INDUCTIVE INFERENCE AS AMPLIATIVE AND NON MONOTONIC REASONING

Isaac Levi*

Abstract

Inductive Inference is reasoning that justifies change from one state of full belief or absolute certainty to another by adding new information to the initial state that is consistent with it but does not entail it. Inductive inference is, therefore, *ampliative*. It is also *non monotonic*. But the deviations from monotonicity differ from those characterized by belief revision according to Alchourrón, Gärdenfors and Makinson or preferential entailment developed by Kraus, Lehmann and Magidor and Lehmann and Magidor. The default reasoning of Reiter resembles ampliative reasoning more closely according to the decision theoretic account of inductive expansion proposed in this essay. The paper seeks to explain the differences and why they occur.

Induction and Abduction:

Inductive inference is a species of justification. What is to be justified is an expansion of the current state of full belief **K** representable by a deductively closed theory or corpus K in a given language L by adding some specific item of information *h* to K and forming the deductive closure. The result is a new state of full belief. The reasoning involved in the justification shows, so I maintain, that the expansion by adding *h* is a best option given the goals of the inquiry in which the deliberation about expansion is embedded and given the expansion strategies available for choice relative to K.

The current state of full belief is the inquirer's standard for distinguishing between logical possibilities that are serious possibilities and logical possibilities that are not. The serious possibilities are those propositions consistent with the inquirer's current state of full belief – that is to say with the background assumptions, beliefs about experimental data and other propositions the inquirer is committed to judging to be certainly true.

Judgments of uncertainty[1] introduce fine-grained distinctions between serious possibilities whose negations are also serious possibilities. Judgments of corrigibility or entrenchment discriminate between absolute certainties and the logical possibilities they rule out as impossibilities. They should not be confused with one another.[2] Someone may be just as certain that the December Tsunami struck the east coast of Africa as he or she is that it devastated Aceh. But where on the east coast it struck may be unknown to him or her. That person may, nonetheless, judge it more likely that it struck the northeastern coast rather than South Africa even though both alternatives are serious possibilities. The same inquirer may be prepared more readily to give up the claim that it struck somewhere in Africa than that it devastated Aceh in the light of countervailing considerations even though in the absence of these considerations, the inquirer is absolutely certain of the truth of both of them. Thus, both judgments of uncertainty and judgments of entrenchment *presuppose* states of full belief. They cannot replace them.

---

* John Dewey Professor of Philosophy Emeritus, Columbia University

[1] There are two types of judgment of uncertainty: judgments of credal or subjective probability and judgments related to degrees of disbelief (potential surprise) and degrees of belief proposed by G.L.S. Shackle (1949).
An inquirer's state of credal probability judgment is representable by a (not necessarily closed) convex set of conditional finitely additive probability functions (the set of *permissible* probabilities) over an algebra of propositions generated by a set of hypotheses exactly one of which is implied to be true according to the inquirer' state of full belief K and each of which is consistent with K and, hence, a serious possibility according to the inquirer. See Levi, 1974, 1980a, 1986, and 1999 for discussion of (determinate and indeterminate) credal probability and its relation to decision making.
Serious possibilities are also evaluated with respect to how well surprising it would be if they turned out to be true or or, how confident the inquirer is that their negations are true. Evaluations of this type were proposed a long time ago by G.L.S. Shackle (1949, 1961) and have been reinvented in one form or another many others (L.J. Cohen, 1977, L. Zadeh, 1978, W. Spohn, 1990, D. Dubois and H. Prade, 1992, P.Gärdenfors and D.Makinson, 1993. My own reconstruction of Shackle's ideas is found in Levi, 1966, 1967a, 1972, 1977, 1979, 1980c and 2002. A brief summary of it is provided later in this paper.

[2] As I believe is done in Spohn 1990 and Gärdenfors and Makinson, 1993.

The conclusion of inductive inference is typically represented as a proposition whose addition to the state of full belief determines what I am calling the inductive expansion of the state of full belief. Formally the inference can then be represented as a transformation K^i of an initial state of full belief K. It is important, however, to distinguish this interpretation of the transformation from other applications.

For example, Gärdenfors and Makinson (1993) take K^i to be an expectation set distinct from the set K of full beliefs.

> "…when we reason, we make use of the information not only that we firmly believe but also the expectations that guide our beliefs without quite being part of them. Such expectations may take diverse forms, of which the two most salient are *expectation sets* and *expectation relations*. (Gärdenfors and Makinson, 1993, p.2.) .

The expectation sets include not only "firm beliefs as limiting cases" but "also propositions that are plausible enough to be used as a basis for inference so long as they do not give rise to inconsistency". (*loc. cit.*) The assessments of plausibility that give rise to expectation sets turn out to have the formal properties of degrees of belief in Shackle's sense. An expectation set, according to Gärdenfors and Makinson as I understand them, is a set of propositions that are plausible to some positive degree or other – i.e., that carry positive degrees of belief in Shackle's sense. The expectation set K^i supported by K, however, does not represent the results of a genuine change in state of full belief but rather a system of judgments of a different kind than full belief.

The expectation set is merely the set of propositions that are judged plausible to some degree or other relative to the unexpanded K. An inquirer X might justify expanding inductively from K to K^i because the items in K^i are "sufficiently plausible" relative to K. But Y might refuse because carrying a positive degree of plausibility is not deemed a sufficient warrant for such expansion. Y might insist on a high degree of belief. Inquirer Z might declare that no level of plausibility is sufficient warrant for inductive expansion. X, Y and Z might, nonetheless, all agree that K^i is the set of positively plausible hypotheses relative to K.

Z's attitude is typical of the sentiments of *anti inductivists* like Wolfgang Spohn (1990) who called the expectation set the set of "plain beliefs". Like Karl Popper (1959), Rudolf Carnap 1950, 1960) and R.C.Jeffrey (1965), Spohn denies the legitimacy of inductive expansion.

I, by way of contrast, am a card-carrying inductivist (Levi, 1967a,1967b, 1991, 1996, 1997, 2001, 2004a.). I think it is important to give some account of how inductive expansions are to be justified.

Induction as understood here ought not to be confused with what Peirce called "hypothesis" at the beginning of his career and later called "abduction". Abduction is conjecturing. It undertakes to identify potential expansion strategies for the inquirer to evaluate when engaged in the task of induction. Its conclusions are as Peirce wrote "in the interrogative."

The conclusions of induction are not conjectures but the coming to be certain that the conclusion is true and to judge it true with certainty. Abduction identifies the potential answers relevant to solving a specific problem or answering a given question. Induction recommends a potential answer to be adopted as the warranted solution. [3]

Inductive Inference and Inductive Behavior.

[3] To be sure, at the beginning of his career, Peirce wrote that both induction and abduction are species of "ampliative" reasoning. This suggested to W.V. Quine that induction is a species of abduction and, hence, is itself a form of conjecturing. But conjecturing is not coming to full belief – to absolute certainty. There is no ampliative inference in conjecturing. There is no change in either in full belief or degrees of belief (whatever these might be). Conjecturing or supposing or entertaining or fantasizing is conflated with believing only at the cost of one's sanity. Inductive expansion is genuinely ampliative. The inquirer X's state of full belief is changed by adding new information to it that is not already implied by the old state. For further discussion of Peirce, see Levi, 1980b and 2004b.

