* R and R' are *incompatible* in a relevant respect.

The two existent alternatives to Kyburg's approach to the Problem of Uninformative Statistics were outlined by Pollock (1990) and Bacchus (1990). Unlike Kyburg who maintained that interval-valued frequency statements are the proper premises for direct inference, Pollock assigned the privileged role to statements of *nomie probability*, and Bacchus assigned the role to statements of *expected frequency*. The accounts of Pollock and Bacchus also differ from the account of Kyburg concerning the conditions under which statistics for a narrower reference class will defeat a direct inference based on a broader reference class. Both Pollock and Bacchus maintain (modulo projectability considerations) that a direct inference is defeated by statistics for a narrower reference class if and only if direct inference based on statistics for the narrower reference class would yield a conclusion that is inconsistent with the one that would have otherwise been drawn using statistics for the broader class.

As it turns out, the incompatibility conditions proposed by Kyburg, Pollock, and Bacchus are each too permissive (since they each allow cases where a direct inference goes undefeated when it should not). On the other hand, each proposal regarding the proper statistical premises for direct inference is in some sense 'workable' as a partial solution to the Problem of Uninformative Statistics.¹² For the moment, I will focus on demonstrating the problem with

¹² The workability of the doctrines of Pollock and Bacchus regarding the sort of statistical statements that are appropriate to serve as premises for direct inference is similar to the doctrine of Kyburg. For all three accounts: (1) statistical statements of the preferred sort may take on values other than *one* and *zero* in the case of unit set reference classes, and (2) the use of known frequencies in the course of direct inference is usually

the incompatibility conditions that have been proposed in the past, and on proposing a new incompatibility condition that delivers the right conclusions. Later on, I will return to briefly argue in favor of the proposal that it is statements of *expected frequency* that properly serve as the statistical premises for direct inference.

3. The ACME Urn Example

A problem with past approaches to the Problem of Uninformative Statistics can be illustrated by a simple example.¹³ Suppose that one is certain that the following propositions are true:

- [1] Many urns exist that were produced by the ACME Urn Company.
- [2] Many of the urns produced by the ACME Urn Company contain balls.
- [3] 51% of all of the balls held in urns produced by the ACME Urn Company are red.
- [4] b is a ball held in an urn produced by the ACME Urn Company.
- [5] The urn, U_b , that contains b contains exactly one hundred balls.

Now make the further assumption that one lacks any additional information about the ACME Urn Company, about the likely distributions of balls of various colors held within urns produced by the ACME Urn Company, and, generally, any information that is relevant to the probability that b is red, that is not already implicit in [1] through [5]. In that case, theories of direct inference prescribe that we assign probability 0.51 to the proposition that b is red.¹⁴ This is the correct conclusion to draw in the present case. It is, of course, unreasonable to think that the relative frequency of red balls among U_b is 0.51. But because it is correct to regard the set of ball in U_b as an unexceptional (one hundred member) subset of the set of balls held in ACME urns (relative to the relative frequency of red balls), it is reasonable to apply our information regarding the frequency of red balls among the full set of balls held in ACME urns in order to conclude that the probability is 0.51 that b is red.

permitted, since point-valued frequency statements usually entail a respective preferred statistical statement of identical (or similar) value.

¹³ The example presented here is adapted from (Stone 1987).

¹⁴ One exception is the theory of Isaac Levi (1982). Levi proposed, roughly, 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 varying results with certain chances, and that the probability we assign to c having a respective target property be identical to the chance of a trial of the respective sort (i.e., a trial which has the relevant reference property) having the respective target property. Since these preconditions are not satisfied in the ACME Urn case, Levi's theory will not permit a direct inference. Levi's theory sets an extremely high threshold to surpass before one is permitted to make a direct inference. I share the view of other advocates of direct inference that Levi sets the bar too high.

Now suppose that one has additional information regarding U_b . In particular, suppose that one is able to inspect the contents of U_b under conditions that allow one to determine, with certainty, the number of balls in U_b that are white. Imagine, for example, that one is permitted to inspect the contents of U_b under unusual lighting conditions which permit one to determine, for each ball, whether or not it is white, and nothing else. As a result, suppose one determines that U_b contains exactly forty-nine white balls. It is thereby correct to conclude that the probability is *zero* that the frequency of red balls among U_b is greater than 0.51.

Given the additional information gained by one's inspection of the elements of U_b , the theories of Kyburg, Pollock, and Bacchus each agree that one's judgment regarding the probability that b is red should not change, and each theory permits one to draw the conclusion that the probability that b is red is 0.51.¹⁵ But that conclusion is unreasonable. Given our new information, it is still *incorrect* to assume that the relative frequency of red balls among U_b is 0.51. Moreover, because there is *zero* probability that the frequency of red balls among U_b is greater than 0.51, it is no longer reasonable to treat U_b as an unexceptional (one hundred member) subset of the set of balls held in ACME urns, and thereby assign probability 0.51 to the proposition that b is red, based on the frequency of red balls among the full set of balls held in ACME urns. Without appealing to that basis, it is unreasonable to assign probability 0.51 to the proposition that b is red.

