

CONTRACTION AND INFORMATIONAL VALUE

**Isaac Levi
Columbia University**

July, 1996

Second version: December, 1996

Third version: February, 1997

Fourth version: March, 1997

Fifth version, August, 1997

Sixth version, February, 1998

Seventh version, April, 1998*

*Thanks are due to Hans Rott and Maurice Pagnucco who have commented on the July draft of this essay and who subsequently sent me a draft of their ideas on contraction. I have sought here to explain how I came to change my mind and favor mild contractions (or severe withdrawals as Rott and Pagnucco call such contractions) and to indicate some of the relations between the approach of Rott and Pagnucco and my ideas that explain our reaching a similar conclusion independently from somewhat different perspectives.

©1998 Isaac Levi

1. Supposition and Belief Change

According to the approach made famous by Alchourrón, Gärdenfors and Makinson (1985), *revision* is a transformation K^*_h of a potential belief state K by adding h yielding another potential belief state.¹ This AGM revision transformation is a composition of two other transformations: *contraction* and *expansion*. $K^*_h = [K_{-h}]^+_h$. This is the expansion by adding h of the contraction K_{-h} of K by removing $\sim h$.²

Among the intended applications of AGM revision, two are especially prominent:

(1) Revisions are supposed to characterize the way states of belief are or ought to be changed when new information h is added.

(2) Revisions are supposed to characterize the conclusions legitimately reached relative to the current state of belief K on the basis of a supposition adopted purely for the sake of the argument.

Supposing that h is neither the same as nor a species of believing that h . To suppose that h entails no change in belief. As a consequence, there is no need to justify supposing that h . By way of contrast, changing one's mind by coming to believe or ceasing to believe is always a fair target for critical scrutiny. Although one may be criticized for reasoning one bases on a supposition, one cannot be faulted for making a supposition except, perhaps, on the grounds of the relevance of the fantasy to the subject under discussion.

I have argued (Levi, 1991) that AGM revision cannot be an adequate basis for an account of legitimate belief change by adding a proposition h . When h is consistent with K , the new belief state should be the expansion of K -- i.e., the weakest potential belief state K^*_h containing the consequences of K and h .³ But when h is inconsistent with K , it is far from obvious that the consistent outcome of adding h and adjusting the result to obtain a consistent belief state should contain h . Perhaps, the adjustment should result in removing h . (See Hansson, 1991 and Levi, 1991, ch.4.) The adjustment, so I argued, should be decomposable into a sequence of legitimate contractions and expansions but the output need not be and, in general, would not be a revision in the sense in which a revision of K by adding h should "successfully" entail h .

¹I shall assume, though Alchourrón, Gärdenfors and Makinson do not, that K is consistent. In the case where K is inconsistent, I take revision to be undefined.

²I take expansion to be defined for all K . I restrict the transformation K^*_h to consistent K in contrast to the practice of Alchourrón, Gärdenfors and Makinson. Although it makes sense to speak of contracting from an inconsistent belief state K_{\perp} by removing h , K_{\perp} and h do not uniquely determine the contraction. See Levi, 1991, ch.4.8.

³A *conceptual framework* is a set K of potential states of full belief (Levi, 1991, ch.2.2). Potential states of full belief are *nonlinguistic attitudinal states* of commitment to having dispositions to linguistic and nonlinguistic behaviors of various kinds. They are not linguistic entities. No internal syntactical structure is attributed to such states. Members of K are partially ordered by a *consequence relation* in a manner satisfying the requirements of a boolean algebra closed under meets and joins of sets of potential states of arbitrary cardinality. The *expansion* of a potential state of belief K by another potential state of full belief K' is the meet of K and K' . In the subsequent discussion (beginning in section 3), I shall focus attention on sets of potential belief states representable in some regimented language \underline{L} . The set of representations is a set of *potential corpora* expressible in \underline{L} . Each potential corpus will be closed under a Tarskian consequence operation for \underline{L} [so that $K \subseteq Cn(K)$, $Cn(K) \subseteq Cn(Cn(K))$ and if $K \subseteq \underline{J}$, $Cn(K) \subseteq Cn(\underline{J})$]. Expansions, contractions and revisions of potential belief states are representable by expansions, contractions and revisions of the potential corpora that represent them. A sentence h in \underline{L} will be equated with the corpus $Cn(\{h\})$ so that the expansion K^*_h of K by adding h is $Cn[K \cup \{h\}]$. The set \underline{K} of potential corpora in \underline{L} is itself partially ordered by set inclusion as a Boolean algebra so that the expansion of a potential corpus \underline{K} in \underline{L} by a sentence h in \underline{L} is well defined. Suppose "p is possible according to X at t" is in corpus \underline{MK} in language \underline{ML} if and only if $\sim p$ is not in X's corpus \underline{K} (in \underline{L}) at t and "p is not possible according X at t" if and only if $\sim p$ is in X's corpus at t. A language \underline{ML} that contains \underline{L} , that contains all truth functional compounds of sentences in \underline{L} with sentences of the type "p is possible according to X at t" (and truth functional compounds of these) and that satisfies the condition just indicated can be used to represent the same potential belief states as can be represented by a corpus in \underline{L} . To each potential corpus \underline{K} in \underline{L} there corresponds a unique corpus \underline{MK} in \underline{ML} . However, the expansion of \underline{K} by adding h conceived as the logical consequences of \underline{K} and h in \underline{L} cannot be represented by the set of logical consequences of \underline{MK} and h in \underline{ML} . The result is not a potential corpus in \underline{ML} .

CONTRACTION AND INFORMATIONAL VALUE

Suppositional reasoning of type (2) differs from genuine belief change. If it is supposed for the sake of the argument that h , the addition of h to \mathbf{K} does not represent a change in the agent's state of belief. In supposing that h , the inquirer does not come to believe that h . Supposition for the sake of the argument is a different propositional attitude than believing either with certainty or with probability. In genuine belief change, adding h when it is inconsistent with \mathbf{K} entails adjustments to make the result consistent; but such adjustments do not always warrant retaining h . In supposition, the situation is different. h is retained in the transformation of \mathbf{K} in virtue of the suppositional stipulation even when it is inconsistent with \mathbf{K} .

For example, if an otherwise reliable reporter testifies that his mother was a virgin at the time of his birth, I do not add his testimony to my beliefs and revise to secure consistency. To the contrary, I would throw out his testimony and cease having confidence in his reliability.

On the other hand, if I were asked to suppose for the sake of the argument that the testimony is true, I would transform my current belief state in fantasy so as to consistently accommodate the claim that his mother was a virgin at birthing time.

AGM revision is inadequate as an account of legitimate belief change. Some of the difficulties with it, however, do not apply when it is used in an account of suppositional reasoning. But even in the setting of suppositional reasoning, AGM revision exhibits some limitations.

As an account of suppositional reasoning, AGM revision insists that if the initial belief state and the supposition are both logically consistent, the transformed suppositional belief state should be consistent as well.. The kinds of contraction relevant to the sort of suppositional reasoning that AGM seeks to capture remove items from the agent's belief state that lead to contradiction with the supposition. There are, however, forms of suppositional reasoning that are not consistency preserving in this way.

When the suppositional reasoning is consensual, the initial contraction is not designed to transform the belief state to something consistent with the supposition but to meet a requirement stipulated independent of the supposition (such as constituting a consensus).. AGM revision and other forms of consistency preserving revision seek to make the belief state hospitable to the supposition.⁴

Even if we focus attention on consistency preserving forms of suppositional reasoning, the sort of revision of \mathbf{K} by adding h suppositionally should not, in general, be an AGM revision but what I call a *Ramsey Revision* (Levi, 1996). The Ramsey revision

⁴Consistency preserving suppositional reasoning is informational value based. In Levi, 1996, ch.2, it is contrasted with consensus based suppositional reasoning. Suppositional reasoning of this kind is not consistency preserving. This is clear in the context of *reductio* argumentation. As I understand such cases, the agent takes his current belief state \mathbf{K} and forms the meet of it with some other belief state \mathbf{K}^c which is the potential belief state that is the target belief state with which the agent seeks a consensus. Sometimes \mathbf{K}^c is, indeed, the belief state of some other agent with whom the agent is having an exchange. For example, Jones and Smith may disagree as to whether Oswald killed Kennedy even though they agree that somebody did. So the consensus corpus \mathbf{K}^c for both Jones and Smith contains the claim that somebody killed Kennedy without containing the claim that Oswald killed Kennedy or its denial. Jones (or Smith) might contract to \mathbf{K}^c and then for the sake of the argument expand by adding "Oswald did not kill Kennedy". In this case, the expansion is consistent and implies that somebody other than Oswald killed Kennedy. So Jones may say, "If Oswald did not kill Kennedy, somebody else did." In other cases, Jones and Smith agree that Her Majesty is not in the vicinity. One of them might expand the consensus corpus \mathbf{K}^c by adding "Her Majesty is present" and getting a contradiction. He might say, "If Her Majesty is present, I'm a monkey's uncle". (By and large, consensus supposition is expressed in English by so called "pure indicatives"). Sometimes $\mathbf{K}^c = \mathbf{K}$. Or it could be a theory that agent has already supposed true for the sake of the argument whose ramifications he is seeking to explore (e.g., the phlogiston theory). Given contraction by consensus, the agent adds his supposition. Now since the contraction was not constrained by the requirement that the contraction be consistent with the supposition, nothing precludes expansion by adding the supposition even when contradiction ensues. In Levi, 1996, ch.2, I contend that consistency preserving informational value based suppositional reasoning is normally expressed in English either by so called "future indicative" or "subjunctive" conditionals. For example, we might say, "Had Oswald not killed Kennedy, Kennedy would have completed his term as President."

of \mathbf{K} by adding h is defined as $[(\mathbf{K}_{\sim h})_{\sim h}]^+_h$. Ramsey revision coincides with the AGM revision in case the *Recovery Postulate* for contraction holds. The Recovery Postulate stipulates that if h is in \mathbf{K} , $[\mathbf{K}_{\sim h}]^+_h = \mathbf{K}$. I have argued against Recovery in both Levi, 1991 and Levi, 1996 and pointed out that when the Recovery postulate is abandoned, the postulates for Ramsey revision no longer coincide with those of AGM revision.. This point is of some importance; for, as Makinson (1987) points out, it makes no difference to the properties of AGM revision whether the Recovery Postulate holds or not.

Although I continue to be skeptical of Recovery and to consider Ramsey Revision a better basis for an account of suppositional reasoning than AGM Revision, the positive account I proposed to replace the AGM account of contraction now seems to me to have been defective. The defect lurks in my account of what I called the problem of how to contract (proposed in Levi, 1991 and reiterated in Levi, 1996). In this essay, I shall recapitulate the proposal I originally favored, why I think it is defective and how I think the defect can be removed. The proposal I would now recommend is the one that Pagnucco and Rott, 1997 call “severe withdrawal”⁵ and I prefer to call “mild contraction” for reasons that shall emerge subsequently.

2. How to Contract as a Decision Problem

I understand the problem of how to contract by removing h from \mathbf{K} to be a decision problem where one is called upon to choose a contraction removing h from \mathbf{K} *from among all the contraction strategies removing h from \mathbf{K} available in the context*. The rational decision-maker engaged in contraction chooses the contraction strategy available to him or her that best promotes the goals of the contraction.⁶ In this respect, the problem of how to contract is a pragmatic issue just as the problem of how to expand is in the context of deliberate inductive inference as spelled out in Levi, 1967; 1984, ch.5; 1980 or 1991.

Such a decision theoretic approach requires that we take two tasks seriously:

1. The specification of the set of contraction strategies that are available options for the agent in the particular setting.
2. The characterization of the goals of contraction in a way that determines how well the available options among the contraction strategies promote these goals.

Let the results of completing these two tasks yield a weak ordering of all of the available options with respect to how well the goals of contraction are promoted. If the set of available contraction strategies is finite, then the inquirer may be urged to choose a contraction strategy removing h from \mathbf{K} that is optimal with respect to this ordering.

⁵My information about this joint work is initially based on an oral report delivered by Hans Rott at a Swedish-German Workshop on Belief Revision held in Leipzig early in May, 1996 and on some notes Rott kindly sent to me. Rott and Pagnucco have developed their ideas in an unpublished (as of December, 1997) manuscript “Severe Withdrawal (and Recovery)”. For reasons that I shall explain in this essay, I came to the conclusion that mild contractions (=severe withdrawals) are to be favored over the contractions I favored previously a month or so before coming to Leipzig. My concern here is to present these reasons which may serve to explain the philosophical perspective from which I approach the questions of contraction and revision. Pagnucco and Rott came to their conclusions via an entirely independent route. Although there seems much common ground between their views and mine, the extent of this consensus remains to be explored. This essay aims to contribute to that exploration.

⁶The focus of this discussion is with giving an account of contraction in decision theoretic terms. I am not concerned to offer a set of axioms for contraction as I understand it that is complete in some sense or other. In comparing the views I favor with the AGM theory and other alternatives, however, I make reference to similarities and differences with respect to satisfaction of various postulates that have been considered and take for granted some familiarity on the part of the reader with the AGM axioms for contraction. For me, a condition of adequacy for any axiomatic account of contraction (or revision or expansion) ought to be an account of how the choice of a contraction (revision, expansion) is rationalized with respect to the goals of the inquirer and the options available.

Even if the set of strategies available is infinite, a weak ordering of the options can yield an optimum.⁷

Thus, it is desirable to spell out the conditions a set of available contraction strategies removing h from \mathbf{K} should meet and the desiderata a weak ordering of these contraction strategies should satisfy to reflect the goals of contraction. Before doing so, however, a review of parallel issues pertaining to inductive expansion might help set the stage.

3. Expansion Strategies

In statistical reasoning, theory choice and other types of inductive or ampliative reasoning involving a deliberate expansion of the agent's initial belief state, the inquirer is called upon to choose between alternative expansion strategies or potential answers to the question under investigation. At least so I have argued elsewhere (e.g., in Levi, 1967, and 1984, ch. 5). Identification of a list of potential answers is the task of *abduction*. The principles of abduction require, for example, that one be in a position to regard suspension of judgment between rival potential answers to be a potential answer.⁸

⁷According to the general account of rational decision making I have proposed elsewhere (Levi, 1986), a rational agent may have value commitments and goals that are in conflict so that he recognizes more than one evaluation of his options as permissible to use. In such cases, a non controversial weak ordering of the options may not exist for the agent and decision making will have to proceed by identifying which options are optimal according to at least one weak ordering that is consonant with the evaluations the agent can make. Such options are *V-admissible*. Rational agents restrict their choices to *V-admissible* options.

This idea can be applied to the question of how to contract. I shall postpone discussion of cases where the inquirer's cognitive goals conflict until section 10 and focus, instead, on the quite likely unrealistic case where a uniquely permissible weak ordering is available.

⁸The task of abduction as Peirce understood it in 1902 is to formulate conjectures that qualify as potential answers to questions under study or potential solutions to problems being confronted. The task of induction is to decide which of the potential answers obtained by abduction at a given stage of inquiry to adopt as the solution to the problem or answer to the question under investigation. Since the aim common to all inquiry is to answer questions and settle problems, there is no point to abduction unless induction is taken seriously. Moreover, the difference between abduction and induction is not a matter of the degree of confidence or evidential support afforded the conclusion of an induction as compared to that provided for abduction. In induction, a conjecture obtained via abduction is assessed relative to the evidence before the conjecture is transformed into a settled answer to the problem under study. On the basis of that evidence, a decision is taken to change the set of settled assumptions, evidence or certainties by adding the hypothesis to them. That is induction. The evidential support such as it is for a conjecture is relative to the evidence *before* the conjecture is converted to a settled assumption -- i.e., relative to the information taken be certain when the conjecture is a conjecture formed by abduction. Conjectures obtained by abduction may differ in the degree to which in some sense or other the evidence supports them; but suggesting that the settled conclusions of induction differ from conjectures in degree of support is misunderstanding the relation between abduction and induction.

Many philosophers and computer scientists have understood Peirce's notion of "abduction" to be inference to the best explanation. Peirce's writings from the 1860's, 70's and '80's do, indeed, afford textual support for such an interpretation of his conception of hypothetic inference. At the same time, however, he apparently thought of hypothetic inference as engaged in the task of proposing conjectures for testing -- a quite different task from testing conjectures and concluding that some are to be eliminated. The latter task is fulfilled by drawing inductive inferences. Around 1902, Peirce explicitly acknowledged that his vision of the distinction between induction and hypothesis (then rebaptised "abduction") could not be "the reasoning by which we are led to adopt a hypothesis". The conclusion of an abductive "inference" is "in the interrogative". It asks whether the conjecture (= the conclusion) is true or not but does not adopt the conclusion as true. Peirce confessed that in his previous work he had confused such conjecturing with adopting a hypothesis as a settled conclusion. He blamed his confusion on his well known early obsession with trying to provide a formal contrast between hypothesis and induction in terms of different permutations of premisses and conclusions of syllogisms. (See Peirce, CP.2.102.) Many contemporary authors have perpetuated the terminological practice that Peirce abandoned for principled philosophical reasons. Unlike Peirce, they either explicitly (as in the case of Popper) or tacitly deny the legitimacy of induction as leading to the adoption of a potential answer as the settled answer to a question at the termination of inquiry. Indeed, Popper's anti inductivism is best understood as recommending conjecturing and testing *without* refutation. Popper's words to the contrary notwithstanding, experiments do not, in general, warrant the rejection of any conjecture. The truth of the reports of outcomes of experiments cannot be conclusively accepted and even if they are, they do not entail the falsity of allegedly refuted conjectures on deductive grounds. Others like Quine and Ullian (1970, pp.58-9) see

Let belief states be represented in a regimented first order language \underline{L} . Any potential belief state \mathbf{K} representable in \underline{L} is representable by a *corpus* \underline{K} which is a set of sentences closed under logical consequence in \underline{L} . Let W be the set of maximally consistent corpora in \underline{L} and $P_{\underline{K}}$ be the set of maximally consistent extensions of \underline{K} .

Partition $P_{\underline{K}}$ into finitely many cells c_1, c_2, \dots, c_n . Form intersections of members of each cell. Suppose that the set of intersections constitute a set of potential expansions of \underline{K} representing strongest consistent potential answers of interest to the inquirer in the context of the given inquiry. In general, the language \underline{L} has means for expressing hypotheses consistent with \underline{K} that are stronger than the elements of this partition. But someone interested in the value of the GNP for the USA in 1995 will not be thinking of the average rainfall in Iraq as part of the information carried by a strongest consistent potential answer to the question under study.