Inductive Inference as I conceive it is an outcome of deliberate choice among options for expansion of a state of full belief. Even so-called induction by simple enumeration as exemplified by the conclusion that all ravens are black on the basis of observation of a sample of the raven population involves a choice between estimates of the percentage of black ravens in the population of ravens sampled. Other more sophisticated versions of parameter estimation, the selection of a theory to model or explain phenomena in a given domain, or predictive tasks all call for coming to be certain that some proposition chosen from a stock of potential answers is true.

Choice of an expansion strategy is, like other forms of deliberate choice, an exercise in optimizing relative to given goals – at least insofar as an optimum can be said to exist. In most activities, however, rational agents do not engage in deliberate decision making in order to promote either cognitive or any kind of practical goals. They respond to the signals they receive from their environment automatically in accordance with some program. We are creatures of custom and habit taught David Hume. Insofar as deliberation is involved at all, it goes into the modification and devising of programs to replace the products of custom and habit. Charles Peirce and, following him, J.Neyman and E.S.Pearson 1932-33), and A. Wald (1950) devised programs for using the data obtained from experiment as input rather than as premises or evidence in predesigned programs for forming estimates of statistical parameters.

The crucial difference between forming beliefs using data as input as a species of inductive behavior and forming beliefs using data as evidence from which the beliefs are inductively inferred may be crudely illustrated by the following example.

X is confronted with an urn. X is certain that the urn contains 99 black balls and one white (H_B) or 99 white balls and one black (H_W). X is afforded the opportunity of observing the color of a ball selected at random from the urn and choosing between a bet where X wins 1000 dollars if H_B is true and losing 1000 dollars if it is false, a bet where X wins 1000 dollars if H_W is true and losing a 1000 dollars if it is false or refusing to bet altogether.. Here the issue is not whether to form a belief but which bet to take. If X had to choose without observation, X should refuse to bet. X has no determinate probability as to whether H_B is true or false. The probability interval [0,1] represents X's state of credal probability judgment. For some values of the probability the first bet is best, for some the second bet and for some refusal to bet. If X then takes security rather than expected utility into account, refusal is recommended.

If X observes the outcome of the draw and then updates X's probability judgment using Bayes theorem when X's prior credal probability state is maximally indeterminate, the same probability interval represents X's probability judgment. X has used the observation of the color as evidence and obtained no benefit from it. The data are useless.

But suppose X adopts a program before sampling for responding to the result of observation. X can reason that adopting a rule dictating acceptance of the first bet if a black is observed or the second bet if a white is observed has a 99% chance of yielding a payoff of 1000 dollars. Choosing this rule is not to be confused with choosing one of the bets. X is letting the gamble X will accept be determined by the outcome of a stochastic process over which X has no control. X receives a sensory stimulus from the environment and responds by accepting a bet.

While implementing this process, X avoids exploiting information about the color of the ball drawn as information to be taken into account in evaluating the expected values of the gambles. Exploiting that information does not change the fact that the chance of obtaining 1000 dollars by following the rule is 0.99. But X can no longer legitimately use that chance to ground X's degree of belief that X will win 1000. Exploiting the information as evidence obliges X to derive X's degree of belief from the chance of following the rule yielding a payoff of 1000 dollars when the ball selected is black. That chance is either 0 or 1. It is not 0.99. Exploiting the information given in the data *as evidence* gives an indeterminate degree of belief. By way of contrast, choosing the rule is adopting a program for *routine decision-making* where the "data" or observations made are used as *input* but not as evidence. In this example, data are used as input in a program whose output is the implementation of an action. *Routine expansion* is routine decision making where the output is the formation of a full belief.

X may not have chosen the program for routine expansion but acquired it as part of X's natural equipment or through education. But we are imagining that X is deliberating as to which such program among the several available to adopt. The expected payoff of the program just described is clearly better than the minimum expectation of any alternative option or mixture of options. X could seek to defend the choice of this program on such a basis. The argument holds only insofar as X does not use the observed color of the ball as part of X's total information but bases the plan and its evaluation on the information prior to observation. The data are used as input but not as evidence. This is the kind of practice that Neyman called inductive behavior even when the output is the formation of a belief or the rejection or acceptance of a statistical hypotheses.

The use of this kind of statistical design is quite legitimate in the devising of plans for estimation and decision making to which the agent is committed prior to data collection. It is entirely illegitimate to use this analysis to evaluate data as evidence retrospectively. In particular, it is unacceptable to use in inductive inference where data need to be marshaled as evidence to be used as premises for or against a hypothesis. Peirce, who recognized that the method could not be used retrospectively, nonetheless, maintained that the approach gives a good account of what he called "quantitative induction". In truth, the idea of precommitment strategies gives a promising account of expansion using sensory inputs or testimony of experts to form beliefs. But Peirce's effort to equate inductive inference with inductive behavior confuses using data as input with using data as evidence in a manner that Peirce's own insights suggested ought to be avoided. Inductive inference is argument where data are used as evidence – i.e., as part of the premises aiming to show that expanding K by adding one potential answer rather than another is admissible for the purpose of promoting the goals of the inquiry relative to all the evidence in K prior to its contemplated expansion.[4]

The Ultimate Partition and the Space of Potential Answers:

Let inquirer X be concerned to answer some question when X's state of full belief is represented linguistically by the corpus (deductively closed theory) K in language *L*. The question may be a question of prediction (e.g., which of several candidates will win an election?, does the bird Tweety fly?), of estimating the value of a parameter (e.g., What is the difference between the mean temperature of the Earth in the past decade and the previous decade?) or deciding between rival theories that provide candidate explanations of certain processes.

The results of the efforts of abduction to articulate potential answers may be represented by a set U_K of sentences in *L* such that (a) K entails the truth of exactly one element of U_K and (b) each element of U_K is consistent with K. U_K is X's ultimate partition. For the purposes of this discussion, U_K is finite.

The elements of ultimate partition U_K express the strongest potential answers recognized by X to the question under investigation relative to X's state of full belief K. They are not maximally specific descriptions of the world. They do not express possible worlds. The assumption of the existence of possible worlds or of maximally specific descriptions of the world is incoherent.

A potential answer relative to K and U_K may be represented as follows:

1. Identify a subset *R* of U_K (the *rejection set*) and the set $\sigma_R = \{U_K/R\}$ of unrejected members of U_K.

2. Let *h* assert that at least one member of σ_R is true.

[4] The discussion of expansion explained here initially was proposed in Levi, 1967a and improved in 1967b. It was combined with an account of indeterminate probability and utility in Levi, 1980a and reiterated in connection with theories of belief change in Levi, 1991 and 2004. The connection with non monotonic reasoning is examined in Levi, 1996 and 2001

3. K^+_h is the potential answer relative to K and U_K. The same potential answer may also be represented by R or by σ_R or by h.