4. Relative Informativeness

We are faced with the problem of determining the sort of *incompatibility* (between a reference class and one of its subsets) that will result in the defeat of a direct inference. According to Kyburg's theory, we have subset defeat only when our frequency information for a subset of a proposed reference class is more informative than it should be. Specifically, our frequency information for a relevant subset of a proposed reference class is deemed *too informative* if and only if the range of values judged to be possible for the broader class is not a subset of the range of values judged to be possible for the narrower class. I concur with Kyburg's idea that subset defeat only occurs when our frequency information for a subset of a proposed reference class is too informative. But in order to extricate ourselves from the problem

¹⁵ According to Kyburg's theory, the direct inference based on the set of all balls held in ACME Urns goes through, since one's frequency information regarding the set of balls held in U_b (namely, that the frequency of red balls among U_b is in $[0,0.51]$) is less precise. The theories of Pollock and Bacchus, respectively, permit us to draw the conclusion that the nomic probability and the expected frequency of red balls among the set of all balls held in ACME Urns is 0.51. Since the theories of Pollock and Bacchus do not permit a direct inference based on the set of balls held in U_b that contradicts the conclusion that the probability that b is red is 0.51, the conclusion that the probability that b is red is 0.51 goes through, according to their theories.

presented by the ACME urn example, I propose that we apply a different test than Kyburg's. In particular: where A is an agent, T a given target class, R a proposed reference class, and R' is a subset of R (where R' contains c , the object of interest), A 's information regarding R' is *too informative* if it is not the case that, for all U and V , A is justified in accepting $\text{PROB}(\text{freq}(T|R') \in U) \in V$ if and only if A would be justified in accepting $\text{PROB}(\text{freq}(T|R^*) \in U) \in V$, in a situation identical to A 's actual situation save that the name " R^* " is introduced to A by a definite description that confers only the information that R^* is a subset of R , and that R^* is the same size as R' .¹⁶

It is intended that the preceding 'informativeness' test be triggered in cases where an agent has information about R' that makes R' an *exceptional* subset of R vis-à-vis the incidence of elements of T (from the agent's perspective). So the test is triggered in cases where an agent has *any* information about the possible values of $\text{freq}(T|R')$ that is not entailed by the agent's judgment that R' is a subset of R , and the agent's judgments regarding the possible values of $\text{freq}(T|R)$, and the possible sizes of R and R' . The condition is also in sync with a natural conception of the justificatory basis of direct inference. In particular, when one makes a direct inference about an object, one assumes that the object is as likely to have a given target property as an object that is drawn *at random* from the proposed reference class.¹⁷ Given such an assumption, direct inference using frequency information for a given reference class is permissible only if the object about which one is reasoning is in relevant respects *indiscernible* from the other elements of the proposed reference class. Corresponding to this conception of the justificatory basis of direct inference, we see that cases where the proposed informativeness test is triggered are cases where an object of interest, c , is relevantly discernible among R , since, in such cases, c 's membership in R' relevantly distinguishes c from the elements of R that are not elements of R' .

In addition to capturing an intuitively correct criterion for when an agent's frequency information for a subset of a set is too informative (relative to a proposed direct inference), the test properly handles the case of unit set reference classes. Indeed, consider any case where it is correct to infer $\text{PROB}(c \in T) = r$, by direct inference from the premises $c \in R$ and $\text{freq}(T|R) = r$. In such cases, our narrowest estimate of the set of possible values for $\text{freq}(T|\{c\})$ and $\text{freq}(T|R^*)$ (where R^* is known only as a one element subset of R) will be identical (i.e., $\{0, 1\}$). We also have $\text{PROB}(\text{freq}(T|\{c\}) = 1) = \text{PROB}(\text{freq}(T|R^*) = 1) = r$,

¹⁶ Suppose, for example, that " R^* " is introduced as a name for the first element of some ordering, ρ , of the subsets of R that contain $|R'|$ elements, where our agent has no information regarding the principle according to which the elements of ρ were ordered.

¹⁷ Here as elsewhere, I will say that c is a random element of R , if c was selected from among the elements of R , by a process that was *equally likely* to yield each element of R .

since, where $R^* = \{c^*\}$, it is correct to infer $\text{PROB}(c^* \in T) = r$, by direct inference from the premises $c^* \in R$ and $\text{freq}(T|R) = r$.