When \underline{L} has for each c_i a sentence d_i expressing the "proposition" $|d_i| = c_i$, the strongest consistent potential expansions of interest to the inquirer can be represented as expansions $\underline{K}_{d_i}^+$ of \underline{K} by adding d_i . The inquirer's *ultimate partition* $U_{\underline{K}}$ relative to \underline{K} and to the inquirer's demands for information consists of the cells in the partition of $P_{\underline{K}}$ into the c_i 's or the corresponding set of sentences d_i . Notice that when there is a set of sentences e_i such that for each d_i , \underline{K} entails the equivalence in truth value between e_i and d_i the ultimate partition relative to \underline{K} can be represented by the set of e_i 's. Of course, for each d_i , $\underline{K}_{e_i}^+ = \underline{K}_{d_i}^+$.

The set of all potential answers (or the potential expansion strategies) of interest to the inquirer in the context of the given inquiry are linguistically representable as expansions of \underline{K} by adding some sentence h in \underline{L} equivalent given \underline{K} to a disjunction of some subset of d_i 's in $U_{\underline{K}}$. The disjunction of the empty set is construed as any sentence whose negation is in \underline{K} .

It is important to keep in mind the trivial point that the sentences d_i in $U_{\underline{K}}$ need not be logically exhaustive but only exhaustive given \underline{K} . For convenience, the d_i 's have been defined so that they must be logically pairwise exclusive. In practice, however, they may be replaced by e_i 's that are equivalent to them given \underline{K} and need not be pairwise exclusive but only pairwise exclusive given \underline{K} .

Moreover, there may be a potential corpus \underline{K}' that is a proper superset of \underline{K} such that $U_{\underline{K}'} = U_{\underline{K}}$. Even though \underline{K}' carries more information than \underline{K} , no element of $U_{\underline{K}}$ is incompatible with the information in \underline{K}' . No new information of interest to the inquirer in the context of the particular question under study has been added to \underline{K} . The set of potential answers remains unaltered. The set of potential expansion strategies relative to \underline{K}' may be representable as expansions by adding disjunctions of the same as the set of cells as the potential expansions relative to \underline{K} . The difference between the two cases may have an impact on which potential expansion strategy the agent should choose because \underline{K}' contains information that alters the probabilities of incurring error in adopting one expansion strategy rather than another. But the extra information does not alter the set of available options.

The elements of $U_{\underline{K}}$ are supposed to represent the strongest consistent potential answers to the question under study of relevance to the demands of the question.

Consider, for example, the question of predicting who will win a given election relative to initial information that either candidate A, B or C will. The agent may be concerned with whether A will win or not in which case, the ultimate partition will consist

induction to be a species of abduction or hypothesis where the latter is the "framing of hypothesis" because if it were true, it would explain some things that the inquirer already believes (p.43). One cannot be sure whether Quine and Ullian reject induction as a way of fixing new beliefs by construing the term 'induction' as singling out certain ways of forming conjectures or whether they, like Peirce confessed his earlier self to have been, are confusing conjecturing with coming to fully believe and taking to be settled by an ampliative inference. Controversy concerns the presuppositions of certain terminological practices rather than terminology itself.

of the two conjectures “A will win”, “A will not win”. Or he may be interested in which of the three will win. Or he may be interested in whether the candidate is from his home town or not so that the ultimate partition has six elements. In the first case, the inquirer contemplates only four potential answers. In the second eight. In the third, sixty-four.

I do not think that there is a standard way of fixing an ultimate partition obligatory on all agents in all settings. There are, to be sure, reasonable suggestions available concerning how agents who disagree in their ultimate partitions and have reason to engage in a common inquiry might come to some agreement as to how to proceed together. They can, in particular, adopt the coarsest common refinement of their partitions (See Levi, 1984, ch.7).

Nonetheless, there are some constraints on what can count as potential answers to a question. Given the choice of an ultimate partition, the inquiring agent is obliged to consider as optional expansion strategies all possible forms of suspense between these expansion strategies and, in addition, the prospect of importing contradiction into his beliefs. That is not to suggest that contradiction is a legitimate mode of expansion; but it is to insist that if it is not, that should be revealed by showing that, given the goals of inductive expansion, contradiction should not be recommended to rational agents.

The choice of an ultimate partition does foreclose consideration of potential expansion strategies that are not representable within the framework of the algebra generated by that partition. This choice reflects how fine-grained our interest in new information is.

Making decisions on this point seems unavoidable. One might seek to duck the issue by urging that we should be as fine-grained as the conceptual resources of \underline{L} or our conceptual framework allow. Each element of the partition in this sense would be a “possible world”. But since there are no conceptual limits on how fine-grained discriminations inquirers can make, there are no possible worlds in this sense. Neither a conceptual framework nor a particular language can be used to standardize the choice of an ultimate partition. To be sure, the language \underline{L} that is to be used allows for the expression of theories that are maximally consistent relative to \underline{L} . But such theories will fail to be maximal relative to languages with richer expressive resources. If an inquirer chooses as his or her ultimate partition, the set of maximally consistent sets of sentences relative to \underline{L} , that choice is as open to controversy as resting content with a coarser ultimate partition. Relative to \underline{L}^* , that partition the elements of that partition are not descriptions of possible worlds. There is no way to escape making a decision as to how fine grained one’s interest in new information is going to be and no reason to insist that it be as specific as a maximally consistent theory in some particular language (or not).

The choice of ultimate partitions can, thus, be a subject of dispute. But as I have suggested, when the choice is disputed, parties to the dispute can reach an agreement as to how to proceed if they care to do so by using coarsest common refinements. To this extent, the choice of an ultimate partition can be brought under critical control. Here then is an example of a way of settling differences by reasoning that does not presuppose that the disagreements concern a question as to what is true.

The closure condition on the set of potential answers I am advocating arises once the issue is settled of how fine-grained an ultimate partition is going to be adopted. I contend that we should not recognize as potential answers to the election question that candidate A will win, that candidate B will win or that candidate C will win without also allowing as potential answers also suspension of judgment between any pair or even all three. If someone insists that definite conclusions are to be recommended over suspense, that person should be required to show why suspense is inferior with respect to the goals of the inquiry. Ruling out suspense as an option by stipulation does not meet this demand.

Finally, it should be noticed that two agents with the same list of potential answers might very well disagree as to which potential answer to adopt owing to differences in their belief states and the probability judgments based on them. Still they agree that exactly one of the three candidates will win.

4. Potential Contractions

The topic of the current essay is not inductive expansion but another cognitive decision problem: How to contract given that we are to remove some specific sentence h from \underline{K} .⁹ The two topics are not totally disconnected. The range of contraction strategies for removing h from \underline{K} is determined by the agent's problem in a manner dual to the case of inductive expansion.

Suppose, to the contrary, that the potential contraction strategies were constrained only by the inquirer's conceptual framework or, perhaps, by the resources of a language \underline{L} used to represent potential states of full belief. Consider, in particular, the set $\{W/P_K\} = I_K$. These are maximally consistent potential corpora in \underline{L} *incompatible* with \underline{K} . If no restrictions other than linguistic ones are imposed on the potential contractions, the potential contractions of \underline{K} removing h would consist of all intersections of \underline{K} with intersections of some subset of $\{W/P_K\}$.

In practice, inquirers do impose restrictions on the set of potential contractions. They immunize themselves (for the time being) against potential contraction to total skepticism by singling out a *minimal contraction* \underline{LK} of \underline{K} such that all potential contractions of \underline{K} to be considered are expansions of \underline{LK} as are all potential expansions of \underline{K} . In previous work, I have taken \underline{LK} to be the weakest potential corpus (the *Urcorpus* \underline{UK} , Levi, 1980, p.7) representing the weakest potential belief state expressible in the language. However, in many settings, it may be the case that some stronger belief set is used as the minimal contraction of \underline{K} . The minimal contraction is the potential corpus representing the weakest potential Belief State expressible in \underline{L} in which the inquirer is interested at the stage of inquiry in which the inquirer is situated.¹⁰

Additional constraints are imposed on contraction by coarse graining. Coarse graining is done by introducing a *basic partition* $V_{LK,K}$ for \underline{K} relative to \underline{LK} (Levi, 1991,

⁹In the setting of belief change, another question needs to be answered in conjunction with the question of how to contract by removing h from \underline{K} -- to wit, whether h is to be removed from \underline{K} . Sometimes contraction is required because the inquirer has inadvertently expanded into inconsistency. According to the view I favor, contraction from inconsistency is not uniquely determined. Whether contraction is called for is not at issue; but what items to remove is very much at issue and deciding this question depends on an account of how to contract by removing specific sentences from consistent corpora in a manner discussed in Levi, 1991, ch.4.8. Sometimes, however, contraction is not coerced by the need to remove inconsistency. In Levi, 1991, ch.4.9, I argue that this can happen when a hypothesis h merits a hearing and this calls for removing $\sim h$ from consistent \underline{K} . In that case, what to remove is clear provided anything is to be removed. One needs to consider, however, whether contraction by removing $\sim h$ is warranted. Once more, it is necessary to consider the problem of how to contract. Finally, contraction removing $\sim h$ may be required in reasoning from the supposition that h for the sake of the argument. Once more, the question of how to contract needs to be addressed. (Levi, 1996.)

¹⁰One setting in which minimal contractions seem useful is in the analysis of future indicative "if" sentences. In such cases, it seems that the suppositional reasoning expressed is prohibited from bringing into doubt any belief about events occurring at the time of utterance or previous to that time. If this view (suggested by V.Dudman) is an accurate reflection of English usage, then in those settings where future indicatives are used in this way \underline{LK} will consist of all those convictions about events occurring in the past and present (as well as theories and generalizations). Notice that sentences in \underline{LK} are immune to removal by contraction as long as \underline{LK} remains fixed. In this sense, the deliberating agent is prohibited from supposing their negations. But I am not suggesting that the deliberating agent misunderstands or fails to entertain the propositions expressed by these sentences. \underline{LK} does not demarcate a realm of conceptual possibility from items that are not entertainable. In a given context, the interests of an agent engaged in supposition are focused in certain directions and not others. Because of those interests, the agent *refuses* to entertain removal of certain convictions from the Belief State even though such removal is conceptually entertainable. \underline{LK} serves to fix the boundary between those items that are entertainably removable from \underline{K} given the inquirer's interests and those that are not. Of course, if my interlocutor insists on entertaining a claim I do not wish to entertain and I wish to engage in joint deliberation with him, I may have to "contract" my \underline{LK} so as to render his claim entertainable by my lights..

p.123-4, 1996, p.26). This partition is an ultimate partition U_{LK} relative to \underline{LK} . The elements of U_{LK} that are consistent with \underline{K} constitute the ultimate partition U_K . As before the sentences representing cells in U_{LK} are exclusive and exhaustive given \underline{LK} and not merely according to logic.

Consider now the cells in $\{U_{LK}/U_K\}$. These are represented by sentences in \underline{L} whose negations are incompatible with \underline{K} . Moreover, these cells represent the strongest hypotheses given the demands of the inquirer's problem incompatible with \underline{K} . I shall call this set of cells U^*_K the *dual ultimate partition* relative to \underline{K} (and \underline{LK}). The set of potential contractions of \underline{K} under consideration may now be represented in the following way.

A *potential contraction of \underline{K} relative to U^*_K* is representable by the corpus that is the intersection of \underline{K} with the intersection of some subset of U^*_K . Potential contractions that replace U^*_K by W/P_K in the definitions that follow constitute a special case of the proposal made here when $U^*_K = \{W/P_K\}$.

A *potential contraction of \underline{K} removing h relative to U^*_K* is a potential contraction of \underline{K} that is obtained by intersecting \underline{K} with a subset of U^*_K containing at least one element of U^*_K that implies $\sim h$ if such exists. Otherwise it is \underline{K} .

Clearly contraction of \underline{K} removing h will be successful (i.e., nondegenerate) only if U^*_K is sufficiently fine grained to contain at least one cell implying $\sim h$. This is not a serious restriction, however, since we may expect that no one would contemplate removing h from \underline{K} in earnest unless he was using a dual ultimate partition that was sufficiently fine grained to have this effect.

We may go even further and assume that a dual ultimate partition U^*_K is *suppositionally decisive* if and only if the following holds:

Given any sentence h in \underline{L} eligible for consideration as an item to be removed "for the sake of the argument" from \underline{K} , all cells in U^*_K either entail h or entail $\sim h$.

In general, the domain of inputs for contraction and the dual ultimate partition will be assumed to satisfy the condition of suppositional decisiveness. All definitions used subsequently will be based on the assumption that U^*_K is suppositionally decisive.

A *maxichoice contraction of \underline{K} relative to U^*_K* is the intersection of \underline{K} with a single element of U^*_K .

A *maxichoice contraction of \underline{K} removing h relative to U^*_K* is the intersection of \underline{K} with a single element of U^*_K that entails $\sim h$.

A *saturatable contraction of \underline{K} removing h relative to U^*_K* is the intersection of a maxichoice contraction of \underline{K} removing h relative to U^*_K with the intersection of a set of elements of U^*_K none of which entail $\sim h$.

Notice that unlike the case with maxichoice contraction, saturatable contraction is defined only for contractions removing some specific sentence from \underline{K} .

Notice also that every potential contraction of \underline{K} removing h relative to U^*_K is the intersection of a subset of saturatable contractions of \underline{K} removing h relative to U^*_K .

By way of contrast, the set of *partial meet contractions* of \underline{K} removing h relative to U^*_K is, in general, a proper subset of the potential contractions of \underline{K} removing h relative to U^*_K . Such partial meet contractions are intersections of subsets of maxichoice contractions removing h from \underline{K} relative to U^*_K . In the AGM account of contraction, the only potential contractions countenanced are the partial meet contractions.¹¹

¹¹Makinson (1987) calls contractions that are not partial meet contractions *withdrawals*.

Ruling out such contractions may be defended on the grounds that it guarantees satisfaction of the Recovery Postulate for contraction. This strategy may appear very attractive at first glance. But the appeal to postulates that seem intuitively compelling at first blush is a risky business. In the case of Recovery, the plausibility of the postulate is highly doubtful as many proffered counterinstances suggest.

Consider, for example, a situation where it is believed that Jones was HIV positive, received a drug treatment and subsequently showed HIV negative. Contract the corpus by giving up "Jones received the drug treatment." The conviction that Jones initially showed HIV positive would be retained. But the judgment that Jones showed HIV negative later on would be abandoned. Moreover, restoring the judgment that Jones received the drug treatment would not resurrect the conviction that Jones subsequently showed HIV negative unless the inquirer had the well entrenched conviction initially that the drug treatment always eliminates the HIV virus. If this belief were not well entrenched or if all that is believed is that the drug treatment is followed by cure in some percentage of cases less than 100%, the Recovery Condition would be violated. (See Levi, 1991, pp.134-5 for a related illustration.)

On a more theoretical level, the merits of postulates for contraction ought to be scrutinized from a decision theoretic viewpoint. If a decision-maker faces a set A of available options and *all* of the options in A are weakly ordered with respect to how well they promote the agent's goals and values, the agent is obliged to restrict choice to one of the optimal options according to that weak ordering. (An option is optimal if it is weakly preferred to every other option in A). If more than one option is optimal, the decision-maker *may* invoke a secondary criterion to decide between the optimal options. Contraction removing h from \underline{K} is a decision problem. No potential contraction should be ruled out as an option until it has been shown to be inferior given the goals of efforts to remove h from \underline{K} .

Such an approach takes a dim view of stipulating that certain available options (potential contraction strategies) be ruled out of court without showing that they are suboptimal means for promoting the goals of contraction.

It may, perhaps, be countered that relativizing contraction to dual ultimate partitions is itself a way of restricting the set of potential contraction strategies. But this relativization is based on the idea that, in contraction as in inductive expansion, the inquirer's concern is not with every issue but only with a restricted set of topics. We can represent the relevant issues by using a coarse grained dual ultimate partition based on a minimal corpus. If a challenge can be brought that the partition is too coarse grained to address the issues under scrutiny or that the minimal contraction \underline{LK} is not minimal enough for the purposes at hand, we can always refine the partition and weaken \underline{LK} . We can then invoke decision theoretic considerations to ascertain whether the initial \underline{LK} and dual ultimate partition yielded recommendations that can no longer be countenanced or not. Any restriction on potential contraction strategies that is due to the choice of a particular \underline{LK} and dual ultimate partition can be called into question.

Challenging the Recovery Postulate requires relaxing the restriction of the options for contraction to partial meet contractions. So there are two ways to protect the Recovery Postulate: (a) by refusing to remove the restriction or (b) by removing the restriction and showing that contraction options violating Recovery are less efficient means for realizing the goals of contraction than contraction strategies satisfying Recovery. Protecting Recovery the first way amounts to a refusal to submit the Recovery postulate to criticism from a decision theoretic point of view and to place it off limits for serious scrutiny. This seems wrong.

Thus, even though (amendable) restrictions on contraction due to coarse graining are permitted, such restrictions are open to challenge. Advocates of the Recovery Postulate seem to regard restriction of partial meet contractions as immune to amendment. I reject this attitude and, for that reason, require that intersections of any

CONTRACTION AND INFORMATIONAL VALUE

subset of saturatable contractions removing h from \underline{K} relative to the dual ultimate partition U^*_K endorsed at the time be taken as options for contraction removing h from \underline{K} at that time.

To illustrate the use of dual ultimate partitions, consider a corpus \underline{K} that contains the claim E that a coin is tossed at time t and the claim H that the coin lands heads at that time. In addition the corpus also contains the general claim L that the coin lands on the surface when and only when it is tossed near the surface. \underline{K} also implies the general claim L' that the coin lands on the surface if and only if it lands heads or lands tails but not both. I shall suppose that the “modeling assumptions” L and L' are implied by \underline{LK} . Clearly \underline{K} also contains the claim O that the coin lands on the surface at time t . The contraction of \underline{K} contemplated is removing E . Consider as the dual ultimate partition the following:

U^*_K : $E\&T$ = the coin is tossed at t and lands tails but not heads. $E\&\sim H\&\sim T$ = the coin is tossed at t , and lands neither tails nor heads (i.e., does not land on the surface) = $E\&\sim O$. $\sim E\&\sim O$ = the coin is not tossed at t and does not land on the surface = $\sim E\&\sim H\&\sim T$.