Notice that when $R = \varnothing$, the potential answer is to remain with the status quo K. When $R = U_K$, the potential answer is to expand into inconsistency.

Seeking valuable information and avoiding error.

I have examined the ramifications of an inquirer deciding between these potential answers when the aim of the inquirer is a compromise between two competing desiderata: Seeking valuable information and shunning error. (Levi, 1967a,b, 1980a, 1997, 2001.)

Seeking valuable information may be characterized as a utility function that orders the potential answers in a manner that satisfies the following *weak positive monotonicity condition*.

If $K^+_h \subset K^+_g$, g carries at least as much informational value as h.

A comparison of potential answers to a given question relative to U_K and K provides a partial ordering with respect to logical strength relative to K. Thus K,g |- h but not vice versa. This assessment of information is converted into a utility function that, among other things, extends the partial ordering to a weak ordering. The conversion is required to obey a weak rather than a strict positive monotonicity condition in the sense that allowance is made for the possibility that although g carries more information than h, the inquirer discounts the value of the extra information. Weak positive monotonicity forbids, however, ranking h higher than g with respect to informational value.

There are many types of utility function that can satisfy weak positive monotonicity. In the context of inductive expansion, I have argued elsewhere for using a function $Cont(x) = 1 - M(x)$ where $M(x)$ is a probability distribution over the elements of U_K that generates a probability on the Boolean algebra generated by U_K. $Cont(x)$ is a measure of *undamped* informational value. We might use this measure or any positive affine transformation of it as the utility function if the sole interest were to maximize informational value.

The utility of avoiding error is representable by a function $T(x,t) = 1$ and $T(x,f) = 0$ (or any positive affine transformation of T.)

The utility of seeking new error free information of value relative to K and U_K should, so I have argued, be equal to a weighted average of these two desiderata. The requirement that the new utility function should be a weighted average flows from the demand that *Cont* and *T* be representations of von Neumann-Morgenstern utility functions over all mixtures or lotteries of the payoffs for the potential answers, that the potential resolution should satisfy the same requirements and that the ranking induced by the potential resolution over the set of mixtures should preserve all shared agreements among the preferences of the two desiderata. So the utility function we seek should have the following form:

$U(x,t) = \alpha T(x,t) + (1-\alpha)(1-M(x,)$.

$U(x,t)$ is a positive affine transformation of the following:

$V(x,t) = T(x,t) - qM(x)$ where $q = (1-\alpha)/\alpha$ for positive α. Since α represents the relative importance of avoiding error as opposed to acquiring informational value, I have argued that $\alpha \geq ½$ (so that $0 \leq q \leq 1$). In that way, committing an error by importing a false belief can never be valued over avoiding an error even if the value of the erroneous information is much greater than the value of the error free information.

Maximizing expected epistemic utility

We are now in a position to consider the potential answer to be recommended relative to K and U_K and the informational value determining probability M_K. I propose to examine the implications of maximizing expected utility. To do this we need a probability distribution Q_K over U_K that represents the inquirer X's state of credal probability judgment.

Where does this probability come from? Two points need to be emphasized:

a) I assume that the inquirer X's probability judgment is an expression X's endorsement of a rule for determining the credal state that should be adopted relative to various bodies of evidence or full belief. The rule is representable by a function $C: K \rightarrow B$ from potential states of full belief to potential credal states. I shall call such rules *confirmational commitments*. I do not assume that there is a single standard confirmational commitment to which all rational agents ought to subscribe. Confirmational commitments are subject to critical review and change just as states of full belief are. States of probability judgment (credal states) may change due to changes in state of full belief (i.e., evidence) or changes in confirmational commitments.

b) The credal states that are the values of the function C are nonempty sets of (conditional) probability functions that satisfy a convexity condition. These are minimal conditions. We may impose other demands as well. Strict Bayesians insist that credal states are singletons. I make no such assumption. The strict Bayesian view is but one special case according to the view I favor. However, for the sake of brevity, I focus on the strict Bayesian case where the credal state for the ultimate partition U_K is representable by a single probability function Q_K.

The expected (epistemic) utility $EV_K(h)$ of expanding K by adding h when U_K is the ultimate partition and Q_K is the credal probability, $M_{\underline{K}}$ is the informational value determining probability and the *degree of boldness* is specified to take the value q is equal to

$$EV_K(h) = Q_K(h) - qM_K(h).$$

When U_K is finite, $EV_K(h) = \sum[Q_K(c_j) - qM_K(c_j)]$. (The sum is taken over the subset σ_R of U_K where R is the rejection set and h asserts that at least one element of σ_R is true.)

Inductive expansion

A recommended expansion should maximize EV_K. To do this, σ_R should satisfy the following two conditions:

(i) If $Q_K(c_j) - qM_K(c_j)$ is positive, $c_j \in \sigma_R$.

(ii) if $Q_K(c_j) - qM_K(c_j)$ is negative, $c_j \notin \sigma_R$.

This leaves the case where $Q_K(c_j) - qM_K(c_j) = 0$ open. Adding these to optimal σ_R does not undermine optimality. If there are such elements of U_K, there can be several optimal expansions. I have proposed that adopting the weakest optimal expansion if such exists should break ties for optimality. When the credal state is a singleton, there always is a weakest optimal expansion. The resulting recommendation runs as follows:

Inductive rejection rule: For finite $U_{K,}$, c_j in U_K should be rejected if and only if $Q_K(c_j) < qM_K(c_j)$. The recommended rejection set $R(U_K, K, Q_K, M_K, q)$ determines the set σ_R and sentence h such that K^{+}_{h} is the recommended inductive expansion K^i of K. Strictly speaking K^i is the inductive expansion of K relative to U_K,K, Q_K, M_K,q.

Notice that the recommended rejection set is also a recommendation of a sentence or proposition to be added to K to yield the recommended inductive expansion. There is symmetry between rejection and acceptance.

Stable inductive expansion by adding e as inductive inference

Suppose that prior to inductive expansion, K is expanded (perhaps via observation or via a side induction) to K^{+}_{e}. U_K, Q_K and M_K are modified accordingly so that the inductive rejection rule can be applied to yield $[K^{+}_{E}]^{i} = K^{+i}_{E}$ or *the inductive expansion of K by adding e*.

Once X has changed from the belief state K to the belief state K^i, X has a new ultimate partition U_Ki, credal probability Q_Ki and informational value determining probability M_Ki. The value of q remains fixed. The inductive rejection or inductive expansion rule can be reapplied to yield $[K^i]^i = K^{2i}$ and the process reiterated until X obtains a result where $K^{ni} = K^{(n+1)i}$. Iteration of inductive expansion to such a fixed point yields a *stable inductive expansion* K^{si}.[5]

> *Stable Inductive Expansion Rule*: The recommended expansion of K relative to U_K, Q_K, M_K and q is K^{si}.

$[K^{+}_{e}]^{si} = K^{+si}_{e}$ is the stable inductive expansion of K^{+}_{e}. It is also the stable inductive expansion of K by adding e.