The proposed condition also properly handles the ACME urn example. In the ACME 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 (after we inspect the elements of U_b) differs, and is more informative than, the lowest upper bound that may be inferred from the given information regarding the number of balls in U_b (the size of the relevant R'), the number of balls held in urns produced by the ACME Urn Company (the size of the relevant R), and the frequency of red balls among balls held in ACME urns. In this case, we are justified in holding that $\text{PROB}(\text{freq}(\text{red-balls}|\text{balls-in-}U_b) > 0.51) = 0$, but where R^* is known only as a one hundred member subset of the set of balls held in ACME urns, we are not justified in holding that $\text{PROB}(\text{freq}(\text{red-balls}|R^*) > 0.51) = 0$.

5. Highly Informative Complement Classes

In the ACME urn example, our frequency information for a relevant subset of a proposed reference class is too informative. The result is that direct inference based on the proposed reference class (the set of balls held in ACME urns) is defeated. A problem related to the one illustrated by the ACME urn example arises in cases where our frequency information for the *relative complement* of a subset of a proposed reference class is too informative. Once again, the problem can be illustrated by a simple example.¹⁸ Suppose that one is certain of the following propositions:

- [1] At least 90 percent of birds are capable of flight.
- [2] There are at least 10 times as many birds as sea tortoises.

In addition, suppose that one has absolutely no information about the frequency with which sea tortoises are able to fly, so that one is only justified in believing that the frequency of sea tortoises that are able to fly is in the interval $[0, 1]$. Despite the absence of information about the proportion of sea tortoises that are able to fly, one may deduce, from [1] and [2], that the frequency of creatures that are able to fly among the set of creatures that are *birds or sea*

¹⁸ The example is adapted from Pollock (1990, p. 84). A second example, from Pollock, that I will not discuss concerns the inference to the conclusion that a given 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 is used to address the example presented here applies equally to the case of the bird with the broken wing.

tortoises is in the interval $[9/11, 1]$. This is a problem for the theories of Kyburg, Pollock, and Bacchus, in the case where one wants to draw a conclusion about the probability that a particular sea tortoise, Herman, is able to fly. Indeed, in the absence of an additional constraint on direct inference, the theories of Kyburg, Pollock, and Bacchus allow one to draw the conclusion that the probability that Herman is able to fly is in the interval $[9/11, 1]$. Kyburg, Pollock, and Bacchus all acknowledge the difficulty that these sorts of case present for their theories. Bacchus postpones treatment of the problem, while Kyburg and Pollock address the problem by claiming that the set of creatures that are *birds or sea tortoises* does not correspond to a projectable predicate/property, so that the set of creatures that are *birds or sea tortoises* cannot be used as a reference class for a direct inference. While Kyburg and Pollock invoke the notion of projectability to deal with such examples, they do not provide criteria for determining when a predicate or property is projectable. Rather the notion of projectability is invoked to deal with counterexamples to their theories in an *ad hoc* manner.

As it turns out, the case of Herman the sea tortoise can be dealt with by a variant of the approach that was used in dealing with the ACME urn example. In the ACME urn example, we found that direct inference about an object, c , based on a reference class, R , may be defeated due to c 's membership in a narrower reference class, R' , in cases where our frequency information for R' is more informative than it should be. In light of the example of Herman the sea tortoise, I propose that a direct inference based on a reference class, R , may also be defeated due to c 's membership in a narrower reference class, R' , in cases where our frequency information for the relative complement of R' (i.e., $R-R'$) is more informative than it should be. More precisely: where A is an agent, T a given target class, R a proposed reference class, and R' is a subset of R (where R' contains c , the object of interest), A 's information regarding R' is *too informative* if it is not the case that, for all U and V , A is justified in accepting $\text{PROB}(\text{freq}(T|R-R') \in U) \in V$ if and only if A would be justified in accepting $\text{PROB}(\text{freq}(T|R-R^*) \in U) \in V$, in a situation identical to A 's actual situation save that the name " R^* " is introduced to A by a definite description that confers only the information that R^* is a subset of R , and that R^* is the same size as R' .

Like the condition used to remedy the ACME urn example, the present informativeness test is triggered in cases where an agent has information about R' that makes R' an exceptional subset of R vis-à-vis the incidence of elements of T (from the agent's perspective). And, once again, the present condition is defensible by appeal to the idea that justified instances of direct inference presuppose that a respective object of interest is *relevantly indiscernible* among a proposed reference class. In the case of Herman, 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 also

members of the set of sea tortoises (relative to the target class, *creatures that are able to fly*). Moreover, the proposed condition applies in the case of Herman, 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 only 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 sea tortoises, and our judgments regarding the frequency of creatures able to fly among the set of creatures that are birds or sea tortoises). Indeed, the narrowest interval in which we can locate the frequency of creatures able to fly among the set of birds ($R-R'$) is $[0.9, 1]$. But if we consider a set R^* , which is known only as a subset of set of creatures that are birds or tortoises, whose size is the same as the set of sea tortoises, then the narrowest interval in which we can locate the frequency of creatures able to fly among $R-R^*$ is $[8/11, 1]$ (so that we are justified in accepting $\text{PROB}(\text{freq}(\text{creatures-able-to-fly}|R-R') \in [0.9, 1]) = 1$, but we are unjustified in accepting that $\text{PROB}(\text{freq}(\text{creatures-able-to-fly}|R-R^*) \in [0.9, 1]) = 1$).