Relative to the modeling assumptions in \underline{LK} , these items are pairwise incompatible and together with $E\&H$ are exhaustive.

Relative to U^*_K , the intersection of \underline{K} with $\sim E\&O$ is the only maxichoice contraction removing E from \underline{K} . The only other contraction removing E from \underline{K} is the intersection of \underline{K} with $\sim E\&O$ and then with $E\&T$.

U^*_K is quite coarse grained. One way to fine grain is to refine $\sim O$ into, say, O_1 = the coin melts, O_2 = the coin is exchanged for a chocolate candy, and O_3 = fails to land on the surface in some other way. Such refinement need not call for modifying \underline{LK} . But it would allow for more maxichoice contractions removing E .

A banal but more interesting possibility is that \underline{LK} is weakened by replacing L' with L'' = the coin lands on the surface if and only if it lands heads, tails or on its edge. If the dual ultimate partition is then refined to allow for the claim that the coin lands on its edge at t = G , we can obtain three saturatable contractions that are not maxichoice.

Observe, however, that \underline{K} still is assumed to contain L' so that adopting a saturatable contraction allowing for the possibility that the coin lands on its edge would call for questioning L' even if L and L'' remain secure. Such a contraction could not then be ruled out in virtue of the stipulated constraints on \underline{LK} due to modeling assumptions. It might be ruled out, however, on the grounds that the loss of information incurred would be too great to be born. Or the loss of informational value incurred by countenancing the possibility may be regarded as sufficiently small to warrant contracting in a way that gives up L' .

Elaborating on this point requires an examination of how potential contraction strategies are to be evaluated relative to the goals of contraction. This topic will be discussed shortly for fixed \underline{LK} and U^*_K . My current concern is to point out that the choice of minimal corpora and dual basic partitions is open to revision. I have already argued, however, the restriction of potential contractions to maxichoice contractions and the meets of subsets of them recommended in AGM ought to be settled by showing that partial meet contractions best promote the goals of contraction and not by stipulation.

As I have already intimated, it is arguable that meets of sets of maxichoice contractions ought almost never to be superior to meets of some sets of saturatable but not maxichoice contractions with respect to promoting the goals of contraction. An examination of the example just used suggests how implausible it is to think that when one comes to doubt that the coin is tossed at t , one moves to suspense between “the coin was not tossed and did not land” and “the coin was tossed and landed heads”. Surely then “the coin was tossed and landed tails” is countenanced as a possibility. This

counterinstance to the Recovery Postulate is by no means exceptional. Any belief state relative to which the outcome of an experiment is known but where in general experiments of the given kind can yield one of several different types of outcome may serve to undermine confidence in the Recovery Postulate. To be sure, the appeal to presystematic judgment or intuition I have invoked here should be buttressed by a rationalization based on an examination of the way contraction strategies ought to be evaluated in a problem where \underline{K} is to be contracted by removing h relative to U^*_K

I have been arguing that the set of available strategies for contracting \underline{K} by removing h relative to U^*_K should consist of all potential contractions removing h from \underline{K} relative to U^*_K and \underline{LK} as defined in this section. Apologists for the Recovery Postulate have followed Alchourrón, Gärdenfors and Makinson (1985) in restricting the options for contraction to a proper subset of such potential contractions -- namely, the partial meet contractions (relative to U^*_K) and proceeding to evaluate the surviving partial meet contractions. This procedure is unacceptable from a decision theoretic point of view because it does not establish the suboptimality of the potential contractions that are not partial meet contractions among *all* potential contractions with respect to the goals of contraction.

The same complaint should be directed against other approaches to evaluating contraction strategies. In Appendix A, base contraction will be discussed from this point of view. Even though base contraction can provide for failures of Recovery, it often does so by restricting the options for contraction to a proper subset of the potential contractions as defined here. This practice is just as objectionable as restricting the options for contraction to partial meet contractions.

5. The Commensuration Requirement

In Levi, 1991, p.65 (see also Levi, 1980a, ch.'s 2-3), I took the position that every legitimate change of belief state from initial state \mathbf{K}_1 to \mathbf{K}_2 should be decomposable into a sequence of contractions and expansions each of which is legitimate. This *commensuration requirement* implies that every justification of a change in state of full belief should be a justification of an expansion or of a contraction. The only kind of changes that are not expansions or contractions are *replacements* where \mathbf{K}_1 entails that h and \mathbf{K}_2 entails that $\sim h$ and *residual shifts* where \mathbf{K}_1 entails that h but leaves the issue of the truth of g open while \mathbf{K}_2 leaves h open while entailing that g . Both of these transformations are decomposable into sequences of contractions and expansions. That is the content of the commensurability thesis (Levi, 1991, p.65).¹² The commensuration requirement holds that no residual shift or replacement is justifiable unless it is decomposable into a sequence of contractions and expansions (in some order or other) each step of which is legitimate.

Why should the commensuration requirement be respected? I claim that an answer is to be found in referring to the *common* features of the goals that *ought* to be pursued in justifying changes in full belief.

The commensuration requirement, therefore, is not dictated by principles of rational decision making alone. It presupposes endorsement of a substantive view of what the common features of the aims of inquiries concerned to justify changes in belief ought to be.

There is no doubt that scientists and other inquirers seek to solve different kinds of problems and settle different kinds of issues. The goals of inquiry are in this sense as diverse as the problems that occasion inquiry. There is no single aim that all inquirers

¹²The commensurability thesis is satisfied as long as the conceptual framework \mathbf{K} is a boolean algebra as I am supposing here.

CONTRACTION AND INFORMATIONAL VALUE

strive to realize in justifying changes in their views. This must be so if for no other reason than the subject matters of inquiries are so diverse.

Even so, there may be some common features that the goals of all properly conducted inquiries share in common. It may be the case that rational pursuit of goals sharing such features entails endorsement of methodologies for fixing beliefs of certain kinds rather than others.

Most authors who agree that beliefs are to be justified by showing how well they satisfy certain goals are skeptical of there being any common features that could serve as the basis for a sensible methodology. Inquirers seek to justify their beliefs by way of reference to goals that focus on moral, political, theoretical, economic, prudential and aesthetic considerations. Agents sometimes fix their beliefs in certain ways because doing so will promote their personal interests or the interests of the group they serve. Relative to the goals they seek to promote, such agents may sometimes succeed in justifying their beliefs in a rationally coherent fashion. In this respect, inquiry displays the same kind of rationality as practical deliberation. There are no autonomous cognitive values to pursue and even if there were, there is no basis for recommending them over other kinds of values.

I contend that inquirers engaged in inquiries that issue in change of view ought to be concerned with seeking new error free information. The quest for such replacement of doubt by true belief is the common feature that I claim ought to be a feature of every serious effort to change beliefs.

I cannot and shall not pretend to be able to demonstrate that one should heed this recommendation. The best defense that can be offered is an exploration of the ramifications of taking it seriously. And one of the ramifications is endorsement of the commensuration requirement.

How might one violate the commensuration requirement? By justifying a change in state of full belief that is neither an expansion nor a contraction relative to one's goals.

Replacement is one such kind of violation. In Replacement, the inquirer shifts from K_1 to K_2 where K_1 contains h and K_2 contains $\sim h$ for some sentence in \underline{L} . According to the view I favor, such a replacement is never justifiable. By hypothesis, the inquirer is committed initially to the truth of every claim in K_1 . To shift to K_2 is to deliberately change one's commitments so that one comes to fully believe what one initially judged to be false. From the prior point of view relative to which the inquirer justifies the change, such a shift cannot be justified *if the inquirer is concerned to avoid error*. Such a shift deliberately replaces sure truth with certain error.

I claim that such a shift is illegitimate because inquirers should seek to avoid error in justifying changes in full belief (Levi, 1980, ch.3, 1991, ch.4.). There are a great many authors who clearly wish to discount a concern to avoid error as a goal of inquiry. My objection to the legitimacy of replacement will not count with them. All I am claiming now is that those who share the view that one ought to avoid error in changing one's full beliefs will have good reason to deny that replacement is ever justifiable.

The other way to violate the commensurability requirement is by changing from K_1 to K_2 where the latter is a *residual shift* from the former. From the initial point of view relative to K_1 no hypothesis is imported that is certainly false. But hypotheses that might from that point of view be false are imported while *at the same time* hypotheses that are certainly true from that point of view are relinquished.

Now an inquiring agent, so I claim, can be justified from the agent's initial point of view in adding some hypothesis that might be false provided that the agent has an incentive for incurring the risk of error. I contend that the incentive is to be found in the *addition in new information* promised by adding the new hypothesis. In a residual shift, however, such an addition in new information is accompanied by a *loss* of information

that, from the initial point of view, is certainly true. There is no logical obstacle to adding the new information without incurring the loss of old information as in an expansion. Expansion would appear to be preferable to a residual shift.

Perhaps, the promise of new information compensates for the risk of error incurred in adding the new information. This would justify an expansion -- not a residual shift. Perhaps, also, there is a way to justify surrendering the information to be given up in a contraction. If so, that would be a good reason for implementing a contraction. Thus, justifying a residual shift appears to decompose into a justification for an expansion and a justification for a contraction.

Thus, a case can be made for claiming that neither replacement nor residual shift is justified by the concern to obtain valuable error free information. To be sure, my argument depends on the assumption that justifications for changes in states of full belief should be with respect to concern to avoid error and obtain valuable information.

This argument for the commensuration requirement might seem to support the conclusion that contraction is also never to be justified. To be sure, contraction never incurs a risk of error from the point of view of the agent prior to contraction. Every item in K_1 is judged true. As long as no new item is added to K_1 there is no risk of error from that quarter. And removing sentences from K_1 cannot incur error from the inquirer's point of view. What such removal does is reduce the information judged to be true with certainty. Insofar as part of the aim is to add new information, contraction appears to be counterproductive.

I have argued elsewhere (Levi, 1980a, ch.3 and 1991, ch.4) that sometimes inquirers inadvertently expand into inconsistency and are compelled to contract because their belief state can no longer serve to distinguish between what is possibly true and not possibly true. In such cases, there is no need to justify incurring a loss of information. The problem becomes one of deciding what to remove in order to remove inconsistency.

On other occasions, agents may deliberately contract K_1 and incur the loss of information because doing so allows them to give a serious hearing to some important conjecture without begging the question either in favor of or against the conjecture or current doctrine. Such deliberate contraction offers some prospect that through further inquiry new information will be added (and added legitimately) that will yield more valuable information than is available in K_1 . (Levi, 1991, ch.4).

My account of the common features of the aims of inquiry concerned to change states of full belief in justifiable ways is thus intimately tied to my advocacy of the commensurability requirement.

In this essay, my primary concern is with contraction. But a brief review of the approach I favor to expansion and, in particular, to the way information and informational value plays a role in the evaluation of potential expansion strategies may help bring some of the questions pertaining to contraction into focus.

6. Deliberate Expansion

In deliberating as to how to expand his or her belief state, an inquirer ought to seek to obtain new error free information. (Levi, 1967; 1984, ch.5; 1980a, ch.2; 1991, ch.3; 1996.) That is to say, all *proximate goals* of specific attempts to justify changes in state of full belief by adopting one of the available expansion strategies generated by an ultimate partition U_K relative to corpus K share in common a concern to avoid error and to increase information. Of course, the proximate aims of different efforts to expand differ in the kind of information being sought. They may differ with respect to how the importance of different types of information is assessed. And inquirers can disagree

about the risks of error that are to be incurred by instituting different changes in point of view. But there are certain features that the proximate goals of all properly conducted inquiries seeking to expand share in common.

By the proximate goals of a decision problem, I mean all of those value considerations that are relevant to assessing the currently available options as better or worse in reaching a decision as to what to do.

I claim that the proximate goal of any effort at deliberate expansion should assess each potential expansion strategy with respect to two desiderata:

(1) Avoidance of error: On the assumption that every consequence of the inquirer's belief state is true (or that all sentences in \underline{K} are true), the inquirer should be concerned to avoid adding new information that is false. From the point of view of the agent endorsing \underline{K} , adopting \underline{K}^+_h avoids error if and only if h is true. That agent takes for granted that all sentences in \underline{K} are true so that all sentences in \underline{K}^+_h are true if h is true and some are false if h is false.

(2) New information: The inquirer is concerned to eliminate elements of U_k and to add the result to his state of full belief in order to produce new information.

The concern to avoid error is manifested in an evaluation of the potential answers generated by U_k with respect to the risk of error incurred. Recall that a potential expansion strategy is represented by a corpus \underline{K}^+_h where h is a disjunction of a subset of members of U_k . The risk of error incurred by adopting that expansion strategy is equal to the degree of credal or belief probability according to $\sim h$ relative to the state of full belief represented by \underline{K} . Let that probability function be $Q_k(\cdot)$.

Another way to understand Q is as the value of a function (the inquirer's *confirmational commitment*) from potential corpora to credal probability functions. (Levi, 1974, 1980a). Confirmational commitments characterize how the inquirer judges potential answers to be probabilistically supported relative to various bodies of information \underline{K} .

In Levi, 1974, 1980a and 1986, I take the position that judgments of credal probability or probabilistic support are better represented by convex sets of Q -functions. Such convex sets then become values of the functions representing confirmational commitments. For the time being I shall simplify the discussion by assuming that the convex sets are singletons.

Obviously if an inquirer were concerned exclusively with avoiding error, the inquirer would refuse to expand except in the degenerate sense that the agent would expand by adopting the disjunction of all elements of U_k . Relative to \underline{K} the risk or probability of error then incurred would be set at 0. The inquirer judges this strategy without risk of error on the assumption that every consequence of the inquirer's state of full belief is true. Claims that are not items of the agent's belief state but are consistent with it are possibly true and possibly false.

If an inquirer is going to expand his or her state of full belief in an "ampliative" fashion, some risk of error is going to be incurred. To justify taking on such a risk, the inquirer must have some incentive. According to the conception of the common features of inquiry I am propounding, incentive is to be found in the new information that is to be obtained. Inquiry aims to relieve doubt as Peirce rightfully observed. In deciding whether to expand his state of full belief by adding such a claim to the initial state, the agent acquires new information but incurs a risk of error and must balance the benefit of the new information against the risk of error.¹³

¹³Throughout this discussion, I am supposing that the justification of a change in belief state is to be based on the information available to the inquirer and the goals the inquirer is seeking to promote. From the inquirer's point of view, all propositions that inquirer fully believes are true are true. It would be incoherent of the inquirer to

Adopting an expansion strategy \underline{K}_h^+ is a choice that has one of two possible consequences of relevance to the concern to obtain new error free information:

- (a) If h is true, the inquirer has increased the value of the information of the state of full belief \underline{K} to the value of the information carried by \underline{K}_h^+ and has done so (on the assumption that no item in \underline{K} is false) without error.
- (b) If h is false, the inquirer has increased the value of the information obtained but has incurred an error in doing so.

Thus, each of the consequences is represented by a two dimensional vector. The first component $T(h,x)$ specifies whether error was avoided ($x = t$) or incurred ($x = f$) and the second spells out the increment $C(h,x)$ in informational value obtained (which shall be assumed to be the same whether $x = t$ or $x = f$. To evaluate the expansion strategy as compared to others generated by the ultimate partition U_k , the values of the two components need to be aggregated into a single value $V(h,x)$.

A common approach to addressing the aggregation of several dimensions of value is to identify some quantitative method of evaluating components in each dimension and then taking some weighted average of the values of the components. In our case, we need to assign utilities to avoiding error and incurring it and to obtaining information carrying certain kind of value. $V(h,x)$ then equals $\alpha T(h,x) + (1-\alpha)C(h,x)$.

Rehearsing this somewhat more carefully, given the inquirer's initial corpus \underline{K} and ultimate partition U_k , a potential expansion strategy \underline{K}_h^+ is restricted to cases where h is a disjunction of some subset of elements of U_k .

Let $x = t$ if and only if all items in \underline{K} and h are true and $x = f$ if and only if all items in \underline{K} are true and h is false.

Let $T(h, t) = a$ and $T(h,f) = b$ where $a > b$. $T(h,x)$ [$x = t$ or $x = f$] is the utility of \underline{K}_h^+ on the assumption that avoidance of error is the sole concern of the inquirer in expansion.

Let $C(h, t) = C(h,f) = c$ the utility of \underline{K}_h^+ on the assumption that the sole concern of the inquirer in expansion is the maximize informational value.

Let $V(h,x) = \alpha T(h,x) + (1-\alpha)C(h,x) = v$ the utility of \underline{K}_h^+ on the assumption that the aim of the inquirer in expansion is to obtain new valuable error free information.

Assigning utilities to the first component $T(h,x)$ is easy enough. One can assign an arbitrary value a to avoiding error and a value $b < a$ to incurring it. It is convenient to let $a = 1$ and $b = 0$. In that way, we can represent an inquirer concerned only with avoiding error as evaluating expansion strategies in terms of expected utility by equating expected utility with the probability that h is true. Maximizing expected utility is maximizing $Q(h)$ or minimizing the risk of error $Q(\sim h)$ as noted before. An inquirer maximizing expected utility where the utility function is sensitive to the concern to avoid error alone will fail to reject any item in U_k and as a consequence the expansion strategy to adopt will be to refuse to expand at all. So now we have:

$$T(h,t) = 1 \text{ and } T(h,f) = 0.$$

The problem we need to face concerns the second component. How is the value or utility of information (informational value) to be assessed? The current

acknowledge that any one of them is false. The inquirer may coherently acknowledge that he or she might change his or her mind in the future; but from the inquirer's point of view, such a change of mind will involve an exchange of true information for false. If the inquirer seeks to avoid error in changing his or her state of full belief, he will do so on the assumption that replacing his current belief that h by belief that $\sim h$ is to incur an error. If the concern to avoid error is to be consistent with an acknowledgement that it is sometimes legitimate to cease to be certain that h , the concern to avoid error should be restricted to avoiding error at the next change of state. Avoidance of error in subsequent changes should be ignored. (See the discussion of Messianic and secular realism in Levi, 1991, ch.4.)

temporary focus is on the value of information in the context of expansion. But the ultimate target is the value of information in the context of contraction.

