Giving up Certainties[1]

Henry E. Kyburg, Jr.

University of Rochester

1. Choosing Among Theories

People have worried for many years — centuries — about how you perform large changes in your body of beliefs. How does the new evidence lead you to replace a geocentric system of planetary motion by a heliocentric system? How do we decide to abandon the principle of the conservation of mass?

The general approach that we will try to defend here is that an assumption, presupposition, framework principle, will be rejected or altered when a large enough number of improbabilities must be accepted on be basis of our experience. If I think that all swans are white, and a student claims to have a counterexample, I will assume that he has made some observational error. I will reject his result, and continue to accept the generalization. When a lot of people claim to have seen counterexamples, I will come around: to continue to accept the generalization would require me to accept too many improbabilities. This is a discontinuous process as we will construe it: it is not a matter of a general statement becoming less probable, while certain reports become more probable. We cannot accept the generalization and even one of the observation reports: that would be a simple inconsistency.

One suggestion, due to Karl Popper, is that we invent Bold Conjectures, and Put Them to the Test. (Popper, 1959) Bold conjecture: the Earth is the Center of the Solar System. Test... what? Bold conjecture: Mass is conserved. Test: weigh a mass of plutonium and its by products before and after. Obviously things are a little more complicated than the slogans suggest.

Alternatively, gather evidence, and accept the hypothesis that is most probable, relative to that evidence. So far, so good (maybe). But then what? How do you change from that hypothesis to one inconsistent with it when the evidence so indicates? For as soon as a hypothesis is accepted, it has probability 1; and as soon as a hypothesis has probability 1, its contraries have probability 0; and as soon as a contrary hypothesis has probability zero, its probability can never leave zero — at least not by Bayes' theorem.

PSA 1990, Volume 2, pp. 333-347

https://doi.org/10.1086/psaprocbienmeetp.1990.2.193078 Published online by Cambridge University Press

A natural response to this observation is to say, as Carnap (1950) did, that "acceptance" is just an approximation to the real truth, and that no hypothesis ever achieves literal acceptance, which would entail its having a probability of 1. What we really have (as opposed to the approximate way we talk) is a probability blanket over a field of empirical propositions, none of which is ever assigned a probability of 0 or 1 unless it is a mathematical or logical truth, or the denial of one.

This latter approach presents us with serious problems. We will consider the problem of assigning prior probabilities to the sentences of a reasonably rich language later, but already we are faced with a difficult computational problem. Gil Harman (1989) has pointed out that in a language win n basic sentences there are 2^n assignments to make. But of course we can get by with wholesale assignments; if we decide that each conjunction of basic sentences or their negations is to have the same measure assigned to it, there is in fact only one assignment to make: one simple algorithm that provides the measure for any sentence.

In general, however, a useful and realistic language will have an infinite number of sentences, and this procedure breaks down. It is still possible to assign measures systematically, without assigning zero to any sentence representing a possibility. The number of sentences in any ordinary language is denumerable, and we can find a denumerable number of finite numbers that add up to 1. But the rationale of the system is hard to find.

It is, at any rate, worth exploring alternatives to either of these approaches to to rational acceptance. One of the first to offer a systematic procedure for this was Isaac Levi. In *Gambling with Truth* and *The Enterprise of Knowledge,* Levi proposes a rule for *adding* to your body of knowledge. Given such a rule, one can obtain a rule for replacing one conjecture, law, theory, hypothesis by another by proposing that when faced with a choice, one simply deletes both candidates from one's body of knowledge, and then adds the one indicated by the application of the rule for addition.

The rule is just this: (1967, p.86) Let U be a set of most specific possible hypotheses — i.e., a set of which exactly one member is true. Let M be an "information determining probability" (1980, p. 48): M(g) represents the informational value of rejecting g, and let p be an expectation forming probability (a degree of belief, a credibility). Let q in [0,1] be an index of caution. The rule (Rule A, of 1967) is to reject all and only those elements g of U such that $p(g) < qM(g)$, and to accept, with deductive closure, the disjunction of the remainder.

Given a rule for acceptance, we can construe contraction as suspending belief in a proposition and then failing to add it back under subsequent expansion; and we can construe replacement as suspending belief in one proposition, and arriving at another on subsequent expansion.

We can accomplish a change of framework of "accepted facts" this way, and we can be sure of maintaining consistency in the process. There are some knotty problems, however. When and how do we decide to suspend belief in a framework proposition? There are clear cases: when observations render our corpus inconsistent, for example. "For the sake of argument," in a friendly social context. In the context of a debate. Levi can afford to be casual about this, since a proposition erroneously deleted will be easily recaptured, and he is not concerned with real time changes. But there are other questions. How should q, the index of caution, be chosen? Where does the information measure M come from? How do we arrive at the credal probability p? More fundamental: How is the "abductive" step — the step in which the ultimate partition U is formed — to be controlled and rationalized?