I have proposed two similar conditions in order to deal with the ACME urn case and the case of Herman the sea tortoise. When either of these conditions hold for a reference class, R , and one of its subsets, R' , I will say that R and R' are *informativeness incompatible* for the agent, A , and the target class, T . The following definition collects the two conditions (which I now state in the negative).

Definition: R and R' are *not* informativeness incompatible for the agent, A , and the target class, T , if and only if for all U and V : (1) A is justified in accepting $\text{PROB}(\text{freq}(T|R') \in U) \in V$ if and only if A would be justified in accepting $\text{PROB}(\text{freq}(T|R^*) \in U) \in V$, and (2) A is justified in accepting $\text{PROB}(\text{freq}(T|R-R') \in U) \in V$ if and only if A would be justified in accepting $\text{PROB}(\text{freq}(T|R-R^*) \in U) \in V$, in a situation identical to A 's present situation save that the name " R^* " is introduced by a definite description that confers only the information that R^* is a subset of R , and that R^* is the same size as R' .

Described as a modification of Reichenbach's theory (as expressed by [RDI] and [RSD]), I propose to amend (iii) of [RSD], so that an instance of direct inference is defeated only if R and R' are *informativeness incompatible*, for the given agent and target class. Modulo considerations that forbid gerrymandered target and reference classes (to be discussed in section 7), the definition of informativeness incompatibility is meant to capture the precise conditions under which our information about a *subset* of a proposed reference class is too informative, where the satisfaction of this condition entails that direct inference based on the

proposed reference class is subject to subset defeat. There is good reason to think that the definition is correct for that purpose. First, the two components of the proposed test (regarding R' and $R-R'$) proceed from the intuitive justificatory foundation of direct inference. This gives us reason to think that the applicability of the definition is relatively general. Second, it is difficult to imagine a more stringent test for whether an agent has any special information regarding a subset of a proposed reference class (relative to the incidence of some target property).

6. Expected Frequencies

The informativeness criterion introduced in the preceding section specifies the conditions under which information about a narrower reference class will defeat an instance of direct inference based on statistics for a broader class. This informativeness criterion provides a partial solution to the Problem of Uninformative Statistics, by insuring that uninformative frequency information does not *in general* result in the (subset) defeat of direct inferences based on informative frequency information for broader reference classes. In order to fully address the Problem of Uninformative Statistics, one must also explain why direct inferences based on uninformative frequency information generally yield conclusions that are consistent with the intuitively correct conclusions based on informative frequency statements for broader reference classes. Within Kyburg's theory, the latter is accomplished by the requirement that the major premises for direct inference be interval-valued frequency statements. While Kyburg's approach achieves the desired effect, the restriction to interval-valued frequency statements is *ad hoc*, inasmuch as the restriction is not adequately motivated. As an alternative to Kyburg's approach, I adopt Bacchus's proposal that it is statements of expected frequency that serve as the proper statistical premises for direct inference.

In probability theory, a random variable is identified with the range of 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 average of the possible values of the random variable weighted 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 the average of the set of possible values of the relative frequency weighted by the probabilities of the respective

values. As a special case, probability statements regarding singular propositions are equivalent to statements of expected relative frequency regarding unit set reference classes.

Before describing the main reason for regarding statements of expected frequency as the proper premises for direct inference, I will explain why using frequency statements as premises for direct inference is a special case of using statements of expected frequency. By demonstrating this connection, I will discharge the demand to explain the manner in which frequency statements are relevant to direct inference. As a corollary, we will see why the use of expected frequency statements as the major premises of direct inference serves as a partial solution to the Problem of Uninformative Statistics.

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. 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 can easily be calculated, by appeal to the following theorem. (Here I use the notation $E[\text{freq}(T|R)]$ to denote the expectation of the relative frequency of T among R.)

