

The value of truth and the value of information: On Isaac Levi's epistemology

Hans Rott

1. A sermon on epistemology

"I preach a sermon on epistemology," says Levi (2004, p. 77). Isaac Levi is an immensely interesting and powerful philosopher. In my opinion, he has erected a most impressive epistemological edifice in the past 40 years. Inspection of current collections on epistemology, however, reveals that there is little interaction between Levi's work and almost any version of mainstream justification-based or reliability-based epistemology.¹ Why is this?

First, Levi resolutely distances himself from many main tenets of mainstream epistemology. He is very decidedly against any sort of "pedigree epistemology" (2004, pp. 11, 232) which holds that whether some belief counts as knowledge is dependent on its origin or its foundation, on where it comes from. He rejects the idea that convergence on the truth is the ultimate aim of inquiry (1980, pp. 70–72). And he is averse to "Parmenidean epistemology" (2004, pp. 10–12) according to which only logical, mathematical or conceptual necessities should be admitted as *full* beliefs, while everything else should get assigned a degree short of "the Permanent One" (2004, p. 10).

For Levi, taking his position just means being true to the pragmatist stance. At any given point of time, a believer doesn't have to justify his or her currently held beliefs, since there is nothing other than the current set of beliefs upon which the evaluation of the believer's mental state could be based (in Levi's terminology: there is no other "standard of serious possibility"). But that does not mean that believers are exempted from any duty of justification. Believers have to justify their *changes of beliefs*. Justification in the pragmatist's sense means justification in terms of a believer's goals and values. Within the pragmatist camp, the *differentia specifica* of Levi's specific position is that justification should be decision-theoretic. As we shall see, truth and information figure most prominently in his own decision-theoretic account.

Secondly, even the term "knowledge" has come to play a minor role in Levi's works. In *The Enterprise of Knowledge*, he has only a short discussion of the justified-true-belief analysis (1980, pp. 1–3, 28–30), and in *The Fixation of Belief and Its Undoing* (1991, p. 45), he rather casually says that knowledge, as he uses the term, is "error-free, full belief", and he dismisses questions of justification as "irrelevant". I have not been able to find any definition or analysis of knowledge in his most recent books, *For the Sake of Argument* (1996) and *Mild Contraction* (2004). To put it provocatively, it begins to seem as if Levi is preaching a sermon on epistemology without knowledge.

A third reason for the unfortunate neglect of Levi's work by epistemologists may lie in the fact that much of Levi's presentation is rather technical. It can be hard to follow his philosophy without attending to a lot of logical and probabilistic niceties. Levi finds large

¹ In the recent *Oxford Handbook of Epistemology*, for instance, there is only one reference to Levi (Kaplan 2002, footnote 25).

audiences among researchers in philosophical logic, AI and knowledge representation, but he has not been equally successful in getting his message across to epistemologists who are less used to dealing with technicalities.

Before venturing a critical evaluation of Levi's work, I shall recapitulate as perspicuously as I can some central elements, paying special attention to their most recently renovated forms. It is mainly the current state of Levi's epistemology as expounded in his new book *Mild Contraction* on which I shall focus.

2. Inquiry

Levi's epistemology focuses on the notion of inquiry. *Inquiry*, according to Levi, consists in the answering of questions and the solution of problems (1996, p. 161; 2004, pp. 39, 56). It has as its proximate cognitive aim the change of bodies of belief or theories with a view to gaining new *error-free information* (1996, p. 165; 2004, pp. 70–71, 76–80). Levi's models can be seen as attempts to implement in a precise and formally explicit manner James's famous double maxim: "Seek valuable information! Shun error!"² As we shall see, one has to be clear about the fact that the two parts of James's maxim give expression to two very different desiderata.

So inquiry can be understood as the process of rational belief change. By *beliefs*, Levi does not refer to conscious occurrences in the inquirer's mind, but to his or her commitments. This reading of 'belief' motivates Levi's imposing the constraint that the set of the inquirer's beliefs be logically closed.

Levi formalizes the process of inquiry. It is not possible to understand fully his intent without going through his formal setting for the analysis of the norms of rational or scientific deliberation. In this section we recapitulate the formalization in its barest backbones.

\underline{K} denotes the class of all potential corpora of beliefs (\underline{K} is sometimes called "conceptual framework", 1996, p. 18; 2004, p. 41). For Levi, a potential belief state \mathbf{K} , abstractly conceived, can be represented by a potential corpus K in \underline{K} . While a belief state is a non-linguistic entity, a corpus is a set of sentences in a regimented language L . Since K is supposed to represent the "full beliefs" of an inquirer, i.e., his or her "commitments" or "standards of serious possibility", K is supposed to be deductively closed (1996, p. 13; 2004, pp. 14, 41–42). Usually, but not always, K is assumed to be consistent.

In order to avoid undue technicalities, I shall in this paper suppose that the set of all L -sentences is partitioned by the relation of logical equivalence into finitely many equivalence classes. As a consequence, each potential corpus K can alternatively be thought of as being represented by a single sentence, viz. the conjunction of representatives of all equivalence classes included in K . I shall take the liberty of jumping between both types of representations of belief states as convenience suggests.

The *ultimate partition* U is a partitioning of the space of all possibilities, where possibilities are identified with maximally consistent sets of L -sentences. Alternatively, U can be thought of as a set of pairwise incompatible and jointly exhaustive L -sentences. Given a corpus K , we mean by the elements of the partition *within* K the classes in U whose maximally consistent sets of L -sentences include K (the first way of thinking of possibilities), or alternatively, the sentences in U that imply K (the alternative way of thinking). The elements *outside* K are just those elements of K that are not within K .

² Levi (1991, p. 81; 2004, pp. 76–80). In James's (1897, p. 24) own words: "*We must know the truth; and we must avoid error* – these are our first and great commandments as would-be knowers." Levi rightly points out that 'truth' here should be replaced by 'information' or 'valuable information', and that James's double maxim was anticipated by Peirce in his *Harvard Lecture X* held in 1865.

For the purposes of this paper, we assume that the content of any potential belief state is expressible as the intersection of the elements of certain sets in U or, alternatively, as the disjunction of certain sentences in U . We also assume that all potential answers to questions the agent may be interested in, all his or her demands for information are expressible as combinations – intersections or, alternatively, disjunctions – of elements of U (1980, p. 45; 1996, p. 163; 2004, p. 49). The ultimate partition determines, among other things, the granularity of currently and potentially held beliefs. Inquirers may alter their ultimate partitions for some reason or other, but such cognitive operations are not main subject of Levi's studies.³

Levi imposes a *commensuration requirement* (1991, p. 65; 2004 p. 16), according to which any legitimate belief change is decomposable into a series of belief expansions and belief contractions. So what inquirers need are principled methods for legitimately expanding and legitimately contracting their corpora.