One can always raise questions, of course. But these questions are disturbing because the rule presupposes a framework (a language, an information measure, a credibility measure, a set of most specific answers), and thus to be not even potentially capable of providing guidance in the choice of a framework. But let us look further.

An approach similar to Levi's has been developed in various ways by Makinson, Alchurron, and Gardenfors (1982, 1985), Gardenfors and Makinson (1988) and Gardenfors (1986, 1988). While Levi approaches the question from a constructive, analytic angle, and seeks to provide formal analysis of what goes on in changes in a corpus of knowledge, Gardenfors and the others approach the question from a logical point of view: they seek to explore axioms that may be taken to characterize the change of a body of knowledge, construed as a set of propositions. Thus, for example, it is clear that if we add the proposition A to our body of knowledge K, then A should belong to that expanded body of knowledge. As is the case with Levi, it is assumed by these writers that a body of knowledge K should be construed as a deductively closed set of propositions.

An excellent examination of these logics of theory change is provided by Gardenfors' book (1988). It is from that source that I take the following axioms. A *belief set* here is construed as a deductively closed set of propositions.

If we denote by K^+_A the expansion of a body of knowledge K by the addition of the consistent proposition A , then we may express the the properties of the expansion of a belief set by the following relatively uncontroversial axioms.

(K^+ 1) K^+_A is a belief set.
(K^+ 2) $K^+_A \supseteq A$
(K^+ 3) If $\sim A \notin K^+_A$, then $K^+_A \supseteq K$
(K^+ 4) If $A \in K$, then $K^+_A = K$
(K^+ 5) If $H \supseteq K$, $H^+_A \supseteq K^+_A$
(K^+ 6) For all belief sets K and all sentences A, K^+_A is the smallest belief set that satisfies (K^+ 1) – (K^+ 5).

What is not so uncontroversial is the question of the principles according to which a body of knowledge should be contracted. This is not a terribly serious question for Levi: any proposition in our body of knowledge can be doubted with relative impunity. It can be doubted with relative impunity, since, if it belongs in our corpus of knowledge, it will be reinstated on reflection. One can thus suspend belief in a proposition A on quite casual grounds.

A serious reason to suspend belief in something arises from the circumstance that our corpus of knowledge is inconsistent. For example, if there are observational routines that warrant our acceptance of the statement that a is a crow and a is not black, then when we practise those routines, we should accept the corresponding statement. (Or proposition.) But if we already accept the generalization that all crows are black, this renders our corpus inconsistent.

With an inconsistent corpus, we are clearly *obligated* to suspend belief in something. Levi says that we should shrink our corpus of knowledge in such a way as to retain the most "information." But it is clear that no simple-minded construal of "information" will lead to the right results. In some sense it is clear that the information

https://doi.org/10.1086/psaprocbienmeetp.1990.2.193078 Published online by Cambridge University Press

content of "all crows are black" is greater than that of "a is a crow and a is not black," but of course on any standard construal of hypothesis testing it is the former that will be suspended and the latter that will be retained.

While Levi offers us no logic of contraction, that is the main concern of Gardenfors et al. Gardenfors offers a number of axioms characterizing the contraction operation, denoted by K_A^-. Most of these axioms are relatively uncontroversial, as in the case of expansion. We have:

(K^- 1) For any sentence A and any belief set K, K_A^- is a belief set.
(K^- 2) $K \supseteq K_A^-$
(K^- 3) If $A \notin K$, then $K_A^- = K$.
(K^- 4) If not $\vdash A$, then $A \notin K_A^-$.
(K^- 5) If $A \in K$ then $K_A^- \supseteq K$.
(K^- 6) If $\vdash A \leftrightarrow B$, then $K_A^- = K_B^-$
(K^- 7) $K_{A\&B}^- \supseteq K_A^- \cap K_B^-$.
(K^- 8) If $A \notin K_{A\&B}^-$, then $K_A^- \supseteq K_{A\&B}^-$

These axioms may be more controversial than those for the expansion of a body of knowledge, but there is still nothing obviously wrong with them. It is possible to provide intuitively plausible axioms for theory replacement, and to show that in general replacement can be construed as a contraction followed by an expansion.

What becomes controversial is the procedure for conducting contractions. The contraction $\supseteq$ is not uniquely determined by these axioms, in contrast to K_A^+ (under the assumption of deductive closure). We must thus consider *how* to perform the contraction. One possibility is the following. Consider a subset of K that is deductively closed, that does not contain A, and that is such that if any other sentence of K is added to it, A will be a consequence of it. The set of all such sets of sentences is denoted by $K \perp A$. Clearly the result of contraction should be a member of this set (if it isn't empty; if A is a theorem, then we can take the contraction of K by A to be K itself. All we need to do is to devise a "selection function" S that will pick one set out of $K \perp A$. But, as Gardenfors shows, this yields contractions that are "too big." If $A \in K$ then this procedure will yield a K_A^- that for any proposition B contains either $A \vee B$ or $A \vee \sim B$.