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

The preceding theorem illustrates the relevance of frequency information to direct inference (assuming that expected frequency statements are the proper statistical premises for direct inference), since 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. Note, for example, the implication between $\text{PROB}(\text{freq}(T|R) = r) = 1$ and $E[\text{freq}(T|R)] = r$. The theorem also illustrates why restricting the major premises of direct inference to statements of expected frequency ensures that direct inference based on uninformative frequency information does not yield conclusions that will contradict the conclusions of direct inference based on informative frequency information. The point becomes clear when one considers the case of unit set reference classes. Consider any case where it is correct to infer $\text{PROB}(c \in T) = r$, by direct inference from the premises $c \in R$ and $\text{freq}(T|R) = r$. In such cases, we invariably know that $\text{PROB}(\text{freq}(T|\{c\}) \in \{0, 1\}) = 1$. Based on this frequency information, we may conclude that $E[\text{freq}(T|\{c\})] \in [0, 1]$ and that $\text{PROB}(c \in T) \in [0, 1]$, which is consistent with the conclusion that $\text{PROB}(c \in T) = r$.

The doctrine that it is statements of expected frequency that are the proper major premises for direct inference serves as a partial solution to the Problem of Uninformative Statistics. Another reason for taking statements of expected frequency as the proper major premises for direct inference is connected to the intuitive justificatory basis of direct inference. When making a direct inference, one assumes that the object about which one is reasoning, c , is as likely to be a member of a 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. Similarly, in cases where one is aware of the correct assignment of probabilities to the values of a given relative frequency, one may calculate the probability of a random element of R being in T by considering the likelihood that $\text{freq}(T|R)$ takes on respective values.¹⁹ For parallel reasons, one is obliged to conclude that the probability that c is in T is equal to the expected frequency of elements of T among R , since the expected value of $\text{freq}(T|R)$ simply encodes a weighting of the possible values of $\text{freq}(T|R)$ according to probability, and since the likelihood that a random 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 assume that proper instances of direct inference proceed from premises of the form $\lceil E[\text{freq}(T|R)] \in U \rceil$ and $\lceil c \in R \rceil$ to conclusions of the form $\lceil \text{PROB}(c \in T) \in U \rceil$. 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 statements, and I will do so for the sake of convenience.

7. The Problem of Relevant Statistics

Consider a case where one is justified in accepting the following propositions:

- [1] 40% of dogs live at least twelve years.
- [2] 70% of dachshunds live at least twelve years.
- [3] Flint is a dachshund.

¹⁹ 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.

On the supposition that the three preceding propositions encapsulate one's knowledge of the factors that are relevant to judging the probability that Flint will live at least twelve years, it seems that one should conclude 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 one to prefer frequency information for narrower reference classes, in cases where one has relevant frequency information for two or more sets. But 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 dachshunds who will not live twelve years. In that case, R_G is narrower than the set of dachshunds, and Flint is an element of R_G . The problem, then, is that the frequency of elements of R_G that will live twelve years is guaranteed to be very near to zero. Indeed, if we suppose that there are only one hundred dachshunds, then $\text{freq}(\text{creatures-that-will-live-twelve-years} \mid R_G) \in \{0, 1/31\}$.

The problem with which we are faced is that of explaining why one is permitted to conclude that the probability that Flint will live twelve years is 0.7, and one is not permitted to conclude that the probability that Flint will live twelve years is in the interval $[0, 1/31]$. Following a suggestion of Bacchus (1990), I call the present problem "the Problem of Relevant Statistics" with the idea being that certain statistical statements (such as the one involving the reference class R_G) are not relevant to 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 one is reasoning, c , is as likely to be a member of a 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* among the other elements of R .

In the case of Flint, the conclusion that Flint is as likely to live twelve years as a randomly selected dachshund is not defeated by the statistical fact that a very high proportion of the 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 superseded, 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 the set of possible values of $\text{freq}(\text{creatures-that-will-live-twelve-years} \mid \{\text{Flint}\})$ is $\{0, 1\}$, while our narrowest estimate of the set of possible values

of $\text{freq}(\text{creatures-that-will-live-twelve-years} \mid R_G - \{\text{Flint}\})$ 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 with 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, γ , of the proposed reference class, where γ is known to contain the object about which we wish to make a direct inference.²⁰ Through the use of such a description, the value of the key statistic is computed via a reference class that is literally *gerrymandered* relative to the given target class. In particular, the value of the statistical statement is computed by appeal to the possible sizes and statistical values for its subsets, where one of the subsets is known to contain the object about which one wishes to make a direct inference. 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 appeal to the possible values for the frequency of creatures that will live twelve years among $R_G - \{\text{Flint}\}$, and by appeal to the possible values for the frequency of creatures that will live twelve years among $\{\text{Flint}\}$. Where “L12” stands for the set of creatures that will live twelve years, and “DH” stands for the set of dachshunds (so that $R_G = (\text{DH} \cap \sim \text{L12}) \cup \{\text{Flint}\}$), the computation proceeds by cases:

Case 1:

If $\text{Flint} \notin \text{L12}$, then

$$\text{freq}(\text{L12} \mid (\text{DH} \cap \sim \text{L12}) - \{\text{Flint}\}) = 0/29,$$

$$|(\text{DH} \cap \sim \text{L12}) - \{\text{Flint}\}| = 29,$$

$$\text{freq}(\text{L12} \mid \{\text{Flint}\}) = 0/1, \text{ and}$$

$$\text{freq}(\text{L12} \mid (\text{DH} \cap \sim \text{L12}) \cup \{\text{Flint}\}) = 0/30.$$

Case 2:

If $\text{Flint} \in \text{L12}$, then

$$\text{freq}(\text{L12} \mid (\text{DH} \cap \sim \text{L12}) - \{\text{Flint}\}) = 0/30,$$

$$|(\text{DH} \cap \sim \text{L12}) - \{\text{Flint}\}| = 30,$$

$$\text{freq}(\text{L12} \mid \{\text{Flint}\}) = 1/1, \text{ and}$$

$$\text{freq}(\text{L12} \mid (\text{DH} \cap \sim \text{L12}) \cup \{\text{Flint}\}) = 1/31.$$

Therefore, $\text{freq}(\text{L12} \mid (\text{DH} \cap \sim \text{L12}) \cup \{\text{Flint}\}) \in \{0/30, 1/31\}$.

²⁰ This feature is also common to all of the examples that I have been able to concoct. Many of these examples are more insidious than the examples found in the literature.

It is easy to see why the preceding sort of gerrymandering violates the indiscernibility condition that is tacitly assumed when one makes a direct inference. In the computation just described, Flint is treated separately from the other elements of the proposed reference class, so that Flint is literally *discerned* from the 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 dependence on certain types of descriptions in the computation of the value of the statistical statements that will be used in direct inference. 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. Rather than concern ourselves with the actual descriptions that an agent employs in the computation of a given statistic, we need only require that the agent could have justified her conclusion through a chain of reasoning that does not rely on a problematic description.

In line with my characterization of the problem in view as the Problem of Relevant Statistics, I will describe cases where a statistic is gerrymandered (in a problematic way) as cases where the statistic is not *relevant* to the probability that a given object is a member of a respective target class.

Definition: $E[\text{freq}(T|R)] \in V$ is (*potentially*) *relevant* to the value of $\text{PROB}(c \in T)$ for an agent, A , if and only if there exists a chain of inference²¹, C , sufficient for justifying A ’s belief that $E[\text{freq}(T|R)] \in V$, where, for all R' : if R' is describable using only vocabulary occurring in the course of C , then

- (i) A is not justified in believing that $R' \subseteq R$,
- (ii) A is not justified in believing that $c \in R'$, or
- (iii) R and R' are not *informativeness incompatible*, relative to A and T .²²

²¹ Assuming that chains of inference resemble proofs in a formal language, it is intended that membership in a set, R' , be characterized by the satisfaction of a first order formula, $\varphi(x)$, with a single free variable x . We may then regard R' as the set of objects that satisfy $\varphi(x)$.

²² In determining the applicability of the present definition, it is assumed that beliefs about the value of an expected frequency cannot be justified on the basis of testimony (or similar means) in cases where one is aware of the basis upon which the testifier formed her judgment. In other words, the more fundamental evidence must be given priority when that evidence is available. (The present proviso is made for the purposes of the definition of *relevance*, and is not proposed as essential to the concept of justification.)

In cases where the present definition fails to apply to a given expected frequency statement (relative to a corresponding single-case probability), the statement is deemed *irrelevant* (to that single-case probability). Note that whether a given statistic is *irrelevant* (to a given single-case probability) may be practically inaccessible in some cases, since a statistic is irrelevant just in case for *every* chain of inference capable of justifying the agent's belief in the statistic, vocabulary is employed that is sufficient to describe some 'suspect' set, R' . But the problem here is not grave, since the conditions under which a given expected frequency statement is *relevant* are (relatively) accessible, and the account of direct inference that I will finally propose requires only that any statistic that serves as a premise for a direct inference, or as a subset defeater for a direct inference, be *relevant*. So an agent may apply the proposed account of direct inference by simply certifying that any statistic she uses (in making or defeating a direct inference) is relevant. In any case, it is often possible to see that a given statistic is irrelevant by observing that the calculation of the statistic (by a given agent) can only proceed by a chain of inference which does involve vocabulary sufficient to describing some suspect set. One such example is the case of Flint and the gerrymandered reference class R_G .