An *expansion* of the corpus K by a hypothesis h is denoted by $K+h$, a *contraction* of K with respect to a hypothesis h is denoted $K\div h$. According to Levi, there are no other legitimate ways of modifying corpora besides expansions and contractions. A *belief revision* by a hypothesis h , for instance, gets reduced to a compound consisting of a contraction followed by an expansion, $K*h = (K\div\sim h)+h$. This equation has become widely known as *the Levi identity*.

3. The aims of inquiry: Truth and information as cognitive values

For Levi, rational belief change has to be justified in decision-theoretic terms. In order to see how this works, we have to supply some more formal means.

If K is the current corpus, then the *credal probability* relative to K is Q_K . All full beliefs, that is, all elements of K , get credal probability 1. The expansion $K+h$ gets credal probability $Q_K(K+h) = Q_K(h)$ relative to K (1996, p. 167). Credal probabilities are also called "*expectation-determining probabilities*", "*belief probabilities*" or "*confirmational commitment*" (1996, pp. 167–169; 2004, pp. 78, 84, 89)

The measure of informational value is similarly functionally dependent on what happens to be the inquirer's current corpus. The basic structure is encoded in an informational-value determining probability function M_K associated with the belief state K . The *informational value* of a sentence h is determined to be $1 - M_K(h)$, where M_K is the *informational-value determining probability function*: $Cont_K(h) = 1 - M_K(h) = M_K(\sim h)$.⁴ All full beliefs get informational-value determining probability 1, i.e. informational value 0.

Levi (2004, pp. 98–101) mentions other content measures such as $1/M_K(h)$ and $-\log M_K(h)$. The content measure chosen by Levi has a number of advantages, $Cont_K$ values for instance have a finite upper and lower bound. More importantly, they satisfy the condition of *Constant Marginal Returns in Informational Value of Rejection* which says that the informational value gained by rejecting an element x of the ultimate partition (thereby expanding the inquirer's corpus) is independent of the set of further elements rejected along

³ We may neglect for our purposes Levi's *basic* or *minimal corpus* LK with the accompanying partition U_{LK} (2004, pp. 49, 57). Levi introduces LK to represent the minimal presuppositions of the inquiry, and U_{LK} to represent the maximally specific relevant answers given these presuppositions (regardless of how K changes). – The ultimate partition is similar in function to Shafer's (1976) *frame of discernment*. Cf. Fioretti (2001).

⁴ Levi (1967, pp. 69 and 107; 1991, p. 84; 1996, pp. 24 and 169; 2004, p. 84). This function has been suggested as a measure of information or content at least since Carnap and Bar-Hillel (1952, p. 237). $Cont(h)$ is independent of the factual truth of h . Measures of informational value are discussed in detail in Levi (1969).

with x .⁵ But the most important advantage is, as we shall see, that this content measure $Cont_K$ can elegantly be weighed against the inquirer's credal probability Q_K .

In Levi's early writings, M_K is assumed to be the traditional "logical probability", so if there are four possibilities between which the agent is to choose, the M_K -value of the hypothesis that a particular one is true would be $\frac{1}{4}$ and its $Cont_K$ -value would be $\frac{3}{4}$. Levi insists that informational-value determining probabilities are in general *not* identical with credal probabilities (1996, p. 22; 2004, p. 88).⁶

There remains an unresolved problem of interpretation: Where do information-determining probabilities come from? Clearly, informational value determining probabilities do not represent subjective beliefs, nor frequencies, nor propensities. I suspect that the only way to understand what they mean is via their intended interpretation as information-value determining. But then it seems a legitimate question to ask whether it would not be clearer to use *informational value* as the primitive notion, without a recourse to uninterpreted probabilities.

We may conveniently summarize what is relevant to inquiry according to Levi, the set of "contextual parameters" (2004, p. 89), in a quadruple of the form $\langle K, U, Q_K, M_K \rangle$, or simply. Levi never introduces such 4-tuples formally as devices for the analysis. For ease of reference, however, let us call them *frameworks for inquiry*. Most of such a framework for inquiry seems subjective, relative to the inquirer in question, but, as I said, the status of informational-value determining probabilities is unclear (to me).

Let us now see how Levi applies frameworks for inquiry in processes of belief change.

4. Aggregating two values in deliberate inductive belief expansions

Levi's account of deliberate inductive expansion is a beautiful piece of theorizing, but it is not easy to understand. Let us develop it carefully.

One of the central problems addressed by Levi is the problem of how to extend a given corpus through inductive reasoning. Such an extension aims at an improving one's doxastic state, without any "external disturbances" through new evidence. The question is which new hypotheses to accept. According to Levi, inquirers should not try to maximize or satisfice probabilities, nor should they look for some tailor-made inductive or non-monotonic logic. They should rather use decision theory. In Levi's view, the traditional notion of a *justified expansion* reduces to (or is to be replaced by) the notion of a *utility-maximizing expansion*. Inductive reasoning is an exercise in decision theory.

Let the framework $\langle K, U, Q_K, M_K \rangle$ be given. Each hypothesis h suitable for strengthening K can be represented as a disjunction $x_1 \vee \dots \vee x_n$ of elements of the ultimate partition U within K .

Levi's principal pragmatist idea is to assess hypotheses according to their cognitive values. The cognitive value of a hypothesis h is a weighted average of *the value of its truth* and *the value of its informational content*.

$$\text{Value}(h) = \alpha \cdot \text{T-value}(h) + (1-\alpha) \cdot \text{I-value}(h)$$

⁵ Levi (2004, pp. 83, 100–101). It is evident from Levi's wording that he is aware that the principle does not transfer well to his theory of belief *contractions*: "The following assumption (proposed in 1967 in Levi, 1984, ch. 5) works very well in the context of expansion" (2004 p. 83, my italics). Also compare Levi's condition of *Extended Weak Positive Monotonicity* (2004, p. 126).

⁶ Levi also discusses the idea that the two probability functions Q_K and M_K associated with the inquirer's current belief state K are obtained by conditionalization from some "master probability functions" Q and M associated with the basic belief state LK (2004, p. 153; cf. 1996, pp. 167–169, 267).

The parameter α expresses the relative importance that is attached to the truth of a hypothesis, and $1-\alpha$ expresses the relative importance of its informational value. α ranges between 0 and 1, and its particular value is chosen by the inquirer. If $\alpha=0$, then truth does not matter at all, only the amount of information conveyed by h counts (so tautologies are regarded to be of minimal value). If $\alpha=1$, then information does not matter at all, only the truth of h counts (so tautologies are regarded to be of maximal value). The more cautious the inquirer, the larger is α ; the bolder the inquirer, the smaller is α . Proceeding from the "plausible assumption that no error is to be preferred to any case of avoiding error" (1967, p. 107; 1996, p. 171; 2004, p. 86), Levi holds that α should not be less than 0.5.