The next idea one might have is to consider the intersection of all the sets of sentences in $K \perp A$. (This is called the "full meet contraction.") This is too small: K_A^- will consist only of the logical consequences of $\sim A$.

Finally, we may consider a selection function S that picks *some* of the members of $K \perp A$, intuitively the most *epistemically entrenched* members, and take K_A^- to be the intersection of these.

But what does epistemic entrenchment come to? That seems to be where the real controversy lies. Levi seeks to preserve information; he can be thought of as construing epistemic entrenchment in terms of information. But the epistemic entrenchment ranking of sets of propositions can plausibly be taken to reflect a system of beliefs, and thus to be sensitive to scientific or conceptual revolutions, whether these be understood in the dramatic Kuhnian sense or not. Gardinfors says that "...the fundamental criterion for determining the epistemic entrenchment of a sentence is how useful it is in inquiry and deliberation." (p.87) (Note that the selection function S is originally defined over sets of sentences, rather than sentences. This reflects a difference that could be exploited.)

One idea for representing such factors is provided by Wolfgang Spohn (1987). Spohn defines an "ordinal conditional function" that maps possible worlds into ordinals. The value of the function represents a degree of implausibility, or a degree of unwillingness to accept, or a degree of potential surprise (Levi, Shackle).

This function is can be extended to propositions in general by taking the value of the function for a proposition, to be the minimum value of the function over the set of worlds in which that proposition is true. Thus, since it is assumed that there is some world with value 0, either $k(A) = 0$ or $k(\sim A) = 0$, and $k(A \cup B) = \min\{k(A), k(B)\}$, where k is Spohn's ordinal conditional function.

Spohn's approach is more general than Gardenfors' since it takes as epistemic input a pair (A,a) consisting of a proposition A and an ordinal a. This yields a new ordinal function on possible worlds, and thus a new ordinal function. In the extreme cases, however, the treatment yields results parallel to those of Gardenfors (p. 73).

2. The Probabilistic Alternative

To be contrasted with this approach in terms of deductively closed sets of propositions, we may consider a purely probabilistic construal of knowledge: We take a statement as acceptable in our knowledge base when it becomes overwhelmingly probable. This is in accord with the nearly universal agreement that when it comes to empirical matters of fact, there is nothing (or almost nothing) that is certain. Almost any of the things we take for granted "could" turn out to be wrong. Nothing is incorrigible. Not even "observation" statements: without knowing how to handle errors of observation, modern science could hardly get off the ground. Of course, very crude observation statements, e.g., "the sun is shining now," are very unlikely to require correction. (They could be wrong: my "observation" may result from post–hypnotic suggestion, rather than the state of the weather.)

One way of dealing with an approach to knowledge that takes nothing empirical to be incorrigible is to become a thoroughgoing Bayesian: Represent knowledge as a probability function defined over the whole algebra of propositions in the language we are using for knowledge representation. Of course, as Carnap observed (1950), we must suppose that all refinements have been made in the language: we cannot introduce new terms without risking having to change our probability function. Then when experience causes us to shift the probability of some proposition, that change in probability propagates through the algebra in accord with some rule of propagation. (One possibility is "Jeffrey conditionalization," Jeffrey, 1965)

This approach to corrigibility has a number of drawbacks. The main one is computational. In language capable of representing some piece of common sense knowledge, or of reasoning about even quite a limited domain, the computational resources needed mount dramatically. The number of possible worlds, describable in even a constrained language, is LARGE. There is also the problem of the source of the original probability measure. Experts? There is the problem of soliciting consistent opinions. Generalize to sets of probability measures? This might be some help, but perhaps not much. There is the problem of updating: No set of probability assessments is likely to be consistent; adjustments will have to be made to achieve conformity with the probability calculus; and one of the items most natural to adjust is the ratio of probabilities $P(A \& B)/P(B)$; but this is just the important probability of A given B. And supposing a collection of agents with a common goal, sharing knowledge: how are disagreements concerning probabilities among these agents to be resolved?

https://doi.org/10.1086/psaprocbienmeetp.1990.2.193078 Published online by Cambridge University Press

These are difficult questions, and while one cannot be certain that plausible answers can't be found, it seems at least worth while to explore an alternative strategy. An alternative that has been explored for some years is that of adopting a purely probabilistic rule of acceptance: In general, "Accept *P* when its probability is high enough." (Kyburg, 1961)

One question rises immediately: how probable is "high enough?" A tentative answer to this question ("It depends on how much is at stake in using the corpus of knowledge in question") has been outlined in (Kyburg, 1988).