Applied to the case of Flint, we see that the expected frequency statement $E[\text{freq}(L12|(DH \cap \sim L12) \cup \{\text{Flint}\})] \in [0/30, 1/31]$ is *irrelevant* to the probability that Flint will live twelve years, since there is no way to compute this statistic (based on the described assumptions) that does not appeal to some variant of the predicate $x = \text{Flint}$, where the set $\{\text{Flint}\}$ satisfies the following three conditions, for any agent, A , whose evidence is as described in the example:

- (i) A is justified in believing that $\{\text{Flint}\} \subseteq (DH \cap \sim L12) \cup \{\text{Flint}\}$,
- (ii) A is justified in believing that $\text{Flint} \in \{\text{Flint}\}$, and
- (iii) $(DH \cap \sim L12) \cup \{\text{Flint}\}$ and $\{\text{Flint}\}$ are *informativeness incompatible*, relative to A and $L12$.

That condition (iii) holds can be seen inasmuch as $\text{freq}(L12|((DH \cap \sim L12) \cup \{\text{Flint}\}) - \{\text{Flint}\}) = 0$, while the smallest set in which we may locate the value of $\text{freq}(L12|((DH \cap \sim L12) \cup \{\text{Flint}\}) - R^*)$ is $\{0/30, 1/30\}$ (where R^* is known only as a subset of $(DH \cap \sim L12) \cup \{\text{Flint}\}$ whose cardinality is one).²³

²³ Because $(DH \cap \sim L12) \cup \{\text{Flint}\}$ and $\{\text{Flint}\}$ are *informativeness incompatible*, it may appear that the application of the definition of informativeness incompatibility (as a criterion for subset defeat) is sufficient to address the Problem of Relevant Statistics. In fact, applications of that definition are not sufficient to address the problem. For one, we must restrict the application of informativeness incompatibility, as a

The prescription that irrelevant statistics not be used as premises for direct inference, or to defeat instances of direct inference via subset defeat, makes sense in light of the precept that correct direct inferences presuppose the relevant indiscernibility of an object of interest from the other elements of a proposed reference class. Indeed, in cases where a statistic is deemed *irrelevant*, we know that the object of interest is discernable from other elements of the proposed reference class, R , in the course of a respective agent's reasoning about the value of statistics for R .

8. Conclusion

With the notions of *relevance* and *informativeness incompatibility* in place, I am in a position to propose some fairly traditional looking principles of direct inference. The principles incorporate three amendments to Reichenbach's theory of direct inference (as expressed by [RDI] and [RSD]). As a remedy to the problem associated with gerrymandered reference classes, I require that only *relevant* statistics play a role in direct inference. As a remedy to the Problem of Uninformative Statistics, I require: (1) that the proper statistical premises for direct inference are statements of expected frequency, and (2) that a direct inference is subject to subset defeat by statistics for a narrower reference class only if the two reference classes are *informativeness incompatible*. These amendments yield the following:

[DI] Direct Inference

If A is justified in believing that $E[\text{freq}(T|R)] \in V$ and $c \in R$, then A has a defeasible reason to believe that $\text{PROB}(c \in T) \in V$, so long as $E[\text{freq}(T|R)] \in V$ is *relevant* to the value of $\text{PROB}(c \in T)$ for A .

[SD] Subset Defeat

A respective instance of [DI] is defeated for an agent, A , if there exists an R' such that:

- (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 *informativeness incompatible*, relative to A and T , and
- (iv) $E[\text{freq}(T|R')] \in U$ is *relevant* to the value of $\text{PROB}(c \in T)$ for A , where U is the narrowest set of values that A is justified in accepting for $E[\text{freq}(T|R')]$.²⁴

criterion for subset defeat, to cases where an agent's statistics for a subset of a proposed reference class are *relevant*.

²⁴ I leave the treatment of reference classes corresponding to reference properties of differing arity for another day (cf. Pust 2011; Thorn *forthcoming*).

Taken together, [DI] and [SD] still allow for the possibility of cases where two instances of [DI] yield reasons for assigning conflicting probabilities to a proposition, where both instances of [DI] are based on relevant statistics and neither of the two inferences is defeated via [SD]. The paradigm example of such cases occurs when an agent has relevant informative statistics, regarding a given target class T, for two overlapping reference classes, but lacks informative statistics, regarding the incidence of T, among the intersection of the two reference classes. I believe that it is sometimes possible to make a reasonable direct inference in such cases, but I will not defend that claim here. In any case, the problem is far from being grave, since we may follow Reichenbach's recommendation and simply suspend judgment in the face of such conflicting reasons. Adherence to Reichenbach's proposal will keep us from forming unjustified beliefs in the face of conflicting reasons for belief, even if there are cases where Reichenbach's proposal is too restrictive.