Now Levi makes a number of important decisions. Both T-value and I-value are "normalized" so that they take values between 0 and 1. The values of truth and falsity are supposed to be *identical* for all possible hypotheses: if h is true, then $T\text{-value}(h) = 1$, if h is false, then $T\text{-value}(h) = 0$.⁷ As already mentioned, the value of information is measured by $Cont_K(h) = M_K(\sim h)$. We will soon see that these particular ways of defining the value of truth and the value of information have important consequences for the formal structure of Levi's theory.

The value of accepting h depends on the chosen value of α (something subjectively chosen) and the truth value of h (something beyond the agent's control). Making explicit this double dependence, we can write down the value of accepting a hypothesis h as

$$\begin{aligned} V_\alpha(h, \text{true}) &= \alpha \cdot T\text{-value}(h, \text{true}) + (1-\alpha) \cdot I\text{-value}(h) = \\ &= \alpha \cdot 1 + (1-\alpha) \cdot (1-M_K(h)) = \alpha + (1-\alpha) \cdot (1-M_K(h)) \end{aligned}$$

$$\begin{aligned} V_\alpha(h, \text{false}) &= \alpha \cdot T\text{-value}(h, \text{false}) + (1-\alpha) \cdot I\text{-value}(h) = \\ &= \alpha \cdot 0 + (1-\alpha) \cdot (1-M_K(h)) = (1-\alpha) \cdot (1-M_K(h)) \end{aligned}$$

(This differs a little from Levi's own notation.) Now the inquirer often does not know whether h is true or false. In this case, she may employ her credal probabilities Q_K about the hypothesis' truth value. Decision theory says that the inquirer ought to maximize her expected utility:

$$\begin{aligned} EV_\alpha(h) &= Q_K(h) \cdot V_\alpha(h, \text{true}) + Q_K(\sim h) \cdot V_\alpha(h, \text{false}) = \\ &= Q_K(h) \cdot (\alpha + (1-\alpha) \cdot (1-M_K(h))) + (1-Q_K(h)) \cdot ((1-\alpha) \cdot (1-M_K(h))) = \\ &= \alpha \cdot Q_K(h) + (1-\alpha) \cdot (1-M_K(h)) = \\ &= (1-\alpha) + \alpha \cdot Q_K(h) - (1-\alpha) M_K(h) = \\ &= (1-\alpha) + \alpha \cdot \sum_{\{x \text{ in } U: x|h\}} Q_K(x) - (1-\alpha) \cdot \sum_{\{x \text{ in } U: x|h\}} M_K(x) = \\ &= (1-\alpha) + \sum_{\{x \text{ in } U: x|h\}} (\alpha \cdot Q_K(x) - (1-\alpha) M_K(x)) \end{aligned}$$

⁷ Levi (1967, p. 106; 2004, p. 80). This of course is a very strong idealization. It neglects all matters of relevance. Even if they are equally likely, the truth that the Dallas Mavericks won their last match matters much less to me than the truth that my daughter got home safely last night. Intuitively, I'd say that these truths differ a lot *in value* for me.

In the transition to the last line, we can see how important it is that Levi chooses the particular definitions of the values of truth and information. Different concepts of T-value or I-value would not have allowed him to break down the value of accepting h into a sum of values of accepting those elements of the ultimate partition that entail both K and h .

What the inquiring subject should do, in Levi's pragmatist picture, is to accept a hypothesis h that maximizes the value of $EV_\alpha(h)$. But such a hypothesis is easy to find. We just have to collect all those x in U within K for which the last term in the formula above is non-negative, that is, for which

$$Q_K(x)/M_K(x) \geq (1-\alpha)/\alpha$$

If this inequality is satisfied, then x is a disjunct of the hypothesis to be accepted. Decision theory is actually silent about whether one should take an x for which the value $Q_K(x)/M_K(x)$ is exactly zero. Levi recommends not to reject zero-valued elements. His *Rule for Ties in Expansions* (1967, p. 84; 1991, p. 93; 1996, 172; 2004, p. 87) instructs the inquirer to take the *weakest* optimal expansion if there is one. And there always is one, namely the one which we obtain by disjunctively conjoining the zero-valued elements to the positively valued elements. We have now arrived at Levi's *Inductive Expansion Principle* (see 1967, p. 86; 1996, p. 172; 2004, pp. 87–88).

Define the *index of boldness* $q = (1-\alpha)/\alpha$. Note that the value of q increases as the inquirer's degree of caution α decreases. Since α is supposed to range from 1/2 to 1, q ranges from 0 to 1. For every element x of the ultimate partition, the chances that x is rejected as a result (as a "conclusion") of an inductive inference increase as the index of boldness q is increased. If q is raised, this means that fewer elements of U within K will be left uneliminated, hence the selected hypothesis h will contain *more* information.⁸

Now let us change perspectives. Instead of taking a certain number α (or equivalently, a certain number q) as given and ask which hypothesis should be accepted, take a certain member of the ultimate partition as given, and vary α (or q). Given the inquirer's credal probability Q_K and information-determining probability M_K , there is for each x in the ultimate partition within K a unique number

$$q(x) = \min \{Q_K(x)/M_K(x), 1\}$$

that is just low enough to make x a disjunct of the hypothesis to be accepted. If the index of boldness q chosen by the inquirer is higher than $q(x)$, then the possibility that x is true is ruled out (that is, $\sim x$ is accepted). If q is lower than or identical with $q(x)$, then x remains a serious possibility. $q(x)$ is "the maximum value of q at which x fails to be rejected" (Levi 1967, p. 137; 1996, p. 185; 2004, pp. 89–90).⁹

That an element x of the ultimate partition is *rejected* means that $\sim x$ is one of the agent's full beliefs (assuming that the elements of an ultimate partition are expressible by sentences of the agent's language). The more elements are rejected, the stronger are the inquirer's beliefs. So if q is higher than $q(x)$, then $\sim x$ gets accepted; if q is less or equal than $q(x)$, then $\sim x$ does not get accepted. The degree of boldness q that the inquirer chooses has to be higher than $q(x)$ in order to make $\sim x$ acceptable. The higher $q(x)$, the more boldness or "mental effort" it takes to find $\sim x$ acceptable. If $q(x)$ is less than 1, it is possible, by a sufficient amount of boldness to accept $\sim x$. But if $q(x)$ equals 1, this is never possible, not

⁸ One could also think of inserting the index q of boldness as a fifth element into a framework of inquiry. But in contrast to the elements we have identified above, q seems to be a matter of the inquirer's free choice, so I prefer to leave it out.