A less immediate question arises when we reflect that probability itself — especially evidential probability — depends on evidence. What is probable depends on what we know; and we are proposing that what we know depends on what is probable. Can we have it both ways? In particular, can evidential probability be serve both functions?

We answer yes. It has been proposed (Kyburg, 1984, Kyburg, 1988, Kyburg, 1974) that having fixed on practical certainty, we can introduce evidential certainty as the square root of practical certainty. (This stems from the fact that, using a probabilistic rule of acceptance, the conjunction of a pair of statements that do not appear conjoined in a higher level corpus will appear in a lower level corpus.)

A purely probabilistic rule of acceptance does not yield what Gardenfors has called "belief sets." The set of accepted statements is not closed under deduction, nor — what comes to the same thing in a logic with compactness — is it closed under conjunction. In general, it is not the case that if *A* and *B* are in our corpus of knowledge, their conjunction will also be in it. Of course it does not follow that the conjunction of a pair of statements in our corpus of knowledge will *not* be in it! There may be large conjunctions of statements whose probability is high enough to qualify for acceptance, and every conjunct of such a set of statements will also be in the corpus. In fact, every logical consequence of *each* statement in our body of knowledge will also be in it.

An immediate consequence is that there is an axiomatic representation of our body of knowledge. That is, there is a (presumably finite) set of statements from which the entire contents of our body of practical knowledge follows. This fact has useful consequences when it comes to talking about revisions of our body of knowledge.

The failure to embody deductive closure is not entirely unintuitive. Our confidence in the conclusion of an argument that involves many premises tends to decrease, even though we cannot put our finger on a specific doubtful premise, as the number of premises decreases. There are good intuitive grounds, even, for thinking that the set of statements that I am well justified in accepting is inconsistent; if it is inconsistent, to apply deductive closure to it would be a disaster. One particularly natural example concerns measurement. Suppose the method *M* yields errors that are distributed approximately normally with a mean of 0 and a variance of .04. Consider a set of applications of that method, from which we infer, in each case, that the length measured lies in the interval r ± .8 (i.e., within four standard deviations of the observed value.) Surely, by any ordinary standard, these results are acceptable. But if we accept a large number *n* of these results, it will also be overwhelmingly probable that at least one of them is wrong — according to the same distribution. The resulting body of knowledge is inconsistent.

The picture we work with so far is this: There are two sets of sentences we use to represent our bodies of knowledge. One, the practical corpus, contains the other, the evidential corpus, as a part. Everything in our evidential corpus is also in the practical corpus, since an item is a member of the practical corpus if and only if the lower bound on its probability (since we are using evidential probability), relative to the evidential corpus, exceeds some fixed probability p.

Statements may come and go, in the practical corpus, according as their probabilities vary with the contents of the evidential corpus. Thus there is no direct problem of revision, expansion, or contraction: all are taken care of by the probabilistic rule of acceptance. This applies to statistical statements, as well as other statements. So we will have such statistical statements in our practical corpus as "about 95% of birds fly," "less than 2% of penguins fly," etc.

Now how about the corpus of evidential certainties? How do statements get in this corpus? By being probable enough, if we are to have a uniform treatment of acceptance and corrigibility. But we can't (for reasons pointed out in Kyburg 1963) just consider simultaneously a sequence of bodies of knowledge. So we must construe a question about the contents of the evidential corpus as shifting context: now we are thinking of a different and higher level as the "evidential" corpus, and what was the evidential corpus as a practical corpus.

3. Probabilistic Inference

Statistical inference is no problem for evidential probability, but there is no ordinary way that empirical generalizations ("All Crows are Black," "Length is additive under collinear juxtaposition," etc) can be given probabilities. And it is just such items of knowledge that we would like to be able to correct. A related fact is that epistemological probability is defined only relative to a fixed language: the definition is syntactical, and depends on the recursive specification of potential reference classes and potential target classes. How do we handle generalization? And how do we deal with the relativization of probability to a language?

The key notion is that of error. We do not suppose that we have a clear cut distinction between "observational" predicates and "non-observational" predicates. We suppose instead that there is a metalinguistic corpus, parallel to our evidential corpus, that contains a representation of our knowledge concerning observational error. For example, it is there that we store the knowledge that method M for measuring length yields errors approximately normally distributed with a mean of 0 and a variance of .04.

The details of this construction are to be found in (Kyburg, 1984). The general idea is that empirical generalizations and theories are construed as features of the language we choose to use. But each of those possible languages will have going along with it, based on a given stock of actual experience, a corpus of knowledge concerning observational error. Good "observational" predicates are those that can be used with little chance of error; "non-observational" predicates will be those that have significant errors associated with them.