7. Undamped Informational Value in Expansion

The information carried by \underline{K}_h^+ can be represented by the set of elements of U_K that are *rejected* in expanding \underline{K} by adding h . One expansion strategy \underline{K}_h^+ carries more information than another \underline{K}_g^+ if and only if the set of elements of U_K rejected by expanding by g is a proper subset of the rejection set for the expansion by h if and only if $\underline{K}_g^+ \subset \underline{K}_h^+$.

The concept of one expansion strategy carrying more information than another partially orders the potential expansion strategies generated by U_K .

The information carried by a potential expansion strategy is one thing. The *informational value* carried by such a strategy is another. Information carried by potential answers partially orders these answers relative to \underline{K} . In assessing informational value, the aim is to construct at the very minimum a weak ordering of the potential expansion strategies. Nonetheless, assessments of informational value should be grounded in judgments of the information carried. This grounding should satisfy the following condition:

Weak Positive Monotonicity: If \underline{K}_h^+ carries more information than \underline{K}_g^+ it carries at least as much informational value.

This condition may be strengthened as follows:

Strong Positive Monotonicity: If \underline{K}_h^+ carries more information than \underline{K}_g^+ it carries more informational value.

Advocates of strong positive monotonicity insist that positive increments of information always should yield a positive gain in informational value. Advocates of weak positive monotonicity insist that sometimes additional information is worthless. Even so, they would insist that if one is concerned with information, the value of new information is not negative.

Assessments of informational value take into account not only how much information potential expansion strategies carry but also such considerations as the explanatory power, simplicity or other attractive (or unattractive) features of potential expansions. A potential expansion that carries more information than another is not required to carry greater informational value for the extra information carried may be of no increased value at all.

The C -function should induce a weak ordering of the potential expansion strategies that satisfies weak monotonicity. But it should be more than a weak ordering. I am supposing here that seeking new error free information requires taking risk of error into account. Maximizing expected utility requires, therefore, that the balance between probability of avoiding error and the informational value that compensates for the risk of error incurred can be given in quantitative terms. This calls for a quantitative representation of informational value. This representation should be a real valued function satisfying the weak ordering that is unique up to a positive affine transformation. That is to say, it should be a utility function.

Moreover, the utility function should be bounded from above and from below. Indeed, if the utility representation for avoidance of error takes two values a and b such that $a > b$, it is useful (but not mandatory) to restrict the values for the C -function to the interval from a to b . In this way, the weight α may be understood to represent the relative importance attached to the concern to avoid error as represented by the T -

function and the concern to maximize informational value as represented by the C-function.

Rejecting an element d_i of U_K increases the informational value of an inquirer's state of information or full belief. Is the increment in informational value obtained the same no matter how many or which other elements of U_K are rejected? Or is there some other relation between the increment and the set of elements of U_K that are also rejected? The issue here concerns the sense in which an inquirer is or should be concerned to obtain new information. The only way to settle this question would be to explore the implications of different proposals. The following assumption (proposed in 1967 in Levi, 1984, ch.5) works very well in the context of expansion:

Constant marginal returns in informational value of rejection: Let the set of elements S of U_K that are rejected in expanding \underline{K} be $R \cup \{d_i\}$. The difference between the informational value obtained by rejecting the elements of S and the elements of R is the same no matter which subset of U_K R happens to be as long as it does not contain d_i .

This increment can be characterized by introducing a measure M that assigns non negative values to the elements of U_K summing up to 1 and such that the M -value of a disjunction of elements of U_K is equal to the sum of their M -values. The increment in informational value in expanding by adding h to \underline{K} is sum of the M -values of elements of U_K that are rejected when the expansion by adding h is implemented. That is to say, it is equal to $M(\sim h) = 1 - M(h) = \text{Cont}(h)$. I shall call this evaluation of the increment in informational value obtained in expansion an assessment of *increment in undamped informational value*.

In Levi, 1967 I proposed a model of deliberate inductive expansion that was a special instance of one where $C(h,x) = \text{Cont}(h)$. In that special case, I required all elements of U_K to be assigned equal M -values.

In Levi, 1984, ch.5 [also first published in 1967], I proposed the condition that $C(h,x) = \text{Cont}(h)$. I took note of the fact that the M -function has the formal properties of a probability function. However, I explicitly denied that the M -function need represent the inquirer's degree of belief function or assessment of how well potential expansions are probabilified by the initial corpus \underline{K} (i.e., by the evidence). The M -function is *informational value determining* and not *expectation determining*.

The C-function is supposed to evaluate the increase in informational value afforded by a deliberate or inductive expansion of corpus \underline{K} as compared to other potential expansions of \underline{K} allowed by the ultimate partition U_K . In the contexts both of expansion and of contraction, it is worthwhile considering the increase in informational value yielded by an expansion of a minimal corpus \underline{LK} relative to the partition U_{LK} . This permits us to assign informational value $C^*(\underline{K})$ to the current corpus \underline{K} . It may be equated with $\text{Cont}^*(\underline{K}) = 1 - M^*(\underline{K})$ where the function M^* is simply the M -function determining increments in informational value from rejecting elements of U_{LK} .

Let h be equivalent given \underline{K} to a disjunction of some subset of U_K . The difference between the informational value of \underline{K} and the informational value of the result \underline{K}^+_h of expanding \underline{K} by adding h and forming the deductive closure is representable by

$$(i) [1 - M^*(\underline{K})] - [1 - M^*(\underline{K}^+_h)] = M^*(\underline{K}^+_h) - M^*(\underline{K}) = 1 - M^*(h)$$

In general, U_K is a proper subset of U_{LK} so that the maximum new informational value to be added to U_K cannot be greater than $M^*(\underline{K})$. Since it is convenient to let the

maximum be equal to 1, divide $1 - M^*(h)$ by $M^*(\underline{K})$ to form the following measure of increment of informational value obtained by adding h to \underline{K} .

(ii) $1 - M(h)$ where $M(h) = M^*(h)/M^*(\underline{K})$ for h that are disjunctions of elements of U_k .

Represented in this way, the assessment of informational value relative to \underline{K} and U_k is derived from an assessment of informational value relative to \underline{LK} and U_{LK} . The advantage of this representation is that it allows for tracing assessment of informational value over a sequence of changes in state of belief where each change brings a state representable as an expansion of \underline{LK} relative to U_{LK} and M^* . As long as these three factors, the minimal contraction \underline{LK} , the basic partition $V_{LK,K} = U_{LK}$ and the informational value determining M -function, M^* relative to U_{LK} are held fixed, for any potential expansion \underline{K} of \underline{LK} , both the ultimate partition U_k and the assessment of increments in informational value by adding information to any expansion \underline{K} allowed by U_k can be determined.

A given inquiry may involve several different expansions in state of full belief. There is no requirement of rationality insisting that evaluations of increments in informational value should be relative to the same \underline{LK} , U_{LK} and M^* . Such coherence in evaluations of increments in undamped informational value is to be expected when (a) the inquirer does not change the range of hypotheses regarded as relevantly entertainable in the given inquiry as defined by \underline{LK} and U_{LK} and (b) holds the M -function relative to \underline{LK} and U_{LK} fixed. In such a case, the inquirer has not altered his or her *demands for information*.

8. Epistemic Utility as the Utility of New Error Free Information:

The V -function representing the utility of new error free information satisfies the following condition:

$$\begin{aligned} V(h,x) &= \alpha T(h,x) + (1-\alpha)Cont(h) \\ &= \alpha T(h,x) + (1-\alpha) - (1-\alpha)M(h). \end{aligned}$$

If we divide the right hand side by α and subtract $q = [1-\alpha]/\alpha$ we obtain the following positive affine transformation of V .

$$V^*(h,x) = T(h,x) - qM(h).$$

Since the ordering of expected utilities of options is not altered by transforming utility functions by a positive affine transformation, we can treat V^* as equivalent to V .

9. Maximizing Expected Epistemic Utility:

Both the V -function and the V^* -function are *epistemic utility functions* representing the cognitive goal of acquiring new error free informational value.

The *expected epistemic utility function* becomes the following:

$$\begin{aligned} EV^*(h) &= Q(h)(1-qM(h)) + Q(\sim h)(0-qM(h)) \\ &= Q(h) - qM(h). \end{aligned}$$

$$= \sum [Q(d_i) - qM(d_i)].$$

The d_i 's are the elements of U_K that are disjuncts in h .

In addition to the assumptions made thus far, a special constraint needs to be imposed on the range of q . I assume that an inquirer seeking new error free information should never prefer to add false hypotheses no matter how informationally valuable rather than avoid error even if doing so adds no new informational value. This condition implies that $\alpha \geq 0.5$ and $0 \leq q \leq 1$. q is called the *index of boldness*.

Assuming that q is restricted to the unit interval in the manner indicated, if there are d_i 's for which $Q(d_i) - qM(d_i) > 0$, then every potential expansion strategy carrying maximum expected utility is a disjunction of all such d_i 's together with some subset of the d_j 's such that $Q(d_j) - qM(d_j) = 0$.

An inquirer concerned to obtain new error free information of value in expansion should, therefore, reject an element d_i of U_K if $Q(d_i) - qM(d_i)$ is negative, refuse to reject it if this difference is positive and may or may not reject it if the difference is 0.

When maximizing expected utility leads to a tie in optimality, a decision-maker is free to choose any one of the optimal options. The decision-maker may, however, undertake to adopt some rule for breaking ties for optimality.

In inductive expansion, it makes sense to adopt a *Rule for Ties recommending the weakest* of the optimal expansions if there is one. As long as the expectation determining probability Q and the informational value determining probability M are determinate, there will be a uniquely weakest option and it can serve as the tie breaker.¹⁴

Maximizing expected utility and the Rule for Ties then leads to the following *rule for inductive expansion*:

Relative to \underline{K} , U_K , Q, M , and q ,

- (1) Reject all and only those elements of U_K such that $Q(d_i) - qM(d_i) < 0$.
- (2) Expand \underline{K} by adding the assertion that one of the unrejected elements of U_K is true.

As already noted, this result was presented for the special case where $M(d_i) = 1/n$ for all n elements of U_K in Levi, 1967 and was generalized in Levi, 1984, ch.5. Further elaborations are found in Levi, 1980 some of which are further discussed in Levi, 1991 and 1996.

10. Stable Inductive Expansion:

Given the expansion yielded at a level of boldness, the rule may be reapplied over and over again until no further elements of U_K can be rejected. The result is the *stable inductive expansion* of the initial \underline{K} at the given level of caution. When $q = 1$, the result of stable inductive expansion is the addition to \underline{K} of the proposition represented by the set of elements of U_K for which $Q(d_i)/M(d_i)$ is a maximum. Stable inductive expansion or acceptance is discussed in Levi, 1980 and 1996.

¹⁴Maurice Pagnucco (1996) has introduced a conception of *abductive expansion* according to which all elements of U_K are evaluated according to some index of explanatory merit. The recommended expansion is the meet of those expansions by adding elements of U_K that rank best. There is no effort to show that the meet or intersection is also optimal since the evaluation of expansion strategies is only of elements of U_K or maximally consistent extensions of \underline{K} . Not only is the proposal I am making one that insists on the optimality of such tie breaking recommendations but, setting this to one side, it is not true that according to my proposal the recommendation is to take the intersection of all expansions by adding elements of U_K carrying maximum EV^* -value. The recommendation is to take the intersection of all such expansions carrying positive EV^* -value.

Notice that stable inductive expansion is defined on the assumption that we may hold U_K fixed and derive the ultimate partitions relative to inductive expansions of \underline{K} as truncations of this one. Not only is this assumed but it is also taken for granted that the assessments of undamped informational value for the iterated expansions are normalizations of the initial M -value. Successive iterated expansions are supposed to cohere in assessments of undamped informational value relative to \underline{K} , U_K and M . Such assumptions seem entirely appropriate make. Inquirers do not alter demands for information without good reason in the context of a given inquiry.

11. Degrees of Surprise and Belief (Plausibility):

Both inductive expansion and stable inductive expansion are heavily context dependent. The legitimacy of an inductive expansion depends on the following contextual parameters:

- (a) The ultimate partition U .
- (b) The information already available in the state of full belief as represented by the corpus \underline{K} .
- (c) The set of potential answers as determined by the ultimate partition U .
- (d) The credal probability function Q defined over the Boolean algebra defined by U , the informational value determining M -function defined over the same domain.
- (e) The index of boldness q (or caution) representing the relative importance attached to the concern to avoid error and to acquire new information of value.

An inquirer who has commitments with respect to \underline{K} , U , Q , and M may wish to appraise the potential answers relative to the various values of q from 0 to 1. Because the values of q have been allowed to take any real value between 0 and 1, for any sentence h that is a Boolean function of the elements of U , if h is rejected for some value of q in this range, we can find a value $q(h)$ such that h is rejected for values of $q > q(h)$. If h is not rejected for any value of q less than or equal to 1, $q(h) = 1$.

We may then offer the following definitions of *degrees of potential surprise* (= degrees of disbelief = degrees of implausibility) and *degrees of belief* (degrees of plausibility).

The degree of potential surprise that h relative to \underline{K} , U , Q , and M
 $= d(h|\underline{K}, U, Q, M) = 1 - q(h)$.

The degree of belief that h relative to \underline{K} , U , Q , and M
 $= b(h|\underline{K}, U, Q, M) = d(\sim h|\underline{K}, U, Q, M)$.

Suppressing the contextual parameters, we can show that potential surprise and degree of belief satisfy the following axioms:

- (1d) If $\underline{K} \vdash g \equiv g'$, $d(g) = d(g')$.
- (1b) If $\underline{K} \vdash g \equiv g'$, $b(g) = b(g')$.
- (2d) $\underline{K} \vdash \sim g$ if and only if $d(g) = 0$.
- (2b) $\underline{K} \vdash g$ if and only if $b(g) = 1$
- (3d) $d(g)$ or $d(\sim g) = 0$
- (3b) $b(g)$ or $b(\sim g) = 0$
- (4d) $d(h \vee g) = \min(d(h), d(g))$.
- (4b) $b(h \wedge g) = \min(d(h), d(g))$.

These axioms for potential surprise (the d - function) were proposed informally in Shackle, 1949 and more formally in the second edition of 1952. Shackle understood these measures to be measures of uncertainty but gave them nothing more than an intuitive interpretation. Other authors have reinvented Shackle's formal measures as measures of degree of belief, support, plausibility or as the dual. Most recently, Gärdenfors and Makinson (1993) have done so. However, none of these authors have seriously improved on Shackle's intuitive interpretation of this kind of measure. The definitions of degrees of potential surprise and belief introduced here were proposed in Levi, 1966 and 1967 and elaborated on in 1980 in Levi, 1984, ch.14. They establish a clear link between potential surprise and degree of belief and inductive expansion and thereby with probabilistic ideas.¹⁵

Many authors think of degrees of belief as subjective probabilities. I do not wish to quarrel over the correct explication of degree of belief; but only to suggest that there is as much presystematic precedent in favor of Shackle's usage as there is for degrees of belief as subjective probability. According to the interpretation of Shackle's idea that I proposed thirty years ago, there is no conflict between using the two ideas. They have different intended interpretations. All we need to do is to avoid confusion due to conflating the two usages. Such confusion is manifest in the tendency of many authors to think of rational inquirers accepting or coming to belief hypotheses when the probability is high enough -- i.e., above a certain threshold. Such authors then proceed to defend violations of deductive closure in one way or another. I insist that this practice is confused. We accept a hypothesis h if $b(h)$ is high enough. The threshold of acceptance is, in effect, an expression of the degree of boldness of the inquiring agent.

12. Probability and Shackle Measures as Indicators of Informational Value:

Assessment of informational value by reference to an informational value determining probability or M -function is motivated by the need to consider a quantitative characterization of the utility or value of information. The task of aggregating informational value and risk of error in a single assessment then calls for the quantitative characterization of informational value.

The idea of using probability based notions of informational value is an old one. Authors like Popper, Carnap and Bar Hillel explored measures of information as inverse functions of probability. Given a probability measure $M(x)$, $1 - M(x)$, $1/M(x)$, $-\log M(x)$ have been suggested as measures of information.

These authors and their epigones equated $M(x)$ with $Q(x)$ -- the credal or expectation determining probability or with the degree of confirmation on the basis of which such credal probabilities are to be derived. The proposal I made in the 1960's departed from this approach by distinguishing between expectation and informational value determining probability. Informational value is probability based in the sense that it is a measure that it is derivable from a measure satisfying structural requirements of the calculus of probability. It does not follow from this that it is derivable from such a probability measure used to represent degrees of credal or conformational probability.

$Cont(x) = 1 - M(x)$ differs from alternative measures of informational value such as $1/M(x)$ or $-\log[M(x)]$ because it satisfies based on the assumption of constant returns in informational value of rejection and has a finite upper and lower bound. These are the features that suit it well for the sense of informational value I am deploying when I claim that inquirers are concerned to acquire valuable new information. As we have seen,

¹⁵L.J. Cohen (1977) has presented the only serious alternative to my own proposal for explicating Shackle measures as degrees of belief. He seeks to derive them from measures of inductive support of a still different kind. I have discussed Cohen's views in Levi, 1984, ch.14.

using this measure together with the T -function representation for the utility for avoiding error has yielded a simple criterion for choosing inductive expansion strategies. I have called assessments of informational value of this kind *undamped assessments*.

In the subsequent discussion, we shall have occasion to consider the role of appraisals of undamped informational value in connection with the evaluation of contraction strategies. In contraction, information is lost and it is desirable to minimize loss of valuable informational value. The question arises as to whether loss of informational value should be assessed using $Cont(x) = 1 - M(x)$ as a measure of informational value. As we shall see, there are reasons in favor of the idea and stronger reasons against. Consequently, we shall want to explore ways of evaluating information alternative to undamped assessments. It turns out that the most promising approach uses measures formally similar to Shackle measures. And just as informational value determining probability ought not to be confused with expectation determining probability so too the use of Shackle measures to represent assessments of losses of *damped informational value* should not be confused with their use in representing degrees of belief and potential surprise. Shackle measures, like probability measures, have many different uses.