⁹ If x is rejected for all (no) values of x , then Levi sets $q(x) = 0$ (resp., $q(x) = 1$).

even by the utmost exertion of one's boldness. In a way, one could say that $q(x)$ is the degree of nonbelief of $\sim x$.

This idea can be generalized to propositions that are not elements of U within K . Hypotheses h that are not maximally specific are identified with disjunctions $x_1 \vee \dots \vee x_n$ of those elements x_1, \dots, x_n of the ultimate partition that imply h . A hypothesis h equivalent with $x_1 \vee \dots \vee x_n$ is *rejected* if all the x_i 's are rejected (1996, p. 185).¹⁰ Thus clearly the maximum level of q at which h fails to be rejected, $q(h)$, equals the maximum of the q -values of the disjuncts x_i :

$$\text{if } h \text{ is equivalent to } x_1 \vee \dots \vee x_n \text{ within } K, \text{ then } q(h) = \max(q(x_i))$$

In my alternative way of speaking, $\sim h$ gets accepted, if the q -value chosen by the inquirer is higher than $q(h)$; otherwise $\sim h$ will not be accepted. Again, it is possible to find $\sim h$ acceptable at some level of boldness if and only if $q(h)$ is less than 1. On the other hand, $q(h)$ is 0 if and only if $\sim h$ is already accepted in K .

The *Shackle degree of belief* of h (relative to a framework of inquiry $\langle K, U, Q_K, M_K \rangle$) is defined to be (Levi 1996, pp. 185–186; 2004, p. 90¹¹)

$$b(h) = 1 - q(\sim h)$$

The more boldness it requires to render h acceptable, the *lower* its degree of belief is. This is why the q -values are subtracted from 1. The b -value of a proposition h is 1 if and only if it is in the inquirer's corpus K . It is 0 if h cannot be found acceptable for any degree of boldness (this is true either for h or $\sim h$, for every hypothesis h).

Levi summarizes his achievement as follows: "we have derived a measure of degree of disbelief and degree of belief exhibiting the formal properties of Shackle measures of degrees of potential surprise or disbelief and of degrees of belief. The derivation involves the use of a family of deductively cogent, caution-dependent, and partition-sensitive criteria for inductive expansion." (1996, p. 186)

It is this construction that most intimately ties Levi's decision-theoretic, pragmatist philosophy together with the work done on the logics of belief revision in the AGM paradigm and related models. An important role in Levi's account of both deliberate expansion and contraction¹² is played by Shackle-like like measures, so-called after the British economist G.L.S. Shackle (1949). Shackle measures assign values to propositions, and Levi presents them in a normalized form with a range between 0 and 1. The characteristic property (1967, p. 133; 1996, pp. 181, 264; 2004, p. 90) of a Shackle measure b is that for all propositions g and h

$$b(g \wedge h) = \min\{b(g), b(h)\}$$

According to Levi (1991, p. 182; 1996, pp. 180, 258; 2004, pp. 90–95), variants of Shackle measures have been rediscovered or reinvented in the last three decades by many researchers,

¹⁰ I find this terminology slightly confusing. On the one hand, h is here thought to be a candidate belief (and not an actual belief), so in a sense, it does not have to be rejected in the first place. On the other hand, the *proposal to expand* the current corpus K by a hypothesis $h = x_1 \vee \dots \vee x_n$ would probably be rejected as soon as *one* (and not *all*) of the disjuncts x_i turned out to be rejectable. So instead of saying that h gets rejected if all x_i 's are rejected, I would find it clearer to say in this case that $\sim h$ is accepted.

¹¹ Actually, Levi (2004) inserts an intermediate concept, viz. the *Shackle degree of potential surprise* or *degree of disbelief* $d(h) = 1 - q(h)$, and then defines $b(h) = d(\sim h)$.

¹² Levi insists that the *kind* of use Shackle measures are put to is very different in deliberate expansion from that in contraction (1996, pp. 267–268; 2002, pp. 136–137; 2004, pp. 40–41). I address contractions in Section 6, but unfortunately, I shall not be able to deal with this question in the present paper.

including L. Jonathan Cohen, Didier Dubois and Henri Prade, Peter Gärdenfors and Wolfgang Spohn.

In my opinion, Levi's decision-reconstruction of belief functions (in Shackle's sense) is a remarkable achievement. Starting with his early masterpiece *Gambling with Truth*, he has developed an ingenious model of combining the demands for truth and information. He has shown how to take the decision-theoretic prescription of utility maximization as the primary principle, how to apply it to the important epistemological problem of inductive expansion, and how to reach an unequivocal decision by the Rule for Ties which recommends to adopt the logically weakest solution of the maximization problem. Moreover, Levi has made clear how the problem of maximizing epistemic utility (by a suitable choice of a hypothesis h) can at the same time be viewed as a satisficing problem with respect to a measure of belief: The agent can accept every possibility above a certain "level of aspiration" or threshold value q .

In a more critical vein, it remains to say that the particular definition of the I-values is not very well motivated. Several questions remain. Where does M_K come from? Why exactly is the content measured by $1-M_K$ rather than, say, by $1/M_K$ or $-\log M_K$? A similar complaint may be raised against the definition of the T-value. What is the justification for assigning truth the constant value 1 and assigning falsity the constant value 0, irrespective of the relevance of the proposition in question?

5. The single value for belief contractions: Informational value

Like belief expansions, belief contractions are considered to be decision-theoretic problems by Levi. Prima facie, the problems presented by contractions, that is, by retractions of belief, are completely dual to the problems presented by expansions. The positive aspect of an expansion is that one gains information, the negative side is that in gaining information one may import error.¹³ We have seen how Levi conceives of the trade-offs between these two factors. The negative aspect of a contraction is that one loses information, the positive side, it would seem, is that in foregoing information one may eliminate error. And it would seem that a similar method for resolving the trade-off problem must be found. Not so, says Levi. The inquirer is committed to treating the full beliefs collected in his or her corpus as infallible, that is, as standards of serious possibility. From the inquirer's own point of view, his or her beliefs *cannot be* erroneous. Thus, since there is no error that could be eliminated, nothing positive can be gained from a contraction. Only one factor, namely loss of information, is to be taken into account. We need informational-value determining probabilities, but we do not need credal probabilities any more. The problem of contraction is unequivocal, there is just one thing to take care of: to minimize the loss of information.¹⁴

But if, from the inquirer's point of view, there is nothing to be gained from a contraction, why should any rational being contract his or her beliefs in the first place?¹⁵ Levi's answer to this question is three-fold. First, inquirers need to contract their beliefs if they are caught in an inconsistent corpus. Second, they may wish to give a new hypothesis a hearing and for this reason withdraw their full belief in its negation. Third, they may want to engage in reasoning for the sake of argument.