Observational error is generated by the interaction of our experience and a language in the following way: We know that error has occurred when we make a set of judgments that cannot all be true. Thus if we were content to live in the flowing sensuous moment, we need never suppose we made an observational error. But our bodies of knowledge would be empty of predictive content, communication would be useless, and language would be impossible. Alternatively, if we were willing to disre-

https://doi.org/10.1086/psaprocbienmeetp.1990.2.193078 Published online by Cambridge University Press

gard experience, we could hold any theories we pleased; observations that conflicted with what our theories led us to could be dismissed as erroneous, and, like the seers of old, we would have achieved TRUTH.

What we need, then, is a way of choosing between candidate languages on the basis of the consequent errors associated the languages. In earlier work (1984, 1990) we approached this question in a very abstract framework, with a view to obtaining treatments of error in both direct and indirect measurement. Here we will adopt the same general standpoint, but examine a variety of replacements of framework assumptions (and expansions and contractions) that are rather more specific.

4. New Observations

There are a number of ways in which new data can impinge on our old body of knowledge. The most common is simply to have new observations added to our body of knowledge. This has an impact on what we believe even when it does not contradict anything we already belief. This impact has two forms. To accept the observation that A is a crow and that A is black entails, in our body of knowledge that A is a bird, since we know that all crows are birds. What is entailed by our background knowledge, and the new observation, becomes part of our background knowledge. (Subject to some caveats we'll get to later: the consequences of long conjunctions of premises may not be in our body of knowledge.)

The other form, more interesting in this context, is the impact that the observation has on our general statistical background knowledge. If we have statistical beliefs concerning the frequency with which A's are B's —e.g., that it is between p and q — and we observe an A that is not a B, that should change our body of knowledge, but not very much. If we had earlier accepted our statistical knowledge on the basis of an observation of n A's, of which m were observed to be B's, we now have, as a basis for our statistical knowledge about A's and B's a sample of $n + 1$, of which m are B's. It is clear that our body of knowledge will change relatively gradually as new observations come in: we will not, in this context, find the discontinuities that we observed earlier.

There is also the possibility that our background knowledge, even though statistical, is based on more than observation. For example, my belief that the chances of a birth being the birth of a male is about in [.50,.52] is based on lore obtained from sources that I regard as reliable. To learn that my daughter just gave birth to a boy will not only have little impact on that statistical generalization: it will have *no* impact. But if my source of knowledge were impugned, that would have a large effect. And it is conceivable that I could myself acquire such a large database of sex observations that my own data would impugn the authority on which I had accepted the conventional interval.

This also applies to the sort of statistical knowledge based on physical principles and assumptions. If a die is well balanced, then the velocities and momenta that characterize its trajectory will lead to its landing on each side with very nearly equal frequency in the long run, in view of the fact that very small changes in these momenta will lead to discontinuously different outcomes. If I roll a die and get a '1', my beliefs concerning its statistical characteristics will be unchanged. (Contrary to the Bayesian view, which would demand a tiny change.) If I roll the die a lot, and get a disproportionate frequency of '1's', then at some point I will question my assumptions — in particular, the assumption that the die is well balanced — and replace (not modify) my belief that the long run relative frequency of 1's is 1/6, by a statistical belief determined by my experience. (This will not be a very exact statistical belief, since I may well make this replacement on the basis of a fairly small sample. Thus I might come to believe that the frequency of 1's is in [.5,1.0].)

Thus even in the case of statistical knowledge, augmented by some more instances, there may be discontinuities. We have continuity (and, strictly, even this is not usually continuity in the mathematical sense) only when our evidential knowledge base contains representations of all the data on which the statistical law in our practical corpus is based, and when, in addition, we obtain additional statistical evidence by a procedure which is evidentially reliable. These conditions are almost invariably met when we are philosophers in our study making up examples out of moonbeams. They are rarely met otherwise.

5. Conflicting Observations

It is useful here to make a distinction between 'observation reports' — what is said to have been observed, and 'observation statements' — what is alleged in the report to have been observed. Observation reports cannot really conflict. If I report the weight of body W on one weighing as 23.654 grams, and on another weighing as 23.655, there need be nothing wrong with my observations, although the observation statements, "W weighs 23.654 grams," and "W weighs 23.655 grams" are inconsistent. This is why the natural and appropriate observation *statement* is rather, "W weighs $23.65 \pm .02$ grams." Note that this statement is not certain: It is acceptable, because the chance of error is negligible, not because the assertion cannot be wrong. On the usual treatment of errors, under which they are treated as normally distributed, an error of any magnitude is *possible.*

We treat the interval statement as *evidence*, however. We take it to be a statement that we can use in designing machinery, in engineering, in prediction, etc. It is not a statement to which we merely assign a high probability.

Even so, it is corrigible. We may weigh W twice again, and conclude (with the same degree of justification as we had before) that it weighs $23.60 \pm .02$ grams. The two interval statements are strictly incompatible. They are contraries. One *replaces* the other.