Shackle style degree of belief

Both the criterion for inductive expansion and the criterion for stable inductive expansion are *boldness dependent*. Inductive expansion depends on other contextual factors as well (K, U_K, Q_K, M_K). However, there may be some interest in avoiding a recommendation of a definite degree of boldness q. Prior to changing X's state of full belief by inductive expansion, X may wish to determine which potential answers would be recommended for various values of q.

Thus far, potential answers have been represented as states of beliefs that are expansions of other states of belief. If K is the initial state, X expands to belief state K^{+}_{x} by adding x to K together with all consequences of K and x. Instead of representing the expansion in that fashion, one might consider taking all the propositions that are believed to a sufficiently high degree of belief relative to the state of belief K prior to expansion and adding these to K. The notion of degree of belief, degree of expectation, degree of plausibility or of "Baconian" probability as it has been variously called is a satisficing notion in the sense that it can be used to formulate inductive expansion rules in the following manner:

> *Satisficing Expansion*: Expand K by adding all propositions that are sufficiently plausible or carry a sufficiently high degree of belief relative to K prior to expansion. After expansion, these propositions are fully believed and maximally certain.

I take it for granted that the notion of plausibility or degree of belief involved in a principle of inductive expansion in the format of satisficing expansion should yield recommendations equivalent to the recommendations of the inductive expansion (or stable inductive expansion) rule at a given degree of boldness. This assumption motivates the following definitions:

> *Degree of unrejectability*: (Levi, 1967a, 1980c). For potential answer *h* that is rejected for some values of q, q(*h*) is the maximum value of q for which *h* fails to be rejected. If *h* is rejected for every value of q in [0,1], q(*h*) = 0.[6]

If *h* is the disjunction of a subset H of U_K, q(*h*) is the maximum value of q which all members of H are rejected. This is the value of q(x) for x in H that is a maximum for members of H. Hence $q(x \vee y)$ must equal max[q(x), q(y)], at least one element of U_K carries q-value 1 and propositions incompatible with K carry q-value 0. These are the properties of possibility measures in the sense of Zadeh (1978) and

---

[5] I first broached the idea of iterating inductive expansion to a fixed point in Levi, 1967a, 151-2. My most recent discussions are Levi, 1996, 2001 and 2004.

[6]n the definition on p.89 of Levi, 2004 (as in Levi, 2002, p.327, degree of unrejectability is defined for *h* that are potential answers that fail to be rejected for some values of but not for others q. The case where *h* is unrejected for all values of q is then cited explicitly as carrying q(*h*) = 1. I failed to make provision for cases where *h* is rejected for all values of q. This is the case where the maximum at which *h* fails to be rejected does not exist. This is an oversight. I intended q(*h*) = 0 for that case as I had explicitly stated in 1967a, p.137, 1980c, p.5-6, 1996, p.185 and other places.

Dubois and Prade (1992). They were derived from boldness dependent expansion rules in Levi, 1967, ch.8 and used to define degrees of disbelief and belief in a manner satisfying the formalism proposed by G.L.S. Shackle (1949, 1961) and reinvented by many others.

d(*h*) = 1 – q(*h*) = degree of potential surprise = degree of disbelief in the sense of Shackle.

b(*h*) = d(~*h*) = degree of belief in the sense of Shackle.

Values for such functions can be derived from either the inductive expansion rule or the stable inductive expansion rule proposed here. They can be derived from other boldness dependent inductive expansion rules as well.

One can specify the degree of belief function b(*h*) relative to the other contextual factors but refrain from inductive expansion because a value for q has not been stipulated. To make such a stipulation is to specify a threshold value such that Boolean combinations of elements of U_K may be added to K.

The set of propositions carrying a degree of belief above a given threshold will yield a deductively closed corpus that expands K at a given level of boldness. If b(*h*) > r, d(~*h*) = 1 – q(~*h*) > r and q(~*h*) < 1-r. The level of boldness associated with q if inductive expansion to such a set is implemented is q = 1-r.

If one refrains from inductive expansion, one can still identify how plausible various Boolean combinations of elements of U_K are. This evaluation of hypotheses with respect to plausibility is thus a way to evaluate inductive or stable inductive expansions of K at various levels of boldness. On the account I am proposing, degrees of belief or plausibility are of value because of their use in assessing the merits of potential expansions of K. If inquirer X has no intention on any occasion of converting the merely plausible propositions to full beliefs on any occasion, I suggest that the judgments of plausibility are a mere epiphenomenon of no value in theoretical inquiry or practical deliberation.[7]

Non-monotonic Reasoning and Inference

Ordinary expansion of K by adding *e* yields the results of reasoning from K and *e* to their deductive consequences. If K, *e* |- *h*, the reasoning from *e* to *h* is appropriately taken to be an inference from K and *e* to *h*. That is because no belief is given up in the series of belief changes beginning with the addition of *e* to K. Moreover, the inference is explicative and it is monotonic. $K \subseteq K^{+}_{e} \subseteq K^{+}_{e\wedge f}$.

Inductive expansion K^{+i}_{e} by adding *e* to K is also reasoning from the addition of *e* to K to some conclusion g. Like explicative inference it also qualifies as inference from K and *e* to *g* in the sense that no belief is given up. However, the inference is *ampliative* and not explicative. It is also *non monotonic*. If the inquirer had expanded K by adding $e\wedge f$ instead of *e*, the inductive expansion $[K^{+}_{e\wedge f}]^{i} = K^{+i}_{e\wedge f}$ of the result need not have been a superset of $[K^{+}_{e}]^{i} = K^{+i}_{e}$. This is so even though K is a subset of K^{+}_{f}. The same is true for stable inductive expansion.

[7] Cohen (1977) introduced a conception of "Baconian Probability" based on a ranking of hypotheses with respect to something akin to boldness and derives Shackle measures from that. This suggests an interpretation close to the one I proposed in 1966 and 1967a and repeated with various elaborations in Levi, 1972, 1977 1996, and 2002. I discuss Cohen's view in Levi, 1979 and in ch. 14 of Levi, 1984. .Zadeh and Dubois and Prade capture the Shackle formalism but do not relate it at all to boldness in inductive expansion. Neither do Gärdenfors and Makinson (1993). Shafer (1976) recognized a formal similarity between his "consonant" support functions and Shackle's measures but claimed to be generalizing this idea by taking dissonance into account. Spohn (1990) does take Shackle's contribution seriously. Like Shackle himself, he thinks of degrees of belief and disbelief as measures of uncertainty with functions similar to probability and criticizes Shackle's failure to come up with an adequate mechanism for updating surprise (although, as Spohn notes, Shackle tried). Interpreting Shackle measures as relevant to inductive expansion according to boldness dependent rules does not fit well with such updating methods. For a discussion of Shafer's theory see Levi 1982 and 1983. I interpret Dempster as more concerned with seeking to address the concerns that led R.A.Fisher to fiducial inference and have discussed his ideas in these terms in Levi, 1980a.