¹³ It does not make much sense to say that one may import truth. It looks as if all one might wish to convey by this phrase is already covered by saying that the inquirer gains information.

¹⁴ It might seem, then, that in Levi's picture (i) expansions are much more interesting and sophisticated change operations than contractions, and that (ii) this is a reversal of the AGM account where expansions are considered to be trivial as compared with contractions. Both parts are wrong. (ii) is superficial because it neglects the fact that AGM simply do not have an operation of inductive expansion (without any input). (i) passes over the fact that finding unique solutions with the help of a "rule of ties" presents a much greater challenge for contractions than it does for expansions. We get back to this point soon.

¹⁵ From a third person's point of view, the answer is trivial: Because inquirers are sometimes mistaken.

The formats of Levi's expansion and contraction operations are different. In deliberate or inductive expansion, the task is to find the right hypothesis to add to the current corpus, subject to the internal constraint that the expected cognitive value be maximized. Inductive expansion is an autonomous process, a purely internal affair as it were. There is no input to the inquirer's belief state from outside.

In contraction, on the other hand, the task the inquirer faces is to discard a particular belief h – and *which* belief to discard is externally determined. As h is given, the task is not to find the right hypothesis to subtract. The options that are being judged with respect to their maximizing the inquirer's cognitive values are target corpora that do not contain h .

So how can we characterize the set of options in a contraction problem? The first answer is that every weakening of K is a potential contraction. More precisely: A *potential contraction of K* relative to the ultimate partition U (of the space of all possibilities) is the disjunction of K with the disjunction of some elements of U outside K .¹⁶ A *potential contraction of K removing h* (relative to U) is obtained by forming the disjunction of K with a subset of U outside K that contains at least one element that does not imply h .¹⁷ These are very general concepts of contraction. I agree with Levi that we should not from the outset restrict the possible options for the process of contracting belief sets.

Levi argues forcefully that hardly any previous author in epistemology has given anything like a decision-theoretic rationale for changes of belief. The following rule for belief contraction is a direct application of this core idea of Levi's:

Rule 1. The Decision-Theoretic Rule

The corpus after a contraction must be optimal, i.e., it must minimize the loss of informational value among all corpora expelling the hypothesis h .

Levi urges that information (the ruling out of some logical possibilities) must not be identified with informational value. The decision-theoretic rule does not stigmatize every loss of information as irrational. There is no objection to the losing worthless information.

The problem is that this rule is not definite enough. There may be many ways to achieve minimum loss of informational value. What is the inquirer supposed to do in the face of such an ambiguity? If two or more options are tied for optimality with respect to the primary decision-theoretic value commitments, Levi recommends to invoke a secondary standard of evaluation in order to break the ties:

Rule 2. The Rule for Ties in Contractions

Given a set of optimal contraction strategies, one should always choose the weakest of them if it exists. (2004, p. 119).

This advice is analogous to the Rule for Ties in Expansions. In contrast to the case of expansions, however, the precondition that there is a unique weakest solution is not readily satisfied in the Rule for Ties for Contractions. There is the danger that in many cases the rule simply cannot be applied. Levi, however, feels strongly that there *should always* be a unique weakest optimal contraction strategy. Combine this desideratum with the fact that the most obvious (and perhaps the only principled) solution to the problem of multiple optima is to

¹⁶ Levi (2004, p. 59). If we stuck to the representation of K and the elements of U as *sets of sentences* rather than sentences, then we would have to mention intersections rather than disjunctions.

¹⁷ If such an element exists; otherwise the only potential contraction of K removing h is set to be K itself (Levi 2004, S. 59). – I have changed Levi's definition slightly. Where Levi requires that $\sim h$ be implied, I only require that h be not implied. My formulation leaves room for hypotheses that are not (yet) expressible in the inquirer's question-and-answer system. But this is a point that we can neglect in the present paper.

settle for the combination – intersection or disjunction – of all optimal contractions. Then we understand the rationale for Levi's installing a third central condition:

Rule 3. The Intersection Equality

If members of a set S of contractions from K are equal in informational value, their intersection is equal in informational value to the informational value of any element of S .¹⁸

Unfortunately, the straightforward way of measuring information does not prove to be suitable for Levi's purposes. Assume that M is an informational-value determining probability function over the class of *all* possibilities (not just those within K).¹⁹ Then the loss of informational value by a shift from K to a potential contraction K' can be determined by $Cont(K) - Cont(K') = (1 - M(K)) - (1 - M(K')) = M(K') - M(K)$ (2004, pp. 84–85, 109). But since the probability of an intersection (or disjunction) of possibilities is in general *more* probable than the possibilities themselves, *more* informational value will be lost by taking intersections (disjunctions) than by taking single possibilities, in violation of the Intersection Equality.

As far as I know, no author except Levi himself has cared to take seriously the ideal of decision-theoretic optimality while at the same time respecting the Rule for Ties.²⁰ We now turn to Levi's techniques of the damping of informational value that are devised precisely to overcome the tension between the decision-theoretic rule and a non-trivial application of the Rule for Ties.

6. Damping informational value

Having established that purely probability-based informational value ought not to be minimized, Levi suggested two different ways of "damping" informational value in such a way that his desiderata for belief contractions can all be simultaneously satisfied (1991, pp. 127–131; 1996, pp. 262–267; 2004, pp. 125–147). Since there are no analogous problems of multiple solutions for inductive expansions, damping is not necessary there.

6.1 Saturatable sets and damping version 1

Levi's first attempt at solving the problem of contraction centers around the notion of a saturatable set. A *saturatable contraction* of a corpus K removing a hypothesis h is the disjunction (intersection) of K with some (possibly empty) set of elements of U outside K that entail h and a single element of U that entails $\sim h$ (cf. 1991, p. 121; 1996, pp. 20–23; 2004, p. 60). Levi points out that every potential contraction of K removing h can be represented as the disjunction (intersection) of a subset of saturatable contractions of K removing h . This is Levi's *Potential Contraction Condition* (2004, p. 61). This is certainly correct, but it does not suffice to accord saturatable contractions removing h an epistemologically distinguished role, since many other types of sets (e.g. maxichoice contractions, saturatable contractions removing $\sim h$) could be mentioned in the Potential Contraction Condition.

¹⁸ This is the strong version of Levi's Intersection Equality (2004, p. 125). The weak version restricts the claim to saturatable contractions removing h from K . As I don't see any reason to bestow a special status upon saturatable contractions, I think it is the strong version that captures the essence of the Intersection Equality.