There are various possibilities. First, we may suppose that we simply have made somewhat unusual errors of measurement. If it is evidentially certain that W weighs between 23.63 grams and 23.67 grams, then W cannot weigh as little as 23.62 grams. But if W can't change weight, the discrepancy *must* be due to errors of measurement. If this is the case, then there are two impacts of our conflicting observations: The observations should be combined; and the discrepancy between the two sets of measurements should be taken as evidence concerning the distribution of errors of measurement for the measuring device(s) involved.

Merely combining the measurements would yield $23.62 \pm .015$, if we assume that all four measurements are simply taken from the same normal population of measurements. But the discrepancy might suggest that we should regard the measurements as coming from two distinct populations (corresponding to two instruments, say), or as coming from a population with a larger variance than we had thought.

In general, the conflict among observation reports must be taken as evidence concerning the reliability of the observer, or of the apparatus, of both. We will find that this is true also in the case of more basic conflicts.

6. Conflict between Observations and an Accepted Framework

This is the most interesting sort of conflict. In the example of weighing just described, it can arise. The conflict of measurements may be taken as evidence throwing suspicion on our framework assumptions: "Who *says* the object can't change weight? This is the sort of conflict that is most likely to be noticed, since we often make relatively local assumptions that we take for granted, act on the basis of, until and unless they lead us into difficulty. Good judgment consists in knowing when to abandon an assumption. But can good judgment be codified, reduced to mechanical rules? In some respects, we will argue, it can.

The simple-minded view of belief change is this: You have a generalization (general assumption) that you have taken for granted that leads you to infer that observational circumstances C will be followed by or accompanied by observational outcome O. You observe C. You observe some contrary of O. You reject your assumption.

But things are almost never this simple. Even when (rarely the case) a qualitative generalization is understood to be strict, to admit of no exceptions, there are alternatives to rejecting the generalization in the face of apparently conflicting observation. We may take the alleged observations to have been in error. Illusion, hallucination, are always available to explain away apparent refutations. And this is not irrational. In fact it has been argued (Kyburg, 1984) that this is the source of our knowledge of the qualitative errors of observation. The identification of an object or observation as belonging to a given *kind* is subject to error. The frequency of such errors is given by two principles: One is the conservation principle:

> We should not attribute more error to our observations than we are obliged to by the model of the world we accept.

The other principle guiding our assessments of error is the distribution principle:

> *Given* the satisfaction of the conservation principle, we should distribute the errors we are obliged to attribute to our observations as evenly as possible among the *kinds* of errors we might have made.

Thus if our model of the world assumes (presupposes) that all crows are black, and we have some observations of blue crows, we would assume that those observations contain errors. And further that the errors (other things being equal) are distributed equally between judgments of blueness and judgments of crowness. The metalinguistic fact that we must assume that we have made these errors of observation provides evidence about the *reliability* with which blueness and crowness can be identified.

7. Quantitative Observations Conflicting with Laws

Suppose in general that we assume the quantitative law, $y = f(x,z)$ in our body of knowledge. Then we observe a series of measurements of the quantities X, Y, and Z. No set of measurements can contradict the law in question, since any measurement is subject to error, and indeed, on the usual theories of measurement error, subject to error that can *possibly* be arbitrarily great. But of course large discrepancies, relative to a body of knowledge that contains the law in question, are extremely improbable.

The same general approach makes sense: The very improbable happens all the time (the particular set of measurements we make would be improbable even if they

https://doi.org/10.1086/psaprocbienmeetp.1990.2.193078 Published online by Cambridge University Press

agreed with our assumed law, and the law *were* true), but if there is an alternative that renders the improbable not so improbable, the observations support that alternative. To put a quantitative measure on this is not trivial. One way, in terms of the framework we have already talked about, is the following: Anomalous observations can have two effects: they can provide new data concerning the errors of observation of a certain sort, or they can be taken at face value, and thus provide grounds for the rejection of general formulas. What we need are principles that will explain why one approach (admit to more errors of observation) or the other (reject a quantititive law) is to be preferred. Note that treating the law as an "approximation" or "idealization" is simply a way of taking it to be false, a way of rejecting it. A full discussion of this would amount to a general discussion of scientific inference. A more detailed treatment will be found in Kyburg (1990).

8. Fundamental Assumptions

Before going on to consider the grounds on which one would choose to give up an assumption in favor of attributing errors to one's observations, it is worth looking at one more extreme cases. This is that of measurement, and has been discussed more fully in (Kyburg, 1984). We suppose that length is additive: that the length of the collinear juxtaposition of two bodies is the sum of their lengths. Our measurements, of course, do not support this supposition; less dramatically: we can maintain the additivity of length only by attributing error to almost all our measurements.