13. Cardinality:

I do not mean to insist that evaluations of consequences of options and of probabilities must be assessed numerically in order to evaluate options in decision problems in general or in assessing inductive expansions strategies in particular. For a variety of reasons (Levi, 1974, 1980, 1986), rational agents may very well not have quantifiable evaluations of consequences, probability judgment or expected utility. According to the view I favor, utilities (and probabilities) should be representable by sets of cardinal utility (probability) functions satisfying a certain convexity condition. (See Levi, 1980, p.78.)

In inductive expansion, this means that the inquirer's *credal state* or state of probability judgment for hypotheses in U_K may be representable by convex set of probability distributions (Q -functions) rather than a singleton. The cognitive or epistemic values of consequences of potential expansion strategies may be represented by convex sets of V -functions. Any EV^* -function obtained from coupling a Q -function in the credal state with a V -function can be used to identify a set of *E-admissible* options that are best according to that EV^* -function). This technique allows for a wider variety of relevant modes of evaluating consequences and probabilities than can be achieved by exclusive focus on *either* cardinal assessments *or* on ordinal assessments.

It is a mistake to think that fewer hostages are committed to fortune by restricting attention to ordinal matters as compared to quantitative issues. A weak ordering of consequences with respect to value is compatible with the claim that the agent's evaluation of consequences is representable by all utility functions preserving the weak ordering and also with the claim that it is representable by some proper (convex) subset of these (including singleton subsets). When uncertainty is taken into account, it turns out that these differences in representability can be relevant to the recommendations as to which option an agent should choose. Someone who insists on ordinal representations is tacitly denying the relevance of the distinctions that are ignored and this may be a quite substantive assumption.

Expansion or contraction in cases where epistemic utility goes indeterminate (as will be common in real life) will not be discussed in this essay. The cardinal representability case may often (but not always) be unrealistic. It is relevant to real life cases, however, because the general constraints proposed for cardinal epistemic utility functions can be construed (legitimately I believe) as constraints on the utility functions that are eligible for membership in a convex set of permissible epistemic utility functions.

In the context of inductive expansion aimed at seeking new error free information, every permissible epistemic utility function should be (a positive affine transformation of) a weighted average of the T -function representing the utility of avoiding error and one of a convex set of permissible utility of information or *cont*-functions. Every permissible expected epistemic utility function should be the weighted average of some permissible credal or expectation determining probability Q and a permissible *cont*-function.

14. Undamped informational value in contraction

At the end of section 5, I suggested that inquirers contract either because they are compelled to do so to remove a conflict in a state of full belief or to give a hearing to some interesting conjecture. And in suppositional reasoning, one may seek to contract for the sake of the argument in order to allow for consistent expansion by adding the supposition. In all of these contexts, we need to consider the following problem. Given that contraction by removing h from \underline{K} is contemplated, what is the best contraction meeting this constraint to adopt?

The need to face this question is especially obvious in suppositional reasoning. Given that one is to suppose that h is true for the sake of the argument, in those contexts where \underline{K} entails $\sim h$, supposition requires that one remove $\sim h$ from \underline{K} and then expand by adding h . But in removing $\sim h$, one must decide among many potential contraction strategies.¹⁶

The same issue arises in coerced contraction. Here the inquirer faces the need to contract from inconsistency. Doing so confronts the inquirer with the following problem: Given a best contraction of \underline{K} removing $\sim h$, should one expand that contraction by adding h , restoring $\sim h$ or by remaining in suspense. In any of these cases, the result can be viewed, of course, as a contraction of the inconsistent corpus \underline{K}_1 . But the contraction of the inconsistent corpus is equivalent to the result of an expansion of a contraction of \underline{K} removing $\sim h$. So once more, we need to consider how to specify what is an optimal contraction removing an item from \underline{K} .

Finally, in seeking to afford a hearing for an interesting conjecture h , $\sim h$ needs to be removed from \underline{K} and once more our question becomes urgent.

In contraction removing $\sim h$ from \underline{K} , avoidance of error cannot be an issue for precisely the same reason that it is an issue in the case of inductive expansion. Since the inquirer is committed to being certain that his initial belief state is free of error, from the inquirer's point of view, the only way a falsehood can be imported is by being added to the agent's belief state. Contraction imports nothing. On the assumption that every item in \underline{K} is true, contraction cannot incur error.¹⁷

In contraction that implements genuine belief change where the agent loses information and in contraction for the sake of the argument, the concern to obtain new

¹⁶ According to the view of consistency preserving supposition I favor, this is true also when the supposition is belief conforming where \underline{K} entails h . In that case, the supposition that h must be removed from \underline{K} just as in the belief contravening case.

¹⁷ Since every item in the initial corpus is judged to be true by the inquirer prior to contraction, removing an item from the corpus cannot, from the inquirer's point of view, incur error. As noted in footnote 8, rationalization of a change of belief by contraction should be based on the inquirer's point of view prior to the contraction whether the change is an expansion or a contraction. Moreover, according to the view expressed in footnote 8, the concern to avoid error is focused on avoiding error at the next step (i.e., contraction or expansion). It could happen in genuine belief change, contraction removing h might lead in subsequent inquiry to coming to full belief that h is false. In rationalizing contraction, such a long term concern to avoid error should not be present -- counter to the messianic realism of Peirce and Popper. A secular realism that insists on a concern to avoid error at the next change but not two or more changes down the sequence of belief changes can consistently endorse the legitimacy of fully believing that h at time t while at the same time recognizing the possibility that h will be given up legitimately at some subsequent stage in inquiry. (Levi, 1991.)

CONTRACTION AND INFORMATIONAL VALUE

error free information is frustrated and the best the inquirer can do is to minimize the loss of valuable information that is incurred.

Whether contraction is part of a genuine change in view or is merely suppositional, therefore, the first order of business is to offer an account of informational value and loss of informational value in a sense elaborating the idea that loss of informational value is to be minimized.

According to the proposal in section 7, the demands for information instituted by an inquirer at a given time are representable by $\underline{L}K$, U_{LK} , and M^* defined over the algebra generated by U_{LK} .

If U_{LK} has n members, there are 2^n potential expansions of $\underline{L}K$ relative to U_{LK} . An informational value determining M -function M^* can be defined for these potential expansions such that for any disjunction x of elements of U_{LK} , $M^*(\underline{K}^+_x)$ is well defined. It represents the increase in informational value accruing from adding x to $\underline{L}K$ relative to U_{LK} and M^* . As noted before, M^* can be used to define the undamped increment in informational value by adding h to \underline{K}^+_x . In that case, however, it is useful to normalize the measure by letting $M(h) = M^*(h)/M^*(\underline{K}^+_x)$.

Potential contractions of \underline{K} are generated from the dual ultimate partition $\underline{U}_K = \{U_{LK}/U_K\}$. Each potential contraction can be evaluated with respect to the undamped informational value relative to U_{LK} with the aid of M^* just as potential expansions of \underline{K} can. Hence, if there is no change in the demands for information in successive expansions and contractions, we should expect a coherence both in evaluations of gains in undamped informational value and losses in undamped informational value.

To be sure, there is no special need to normalize the M -values of potential contractions to values between 0 and 1. Moreover, if we wish to compare potential contractions of \underline{K} with potential expansions, it is useful to utilize the function M^* rather than a normalization for the purpose of evaluating gains in undamped informational value.

Although each potential contraction of \underline{K} has an informational value equal to its undamped informational value relative to $\underline{L}K$ and U_{LK} and this is equal to $cont^*(x)$ where the contraction is the expansion of $\underline{L}K$ by adding x , the *loss of undamped informational value* incurred by shifting from \underline{K} to a potential contraction \underline{K}' of it is equal to $cont^*(\underline{K}) - cont^*(\underline{K}') = M^*(\underline{K}') - M^*(\underline{K})$. This difference is equal to the sum of the M -values of the cells in U_{LK} that are eliminated in moving from \underline{K}' to \underline{K} .

The assessment of losses in informational value incurred in potential contractions of \underline{K} like assessments of gains in informational values yielded by potential expansions of \underline{K} is derived from an evaluation of content or undamped informational values of potential expansions of $\underline{L}K$ generated relative to U_{LK} . Both the assessments of losses in contraction and of gains in expansions have been represented quantitatively.

Quantitative assessments in evaluating expansion strategies are needed because possibility of error has to be taken account so that the problem of deciding how to expand becomes a problem of decision making under risk. The problem of contraction differs in this respect. In choosing among potential contractions, there is no risk of error to be incurred. The relevant consequences of a contraction strategy (the information given up) are known with certainty. In such a setting, quantitative assessment appears to be gratuitous. We could rest content with an ordinal evaluation of losses of informational value.

Nonetheless, as long as the demands for information remain constant, we should expect that the evaluation of losses in undamped informational value incurred in contraction from \underline{K} should be derived from the same $\underline{L}K$, U_{LK} , and M^* as is the M -function relative to \underline{K} in the context of expansion. It is true that the inquirer needs to take into account only the weak ordering of contraction strategies determined by the

quantitative assessment of losses in undamped informational value. But the weak ordering should cohere with the demands for information by being derivable from \underline{LK} , U_{LK} and M^* .

In Levi, 1980, pp.61-62, I asserted that the aim of contraction should be “to minimize the loss of informational value resulting from contraction subject to the constraint that the need occasioning the demand for contraction be satisfied.” Here the constraint can be that h be removed from \underline{K} . And the understanding of informational value used was that explained in Levi, 1980, ch.2.4. It is the understanding of informational value relevant to expansion given above. Relative to an ultimate partition, a probability measure M is defined over the cells and their Boolean compounds. The informational value of expanding \underline{K} by adding h where h is a disjunction of elements of U_K is $Cont(h) = 1 - M(h)$ or $M(\sim h)$ which, given the truth of \underline{K} is equal to the total probability of the rejected elements of U_K .

This is the understanding of informational value I deployed in Levi, 1980. I called the 1980 conception “probability based” informational value in Levi, 1991 and “undamped informational value” in Levi, 1996. It was intended to obey what I am now calling the undamped coherence constraint on losses of informational value. As I have explained the idea here, the losses in undamped informational value by removing h from \underline{K} are to cohere with assessments of gains in undamped informational value relative to some minimal contraction \underline{LK} , ultimate partition U_{LK} relative to \underline{LK} and probability function M^* defined over the Boolean algebra generated by U_{LK} . The function M used in Levi, 1980, ch.2.4 for the algebra generated by U_K is the normalization of M^* through division by $M^*(K)$.

Undamped informational value so conceived fits in well with a conception of the *information* carried by a belief state. The belief state represented by corpus \underline{K}_1 carries at least as much information as the belief state represented by corpus \underline{K}_2 if and only if $\underline{K}_2 \subseteq \underline{K}_1$. Potential states of full belief represented in \underline{L} can be partially ordered by set inclusion and, hence, with respect to information carried. Many potential states will, of course, be noncomparable with respect to the information carried. An assessment of *informational value* extends the partial ordering with respect to information to a total ordering by introducing a way of comparing noncomparable states. If the focus of concern is an assessment of the informational values of potential expansions of \underline{K} relative to an ultimate partition, this extension will involve making evaluations of the cells of the ultimate partition. Such comparisons might involve taking into account the simplicity in some sense of the elements of ultimate partition, explanatory power in some sense or any other desiderata that reflect the *demands for information* of the inquirer. These comparisons can be made quantitative by assigning to the elements of U_{LK} numerical values (M -values) summing up to 1 represented by the function M^* and assigning all other potential answers M -values equal to the sums of the M -values of elements of U_{LK} that entail them. The undamped informational value of a potential expansion \underline{LK}^+_h is then equal to $Cont^*(h) = 1 - M^*(h)$. In the context of contraction, a potential contraction of \underline{K} relative to U^*_K carries undamped informational value equal to the undamped informational value of that corpus relative to \underline{LK} and U_{LK} just as the undamped constraint on losses of informational value requires.

Measures of undamped informational value preserve the partial ordering of corpora or belief states with respect to information in the following weak sense already enshrined in the weak monotonicity condition. If \underline{K}_1 carries more information than \underline{K}_2 , \underline{K}_2 does not carry more undamped informational value than \underline{K}_1 . This requirement is, in effect, the weak positive monotonicity principle of section 7.

In Levi, 1991 and 1996, I did not require that \underline{K}_1 carry more undamped informational value than \underline{K}_2 . I allowed for the possibility that the increments in informational value obtained by rejecting some cell in U_{LK} may be null even though there is an increase in information. This can happen when, according to the inquirer's

demands for information, the increment in information is of no value (and, hence, the cell in question carries 0 M -value). Allowing for this possibility is surely coherent; but it now seems to me that it goes against the spirit of the overall approach to inductive expansion I favor.

Suppose that d is a cell in U_{LK} such that $M^*(d) = 0$. This implies that for any expansion \underline{K} of \underline{LK} relative to which the truncation U_K of U_{LK} contains d , it is not possible to reject d via inductive expansion no matter what $Q(d)$ might be. An inquirer who endorsed assessments of informational value of this kind would have rendered d immune to rejection relative to all probability distributions over U_K and would in this sense have rendered the testing of d pointless. It is true that if d were a hypothesis specifying the precise real value of some parameter and U_K consisted of d and its negation, $M(d)$ would plausibly be assigned the value 0 in many contexts. However, in investigations concerning the values of such parameters where this is so, the ultimate partition relative to \underline{K} should be considered as containing infinitely and, indeed, noncountably many cells. Such problems call for special treatment of the sort sketched in sections 23-25. In this discussion, I have restricted attention to cases where both U_{LK} and U_K are finite. In those cases, it seems plausible to require $M^*(d)$ and, hence, $M(d)$ to be positive.

This objection to allowing $M(d)$ to be 0 is based on the ramifications of doing so for inductive expansion. According to the approach I am developing here and already endorsed in Levi, 1991 and 1996, if d is in the dual ultimate partition $U^*_{K'}$ of an expansion $\underline{K'}$ of \underline{LK} , it should follow that $M^*(d)$ be positive. Yet, this requirement stands at odds with the view I endorsed in Levi, 1991 and 1996 in order to give an account of contraction as violating the Recovery Condition. I still think requiring the Recovery Condition is mistaken; but my previous theoretical rationale for this view is also mistaken.

Those who insist that all cells carry positive M -value adopt a *hyper-undamped* assessment of informational value. Undamped informational value whether hyper-undamped or not does preserve the partial ordering with respect to information in the weak sense just stated. In this weak sense the conception of undamped informational value "fits in" with the conception of information. In order to explain the views I took in Levi, 1991 and 1996, I shall, for the time being *avoid* requiring that undamped assessments of informational value be hyper-undamped even though allowing assessments of informational value to be undamped without being hyper-undamped now seems to me to be implausible.

15. Why losses in undamped informational value ought not to be minimized.

As already stated, in Levi, 1980, I adopted an evaluation of undamped informational value for the purpose of assessing contraction strategies.

In the later essays, I *abandoned* my endorsement of undamped informational value on the grounds that minimizing loss of undamped informational value leads to recommending a saturatable contraction removing h from \underline{K} .

The objection to doing so is that if one subsequently expands by adding $\sim h$ to the contraction, the result is going to be expanding to a cell in the dual ultimate partition $U^*_{K'}$. Mandating this kind of result in every case is clearly absurd. For those who share my view that this result is unacceptable (and this includes the authors of AGM), minimizing loss of informational value in contraction cannot be minimizing loss of probability based or undamped informational value.¹⁸

¹⁸ I had not thought through the implications of my 1980 proposal at that time and was only prompted to do so when I saw a prepublication version of the classical AGM paper of 1985 and the several papers of Alchourrón and

Since the issue is relevant to understanding alternative ways of understanding the injunction to minimize loss of informational value, it will be well to elaborate on details somewhat more explicitly.

Consider the minimal contraction $\underline{L}\underline{K}$ and that ultimate partition $U_{\underline{L}\underline{K}}$ relative to $\underline{L}\underline{K}$. I shall assume that $U_{\underline{L}\underline{K}}$ is finite here. Let us be given an M function, M^* defined for Boolean combinations of elements of $U_{\underline{L}\underline{K}}$. The corpus \underline{K} is an expansion of $\underline{L}\underline{K}$ which, as noted before, may contain sentences additional to those in $\underline{L}\underline{K}$ are not expressible as Boolean combinations of elements of $U_{\underline{L}\underline{K}}$. However, in the context in which contraction of \underline{K} by removing h is being assessed, the probability based or undamped informational value of \underline{K} is equal to the M value of the *disjunction* of those elements of $U_{\underline{L}\underline{K}}$ that are *rejected* when \underline{K} is endorsed. That is to say, it is the undamped informational value of the disjunction of all elements the dual ultimate partition $U_{\underline{K}}^*$. Any potential contraction strategy removing h from \underline{K} will be tantamount to shifting to a corpus \underline{K}' whose ultimate partition $U_{\underline{K}'}$ is a superset of the ultimate partition $U_{\underline{K}}$ for \underline{K} and where at least one element of $U_{\underline{K}'}$ implies $\sim h$.

The loss of undamped informational value incurred by contracting from \underline{K} to \underline{K}' is going to be $\text{cont}^*(\underline{K}) - \text{cont}^*(\underline{K}') = 1 - \sum_{d \in U_{\underline{K}}} M^*(d) - 1 + \sum_{d \in U_{\underline{K}'}} M^*(d)$. This may easily be seen to be equal to the sum of the M -values those elements of the dual ultimate partition $U_{\underline{K}}^*$ for \underline{K} that are shifted to the ultimate partition $U_{\underline{K}'}$ for \underline{K}' . The greater this sum, the greater the loss incurred.

Let us now consider the ramifications of seeking to minimize losses of undamped informational value. I shall consider a series of special cases first.

Special Assumption 1 about M^ :* All elements of $U_{\underline{K}}^*$ carry positive M -value.

An immediate corollary is, of course that all elements of $U_{\underline{K}}^*$ carry positive M -value. But then it follows that no more than one element of $U_{\underline{K}}^*$ should be shifted to $U_{\underline{K}'}$ if shifting from \underline{K} to \underline{K}' minimizes loss of informational value among contraction strategies. Moreover, the element shifted must be one entailing h . Hence we have the following:

Observation 1: If we are to minimize loss of undamped informational value in removing h from \underline{K} and if the special assumption 1 about M^* is applicable, all optimal contractions removing h from \underline{K} must be maxichoice and, hence, saturatable.