¹⁹ The critical questions about the inquirer's probability function M_K raised at the end of Section 4 transfer and are indeed aggravated for the impersonal probability function M .

²⁰ But compare the careful discussion in Sandqvist (2000).

Damped informational value version 1 is then defined in two steps (2004, p. 131).²¹ The loss of damped informational value incurred by shifting from the corpus K to a saturatable contraction K' removing h is equal to the loss of undamped informational value, i.e., $Cont(K) - Cont(K') = M(K') - M(K)$. The loss of informational value incurred by shifting from the corpus K to a disjunction (intersection) of some set of saturatable contractions removing h is defined to be the largest loss incurred by any element in that set, i.e., $Cont(K) - Cont(\bigvee K_i) = \max_i M(K_i) - M(K)$.

Levi chose this definition in order to make sure that the informational value of the disjunction (intersection) of two contractions removing h is equal to the minimum of their informational values. Now what Levi recommends (or rather: recommended) comes down to saying that *the right contraction* removing h from K is the disjunction (intersection) of all saturatable contractions K' removing h that minimize the loss of damped informational value, that is, for which $M(K')$ is minimal (1991, p. 130; 1996, p. 263; 2004, pp. 132–133). It follows from the definition of damping version 1 that any disjunction (intersection) of contractions minimizing loss of informational value²² minimizes loss of informational value itself. So the Rule for Ties for contractions does not run counter the decision-theoretic rule.

Unfortunately, the definition as it stands is not well-defined. We show this by giving an example. Consider the eight-cell partition U that is generated by the truth-value combinations of the three atomic sentences p , q and r . Suppose that M assigns probability 0.2 to the two cells with p and q being false, and probability 0.1 to the six remaining cells. Let K be $p \wedge q \wedge r$ (we identify a theory with its generating sentence), and suppose we want to contract K by p . Consider the potential contraction $K' = p \vee q$. Clearly K' is not saturatable as a contraction that removes p . But K' can be represented as the disjunction $K_1 \vee K_2$ of two saturatable contractions K_1 and K_2 removing p , where $K_1 = p \vee (q \wedge r)$ and $K_2 = p \vee (q \wedge \sim r)$. Since the M -value of both K_1 and K_2 is 0.5, their informational content $Cont$ is 0.5. So by damping version 1, the content of K' is 0.5, too. However, K' can also be represented as the disjunction $K_3 \vee K_4$ of the saturatable contractions K_3 and K_4 removing p , where $K_3 = (p \wedge \sim r) \vee (q \wedge r)$ and $K_4 = (p \wedge r) \vee (q \wedge \sim r)$. The M -value of both K_3 and K_4 is 0.4, so their informational content $Cont$ is 0.6. So by damping version 1, the content of K' must be 0.6, too. But this contradicts the value we have calculated before.

Usually, this ill-definedness does not do any harm since one has to look at maxichoice (rather than just saturatable) contractions anyway. Levi's first damping construction recommends as the right contraction of K removing p the disjunction (intersection) of the saturatable contractions removing p that minimize loss of damped informational value version 1. In our example, this is $K \dot{\div} p = (q \wedge r) \vee (q \wedge (p \leftrightarrow r)) = q \wedge (\sim p \vee r)$, the damped informational value version 1 of which is $\min(Cont(q \wedge r), Cont(q \wedge (p \leftrightarrow r))) = 0.8$. Saturatable sets that are not maxichoice come into play only if there are cells of zero M -value outside K that entail the hypothesis to be removed.

Another point of criticism is more substantial. Damping version 1 obviously bestows a privileged status onto saturatable subsets of K – these are the only sets for which, in Levi's terminology, damped equals undamped informational value. For such a privileged status I can see no good reason. Advocating doxastic conservatism, AGM had begun by considering maximal-nonimplying subsets of K ("maxichoice contractions") as the only options for contraction. Levi is correct in emphasizing that this restriction cannot be justified. As he points out, agents may sometimes turn to logically weaker belief sets without incurring any additional loss in informational value (if the relevant additional possibilities bear zero M -value). However, if this line of reasoning is right, there is no motivation any more for

²¹ A similar definition is given in one step in Levi (1996, p. 23).

²² These are the ones for which the $\sim h$ -cell is minimally M -probable among the $\sim h$ -cells outside K , and for which the h -cells outside K bear zero M -value.

insisting on the property of saturatability which is a remnant of strict conservatism. In sum I think that Levi's (2004, pp. 134–135) recent decision to give up his earlier "version 1" damping of informational value was a good one.

6.2 Damping version 2

Levi's new theory centers around the old AGM notion of maxichoice contractions. A *maxichoice contraction* of a corpus K is the disjunction (alternatively, intersection) of K with a single element of U outside K (cf. 1996, pp. 20, 262; 2004, p. 60).

Damped informational value version 2 can also be defined in two steps (2004, p. 141). The loss of informational value incurred by shifting from the corpus K to a maxichoice contraction K' is equal to the loss of undamped informational value, i.e., $Cont(K) - Cont(K') = M(K') - M(K)$. The loss of informational value incurred by shifting from the corpus K to an disjunction (intersection) of some set of maxichoice contractions removing h is the largest loss incurred by any element in that set, i.e., $Cont(K) - Cont(\bigvee K_i) = \max_i M(K_i) - M(K)$. Since maxichoice contractions are more definite than saturatable contraction, no problem of well-definedness arises here.

This definition makes sure that the informational value of the disjunction (intersection) of two contractions removing h is equal to the minimum of their informational values. What Levi now recommends comes down to saying that *the right contraction* removing h from K is the disjunction (intersection) of all maxichoice contractions K' that are at least as informationally valuable as a maxichoice contraction removing h that minimizes the loss of damped informational value version 2. This results in the disjunction of K with all the lowest M -valued $\sim h$ -cells and with all h -cells outside K carrying no higher M -value than these (2004, pp. 142–143). By the definition of damping version 2, this disjunction (intersection) of maxichoice contractions is exactly as informationally valuable as any optimal maxichoice contraction removing h . And it is clearly the weakest one among the maximally informative contractions removing h . Thus there is no problem to apply the Rule for Ties for contractions, and to combine it with the decision-theoretic rule. Levi calls this method of removing a hypothesis from a corpus *mild contraction*.

I like Levi's idea of refocusing on AGM's maxichoice contractions rather than saturatable contractions, but I find his concept of damped informational value version 2 counterintuitive.