Is this the alternative? To suppose that we can measure accurately, but that length is not additive, on the one hand, or, on the other, to suppose that length is additive, but that all our measurements are infected with error? Put this way it seems odd that one would ever opt for the second alternative. But we do.

Here is a possible explanation. The errors of measurement we need to introduce are very rarely large. They therefore do not deprive us of much useful knowledge. But the additivity of length is an enormously powerful predictive device. Knowing the length of two rigid bodies, we know, without even measuring, the *approximate* length of their collinear juxtaposition.

The choice between attributing error to observations and maintaining a generalization, as opposed to taking observations to be accurate and to refute the generalization, lies in the predictive observational content of the whole body of knowledge involved.

9. Choosing Between Assumptions and Errors

Suppose we consider two bodies of knowledge, one that embodies among its evidential certainties (among other things) the assumption A, the other of which does not. We make a set of observations (add to our evidential certainties a set of observation reports). We have in our background knowledge statistical information about errors in observations of this sort. Given the assumption A, the observation reports must be taken to embody unusually (improbably) large errors. These errors are not without observational consequences. They render observational predictions less dependable, since the correspondence between what is predicted and what *probably* going to be observed is only approximate, and reflects our knowledge of errors of observation.

How do we weight the advantages of one choice or the other? In order to have an actual measure that will yield an answer in these cases, we must focus on a class of predictive statements — that is, a class of statements that is of interest to us in the circumstances at hand. It is in this class that the predictions of the two cases are to be found. Let this class be **C**. We also need a measure of the precision of the predictions: thus if a prediction has the form "Bird *B* is Blue," the amount of content of that prediction should reflect the chance of an error in the observation that would test that prediction. If we can't accurately tell blue things, there is less content to the prediction that something is blue. If the prediction has the form, "Object *O* will be observed at an angle between a – d and a + d," then its content will reflect the distribution of errors of observation of angle in the circumstances under consideration.

The class **C** of predictive statements about which we are concerned should be finite. It can be large, but we want to ensure that ratios are well defined in it. Next we need a measure **m** of accuracy, or predictive usefulness. What we need are only:

(i) a (finite) set of sentences **C** that include all those that may be of predictive interest in a given context, and
(ii) a measure **m** of how important errors of various kinds are.

We get the frequencies of error from our background knowledge of the observation reports we have had, together with assumptions of our body of knowledge. When we change the assumptions (or eliminate one) we change the statistical representation of these errors that we have reason to accept. If, for example, we eliminate an assumption, we can replace a number of predictions (those that stemmed from that assumption) by no predictions. If we increase the error of a certain kind of observation, we decrease the value of what we can predict.

Let **B** be a set of atomic sentences from **C** reflecting their historical proportion in our experience. Add **B** to our body of knowledge. Let **P** be the set of sentences that then become newly practically certain. The *predictive content* of a body of knowledge, relative to **C** and **m**, might be measured by the sheer number of predictions in **P**, each weighted by its reliability. This is a crude measure for determining the replacement of one general scientific theory by another, but for many purposes, it might be illuminating. In limited circumstances, we can find a class **C** that includes the statements that concern us; what is at issue is itself relatively straight-forward (Shall we assume that instrument *I* is working correctly, or shall we assume that it is broken?); and in these cases predictive content provides an appropriate criterion.

When we have quantitative statements in **P** , the natural measure of predictive content is $1/|u - l|$— the measure of the precision which which we can confidently (with evidential certainty) predict. Of course the interval reflects the scale on which the quantity is measured. To alleviate the problem of artificial changes of scale, we can take the maxium and minimum values of a quantity in **B**, normalized to [0,1], to determine the scale. The measure **m** will then consist of the sum over **P** of the lengths of all the predictive intervals, reduced to that common scale.

10. Relation to Other Procedures

Can we relate this approach to replacement to other replacement formalisms that have been used? There is no direct reduction, obiously, since we are looking here at only a small segment of the statements of the language. Furthermore, these statements are not even statements that we have (now) reason to believe: what we have is a set **C** of statements that we are using a a *test instrument* for determining the relative desirability of two alternative frameworks. The procedure offered here is far less global than the procedures offered by Levi, Gardenfors, Spohn, and the others. It only takes into account the usefulness of one alternative, compared to another, when they make a difference to the test consequences **C** under the information measure **m**.

https://doi.org/10.1086/psaprocbienmeetp.1990.2.193078 Published online by Cambridge University Press

This seems quite natural (and perhaps even useful) when it comes to weighing relatively local assumptions. But it is not clear how far it can be extended, and how generally plausible the procedure can be made.

On the other hand, the procedure outlined is more general than the global replacement schemes previously suggested, since it allows us to compare two quite different languages, so long as they have the same test consequences **C**. We do not need an information measure on the languages themselves, nor do we need a global probability measure. Straight forward evidential probabilities will suffice.