If two or more maxichoice contractions are optimal as is possible, we cannot under special assumption 1 recommend using the meet of these maxichoice contractions. *The meet of the optimal maxichoice contractions must be suboptimal.* Given the goals of contraction as stated, partial meet contractions in the sense of AGM are forbidden as forbidden.

Special Assumption 2 about M^ :* All elements of $U_{\underline{K}}^*$ entailing $\sim h$ carry positive M -value. Some entailing h carry 0 probability.

Observation 2: If special assumption 2 about M holds, then all optimal contractions are either maxichoice or intersections of a single optimal maxichoice contraction with elements of $U_{\underline{K}}^*$ that entail h and carry 0 M -value. The latter type of contraction is saturatable.

Special Assumption 3 about M^ :* Some elements of $U_{\underline{K}}^*$ entailing $\sim h$ carry 0 M -value and some entailing h do so as well.

Observation 3: If special assumption 3 about M^* holds, the optimal maxichoice contractions are precisely those shifting single cells entailing $\sim h$ and carrying 0 M -value from $U_{\underline{K}}^*$ to $U_{\underline{K}'}$. Optimal saturatable contractions exist consisting of the

Makinson of that period. My objection to the 1980 proposal has nothing to do with its being a quantitative notion of informational value but rather with the structural properties of that measure.

CONTRACTION AND INFORMATIONAL VALUE

intersections of \underline{K} and such singletons and intersections of cells from U^*_K entailing h and carrying 0 M -value. Any intersection of \underline{K} with a subset of such optimal saturatable contractions is also optimal.

Special Assumption 4 about M^ :* Some elements of U^*_K entailing $\sim h$ but none entailing h carry 0 M -value.

Observation 4: Under special assumption 4, all optimal contractions are either maxichoice or intersections of sets of optimal maxichoice contractions (that is to say, partial meet contractions).

In Levi, 1991, I restricted attention to cases characterized by special assumptions 1 and 2 about M^* (pp.124-5) and for those cases concluded correctly that minimizing probability based informational value requires choosing a saturatable contraction while forbidding intersections of optimal saturatable contractions. This seems as absurd as mandating maxichoice contractions as special assumption 1 alone does.

I neglected special assumptions 3 and 4 where, indeed, intersections of optimal saturatable contractions can be optimal and in the case of 4, partial meets of optimal maxichoice contractions can. At the time, it seemed to me that the only interesting cases of contraction from \underline{K} removing h would be those where removing h incurred a positive loss of undamped informational value as under special assumptions 1 and 2. But strictly speaking, the implications of assumptions 3 and 4 ought to be taken into account. Of course, if we require assessments of informational value that are hyperundamped as seems desirable in inductive expansion from finite ultimate partitions, only special assumption 1 need be considered.

Levi, 1980, p.62 argued that if two or more contraction strategies are optimal, they should "if feasible" be implemented jointly. The qualification "if feasible" should have been if the "joint implementation" is optimal or "admissible" in the sense of not ruled out for choice relative to the goals of the decision problem. The recommendation is to adopt the contraction that is the intersection of all optimal contractions -- i.e., those minimizing undamped informational value.

When two or more options are tied for optimality with respect to the goals of the decision maker, the decision maker can, without betraying his aims, invoke a secondary standard of evaluation to "break ties". Some think of principles recommending giving preferential treatment to minorities in hiring for jobs in just these terms. If two candidates for a job are equally qualified but one comes from a disadvantaged group, the principle recommends breaking ties in favor of the disadvantaged. The choice of such secondary value commitments (there could be tertiary, etc. principles) need not be made. And if it is, it is an issue of values and, perhaps, of morals and politics. However, so is the agent's choice of his primary goals and value commitments.

I am supposing as a characterization of common features of the goals of inquiries focused on expansion that they seek to obtain new error free information. The common features of contraction focus on minimizing loss of informational value.

In both expansion and contraction, there is one principle I have always favored as a secondary criterion. When two or more options tie for optimality one should adopt the intersection of all of them. In Levi, 1980, I endorsed that principle in a form that did not insist that the "tie breaking" option be optimal as well against what were otherwise my own principles. Thanks to the entirely just admonitions of Teddy Seidenfeld, I subsequently (Levi, 1986) came to insist that tie breaking criteria must recommend optimal or, more generally, E-admissible options.

The motivation for urging suspense whenever otherwise admissible as a secondary criterion derives from the idea that when the decision maker cannot rule out two or more cognitive options in expansion and in contraction, the agent should move to

a state of doubt with respect to the unsettled issues. In Levi, 1967 in the context of expansion, I called such a principle a “Rule for Ties”. I shall use the same nomenclature in the case of contraction as well.

The Rule for Ties is a proposal for a secondary criterion for breaking ties in cognitive decision problems. This secondary criterion is a value judgment. It reflects an attitude I think important in efforts to adopt a sensible balance between skepticism and the anxiety to relieve doubt. When all other relevant cognitive decisions cannot adjudicate between alternative cognitive options, suspension of judgment is to be recommended *provided* suspension of judgment is an optimal or admissible option given the inquirer’s primary cognitive goals.

But as I have been going to great lengths to emphasize, one cannot and should not invoke the Rule for Ties unless the intersection of optimal cognitive options is optimal (or, more generally) is E-admissible.

AGM agrees with the idea that one should adopt a “skeptical” stance vis a vis potential maxichoice contractions that tie for optimality. And Gärdenfors (1988) registers some sympathy for the idea that the value of maxichoice contractions resides in its informational value. But AGM fails to offer an account of how intersections of maxichoice contractions are to be evaluated and registers no concern to insure that the intersection or “meet” of optimal contractions that they recommend is itself optimal. To secure this, it is necessary to extend the “preference” over potential maximal contractions to intersections of such contractions. AGM fails to do this. In this respect, AGM is not sufficiently committed to a decision theoretic approach to contraction.

Lindström (1990) and Rott (1993) show how “preferences” over maxichoice contractions that are candidates for being ingredients in partial meet contractions can be characterized within the framework of revealed preference theory. However, they do not include intersections of maxichoice contractions among the options in the domain over which choice functions are defined. Yet, both Lindström and Rott claim that according to the versions of AGM they are exploring, partial meet contractions are recommended. As a consequence, partial meet contractions ought to be considered options in AGM theory. If that is so, the preference relation ought to be defined for partial meet contractions. Lindström and Rott fail to extend their choice function based notion of preference to the domain of partial meet contractions. As I have argued, the option set should be extended even further so that intersections of saturatable contractions are included in the preference ranking.

The same neglect is revealed in the use of Grove models that do well in representing a weak ordering of states or worlds and representing various types of contractions in a geometric picture but fail totally to incorporate sets of points, states or worlds into the preferential ordering.

The idea of using a probability based measure of undamped informational value to achieve this end is the first that comes to mind. If we understand attempts to minimize loss of informational value as attempts to minimize loss of probability based or undamped informational value and combine this recommendation with Rule of Ties, we get the following:

Main Claim A: Attempts to minimize loss of undamped informational value in contraction of \underline{K} by removing h supplemented by the Rule for Ties have the following consequences:

Under special assumption 1, the contraction *must be* a maxichoice contraction even if there are several optimal maxichoice contractions. Recovery is mandated.

Under special assumption 2, the contraction *must be* a saturatable contraction even if there are several optimal saturatable contractions. Recovery is violated.

Under special assumption 3, the contraction *must be* the intersection of all optimal saturatable contractions and must incur 0 loss of undamped informational value. Recovery is violated.

Under special assumption 4, the contraction *must be* the partial meet contraction that is the intersection of all maxichoice contractions carrying loss of undamped informational value and must itself incur such 0 loss. Recovery is mandated.

In sum, if a positive loss of undamped informational value is incurred by an optimal contraction removing h from \underline{K} (so that either special assumption 1 or 2 obtains), minimizing loss of undamped informational value combined with the Rule for Ties entails choosing an optimal saturatable contraction. If a 0 loss of undamped informational value is incurred (so that special assumptions 3 or 4 obtain), intersections of all saturatable contractions incurring 0 loss in undamped informational value is mandated.

So it turns out that the only way we can require conformity with the Recovery Condition without recommending maxichoice contractions is by enforcing special assumption 4. But given that the loss of informational value is loss of undamped informational value derivable from \underline{LK} , U_{LK} and M^* , special assumption 4 cannot hold for all \underline{K} that are expansions of \underline{LK} . Some such \underline{K} , \underline{K}_1 , is the intersection of all elements of U_{LK} that entail h . Let \underline{K}_2 be the intersection of all elements of U_{LK} that entail $\sim h$. $U_{K_1} = U_{K_2}^*$ and $U_{K_2} = U_{K_1}^*$. Special assumption 4 cannot apply both to removing h from \underline{K}_1 and to removing $\sim h$ from \underline{K}_2 . Requiring partial meet contractions that are not maxichoice for contractions of all potential expansions of \underline{LK} and are to satisfy the undamped constraint on informational value relative to \underline{LK} , U_{LK} and M^* is not feasible.

There is another more direct way to reach the conclusion that enforcing partial meet contractions is untenable given the injunction to minimize undamped informational value. For any expansion \underline{K} of \underline{LK} , U_{K}^* is a subset of $U_{LK} = V_{LK,K}$. It is the part of U_{LK} consisting of elements inconsistent with \underline{K} . Relative to \underline{LK} , at least some elements of U_{LK} must carry positive M -value when U_{LK} is finite. Hence, there must be an expansion of \underline{LK} , \underline{K}'' such that all elements of $U_{K}''^*$ carry positive M -value. For such cases, contractions removing any extralogical sentence from \underline{K}'' whose negation is entailed by at least one element of $U_{K}''^*$ must incur positive loss of undamped informational value and, indeed, must do so under the conditions of special assumption 1. In those cases, any such contraction must be maxichoice. That is sufficient reason, so it seems, to call into question minimizing loss of undamped informational value.

16. Informational Economy and Plausibility.

In spite of this argument, it remains the case that in non-degenerate contraction, one shifts from a corpus \underline{K} to another less informative corpus \underline{K}' in the sense that \underline{K}' is a proper subset of \underline{K} . In keeping with the Peircean observation that in inquiry we seek to remove doubt, such loss of information should, in general, be seen as incurring a loss or, at least, as yielding no gain. Moreover, the extent of the loss should be assessed as a loss of valuable information. In contraction one ought to minimize the loss of *informational value* incurred. For reasons already explained, in Levi, 1991 and 1996, I abandoned the recommendation to minimize undamped informational value in contraction that I favored in 1980. But I did not give up the view that the goal being promoted in choosing between potential contraction strategies is minimizing loss of informational value in a sense other than minimizing loss of undamped informational value. The problem was to identify the sense.

Even without complete characterization, the notion of informational value has some import. Relative to a dual ultimate partition U^*_K , the set of potential contractions removing h from \underline{K} is well defined. Such a set of potential contraction strategies can always be partially ordered with respect to loss of information. That partial ordering stipulates that \underline{K}_1 carries less information than \underline{K}_2 if and only if $\underline{K}_1 \subset \underline{K}_2$.

A weak ordering of potential contractions removing h from \underline{K} with respect to loss of *informational value* as an extension of such a partial ordering of the same set of potential contractions with respect to loss of *information*. There are two interpretations of an extension of a partial ordering to a weak ordering relevant here.

Strict positive monotonicity requires that potential contraction \underline{K}_1 incur a greater loss of informational value than potential contraction \underline{K}_2 if $\underline{K}_1 \subset \underline{K}_2$. Any assessment of informational value meeting this requirement is equivalent to an assessment of undamped informational value under special assumption 1. Minimizing losses of informational value so conceived must recommend maxichoice contractions.

I understand a weak ordering of the potential contractions from \underline{K} to be an assessment of losses of informational value provided it satisfies a weak positive monotonicity condition:

\underline{K}_1 incurs at least as great a loss as \underline{K}_2 if $\underline{K}_1 \subset \underline{K}_2$.¹⁹

Assessments of losses of undamped informational value satisfy weak positive monotonicity. But they are not the only weak orderings of potential contractions of \underline{K} that do so.

I have also defended, since 1967 in the case of expansion and since 1980 in the case of contraction, a Rule for Ties which states that given the set of optimal expansion (contraction) strategies, one should always choose the weakest of them if it exists. Such a Rule for Ties does not evaluate the weakest expansion (contraction) strategy as better than the others.

Because the weakest expansion (contraction) is tied for optimality with the other optimal strategies, it is equiprefered to them. The Rule for Ties assesses the optimal strategies for some property other than the loss of informational value it incurs. When two or more expansion (contraction) strategies are optimal, we should choose an optimal strategy in a manner that minimizes controversy. We should choose the weakest.

Minimizing controversy is not a primary aim of expansion (contraction). But when the primary aim fails to render a verdict, we are free to choose among the optimal options according to some other standard. I am urging adoption of a secondary tie breaking standard that minimizes controversy. The Rule for Ties captures that idea.

In the case of contraction, the Rule for Ties recommends choosing the weakest potential contraction minimizing loss of informational value (in the sense as yet to be explained). Given that assessments of informational value must satisfy the condition of weak positive monotonicity, we know that the weakest potential answer in this sense cannot carry more informational value than other potential contractions minimizing loss of informational value unless there are no others.

A weak ordering of potential contractions satisfying weak positive monotonicity has embedded in it a weak ordering of the cells of the dual ultimate partition U^*_K or of the maxichoice contractions obtained with their aid. We may seek a rule for deriving a weak ordering of all the potential contractions of \underline{K} removing h or a real valued evaluation representing that weak ordering from the weak order of maxichoice

¹⁹ Section 7 discusses weak positive monotonicity in the context of inductive expansion. Section 14, p.26 discusses it in the context of contraction. The weak monotonicity condition is introduced as a constraint on assessments of informational value in Levi, 1991, p.122.

CONTRACTION AND INFORMATIONAL VALUE

contractions. Or we may seek to derive a numerical representation of the ordering of the set of potential contractions from a numerical representation of the ordering of the maxichoice contractions.

Thus, if we assign M^* -values to the elements of U^*_K or the associated maxichoice contractions, we can determine M^* -values for all potential contractions and, hence, order potential contractions with respect to loss of informational value. The result is the discredited undamped measure of losses of informational value.

Pagnucco and Rott (1997) begin with a weak order of the maxichoice contractions. However, they are not interested in deriving a quantitative evaluation or even a weak ordering of all potential contractions. They aim to identify or "select" a recommended contraction strategy. So they do not urge inquirers to minimize loss of informational value.

As shall emerge, the recommendation they make is identical with the recommendation favored when one seeks to minimize loss of damped informational value of type 2 as explained in section 18 and then applies the Rule for Ties. However, they are not interested in showing that their recommendation is one that is best among all the available options relative to a specified goal or evaluation of potential contraction strategies as better or worse. So the principles they invoke for recommending a contraction strategy on the basis of a weak ordering of the maxichoice contractions differ from and are motivated differently from the principles I use in reaching the same recommendation.

Pagnucco and Rott survey several methods for deriving recommendations of potential contractions from evaluations of maxichoice contractions. They begin with the Principle of Economy:

Keep Loss to a Minimum. (p.4)

Pagnucco and Rott fail to explain what they mean by a loss. I conjecture that they mean that in choosing a contraction removing h from \underline{K} , one should not adopt a contraction that allows more cells to be intersected with \underline{K} than is necessary in order to obtain such a contraction. One should minimize loss of *information*. Comparison of potential contractions with respect to loss of information leads to a partial ordering of the potential contractions. No potential contraction incurs minimum loss of information. I take it that Pagnucco and Rott mean to recommend choosing any potential contraction that incurs no loss of information greater than the loss incurred by some other potential contraction. That is to say, one is to choose a maxichoice contraction.

As it stands, the Principle of Economy draws no comparisons between different maxichoice contractions. But Pagnucco and Rott do invoke a Principle of Strict Preference where the preference concerns the evaluation of maxichoice contractions. As I understand this principle, it requires that the choice of a potential contraction removing h from \underline{K} be restricted to those that are intersections of subsets of most preferred maxichoice contractions.

If one adopts this version of the Principle of Strict Preference and applies the Principle of Economy only to contractions that are intersections of subsets of most preferred maxichoice contractions, the recommendation becomes to choose a most preferred maxichoice contraction.

The Principle of Economy modified to accommodate the Principle of Strict Preference recommends the same contractions as those prescribed by an injunction to minimize undamped informational value under special assumption 1. But there is no effort to identify the assessment of undamped informational value for all potential contractions on the basis of the evaluation of the maxichoice contractions.

Pagnucco and Rott recognize the dubious implications of this modified Principle of Economy. They contend that the recommendation of partial meet contractions in AGM already compromises the Principle of Economy. Contractions that are intersections of optimal maxichoice contractions are recommended.

Pagnucco and Rott seem to think that the operation of the modified Principle of Economy is compromised in AGM by an appeal to a Principle of Indifference. If elements of a set of maxichoice contractions are optimal and equiprefered, they should be "treated equally".

Equal treatment here is glossed as taking the intersection of the optimal maxichoice contractions removing h . The Principle of Economy requires us to incur no more loss in information than is necessary to contract K by removing h , have a deductively closed theory and obey the Principle of Strict Preference. The Principle of Indifference requires us to violate that injunction. When its use is combined with the Principle of Strict Preference, the recommendation is to choose the intersection of those maxichoice contractions that are best in the preference ranking of maxichoice contractions. When there are several optimal maxichoice contractions, the intersection does not minimize loss of information.

Pagnucco and Rott then point out that to apply the Principle of Indifference strictly, we should include in the intersection or meet of maxichoice contractions not only the best that imply $\sim h$ but also those that imply h as well.

This reasoning yields the same result as minimizing the loss of damped informational value of type 2 and then breaking ties by taking the intersection of the set of optimal contraction strategies. This strategy itself is optimal as it should be.

Pagnucco and Rott do not explicitly say that their use of the principles of Economy, Indifference and Strict Preference are to be understood as developing a principled method for deriving a choice of a potential contraction from the an evaluation of maxichoice contractions. But their examples especially in connection with the discussion of Grove modeling suggest that they have something like this in mind.

In any case, Pagnucco and Rott do not derive a weak ordering of the potential contraction strategies on the basis of which it can be argued that the contraction recommended is optimal.