In order to see this, let us first have a fresh look at the example of the previous section. The maxichoice contractions removing p with the smallest loss of information are $K \vee (\sim p \wedge q \wedge r) = q \wedge r$ and $K \vee (\sim p \wedge q \wedge \sim r) = q \wedge (p \leftrightarrow r)$, each bearing M -value 0.2 and $Cont$ -value 0.8. But all the maxichoice contractions that do *not* remove p have the same values. So what Levi recommends is actually the disjunction (intersection) of six maxichoice contractions with these values. It is easily seen that this disjunction is $K' = p \vee q$, the damped informational value version 2 of which is the minimum of the $Cont$ -values of the maxichoice contractions involved, that is, 0.8.

We have finally found that according to damping version 2 the belief state represented by $K' = p \vee q$ has the same informational value as any one of the six maxichoice contractions that are used for the construction of K' . This is strange, since K' is obviously *much* weaker than the latter contractions (it comprises 6 as opposed to only 2 cells of the ultimate partition). What could be the justification for this deviation from our ordinary intuition of informational value? I cannot think of one.

More generally, a typical situation is this. Suppose K is a corpus and K_1 and K_2 are two different proper subsets of K that both minimize the loss of information, subject to the constraint that h be removed from K . Let us suppose that K_i is the disjunction (intersection) of

K and a single cell x_i of U outside K ($i = 1,2$), with x_1 different from x_2 . That is, K_1 and K_2 are maxichoice relative to U . Suppose further that both K_1 and K_2 incur a non-zero loss of informational value. This means that if the inquirer starts from the corpus K , admitting the possibility x_1 loses some informational value, and admitting the possibility x_2 loses some informational value as well. My thesis now is that on any natural account of informational value, admitting the possibility x_1 also incurs a non-zero loss of information if the inquirer were to start from K_2 , and admitting the possibility x_2 also incurs a non-zero loss of information if the inquirer were to start from K_1 . The *amount* of information lost may vary from context to context, but the fact that *some* information is lost if the inquirer ceases to be able to rule out a possibility seems indisputable, given that this fact means a loss of informational value when the inquirer sets out from K . The particular corpus from which the inquirer starts should not make *that* much of a difference. But if this is right, it follows that $K_1 \vee K_2$ which rules out neither x_1 nor x_2 has *less* informational value than either of K_1 and K_2 , and hence its informational value is *lower* than the minimum of the values of K_1 and K_2 .

This consideration is very close to the Principle of Constant Marginal Returns that Levi endorses in the context of deliberate expansions (see Section 3 above). Unfortunately he does not comment on why he refrains from employing the same principle for belief contraction.

Levi (2004, pp. 181–186) points out that his model of mild contraction is formally identical with a model studied under the name '*severe withdrawal*' by Pagnucco and Rott (1999). In this paper, the model was justified in terms of Principles of Preference and Indifference, and this still seems more convincing to me than Levi's justification in terms of damped informational value version 2. Levi is exactly right in saying that we all (Levi, Pagnucco and Rott) agree that mild contraction *alias* severe withdrawal should be "taken seriously" (2004, p. 147). But I personally think that neither Levi's nor Pagnucco's and my justification of mild contraction is strong enough to warrant its endorsement as *the* distinguished legitimate way of contracting corpora of belief.²³ Levi's mild contraction surely deserves to be taken seriously, but it is very severe indeed, much more severe than AGM contraction and his own contraction based on damped informational value version 1.

In the past couple of years, Levi has come to advocate strongly version 2 of damped informational value. I think that this notion is motivated by his wish to construct a measure of information that conforms both to the Decision-theoretic Rule and the Rule for Ties at the same time.²⁴ I am not convinced that this project is on the right track. I rather think one should acknowledge that the desiderata expressed by both rules tend to require genuinely different kinds of contraction behaviour: the former requires informational economy (or 'thrift'), the latter requires the equal consideration of multiple solutions (or 'fairness').

Pagnucco and Rott (1999) argued that the AGM model of partial meet contraction is above all committed to the Rule for Ties, since partial meet contraction effectively makes considerations of fairness override considerations of minimal change (or 'maximum information'). The strict idea of minimal change would indeed insist on maxichoice contraction. For AGM, the Rule for Ties is not, as Levi's picture suggests, a secondary value commitment that comes after the idea of minimizing loss of informational value (pp. 119, 150–151). It is a primary value commitment. This is not in conflict with Levi's own ideas, since, as we saw, he claims that we can fully satisfy the idea of fairness, provided that we use

²³ Hansson's (1999, Observation 2.52) criticism that mild contraction is too "expulsive" still stands unanswered: For any two hypotheses h and h' , the inquirer either loses h' when removing h , or she loses h when removing h' from her corpus – even if the contents of h and h' are in no way related.

²⁴ See Levi (2004, p. 125): "When two or more saturatable contractions removing h minimize loss of informational value, we *want to* be in a position to recommend adopting the 'skeptical' contraction that is the intersection of these optimal contractions. And we *want to* be in a position to do so while still claiming that informational value is being minimized." (My italics)

the right notion of informational value: damped informational value version 2. But he goes farther than AGM in claiming that he can at the same time be true to the idea of minimizing the loss of informational economy.

To me it seems that the *only* support for Levi's notion of damped informational value version 2 is that it represents a way of simultaneously satisfying the two desiderata of informational economy and fairness. But the whole strategy itself seems counterintuitive to me. Rule 1 and Rule 2 give expression to conflicting desiderata, and I see no reason why one should not say so. As we all know, economy and fairness just don't always agree. Agents have to rank or to weigh them in order to reach principled and satisfactory decisions.²⁵ It is somewhat ironic that Levi who has said so many insightful things about *Hard Choices* has here refused to acknowledge the existence of conflicting desiderata in belief contraction.

7. Conclusion

I have tried to present Levi's basic decision-theoretic models of inquiry as perspicuously as possible. His account of deliberate or inductive expansion (without inputs) combines the concerns for truth and information in a very interesting and elegant manner, successfully linking inductive reasoning to decision theory. On the other hand, I find Levi's account of belief contraction (with regard to some preselected hypothesis) not quite as convincing. His aim is to obey two rules for belief contraction at the same time: The Decision-theoretic Rule of minimizing the loss of informational value and the Rule for Ties. In order to achieve this aim, Levi has invented two notions of damped informational value that satisfy the Intersection Equality. I have argued that both notions lack independent motivation and thus fail to render compatible the two desiderata that Levi set out to meet. These desiderata just pull in different directions. It would be nice if Levi were right, but I think we have to put aside our aspirations toward theoretical elegance and admit that a compromise between (i.e., intersection or disjunction of) optimal solutions need not itself be optimal in the same sense. Obeying the tie-breaking rule means sacrificing some informational value.