Furthermore, our approach opens up the feasibility of treating framework assumptions as if they were infallible — that is, as genuine assumptions. The problem that has always surrounded talk of "presuppositions," "local assumptions," "framework assumptions," and the like has been their imunity from critical control. How do you weigh one against another? How do you tell when an assumption is dumb, compared to another that might have been made? It is hoped that we have offered an approach which allows these comparisons to be made in a rational way.

11. Summary

Global approaches to replacing one theory by another require relatively universal conventions: an ordering of all the sets of sentences in a formal language, for example, as well as, a Bayesian probability measure over all the sentences in the language. Approaches to eschewing acceptance, and therefore replacement, such as proposed by "Bayesian probabilists" tend to be impractical for simpler reasons: too much computation is devoted to issues that are at best peripheral to the often relatively simple question at hand, e.g. "Should we assume that instrument *I* is operating correctly?"

We have proposed instead an approach characterized by a set of sentences **C** (sentences that could, in principle, be construed as predictive observational sentences in the sense characterized above), and also by a measure of informational value **m** determined by a distribution of errors for these sentences. Suppose we are given a pair **(C,m)** consisting of a set of sentences and a measure of the importance of errors. Suppose we are given a body of knowledge. Then the relative value, in the face of a given body of observation reports, of two assumptions, or of one assumption as opposed to none, is determined. It is determined by machinery of evidential probability that we already have in hand.

There is, of course, the problem of determining the pair **(C,m)** to fit a given context. We have not yet dealt with this problem. We observe only that it is a far less overwhelming problem than that of determining informational content of all the sentences of a language (Levi) or of associating with each sentence of the language an ordinal number (Spohn). It can be done for a specific class of circumstances when certain kinds of predictions or anticipations are the kinds at issue. When the "assumptions" about which we are talking are relatively limited in scope ("Instrument 47 is working correctly"), it is not at all unreasonable to suppose that in fact we can isolate such a useful set of sentences. The question of deriving such a set of sentences from our concerns in a given context, and the question of deriving the importance of various kinds of error from the utilities of the outcomes possible in a given context, are questions that must be reserved for another time.

Note

[1]Research on which this work is based was supported by the Signals Warfare Center of the United States Army.

References

Alchourron, C.E., Gardenfors, P., and Makinson, D. (1985), "On the Logic of Theory Change: Partial Meet Functions for Contraction and Revision," *Journal of Symbolic Logic* 50: 510-530.

________________________________. (1982), "On the Logic of Theory Change: Contraction Functions and Their Associated Revision Functions," *Theoria* 48: 14-37.

Carnap, R. (1950), *The Logical Foundations of Probability* , University of Chicago Press, Chicago.

Fisher, R.A. (1956), *Statistical Methods and Scientific Inference*, Hafner Publishing Co, New York.

Gardenfors, P. (1986), "The Dynamics of Belief: Contractions and Revisions of Probability Functions," *Topoi* 5: 29-37.

_________, (1986), "The Dynamics of Belief: Contractions and Revisions of Probability Functions," *Topoi* 5: 29-37.

_________, and Makinson, D. (1988), "Revisions of Knowledge Systems Using Epistemic Entrenchment," *Proceedings of the Second Conference on Theoretical Aspects of Reasoning About Knowledge*, M. Vardi (ed) Morgan Kaufman, Los Altos: 83-95.

Harman, G. (1989), *Change in View*, Bradford Books, Cambridge.

Jeffrey, R.C. (1965), *The Logic of Decision* , McGraw-Hill, New York.

Kyburg, H.E. Jr. (1963), "A Further Note on Rationality and Consistency," *Journal of Philosophy* 60: 463-465.

__________. (1988), "Full Belief," *Theory and Decision* 25: 137-162.

__________. (1961), *Probability and the Logic of Rational Belief*, Wesleyan University Press, Middletown, Ct.

__________. (1990), *Science and Reason*, Oxford University Press.

__________. (1974), *The Logical Foundations of Statistical Inference* , Reidel, Dordrecht.

__________. (1984), *Theory and Measurement* , Cambridge University Press, Cambridge.

https://doi.org/10.1086/psaprocbienmeetp.1990.2.193078 Published online by Cambridge University Press

Levi, I. (1967), *Gambling with Truth*, Knopf, New York.

_ _ _ _. (1980),*The Enterprise of Knowledge*, MIT Press, Cambridge.

Popper, K.R. (1959), *The Logic of Scientific Discovery* , Hutchinson and Co., London. First German Ed. 1934.

Spohn, W. (1987), "Ordinal Conditional Functions: A Dynamic Theory of Epistemic States," in *Causation in Decision, Belief Change, and Statistics*, W. Harper and B. Skyrms (eds) Reidel, Dordrecht, pp. 105-134.

https://doi.org/10.1086/psaprocbienmeetp.1990.2.193078 Published online by Cambridge University Press