I propose the following non-monotonic consequence relation to codify such ampliative inferences:

e ||~ *h* @ K if and only if $h \in K^{+i}_e$ (or K^{+si}_e) where '@K' is meant to make explicit the relativity of the inference from *e* to *h* to K (and, if need be to other contextual parameters U_K, Q_K, M_K,q).

The trite and true example of poor Tweety will serve to illustrate. If inquirer X in belief state K comes to believe that Tweety is a bird and then concludes that Tweety is able to fly, X has drawn an ampliative inference from Tweety's avian credentials. There has been no belief contravening belief change at any step. This is so even though X knew prior to making the inductive leap that no penguins fly and that it is compatible with K and that Tweety is a bird that Tweety is a penguin. In jumping to the conclusion that Tweety can fly X comes to rule out the prospect of Tweety's being a penguin as not a serious possibility.

X's inductive leap is non-monotonic because had X recognized that Tweety is not only a bird but a penguin *prior to inductive expansion*, X would not have been entitled to conclude that Tweety can fly. K^{+i}_e need not be a subset of $K^{+i}_{e \wedge f}$ even though K^{+}_e is a subset of $K^{+}_{e \wedge f}$. But neither inductive expansion being compared is belief contravening.

But suppose X had found out that Tweety is a penguin *after* inductive expansion. That is to say, X expanded X's state of full belief implying that Tweety is not a penguin with an observation that Tweety is a penguin or some other compelling testimony. The new information is inconsistent with the information in the initial inductive expansion. To save consistency X will either have to give up the claim that Tweety is not a penguin or that Tweety is a penguin or both. The conclusion of inductive expansion (such as Tweety is not a penguin) is a full belief. It is not automatic that this full belief should be contracted in the face of recalcitrant experience.

Suppose that Tweety is not a penguin is given up. Belief contravention has taken place. The result looks like an AGM revision of the result of the initial inductive expansion by adding Tweety is a bird.[8] The form looks like this: $[K^{+i}_e]^*_f$. AGM revision is not inference from *f* and K^{+i}_e by adding *f*. This is because ~*f* and other items need to be removed from K^{+i}_e in order that *f* can be added consistently.

AGM revision is non-monotonic. K^*_x need not be a subset of L^*_x even though K is a subset of L. As we have just noticed, inductive expansion is also non-monotonic. K^{+}_f is a subset of $K^{+}_{e \wedge f}$. Yet K^{+i}_e need not be a subset of $K^{+i}_{e \wedge f}$. AGM revision supports a form of non-monotonic reasoning that is not ampliative and is not inference. Inductive expansion supports a form of non-monotonic reasoning that is ampliative and is inference.

Consider the non-monotonic consequence relation built on AGM belief revision (a modified version of Makinson and Gärdenfors, 1991, 189):

e |~* *h* @K if and only if $h \in K^*_e$.

In cases where K entails ~*e*, to consider this to be codifying inferences from K and *e* to *h* is a confusion precisely because the "inference" requires giving up ~*e* and other items in K. Makinson and Gärdenfors (1991, p.188) take note of this by saying that the non-monotonic inference is not from K and *e* but from *e* rather than K.

Their notion of inference is different from mine. According to my terminological practice, a genuine inference should be supported by all the assumptions taken for granted including both the new data and the background information. I think my usage is in closer conformity with customary practice; but terminology is not the central issue here. Makinson and Gärdenfors and I agree that the reasonings supported by AGM revisions are non-monotonic and are not inferences in the sense I adopt. The crucial issue is whether there are ampliative inferences in the sense I favor that are genuinely non-monotonic.

[8] AGM as an acronym for Alchourrón, Gärdenfors and Makinson (1985).

As I mentioned previously, Gärdenfors and Makinson (1993) maintain that relative to K one can determine an "expectation set" that may be derived from a measure of degree of belief formally like degree of belief in the sense of Shackle.

The expectation sets include not only "firm beliefs as limiting cases" but "also propositions that are plausible enough to be used as a basis for inference so long as they do not give rise to inconsistency". (*loc. cit.*) An expectation set, according to Gärdenfors and Makinson, is, as I understand them, a set of propositions that are plausible to some positive degree or other – i.e., that carry positive degrees of belief in Shackle's sense. If Shackle measures are derived from inductive expansion rules along the lines I suggest, this implies that the expectation set is the set of propositions that would be fully believed if an inquirer were to engage in inductive expansion at the maximum degree of boldness q = 1. Prior to expansion, the members of the expectation set that are not fully believed are merely plausible.

Gärdenfors and Makinson do not consider boldness-dependent inductive expansion rules. Because they do not consider inductive expansion, Gärdenfors and Makinson cannot provide the kind of application for expectation sets that the account I have offered can. They suggest that we "make use" of the positively expected propositions without quite explaining what that use is except that it is something akin to but different from the use of full belief.

Thus, according to Gärdenfors and Makinson, the transformation of K into K^{+i} or into K^{+si} and the associated transformations of K^{+}_{e} determine expectation sets without representing changes in states of full belief by ampliative but non-monotonic inference. It is unclear whether such expectation sets are ever legitimately converted into full beliefs.

Keeping this in mind, let x be consistent with the firm beliefs K. Add x to K and then derive the expectation set from K^{+}_{x}. The result is K^{+i}_{x} or K^{+si}_{x}. K^{+i}_{x} (K^{+si}_{x}) taken as a transformation of K by x to another potential state of full belief will now be compared with K^{*}_{x} in bare outline. See Levi, 1996, ch.5 for extended discussion.

Let the order of the operations be permuted. An inductive expansion transformation is first performed on K. Suppose that K^{+si} entails ~x. Subsequently x is added to the firm beliefs by observation or some other independent manner. Because ~x is in the expectation set and not the set of full or firm beliefs, it is automatically given up as Gärdenfors and Makinson explicitly state (1993, 220). And the old expectation set is replaced by the expectation set for K^{+}_{x}. So the end result must be the same. And the expectation rankings of the old and the new expectation sets should agree in their intersection. We may conclude that the following two claims must hold according to Gärdenfors and Makinson for e consistent with K:

A. If e is consistent with K, e ||~ f @ K if and only if e |~* f @ K^{i}.

B. If e is consistent with K, $K^{+i}_{e} = [K^{i}]^{*}_{e}$.

Notice that when consistent e is inconsistent with K, $K^{+i}_{e} = K_{\perp}$ in violation of the analogue of K*5 for consistency preservation imposed on AGM revisions. However, $[K^{i}]^{*}_{e}$ is consistent. The consequents of A and B fail in this case.