For reasons of space, I have refrained from making an issue of another, more fundamental point. Levi puts great emphasis on an important asymmetry between belief expansion and belief contraction. The notion of truth – or more precisely, of probability of truth – is crucial in belief expansion, but is completely absent in belief contraction. The reason is that Levi's inquirer is conscious of the possibility of importing error through expansion, but completely rules out the possibility of having an error contained in his or her current corpus of full beliefs:

If an inquirer's belief state is K , any hypothesis h incompatible with K is not a serious possibility. The coherent inquirer must regard every such hypothesis as certainly false *and maximally implausible*. ... as long as long as the inquirer is in the state of full belief represented by K , the inquirer cannot coherently acknowledge the serious possibility of being mistaken in this belief. ... Each and every cell in U^*_K [in U outside K , H.R.] is maximally and equally implausible. There can be no distinction between hypotheses incompatible with K with respect to plausibility. (Levi 2004, p. 174)

²⁵ Many people will be inclined to think that the pair of payoffs $\langle 60,40 \rangle$ is "better" than both of the payoff pairs $\langle 45,45 \rangle$ and $\langle 80,30 \rangle$. The latter is too unjust, the former too wasteful.

This dogmatic insistence of the inquirer that all of his or her current beliefs must be true is a central feature of Levi's brand of pragmatism. In his terminology, inquirers are not *incorrigible* (otherwise there would be no contractions), but they are bound to consider themselves *infallible*. All beliefs in the corpus K have maximal degree of belief.²⁶ This is why credal probabilities need not be considered in Levi's contraction, and only informational-value determining probabilities matter.

My feeling is that this (temporary) dogmatism of the inquirer about her beliefs might overtax pragmatism to a point where it becomes flatly implausible. First, it is one of Levi's ideas that contractions are performed in order to give alternative hypotheses a hearing. Why should an inquirer do this if she is certain that she is right anyway? She must be having a faint idea that there might be something wrong with her beliefs. Second, at the time an inquirer incorporates new beliefs into her corpus by expansion, she must be aware that inductively expanding her beliefs is risky business. Why should she be committed to forgetting all about the previously uncertain status of her newly promoted full beliefs? In so far as these questions are to the point, inquirers are and should be concerned about truth not only in expansions but also in contractions. Error cannot, of course, be imported in contractions, but it can be expelled by contractions. The challenge is to come up with a model that reverses Levi's model for expansions and uses sensible, well-motivated credal and informational-value determining probability functions.

Usual notions of justification play no role in Levi's epistemology which thus seems orthogonal to the traditional concerns of mainstream theories of knowledge.²⁷ Levi's pragmatist attitude opposes pedigree epistemology, holds that it is not beliefs but changes of belief that are in need of justification, and gives such justification in decision-theoretic terms. At the time of writing, Levi's immensely important approach is still only loosely connected with mainstream epistemology. This is a regrettable state of affairs, and one I hope will be alleviated soon.

References

- Carnap, Rudolf, and Yehoshua Bar-Hillel, "An Outline of a Theory of Semantic Information". Technical Report 247, Research Laboratory of Electronics MIT, 1952. Reprinted in Yehoshua Bar-Hillel, *Language and Information*. Reading, Mass., London: Addison-Wesley, 1964, pp. 221–274.
- Fioretti, Guido, "A Mathematical Theory of Evidence for G. L. S. Shackle." *Mind and Society. A Journal of Cognitive and Epistemological Studies on Economics and Social Sciences* 2 (2001): 77–98. (<http://www.icer.it/docs/wp2001/Fioretti3-01.pdf>)
- Hansson, Sven Ove. *A Textbook of Belief Dynamics: Theory Change and Database Updating*, Dordrecht: Kluwer, 1999.
- Harman, Gilbert. *Change in View*, Cambridge, Mass.: Bradford Books, MIT Press, 1986.
- James, William, "The Will to Believe." In: F.H. Burkhardt, F. Bowers and I.K. Skrupskelis, eds., *The Will to Believe and Other Essays in Popular Philosophy – The Works of William James, Vol. 6*, Cambridge, Mass., and London: Harvard University Press, 1979, pp. 13–33. Originally published 1897.

²⁶ According to Shackle's way of determining degrees of *belief*. For those who Levi calls "Parmenidean epistemologists", including Gärdenfors and Spohn, only logical truths have maximal degree of *invulnerability*. – In essence, Shackle belief functions as used by Levi order non-beliefs, while invulnerability functions order beliefs. Shackle functions are useful for inductively expanding K into some decision-theoretically recommended $K+h$, invulnerability functions are useful for contracting K into some decision-theoretically recommended $K-h$. Shackle measures are dependent on K , U , Q_K and M_K , while invulnerability measures are dependent on K , U and M and, possibly, on the sentence to be retracted (Levi 1996, p. 267, 2004, pp. 192, 196).

²⁷ A still inspiring intermediary can be found in Harman (1986).

- Kaplan, Mark, "Decision Theory and Epistemology." In: Paul K. Moser (ed.). *The Oxford Handbook of Epistemology*. Oxford: Oxford University Press, 2000, pp. 434–462.
- Levi, Isaac. *Gambling With Truth: An Essay on Induction and the Aims of Science*. New York: Alfred Knopf, 1967.
- Levi, Isaac, "Information and Inference." *Synthese* 17 (1969): 369-391. (Reprinted in Isaac Levi. *Decisions and Revisions*. Cambridge: Cambridge University Press, 1984, pp. 51–69.
- Levi, Isaac. *The Enterprise of Knowledge: An Essay on Knowledge, Credal Probability, and Chance*. Cambridge, Mass.: MIT Press, 1980.
- Levi, Isaac. *The Fixation of Belief and Its Undoing: Changing Beliefs Through Inquiry*. Cambridge: Cambridge University Press, 1991.
- Levi, Isaac. *For the Sake of the Argument: Ramsey Test Conditionals, Inductive Inference and Nonmonotonic Reasoning*. Cambridge: Cambridge University Press, 1996.
- Levi, Isaac, "Seeking Truth." In: Wolfram Hinzen and Hans Rott (eds.). *Belief and Meaning: Essays at the Interface*. Frankfurt: Hänsel-Hohenhausen, 2002, pp.119–138.
- Levi, Isaac. *Mild Contraction: Evaluating Loss of Information due to Loss of Belief*. Oxford: Oxford University Press, 2004
- Pagnucco, Maurice, and Hans Rott, "Severe Withdrawal (and Recovery)." *Journal of Philosophical Logic* 28 (1999): 501–547.
- Sandqvist, Tor, "On Why the Best Should Always Meet." *Economics and Philosophy* 16 (2000): 287–313.
- Shackle, George L.S., *Expectation in Economics*, Cambridge: Cambridge University Press, 1949, second edition 1952.
- Shafer, Glenn, *A Mathematical Theory of Evidence*, Princeton: Princeton University Press, 1976.