Acknowledgements This work was supported by the LogiCCC EUROCORES program of the ESF and DFG. For helpful comments on earlier presentations of this paper, I thank audiences at the University of Arizona, the University of Düsseldorf, and the Third Formal Epistemology Festival at the University of Toronto. I also thank Terry Horgan, Shaughan Lavine, Gerhard Schurz, and especially John Pollock and two anonymous referees for *Erkenntnis* for helpful comments on earlier drafts of the paper.

References

- Bacchus, F. (1990). *Representing and Reasoning with Probabilistic Knowledge*. (MIT Press, Cambridge, Massachusetts.)
- Bacchus, F., Grove, A., Halpern, J., and Koller, D. (1996). From statistical knowledge bases to degrees of belief. *Artificial Intelligence*, 87(1-2), 75-143.
- Colyvan, M., and Regan, H. (2001). Is it a Crime to Belong to a Reference Class? *Journal of Political Philosophy*, 9(2), 168-181.
- Colyvan, M., Regan, H., and Ferson, S. (2007). Legal decisions and the reference class problem. *International Journal of Evidence and Proof*, 11(4), 274-286.

Two Problems of Direct Inference

- Fetzer, J. (1977). Reichenbach, Reference Classes, and Single Case 'Probabilities'. *Synthese*, 34, 185-217.
- Fitelson, B., Hájek, A., and Hall, N. (2005). Probability. (In , S. Sarkar and J. Pfeifer (Eds.) *Philosophy of Science: An Encyclopedia*. Oxford: Routledge.)
- Geffner, H., and Pearl, J. (1992). Conditional entailment: bridging two approaches to default reasoning. *Artificial Intelligence*, 53(2-3), 209-244.
- Goodman, N. (1955). *Fact, Fiction, and Forecast*. (Cambridge: Harvard University Press)
- Hájek, A. (2007). The Reference Class Problem is Your Problem Too. *Synthese*, 156(3), 563-585.
- Halpern, J. (2003). *Reasoning About Uncertainty*. (Cambridge: Harvard University Press)
- Hempel, C. (1968). Lawlikeness and Maximal Specificity in Probabilistic Explanation. *Philosophy of Science*, 35(2), 116-133.
- Howson, C., and Urbach, P. (1989). *Scientific reasoning: The Bayesian Approach*. (Chicago: La Salle and Open Court)
- Horty, J., Thomason, R., and Touretzky, D. (1990). A sceptical theory of inheritance in nonmonotonic semantic networks. *Artificial Intelligence*, 42(2-3), 311-348.
- Kraus, S., Lehmann, D., and Magidor, M. (1990). Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44(1-2), 167-207.
- Kyburg, H. (1956). The Justification of Induction. *Journal of Philosophy*, Vol. 53, 394-400.
- Kyburg, H. (1961). *Probability and the Logic of Rational Belief*. (Middleton , Connecticut: Wesleyan University Press)
- Kyburg, H. (1974). *The Logical Foundations of Statistical Inference*. (Dordrecht-Holland: Reidel Publishing Company)
- Kyburg, H., and Teng, C. (2001). *Uncertain Inference*. (Cambridge: Cambridge University Press)

Two Problems of Direct Inference

- Levi, I. (1982). Direct Inference. *Journal of Philosophy*, 74, 5-29.
- Pollock, J. (1990). *Nomic Probability and the Foundations of Induction*. (Oxford: Oxford University Press)
- Pollock, J. (2007). The Y-function. (In Wheeler, G. and Harper, B. (Eds.), *Probability and Evidence: Essays in Honour of Henry E. Kyburg Jr.* London: College Publications.)
- Pollock, J. (*unpublished*). Probable Probabilities.
- Pust, J. (2011). Sleeping Beauty and Direct Inference. *Analysis*.
- Reichenbach, H. (1949). *A Theory of Probability*. (Berkeley: Berkeley University Press)
- Rhee, R. (2007). Probability, Policy and the Problem of the Reference Class. *International Journal of Evidence and Proof*, 11(4), 286-292.
- Salmon, W. (1971). Statistical Explanation. (In Salmon, W. (Ed.), *Statistical Explanation and Statistical Relevance*. Pittsburgh: University of Pittsburgh Press.)
- Salmon, W. (1977). Objectively Homogeneous Reference Classes. *Synthese*, 36, 399-414.
- Salmon, W. (1984). *Scientific Explanation and the Causal Structure of the World*. (Princeton: Princeton University Press)
- Stone, M. (1987). Kyburg, Levi, and Petersen. *Philosophy of Science*, 54(2), 244-255.
- Thorn, P. (2007). *Three Problems of Direct Inference*. Dissertation, University of Arizona.
- Thorn, P. (forthcoming). Undercutting defeat via reference properties of differing arity: a reply to Pust. *Analysis*.
- Venn, J. (1866). *The Logic of Chance*. (New York: Chelsea Publishing Company)