This argument for A and B is based on the Gärdenfors-Makinson distinction between sets of full beliefs and sets of expectations. What I call inductive expansion are strictly speaking arguments for the plausibility of hypotheses relative to a set of firm (i.e., full) beliefs. The non-monotonic consequence relation '||~' codifies the derivation of a plausibility or expectation set (not a state of full belief) from the state of full belief K^{+}_{e}. Such a derivation is neither an inference nor is it inductive. Whether it qualifies as a form of reasoning and not a mere syntactic construction depends upon what interpretation can be given to plausibility or expectation sets. On my preferred interpretation of the formalism of Gärdenfors and Makinson, the expectation set derived from a state of full belief K is the recommended inductive expansion K^{i} of K to a new state of full belief when the inquirer is maximally bold. Assume as I do, that inductive expansion is sometimes legitimate and judgments of expectations have a use in determining when and how to engage in inductive expansion. It is possible to identify assumptions that secure the correctness of theses A and B for changes in states of full belief. I shall argue that these assumptions, however, are clearly unacceptable. [9]

If A and B were true, genuinely inductive expansion, ampliative but non-monotonic inference from K^{+}_{e} would be equivalent to non-ampliative and possibly belief contravening, non-monotonic reasoning from K^{i} and e except for the case where e is inconsistent with K. In that case, inductive expansion violates the analogue of K*5. The associated inference relation '||~@K' violates the corresponding version of consistency preservation and obeys all the other conditions of rational preferential entailment relations according to Lehmann and Magidor (1992). Revision of K^{i} by e remains consistency preserving. The non-monotonic consequence relation '|~*@K^{i}' is also consistency preserving but like '||~@K' obeys the conditions of rational preferential entailment otherwise.

Observation 1: Conditions A and B hold only if the inductive expansion operator satisfies the *permutability condition*: – i.e., where $K^{+i}_{e} = [K^{i}]^{+}_{e}$ when e is consistent with K^{i}.

Observation 2: If inductive expansion with q = 1 and e is consistent with K^{si}, the permutability condition, $K^{+si}_{e} = [K^{si}]^{+}_{e}$ holds.

Sketch of Proof: When q = 1, the only elements of finite U_K that survive rejection and, hence belong in σ_R when inductive rejection is stable rejection are those elements x such that $Q(x)/M(x)$ are a maximum. Hence K^{si} entails that exactly one of the members of this set $\sigma_R = U_K$si is true. If e is consistent with K^{si} it is consistent with this claim. Hence, $[K^{si}]^{+}_{e}$ implies that a subset of $U_K$si is true. All of these elements carry the same ratio of credal probability to informational value probability given K^{+}_{e} so that inductive expansion does not lead to any element being rejected. So $[K^{si}]^{+}_{e} = [K^{+}_{e}]^{si}$ (= K^{+si}_{e}) as permutability requires.

If q < 1, on the other hand, permutability can fail because σ_R may contain elements of U_K for which $Q(x)/M(x)$ is not a maximum.

To show this Permutability will be decomposed into two conditions:

Importability: $K^{+i}_{x} \subseteq [K^{i}]^{+}_{x}$

Exportability: For every x consistent with K^{i}, $[K^{+i}]^{+}_{x} \subseteq K^{+i}_{x}$.

Observation 3: If stable induction is implemented with q < 1 or if induction is not stable, importability may fail.

Proof: Let x rule out just those elements of U_K carrying maximum Q/M ratios. Once these are eliminated, it could happen at the given level of q that no other elements of U_K can be eliminated inductively. However, at the same level of boldness, K^{i} not only fails to eliminate the maximum Q/M ratios but the next level as well. But it eliminates the rest. When x is then added to the result, $[K^{+i}]^{+}_{x}$ is a proper subset of K^{+i}_{x}. This could happen whether or not induction is stable. And it violates importability. Example: U_K = {a,b,c,d} with equal *M* values and *Q*-values of <0.5, 0.2,0.2, 0.1>. Let q = 0.5. x = ~a. K^{i} = {a,b,c}. $[K^{i}]^{+}_{x}$ = {b,c} K^{+}_{x} = {b,c,d} = K^{+i}_{x}.

9 Arló Costa (1995) points out that the transformation of K supporting the non-monotonic relation satisfying requirements of consistency preserving rational preferential entailmnent could be a revision not of K itself but of a set of expectations such as K^{i}. Substituting K^{+i}_{e} for K^{*}_{e} in K*3 does not insure satisfaction of the result. Arló Costa had insisted to me that K*3 could be undermined under some interpretations. But it was only when I appreciated that it could be called into question in some versions in inductive expansion, that I appreciated Arló Costa's insight as having relevance to epistemology. (See Levi, 1996.) However, in 1995, Arló Costa did not think of the expectation set as determined by stable inductive expansion rules at maximum boldness as I do here and did in Levi 1996 as candidates for addition to the state of full belief. At best, Arló Costa would acknowledge my proposal as one of several possible models of non-mmonotonic consequence relations obeying Rational Monotonity. I am more concerned with which of these models have useful application in inquiry.

As a consequence of the failure of importability, the inductively extended version of conditionalization fails:

Inductively Extended Conditionalization: If h∧f ||~ *g*, *h* ||~ *f* ⊃ *g*.[10]

Observation 4: If stable induction is implemented with q < 1 or if induction is not stable, exportability may fail.

Proof: Take the same example as before but let x be ~b. Let q = 0.4. K^i = {a,b,c,d}. $[K^i]^+_x$ = {a,c,d}, K^+_x={a,c,d}. K^{+i}_x = {a, c}. These results hold for both inductive and stable inductive expansion and violate importability. As a consequence, inductively extended rational monotony fails. (Levi, 1996, 152-3/)

Inductively extended Rational Monotony: If not-(*h* ||~ ~*f*) and *h* ||~*g*, *h*∧f ||~ g

So does its consequence, inductively extended cautious monotony.

Inductively extended Cautious Monotony: If *h* ||~ *f* and *h* ||~ *g*, *h*∧*f* ||~ *g*.

Observation 5: Inference relations codifying stable induction at q < 1 satisfy Cut.

Inductively extended Cut: If *h* ||~ *f* and *h*∧*f* ||~ *g*, *h* ||~ *g*.

Proof: Let $f \in K^{+i}_h$ and $g \in K^{+i}_{h\wedge f}$. ~*f* is already rejected by the stable inductive expansion of K^+_h, So stable inductive expansion by adding *h*∧*f* to K cannot reject more elements of U_K. (This is not true if inductive expansion is not stable.) If stable inductive expansion by adding *h*∧*f* to K implies g, so must stable inductive expansion by adding *h* to K.

Observation 6: Inductive expansion whether stable or not and for any value of q > 0 defines a non-monotonic consequence relation satisfying inclusion, idempotence, supraclassicality and left absorption.

Permutability is necessary for theses A and B. And stable induction at q = 1 is necessary for permutability. How tenable is it to endorse stable induction at q = 1?

In many contexts, the assessment of informational value might well assign equal informational value to all elements of U_K. In such cases, the preference relation for inductive expansion will be credal probability. If q = 1, the recommended expansion will favor rejecting all elements of U_K except those carrying maximum probability. Thus, if one were to be predicting the outcome of a thousand tosses of a coin in relative frequencies, the ultimate partition would have 1001 equally informative members. All hypotheses except the one predicting that the coin will land heads exactly 500 times will be rejected. This recommendation is not inconsistent but it seems absurd. A sensible agent should predict that the coin will land heads in the neighborhood of 500 times. How wide that neighborhood might be will depend on the degree of boldness exercised. But exercising maximum allowable boldness of seems mistaken.