Consider the set of optimal maxichoice contractions removing h or better yet the set of all maxichoice contractions as good as the optimal maxichoice contractions removing h . Does the Principle of Indifference state whether the intersection of these maxichoice contractions is equiprefered to each of the maxichoice contractions in the set or is strictly preferred to them? There is no way of telling. Worse yet, any answer we may provide is unsatisfactory.

Suppose the intersection is indifferent to each of optimal maxichoice contractions. Then no recommendation is forthcoming as to whether to choose the intersection or one of the maxichoice contractions (or the intersection of a subset of these).

Suppose the intersection is better than any of the maxichoice contractions constituent in it. Then the ordering of potential contractions violates the weak positive monotony condition. It allows a weaker less informative contraction to be favored over a stronger one.

If Peirce is correct that inquirers seek to relieve doubt, then in contraction where doubt is increased we should be seeking to keep the increase to a minimum. Consequently, in contraction, one should never favor a weaker, less informative potential contraction over a stronger, more informative one. We should seek to minimize loss of informational value in a sense satisfying the weak positive monotonicity condition. In inquiry we seek to relieve doubts. That is to say we seek to obtain valuable

information. And we seek to avoid losing valuable information. Pagnucco and Rott do not do justice to this important point.

Pagnucco and Rott seek to derive a recommendation of a contraction removing h from \underline{K} from an evaluation of the maxichoice contractions removing h from \underline{K} . If this is to be done in a manner that shows that loss of informational value is minimized, it must be admitted at the outset that more than one potential contraction strategy may minimize such loss. So deriving a recommendation should first identify the set of loss minimizing contraction strategies and then recommend a way of breaking the tie for optimality.

The Pagnucco-Rott Principle of Indifference tries to do both tasks in one fell swoop. As a consequence, it does neither well.

There is another aspect of the Pagnucco-Rott argument that merits some scrutiny. The partial ordering of the elements of U^*_K by subset inclusion is compatible with any weak ordering of the elements of U^*_K . To satisfy the weak positive monotonicity requirement, a partial ordering of all of the potential contractions removing h from \underline{K} is needed. Some weak orderings of all the potential contractions will extend this partial ordering in the sense of weak positive monotonicity. Others will not. Pagnucco and Rott do not consider weak orderings or quantitative representations of all potential contractions satisfying weak positive monotonicity. The assessment of loss of informational value does not even arise if we restrict attention to assessments of members of U^*_K and not to subsets of U^*_K as well.

Whether or not this is a motivation for refusing to consider orderings of the maxichoice contractions as assessments of informational value, it is a motivation. And Pagnucco and Rott do explicitly evaluate maxichoice contractions in terms of *plausibility*. (Pagnucco and Rott, 1997, section 3, p.8.)

What does plausibility mean here? Here is what Pagnucco and Rott write when illustrating how their method works with Grove models. Interpret possible worlds to be cells in U^*_K .

Essentially, a system of spheres centred on $[\underline{K}]$ orders those worlds inconsistent with the agent's epistemic state \underline{K} . Intuitively, the agent believes the actual world to be one of the \underline{K} -worlds but does not have sufficient information to establish which one. However, the agent may be mistaken, in which case it believes that the actual world is most likely to be one of those in the next greater sphere and so on. As such, a system of spheres can be considered an ordering of plausibility over worlds; the more plausible worlds lying further towards the centre of the system of spheres. (p.8.)

Plausibility is clearly an evaluation with respect to some notion of likelihood. However, if a hypothesis is plausible, it is, at a minimum, possibly true in the sense of serious possibility. If an inquirer's belief state is \underline{K} , any hypothesis h incompatible with \underline{K} is not a serious possibility. The coherent inquirer must regard every such hypothesis as certainly false and *maximally implausible*.

Pagnucco and Rott point out that the inquirer "may be mistaken" in believing that the actual world is one of the \underline{K} -worlds. It is, of course, logically possible that this is so. Another inquirer can judge it seriously possible that the inquirer in question is mistaken. The inquirer him or her self may acknowledge that he or she may cease being certain in the future and that his or her future (and past beliefs) may be (and often are) in error. In any one of these senses, the inquirer may be mistaken in believing that the actual world is one of the \underline{K} -worlds. However, as long as the inquirer is in the state of full belief represented by \underline{K} , the inquirer cannot coherently acknowledge the serious possibility of being mistaken in this belief.

This means that every element in the dual ultimate partition U^*_K (every not- \underline{K} -world) is maximally and equally implausible. There can be no distinction between hypotheses incompatible with \underline{K} with respect to plausibility. When plausibility is understood in this way, the principle of indifference seems to favor treating all potential contractions alike. Either everything \underline{K} must be given up or everything retained.

Many philosophers respond to this kind of observation by suggesting that no reasonable inquirer is certain of any extralogical or a posteriori truth. Items incompatible with \underline{K} are not judged to be false with maximum certainty but only with more or less confidence. Whether such confidence is measured by probabilities or by generalizations of Shackle measures, as authors such as Spohn (1988) have done, \underline{K} no longer represents a state of full belief. That is to say, \underline{K} is no longer a state of information or evidence cum background information relative to which probability judgments or other judgments of uncertainty such as degrees of belief in the Shackle sense may be made.

The only state of full belief allowed is the state of total ignorance where the corpus is the minimal urcorpus consisting of logical truths and whatever passes for conceptually necessary truth. What happens to the study of belief change under these circumstances?

I suggest that we remove the incoherence from the proposal of Pagnucco and Rott by reconstruing plausibility to be a measure of informational value. That is to say, the weak ordering of elements of U^*_K should be used to extend the partial order of the potential contraction strategies removing h from \underline{K} to a weak ordering with respect to informational value satisfying the weak positive monotonicity condition. But how is this to be done?

In Levi, 1991 and 1996, I made a new proposal for assessing losses of informational value that furnishes a weak ordering of the contraction strategies satisfying weak positive monotonicity. I shall first restate the proposal I made in 1991 and reiterated in 1996 and explain why I now think it too is doubtful. I shall then offer a new proposal that I think is acceptable and continues to protect the derivability of such losses from demands for information.

Minimizing loss of informational value in this last sense (damped informational value of type 2) and invoking a rule for ties that recommends choosing the weakest of all the optimal potential contractions if there is one yields the same potential contraction that Pagnucco and Rott (1996) obtain. But the incoherent evaluation of hypotheses that are not seriously possible with respect to plausibility is discarded.

And the clash between the Principle of Indifference advocated by Pagnucco and Rott and the requirement sanctioned by the weak positive monotonicity condition on informational value that a weaker potential contraction should never be preferred to a stronger one is avoided. The Principle of Indifference is abandoned. Minimizing loss of damped informational value, version 2 identifies a set of optimal potential contractions. These include the intersection of all optimal maxichoice contractions just as one construal of the Principle of Indifference requires. The Rule for Ties then is invoked to select the intersection of all optimal maxichoice contractions as the second construal of the Principle of Indifference insists without implying that this contraction is better than the others in minimizing loss.

17.Damped Informational Value: Version 1

In Levi, 1991, I proposed a method of evaluating loss of informational value that I called loss of *damped informational value*. I reaffirmed my endorsement of its use in Levi, 1996 as well. Damped informational value as defined in Levi, 1991, ch.4.4 and in Levi, 1996, p.263 is understood as follows:

Consider all potential contractions of \underline{K} removing h as defined relative to \underline{LK} , $V_{LK,K} = U_{LK}$ and M^* .

The version 1 loss in damped informational value incurred by shifting from \underline{K} to a saturatable contraction removing h from \underline{K} is equal to the loss in undamped informational value thereby incurred as represented by M^* .

CONTRACTION AND INFORMATIONAL VALUE

The version 1 loss in undamped informational value incurred by shifting from \underline{K} to an intersection of saturatable contractions removing h from \underline{K} is the largest loss incurred by any saturatable contraction in the set.

Thus, loss in version 1 damped informational value is derived from losses in undamped informational value derived from \underline{LK} , U_{LK} and M^* according to *the version 1 damped constraint on informational value*. Given the dual ultimate partition U^*_K , the first clause of the definition combined with the injunction to minimize loss of informational value tells us to consider those maxichoice contractions that minimize loss of undamped informational value in removing h from \underline{K} . They consist of intersections of U_K with each of those d_i 's in U^*_K entailing $\sim h$ for which the M^* -value is a minimum. For each such maxichoice contraction, we can construct the saturatable contraction obtained by intersection with all d_i 's entailing h that carry 0 M^* -value. The resulting saturatable contraction will carry the same undamped informational value as the maxichoice contraction corresponding to it. Thus, the first clause restricts choice to a specific subset of saturatable contractions and intersections of members of subsets of that set.

The second clause stipulates that preference between intersections of two distinct sets of such saturatable contractions should be determined by the saturatable contractions in each set incurring greatest loss of undamped informational value. We are thus not concerned to minimize loss of undamped informational value any longer. What we wish to minimize is greatest loss in undamped informational value among the saturatable contractions

The consequence of adopting this index is that some of the grosser deviations from what I, at any rate, regard as near common sense that are thrust on us by the injunction to maximize undamped informational value are avoided.

Observation 1': Under special assumption 1, minimizing loss of version 1 of damped informational value and using the rule for ties requires choice of a partial meet contraction -- to wit, using the intersection of all maxichoice contractions removing h carrying minimum M^* -value. Recovery is satisfied.

Observation 2': Under special assumption 2, minimizing loss of version 1 of damped informational value and using the Rule for Ties recommends choosing the intersection of all saturatable contractions removing h from \underline{K} where each saturatable contraction is produced by forming the intersection of a maxichoice contraction incurring minimum (but positive) loss of informational value with the intersection of all those cells of U^*_K that carry 0 M^* -value and entail h . Recovery fails

Observation 3': Under special assumption 3, minimizing loss of version 1 of damped informational value and using the Rule for Ties recommends the intersection of all saturatable contractions removing h from \underline{K} incurring 0 loss of undamped informational value. Recovery fails.

Observation 4': Under special assumption 4, minimizing loss of version 1 of damped informational value and using the Rule for Ties recommends the intersection of all maxichoice contractions incurring 0 loss of undamped informational value. Recovery is satisfied.

Under none of the special assumptions does minimizing loss of version 1 damped informational value require choice of a maxichoice or even a saturatable contraction (although such choices may be permitted sometimes).

In Levi, 1991 and 1996, I sought to show how Recovery might fail. And under special assumptions 2 and 3 it does indeed fail. The argument I offered for using version 1 of damped informational value, however, depended on an appeal to examples where it seems clear presystematically that Recovery does fail and where I claimed that minimizing loss of damped informational value and using the Rule for Ties rationalizes

such failure. I considered only situations under special assumptions 1 and 2 without explicit argument for doing so. Arguments have been offered here as to why at least sometimes we want to avoid requiring maxichoice contractions under special assumption 1. But, as just noted, under special assumption 1, meet contractions are required and the Recovery Condition is satisfied. Recovery fails under special assumption 2. The question I did not address is whether the counterinstances to Recovery can all be plausibly construed as satisfying special assumption 2 or whether it is desirable to guarantee that Recovery sometimes fails for special assumption 1 as well. I now think that it must. The case where it is known that a given coin has been tossed and landed heads (E&H) may serve to illustrate this.

Let \underline{L} imply that the coin lands on the surface if and only if it is tossed (an instance of L) and the coin lands on the surface if and only if it lands heads or lands tails but not both (an instance of L'). It is not given in \underline{L} , however, whether the coin is tossed at the given time (E). The basic partition $V_{L,K} = U_{L,K}$ consists of the hypotheses E&H, E&T and $\sim E$.

Relative to this information, three assumptions about the undamped informational value of these three alternatives seem plausible:

- (1) The undamped informational value of $\sim E$ is less than the undamped informational value of E so that the M^* -value of E is lower than that of $\sim E$. Since $M^*(E) = M^*(E\&H) + M^*(E\&T)$, each of the summands must be less than $M^*(\sim E)$.
- (2) The gain in informational value by adding E&H to \underline{L} to form \underline{K} is positive so that $M^*(E\&H) > 0$.
- (3) The undamped informational values of E&H and E&T are equal.

From these three entirely plausible assumptions about such a situation, it follows that $M^*(E\&T) > 0$ as under special assumption 1. If we contract \underline{K} by removing E&H, the dual ultimate partition U^*_K contains two elements: $\sim E$ and E&T. According to the injunction to minimize loss of version 1 damped informational value, we should adopt the maxichoice contraction that suspends judgment between $\sim E$ and E&H. That is to say, we should judge that either the coin was not tossed or was tossed and landed heads. We cannot include as a possibility that it was tossed and landed tails (E&T) because $M^*(E\&T)$ carries positive M^* -value. Yet presystematic judgment clearly indicates that E&T would be included as a possibility and that restoring E should *not* return H as Recovery requires.

This toy example can be mirrored in countless cases where one knows that an experiment has been conducted and a definite outcome realized from among the "points" that represent possible outcomes of the experiment (the "sample space"). I do not see how it can be denied that in most such cases, special assumption 1 holds but Recovery fails. Minimizing damped informational according to version 1 fails to secure this result.

This observation is closely related to the more general observation that when the basic ultimate partition $U_{L,K}$ is finite, all its cells should carry positive M^* - value as is required of hyper-undamped assessments of informational value. Once more we are driven to focus on special assumption 1.

There is a third consideration arguing that we should look to an alternative to minimizing damped informational value according to version 1. According to version 1, the loss in damped informational value of a saturatable contraction that is not maxichoice is equal to its undamped informational value. This is so even though more than one element of U^*_K is shifted out in forming the contraction. (More than one element must be removed for one must be an element entailing $\sim h$ and one an element entailing h.) Yet, the intersection of two or more saturatable contractions that does not

yield a saturatable contraction incurs a loss of damped informational value that does not, in general, equal the loss in undamped informational value but could incur a greater loss. In both cases, the resulting contractions are not maxichoice. More than one element is removed from U^*_K . According to version 1 standards for evaluating damped informational value, in one case, damped informational value equals undamped informational value yet in the other case, damped and undamped informational value come apart. Should we not use the same standards for evaluating losses in informational value in both cases?

One way to avoid a double standard, of course, is to revert to the standards for undamped informational value; but we have seen why that will not do. The other alternative is to use the same “damping” method for assessing loss of informational value incurred by a saturatable contraction as one does for the intersection of a set of such contractions. I shall now consider this procedure.

18.Damped Informational Value: Version 2

The loss of *damped informational value according to version 2* incurred by a contraction removing h from \underline{K} is equal to the largest M^* -value belonging to an element of U^*_K in the subset of such elements whose intersection with \underline{K} is the contraction in question.

Minimizing version 2 loss of damped informational value satisfies *the version 2 damped constraint* on deriving loss of informational value from \underline{LK} , U_{LK} and M^* .

Calculating damped informational value of type 2 begins by looking at the undamped informational values of all maxichoice contractions relative to \underline{K} and U^*_K and then identifying all those maxichoice contractions removing h from \underline{K} that minimize loss of undamped informational value. The saturatable contractions that are intersections of these maxichoice contractions with members of U^*_K entailing h are evaluated. These incur losses of informational value equal to the largest M^* -value carried by any element of U^*_K in the intersection. As long as these largest M^* -values are no greater than the M^* -value of the associated maxichoice contraction removing h (entailing $\sim h$), the saturatable contraction minimizes loss of damped informational value according to version 2. And so do intersections of all such saturatable contractions as before.

This procedure yields an ordering of all contraction strategies removing h from \underline{K} . The weak ordering guarantees that x is an optimal contraction removing h from \underline{K} if and only if and only if x is an intersection of a set of maxichoice contractions all of which carry no greater M^* -value than the minimum M^* -value carried by maxichoice contractions removing h from \underline{K} and x is itself a contraction removing h from \underline{K} .

The Rule of Ties then stipulates that the weakest of these optimal contractions should be chosen. That is to say, the intersection of all optimal maxichoice contractions removing h from \underline{K} should be chosen.

This contraction strategy is the intersection of the set of all maxichoice contractions carrying no greater M^* -values than the minimum M^* -value carried by maxichoice contractions removing h from \underline{K} .

Since the M^* -values of maxichoice contractions weakly order them, this latter recommendation is easily seen to be identical to the recommendation Pagnucco and Rott (1997) derive from a weak ordering of the maxichoice contractions by their Principles of Strict Preference and Indifference.

Main Result B: If the contraction recommended for removing h from \underline{K} is the result of minimizing loss of damped informational value according to version 2 and the Rule of Ties, the contraction is the intersection of the partial meet contraction removing all elements of U^*_K entailing $\sim h$ that have minimum M^* -

value among all elements of U^*_K entailing $\sim h$ and the intersection of all elements of U^*_K entailing h that carry M^* -value less than or equal to that minimum.

Observation 1": Under special assumption 1, minimizing loss of damped informational value according to version 2 is an intersection of saturatable (but not always maxichoice) contractions removing h from \underline{K} . Recovery Fails.

Observations 2" and 3" make the same claim for special assumptions 2 and 3. Observation 4" and that alone strengthens the claim to require partial meet contractions.

Not only does version 2 of damped informational have fewer epicycles than version 1, it also handles sample space examples such as the case of the coin that is known to have landed heads in precisely the way I claim presystematic judgment requires. In particular, it secures failures of Recovery under special assumption 1 and removes, thereby, an obstacle to assigning all cells in finite ultimate partitions positive M -value.

Finally, it offers a philosophically compelling way of arguing for the recommendations that Pagnucco and Rott obtain. Contraction removing h from \underline{K} is a decision problem where the aim should be to minimize the loss of damped informational value according to version 2. Assessments of loss of version 2 damped informational value weakly orders the potential contractions in a manner satisfying the weak positive monotonicity condition. And the weak ordering over all the potential contractions is readily derivable from the weak ordering of the maxichoice contractions of \underline{K} .

That weak ordering singles out the best maxichoice contractions removing h from \underline{K} . The intersection of all such maxichoice contractions is also a best contraction removing h from \underline{K} . So is the intersection of this with the set of maxichoice contractions entailing h . This latter contraction is not merely tied for optimality with its maxichoice constituents, it is recommended for choice in virtue of the tie breaking rule to choose the weakest of the optimal options.