This toy example illustrates an obvious fact – to wit, that in many contexts, adopting the modal point estimate as a full belief is too bold. If this is right, stable inductive expansion with q = 1 is unacceptable. In general, non-monotonic inference relations meeting condition A will be violated as will condition B.

But even if this objection to q =1, is waived, we are not out of the woods. Permutability is necessary but not sufficient for theses A and B. The second assumption needed to guarantee theses A and B is the following.

[10] For proof, see Levi, 1996, 5.7

Sameness of Preference for Inductive Expansion and Contraction:: The values for q(x) of the members of the ultimate partition U_K consistent with *e* in inductive expansion of K^+_e must also rank the maxichoice contractions removing ~*e* from K^i for contraction.

In formulating the sameness of preference assumption, I have not required that the ranking q(x) conform to the values of $Q_K(x)/M_K(x)$. This is, however, the ranking I shall deploy.

In case, *e* is incompatible with K^{si}, even with q = 1, it will be necessary to remove ~e from K^{si}. Moreover, it will be necessary to do so in a manner that guarantees that the Q/M ratios of the elements of U_K that are reinstated as serious possibilities all carry maximum Q/M ratios among all elements of $U_K$consistent with e. This will require that the ranking of elements of U_K eligible for reinstatement must agree with the Q/M ranking of elements of U_K.

The sameness of preference assumption makes sense if inductive expansion to full belief is not recognized but only inductive expansion as determining plausibility or expectation sets. We have already seen that adding x to K^{+si} is, under this interpretation, tantamount to expanding K by adding x and then determining an expectation set. ~x is then ruled out of K^+_x and, hence, K^{+si}_x. There is no need for contraction.

But if inductive expansion is taken to be inductive expansion to full belief, the sameness of preference assumption is very doubtful. The reason is this. In inductive expansion, the preferences are a function of credal probability as represented by Q_K and of informational value determining probability M_K. At least that is what I would hold. I do so because inductive expansion involves a choice between options for expansion that should be compared with respect to the risk of error they incur and the value of the additional information they provide.

Contraction of a state of full belief is an entirely different affair. If the state of full belief is K^{+si}, the elements of U_K entailing x are ruled out as not possibly true. Relative to that state of full belief, the inquirer assigns 0 *Q*-value to such elements of U_K. So the *Q*/*M*-ratios must be 0 as well..

On the other hand, elements of U_K that were ruled out as false are from the inquirer's point of view distinguishable with respect to the value of the information they carry as represented by the *Cont* –function or the *M*-function. Even if an inquirer is certain that a proposition or theory is false, he or she may appreciate the proposition because of its explanatory virtues, simplicity or whatever and, hence, differentiate between relatively important good and bad falsehoods. I claim that that the degree of vulnerability to being given up possessed by a full belief increases with decrease in the loss of informational value incurred by giving it up in contraction. And that loss depends on the *M*-function alone. (Levi, 1991, 2004a.)

For this reason, the sameness of preference condition should be rejected

Conditions A and B do not appear to be acceptable if they are intended to regulate ampliative inference.

Conclusion

The large moral to this story is that ampliative inference should be taken seriously, taking it seriously calls for providing it with a decision theoretic rationale and that doing so will shed light on the understanding of non-monotonic consequence relations relevant to ampliative or inductive inference.

The non-monotonic consequence relations supported by AGM revisions are useful in evaluating conditional modal judgments according to the famous Ramsey Test. They also codify non-monotonic inference relations when belief contravention takes place. The same is true for non-monotonic inference relations characterized by the preferential models of Kraus, Lehmann and Magidor (1990) and Lehmann and Magidor (1992). The conditional judgments and non-monotonic consequences supported by AGM satisfy the requirements of rational preferential entailment according to Lehmann and Magidor supplemented by consistency preservation.

The non-monotonic consequence relations supported by stable inductive expansions are distinguished from those characterized by AGM by the fact that they are (a) strictly ampliative and violate consistency preservation and (b) from both AGM and KLM because the consequence relations are parameterized by the degree of boldness exercised. And only when $q = 1$ are the properties of AGM aside from consistency preservation and the consequence relations obeying rational monotony secured.

When $q < 1$, the non-monotonic consequence relation fails to satisfy either the inductive variant of conditionalization or cumulative rationality. Cut, however, is satisfied. The consequence relation for stable inductive inference shares the same minimal properties possessed by the consequence relation for Reiter (1980) defaults. If, as I think, q should be less than 1, the failure of Reiter non-monotonicity to measure up to the standards seemingly favored by many logicians is far from being a defect. It displays sensitivity to the importance of ampliative inductive inference.

The conception of boldness has played an important role in this discussion. Boldness as understood here represents the degree to which the inquirer is willing to incur risk of error for the value of the information acquired. And this vision of boldness is explicated within a decision theoretic framework. Within this framework, it is possible to assign a role not only to the value of information but also to the risk of error and through this idea to the notion of credal probability. [11]

In this discussion, I have pretended that both the value of information and credal probability may be represented in numerically determinate fashion. This is but a pretense. It is important to extend the ideas outlined here to cases where both credal probability and utility go indeterminate. Since the early 1970's, I have proposed a general approach to decision-making that insists on sticking as close as possible to expected utility theory as possible while taking seriously indeterminacy in probability and utility judgment. Unlike many other proposals concerned with these matters, the proposals I have made give up the ambition of ordering the available options in favor of an alternative approach to making recommendations for rational choice. I believe this approach can be brought to bear on inductive expansion. But that is a subject for another time.

The same is true for the question of inductive expansion when the ultimate partition is countably infinite or when it is non-countable as when making estimates of parameters in n-dimensional Euclidean space.[12]

Even so, treatment of the finite case with determinacy in place indicates how fruitful a decision theoretic approach to ampliative expansion can be. To my way of thinking it suggests that the widely popular prejudice in favor of anti inductivism ought to be reconsidered.

[11] In Levi, 1996 I consider some alternative inductive expansion rules. In particular, I consider the idea of rejecting element x of U_K if the ratio $[Q(x)/Q^*][M^*/M(x)]$ where $Q^*/M^*$is the maximum Q over M ratio assigned an element of U_K is less than a value k between 0 and 1. The resulting rejection rule is obviously stable. For $k = 1$ it yields the same results as the proposals I make for stable expansion with $q = 1$. I mistakenly thought that although it satisfied Cut it would lack the property of importability that assures inductively extended conditionalization. I recognized that it satisfied importability in Levi, 2001 but failed to state explicitly that this assures inductively extended conditionalization although this observation is made in Levi, 1996 , p.150. The relevance of this point is that if one can find an acceptable rationale for this ratio rejection rule for $k < 1$, the corresponding non-monotonic consequence relation would satisfy Cut and inductively extended conditionalization. This is no reason for embracing the ratio rule. The inductive expansion rule based on seeking new, valuable error free information is rationalized in a straightforward and intelligible manner. The Ratio Rejection Rule lacks such a decision theoretic rationale. Pending the provision of such a rationale, its status remains questionable.

[12] (See Levi, 1980, 1996, 2004)

References

Alchourrón, C. Gärdenfors, P. and Makinson, D. (1985), "On the Logic of Theory Change: Partial Meet Functions for Contraction and Revision,: *Journal of Symbolic LogicI 50-510-30.*

Arló Costa, H. (1995), "Epistemic Conditionals, Snakes and Stars," in I.Farinas dell Cerro, G.Crocco, A Herzig (eds.) *Conditionals, from Philosophy to Computer Science*, Oxford: Oxford University Press, 193-239.

Carnap, R. (1950), *The Logical Foundations of Probability*, Chicago: U.of Chicago Press, 2nd ed. 1962.

Carnap, R. (1960), "The Aim of Inductive Logic", *Logic, Methodology and Philosophy of Science* ed. by E.Nagel, P.Suppes and A.Tarski,. Palo Alto: Stanford University Press, 302-18.

Cohen, L.J. (1977), *The Probable and the Provable*, Oxford: Clarendon Press.

Dempster, A.P. (1967), "Upper and Lower Probabilities Induced by a Multivalued Mapping, " *Annals of Mathematical Statistics* 38, 325-39.

Dubois, D. and Prade, H. (1992), "Belief Change and Possibility Theory," *Belief Revision* ed. by P.Gärdenfors, Cambridge: Cambridge University Press, 142-82.

Gärdenfors, P. and Makinson, D. (1993), "Nonmonotonic Inference based on Expectations," *Artificial Intelligence* 65, 197-246.

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

Kraus, S., Lehmann, D. and Magidor, M. (1990), "Nonmontonic Reasoning, Preferential Models and Cumulative Logics," *Artificial Intelligence* 44, 167-201.

Lehmann, D. and Magidor, M. (1992), "What Does a Conditional Knowledge Base Entail?, *Artificial Intelligence* 55, 1-60.

Levi, I. (1966), "On Potential Surprise," *Ratio* 8. 107-29.

Levi, I. (1967a), *Gambling with Truth*, New York: A.Knopf reissued in paperback without revision in 1973 by MIT press.

Levi, I.(1967b), "Information and Inference," *Synthese* 17, 369-91 reprinted in Levi, 1983, ch.5/

Levi, I. (1972), "Potential Surprise in the Context of Inquiry," in *Uncertainty and Expectation in Economics: Essays in Honoor of G.L.S.Shackle* ed. by C.F.Carter and J.L.Ford, Oxford: Blackwell, 213-36.

Levi, I. (1974), "On Indeterminate Probabilities," *The Journal of Philosophy* 71, 391-418 reprinted in Levi, 1997, ch.6.

Levi, I. (1977), "Four Types of Ignorance," *Social Research* 44, 745-56.

Levi, I. (1979), "Support and Surprise,", L.J. Cohen's View of Inductive Probability." *British Journal for the Philosophy of Science* 30, 279a-92.

Levi, I. (1980a), *The Enterprise of Knowledge*, *Credal Probability and Chance*, Cambridge, Mass,: MIT Press reissued with slight changes in paperback in 1984.

Levi, I. (1980b), "Induction as Self Correcting According to Peirce," *Science, Belief and Behaviour*: *Essays in honour of R.B.Braithwaite* ed. by D.H.Mellor, Cambridge: Cambridge University Press, 127-140.

Levi, I. (1980c), "Potential Surprise: Its Role in Inference and Decision Making," *Applications of Inductive Logic*, ed. by L.J.Cohen and M.Hesse, Oxford: Clarendon, 1-27.

Levi, I. (1982), "Dissonance and Consonance according to Shackle and Shafer," *PSA 1978* 2, 466-75.

Levi, I. (1983), "Consonance, Dissonance and Evidentiary Mechanisms," in *Evidentiary Value, Philosophical, Judicial and Psychological Aspects of a Theory* ed. by P.Gärdenfors, B.Hansson and N.-E.Sahlin, Lund: Gleerups, 27-43.

Levi, I. (1986), *Hard Choices*, Cambridge: Cambridge University Press.

Levi, I. (1991), *The Fixation of Belief and Its Undoing: Changing Beliefs Through Inquiry*, Cambridge: Cambridge University Press.

Levi, I. (1996), *For the Sake of the Argument*: *Ramsey Test Conditionals,, Inductive Inference and Nonmonotonic Reasoning*, Cambridge: Cambridge University Press.

Levi, I. (1997), *The Covenant of Reason*: *Rationality and the Commitments of Thought*, Cambridge: Cambridge University Press.

Levi, I. (1999), "Value Commitments, Value Conflict, and the Separability of Belief and Value," *Philosophy of Science* 66, 509-33.

Levi, I. (2001), "Inductive Expansion and Nonmonotonic Reasoning," *Frontiers in Belief Revision* edited by Mary-Anne Williams and Hans Rott, Dordrecht: Kluwer, 7-56.

Levi, I. (2002), "Maximizing and Satisficing Measures of Evidential Support," *Reading NaturalPhilosophy : Essays in the History and Philosophy of Science and Mathematics*, edited by David Malament, Chicago: Open Court, 315-333.

Levi, I. (2004a), *Mild Contraction*: *Evaluating loss of information due to loss of belief*. Oxford: Oxford University Press.

Levi, I. (2004b), "Beware of Syllogism: Statistical Reasoning and Conjecturing According to Peirce," *The Cambridge Companion to Peirce* ed. by C. Misak, Cambridge: Cambridge University Press.

Makinson, D. and Gärdenfors, P. (1991), "Relations between the Logic of Theory Change and Nonmonotonic Logic," in *The Logic of Theory Change* ed. by A.Fuhrmann and M.Morreau, Berlin: Springer Verlag, 185-205.

Neyman, J. and Pearson, E.S. (1932-3), "The Testing of Statistical Hypotheses in Relations to Probabilities a priori", *Proceedings of the Cambridge Philosophical Society* 29, 492-510.

Popper,K.R. (1959), *The Logic of Scientific Discovery*, London: Hutchinson.

Reiter, R. (1980), "A Logic for Default Reasoning," *Artificial Intelligence* 13, 81-132.

Shackle, G.L.S. (1949), *Expectations in Economics*, Cambridge: Cambridge University Press. 2nd ed. 1952.

Shackle, G.L.S. (1961), *Decision, Order and Time*, Cambridge: Cambridge University Press. 2nd ed. 1969.

Shafer, G. (1976), *A Mathematical Theory of Evidence*, Princeton: Princeton University Press.

Spohn, W (1990), "A General non Probabilistic Theory of Inductive Reasoning," *Uncertainty in Artificial Intelligence,*" 4, Schechter *et. al*. North Holland.

Wald, A. (1950), *Statistical Decision Functions*, New York: Wiley.

Zadeh, L. (1978), "Fuzzy Sets as a Basis for a Theory of Possibility," *FuzzySets and Systems* 1: 3-28.