19. Severe Withdrawal = Mild Contraction

As noted before, Pagnucco and Rott (1997) have also proposed quite independently and utilizing different principles that contractions of the sort described in Main Conclusion B should be taken seriously. They call such contractions "severe withdrawals". I prefer another epithet "mild contractions". Two reasons motivate my terminological predilections:

The first is a stubborn insistence that the problem of how to contract should look at all potential contraction strategies and not merely those that are partial meet contractions. Discrimination between potential contraction strategies that are partial meet contractions and the others by restricting the application of the epithet "contraction" to partial meet contractions and calling the others "withdrawals" suggests even if it does not entail that contractions that fail to violate the Recovery Condition are not true contractions. I wish to resist using terminology that even suggests this view.

Second, contractions often give birth to revisions so that an obstetric metaphor rather than one from the drug culture is not totally out of place. The only issue is whether the contractions are mild or severe. To answer this question, we must ask what sorts of values are of concern in contraction. I have been insisting that we ought to be concerned to minimize loss of informational value. Now if the informational value is undamped informational value, the loss incurred would be considerable. Thinking of the contractions under consideration as severe seems entirely reasonable according to that reckoning.

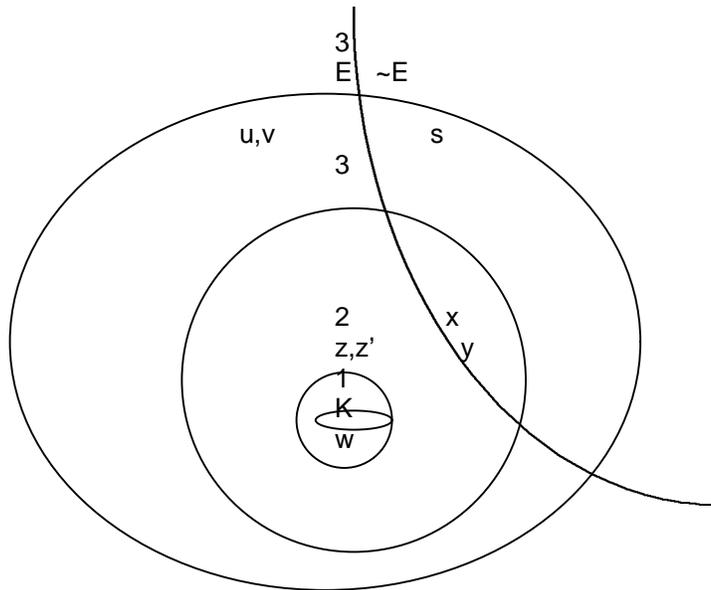
Furthermore, assessing loss of informational value along the lines of version 1 of loss of damped informational value also implies that the contraction would be severe.

CONTRACTION AND INFORMATIONAL VALUE

On the other hand, assessing losses of damped informational value according to version 2 yields the result that the contractions are mild. There is no more loss in informational value of this kind by adopting the severe withdrawal/mild contraction than there is in adopting the corresponding partial meet contraction.

To be sure, using the geometry of Grove modeling to represent alternative contractions makes it appear that the contractions are severe. But this seems to me to be one of the misleading aspects of Grove-think.

Consider the following Grove diagram of contraction of \underline{K} by removing E .



Here the innermost circle represents the current corpus \underline{K} . We shall suppose that the lower case roman letters represent elements of U^*_K . Circle 1 minus the innermost circle consists of worlds carrying 0 M -value. Circle 2 minus circle 1 includes worlds carrying positive M values no greater than worlds x and y which are worlds according to which E is false. Circle 3 minus circle 2 contains worlds carrying positive M -values greater than the M -values of x and y .

Minimizing loss of undamped informational value calls for contracting by choosing one of the two maxichoice contractions obtaining by forming the union of $|\underline{K}|$ and x or of $|\underline{K}|$ and y . Both of these incur a positive loss of undamped informational value. Hence, it is clear that the union of the two sets of points increases the loss. The partial meet contraction cannot be recommended.

If we calculate minimum loss in informational value incurred by removing E from \underline{K} using undamped informational value of the first kind, the optimal saturatable

contractions are the unions of $|\underline{K}|$ and $\{x\}$ and $\{w\}$ or $|\underline{K}|$ and $\{y\}$ and $\{w\}$. Recovery is violated.

If the Rule for Ties is applied to the set of contractions minimizing damped informational value of the first kind, the union of the two saturatable contractions mentioned above is recommended. The loss in undamped informational value is equal to the total M -value assigned points in the intersection of the set of points $\sim E$ and sphere 2 and to points in sphere 1. But the total M -value of points in sphere 1 is 0. Hence, the loss of damped informational value of the first kind incurred by the recommended contraction is equal to the loss of damped informational value of the first kind incurred by the corresponding partial meet contraction.

Recovery obviously fails as long as there are some points of U^*_K in sphere 1. But loss of informational value incurred by violating Recovery as compared to partial meet contraction is 0. The Rule for Ties clearly then recommends violating Recovery.

The geometry of the Grove model, however, suggests even if it does not imply that a genuine loss is incurred by violating Recovery; for it invites you to think of items in the first circle as incurring a genuine loss. We could, however amend the modeling so that the items in circle 1 fell on the boundary of the sphere for \underline{K} where \underline{K} is now taken to be an open set. It then becomes clear that the difference in loss of both undamped and damped informational value of the first kind incurred by partial meet contraction and by minimizing loss of damped informational value of the first kind is 0. Let us call contractions recommended by the injunction to minimize loss of damped informational value *impalpable contractions*.

If we calculate damped informational value of the second kind and use the Rule for Ties, we form the union of the impalpable contractions just described with $\{z, z'\}$. The result is what Pagnucco and Rott call a severe withdrawal. Looking at the geometry of the Grove model it certainly looks severe. For it suggests an additional loss of informational value equal to the informational value of the contents of sphere 2 which would be positive. If we are measuring loss in undamped informational value, this claim is entirely correct.

However, there is no such loss in damped informational value of the second kind. There is no more loss incurred than if the maxichoice contraction represented by the union of $|\underline{K}|$ with $\{x\}$ were used.

Whether the contraction is mild or severe is clearly going to depend on the kind of loss of concern to the inquirer. This is the important philosophical point that I mean to belabor.

In the case of impalpable contractions, the loss (as compared to that incurred by the partial meet contraction) is nonexistent whether the informational value is undamped, damped of the first kind or damped of the second kind. In the case of mild contractions, a positive loss is incurred with respect to undamped and damped informational value of the first kind. So the contraction is not impalpable from all perspectives of interest. But there is no loss in undamped informational value of the second kind. From that perspective, the contraction is mild. It can only be taken as severe when we seek to minimize loss of undamped informational value or damped informational value of the first kind.

The Grove model diagram cannot reveal this for the simple reason that such diagrams represent the ordering of cells in the dual ultimate partition with respect to how "severe" the maxichoice contraction that is the union of $|\underline{K}|$ and $\{x\}$, where x is a cell, happens to be. The Grove model diagram says nothing about how severe contractions are supposed to be that are unions of $|\underline{K}|$ with sets that are not such singletons.

If one thinks that somehow the more inclusive a sphere, the more severe the contraction will be, that is, indeed, a tacit commitment to undamped informational value -- a commitment, so I have argued, no one should want to make.

Grove models rest silent with respect to this proposal as well as with respect to alternatives. And that is why Grove models are in the final analysis useless in revealing the structure of acceptable contractions.

Grove models can distinguish between maxichoice contractions and nonmaxichoice but saturatable contractions. Within the latter category, Grove models can be used to distinguish between saturatable contractions removing E from \underline{K} where the $\sim E$ states are all in sphere 1, where the $\sim E$ states are in sphere 1 or sphere 2 and where the $\sim E$ states are all in sphere 3 and subspheres. We can even represent intersections of saturatable contractions. But there is no way to indicate how potential contractions of \underline{K} other than maxichoice contractions are to be compared. The absence of such a standard seems to suggest to many that no systematic approach to critically examining such standards is available. I believe that by thinking of the goal of a contraction as minimizing loss of informational value of certain kind and then exploring what kind it should be, we can add to the resources we have for coming to some nonarbitrary adjudication of the controversies concerning the character of contraction.

20. Revision and Recovery

As Makinson has correctly observed, whether contraction satisfies the Recovery condition or not makes no difference to AGM revision. For purposes of characterizing AGM revision, therefore, one might as well deploy partial meet contraction in defining contraction as an expansion of a contraction as any contraction transformation corresponding to the given partial meet contraction. The results will be the same.

In belief change, this does not diminish the importance of studying types of contraction failing to satisfy Recovery; for not every contraction is followed by an expansion. But in suppositional reasoning where the success of a revision is stipulated, it may appear that questions about Recovery are of lesser importance.

This is not so. In suppositional or conditional reasoning, the revision transformation ought not to be AGM revision but rather what I have called Ramsey Revision. (Levi, 1996)

In a situation where the agent is convinced that the coin was tossed and landed heads (E&H) so that E&H is in \underline{K} , reasoning on the supposition that E requires that E be contracted from \underline{K} and then restored before inferences from E are drawn. We want to consider whether on the supposition that the coin had been tossed, the coin would have landed heads. Presystematically, it is clear that in most contexts we are prepared to judge that because the coin had been tossed, it would have landed on the surface. We are not prepared to judge, however, that because the coin had been tossed, it would have landed heads. We do not think that the tossing explains why the coin landed heads.

AGM delivers an opposite verdict. In my judgment, this constitutes an argument from presystematic practice supporting the idea that revision in suppositional reasoning is not AGM revision but Ramsey revision. Ramsey revision requires removing $\sim h$ when $\sim h$ is in \underline{K} and removing h when h is in \underline{K} where h is the supposition that is to be added subsequent to the contraction. Ramsey revision \underline{K}^r_h coincides with AGM revision \underline{K}^*r_h when the Recovery condition on contraction is satisfied; but it does not otherwise. Suppositional reason in "sample space" examples illustrated by the coin requires that contraction violate Recovery. If this is right, not only is failure of Recovery manifest in belief change but also in suppositional reasoning and in the logic of conditionals.

Some of the ramifications of this point were explored in Levi, 1996. Both the postulates (K*4) and (K*8) for AGM revision were weakened to characterize Ramsey revision (p.40). However, the approach adopted there was predicated on the idea that in contraction loss in damped informational value of type 1 was being minimized and that, as a consequence, the contractions involved that violated Recovery were impalpable contractions. As Hansson and Olsson (1995) have correctly shown, all the postulates for AGM contraction except (K-5), which, in the light of the others, entails Recovery, are satisfied.²⁰ This is true for mild contraction (alias severe withdrawal) as well.

AGM (1985 observation 6.5) shows that (K-7) and (K-8) in the light of the other postulates for contraction except (K-5) entail the following:

$$(X) \underline{K}_{h\&f} = \underline{K}_h \text{ or } \underline{K}_{h\&f} = \underline{K}_f \text{ or } \underline{K}_{h\&f} = \underline{K}_h \cap \underline{K}_f$$

Without (K-5), (X) implies (K-7) in the light of (K-1) -(K-4) and (K-6). (K-8) is not implied although with the aid of (K-5) it is.

(X) holds for Ramsey as well as AGM contraction. For Ramsey revision, the following holds for cases where h&f is in \underline{K} *even though none of these revisions are identical with \underline{K} as they would be if AGM revision obtained and whether or not Recovery with Ramsey revision held.*

$$(a) \underline{K}_{h\&f}^{*r} = \underline{K}_h^{*r} \text{ or } \underline{K}_{h\&f}^{*r} = \underline{K}_f^{*r} \text{ or } \underline{K}_{h\&f}^{*r} = \underline{K}_h^{*r} \cap \underline{K}_f^{*r}$$

Rott and Pagnucco, 1996 point out that for mild contractions, condition (X) can be strengthened as follows:

$$(\text{strong X}) \underline{K}_{h\&f} = \underline{K}_h \text{ or } \underline{K}_{h\&f} = \underline{K}_f.$$

This means that when damped informational value of type 2 is minimized in contraction and Ramsey revision is used, in cases where \underline{K} contains h&f, we have the following result.

$$(b) \underline{K}_{h\&f}^{*r} = \underline{K}_h^{*r} \text{ or } \underline{K}_{h\&f}^{*r} = \underline{K}_f^{*r}.$$

Hence, even though we cannot say, as in the case of AGM revision that if f is not in \underline{K}_h , that $[\underline{K}_h]_f^+$ is a subset of $\underline{K}_{h\&f}$, we can say that either $[\underline{K}_h]_f^+$ or $[\underline{K}_f]_h^+$ is a subset of $\underline{K}_{h\&f}$.

Thus, whether under the sobriquet “mild contraction” or “severe withdrawal”, contractions minimizing loss of damped informational value type 2 in conformity with the Rule for Ties do have a non-negligible effect on the formal properties of Ramsey revision. Because (for reasons already given) I now recommend mild contraction over impalpable contraction, I am prepared to strengthen the postulates for Ramsey revision to the extent just indicated.

In this essay, I am arguing that the case for seeking to minimize loss of damped informational value of type 2 is also a case for violating Recovery by mandating mild contractions. Pagnucco and Rott reach this conclusion by a different argument. I have suggested that their argument is not sufficiently loyal to decision theoretic requirements. In particular, Pagnucco and Rott fail to show that the severe withdrawal (or mild contraction) is optimal because their principles do not allow for evaluating intersections of sets of maxichoice contractions with respect to loss of informational value. When details are filled in by introducing version 2 of damped informational value and the Rule for Ties is invoked to arbitrate between optimal contraction strategies, their

²⁰ Hansson and Olson (1995, p.108) correctly point out the need for a special postulate stipulating that if \perp -h, $\underline{K}_h = \underline{K}$.

recommendations are obtained on a foundation more in accord with decision theoretic principles.²¹

Pagnucco and Rott fail to recognize the importance of Ramsey revision as compared to AGM revision. For this reason, they do not recognize how significant the difference between severe withdrawal and partial meet contraction actually is in the context of suppositional reasoning and the logic of conditionals. In the absence of a decision theoretic argument, the attractiveness of some of the properties of severe withdrawals looks like a matter of taste about which one may dispute without much argument. The appeal to Ramsey revision and the decision theoretic rationale strengthens the argument, in my judgment, quite substantially.

21. Entrenchment and Incorrigoibility

In minimizing loss of informational value incurred by removing h from \underline{K} , the concern is to identify those potential contractions removing h from \underline{K} that carry maximum informational value. As we have seen, this calls for an evaluation of potential contractions of \underline{K} with respect to informational value in some reasonable sense or other.

This evaluation carries with it another sort of assessment. Both \underline{K} and all potential contractions removing h have other belief states as consequences. Indeed, every potential contraction of \underline{K} in the sense of a deductively closed subset of \underline{K} containing the logical truths is such a consequence and all finitely axiomatisable subsets of this type are representable by sentences in \underline{K} . Implementing a contraction removing h from \underline{K} involves giving up some consequences of \underline{K} and retaining others and this can be characterized as giving up some sentences in \underline{K} while retaining others.

From this consideration, it becomes apparent that an assessment of potential contractions of \underline{K} removing h also determines an assessment of which consequences of \underline{K} (or the sentences that represent them) are more vulnerable to being given up in removing h .

In Levi, 1980, I contended that although sentences in an inquirer's corpus \underline{K} at time t are all maximally certain from the agent's point of view at that time, they may, nonetheless, differ from one another with respect to their degrees of incorrigoibility (or degrees of corrigoibility). I claimed that discrimination with respect to incorrigoibility depended on the losses in informational value that are incurred. (p.62) Degrees of incorrigoibility or vulnerability to being given up are not degrees of confidence, degrees of expectation, degrees of support by evidence but are judgments of informational value -- a species of epistemic utility. Such judgments reflect the research project or other demands for information of the inquirer in the given situation.

²¹In a recent interesting paper, Meyer, Labuschagne and Heidema (1998) have proposed an account of contraction they call "systematic withdrawal". Like Pagnucco and Rott, Meyer, Labuschagne and Heidema begin with an ordering of cells in U^*_K and seek to derive a recommended contraction. The partial meet contraction favored by AGM is the intersection of \underline{K} with the cells entailing $\neg h$ that are "nearest" \underline{K} . Meyer, Labuschagne and Heidema argue that the intersection ought to be extended to cover cells in U^*_K that entail h that are nearer to \underline{K} than the nearest cells entailing $\neg h$. (1998, pp.15-17.) Recovery is violated. The contraction is, in general weaker than that yielded by minimizing damped informational value of type 1 but not of damped informational value of type 2. Let the ordering of elements of U^*_K be a "K-faithful" a weak ordering (total preorder). According to the decision theoretic approach I favor, the ordering should be extended to all potential contractions of \underline{K} . The contraction of \underline{K} by removing h to be recommended is a best such contraction according to the weak ordering that satisfies the Rule for Ties. One might achieve the desired result for contraction of \underline{K} by removing h ; but I do not see how this ordering can be made to cohere with contraction of \underline{K} by removing f . Meyer et al. do not endorse my decision theoretic program any more than Pagnucco and Rott do. As a consequence, the merits of the Meyer, Labuschagne and Heidema proposal for systematic withdrawal as compared with severe withdrawal can be compared only by consulting intuitions about the implications of the two proposals. I suspect such comparisons will prove inconclusive. But the fact that mild contraction or severe withdrawal has a decision theoretic rationale while systematic withdrawal does not seems to offer a strong argument for the former over the latter.

As already indicated, in Levi, 1980, I assumed that losses in informational value incurred in contraction are to be assessed as losses in undamped informational value. I have explained why I now think this is a bad idea. In any case, in 1980, I did not elaborate in detail on how grades of incorrigibility are to be derived from such assessments.

I sought to elaborate the idea of incorrigibility as derived from losses of damped informational value of type 1 and to compare it with Gärdenfors's notion of entrenchment in Levi, 1991 and 1996. Here I shall reformulate the conceptions of incorrigibility and entrenchment in a manner that is neutral with respect to whether losses being assessed are losses of type 1 or type 2 damped informational value. I shall then consider the specializations of incorrigibility and entrenchment relative to losses of each of these two types.

Consider the subset of potential contraction strategies removing h from \underline{K} (relative to \underline{LK} and U_{LK}) that also remove g . Whether assessments of damped informational value are of type 1 or type 2, there will be at least one potential contraction strategy of this kind that carries maximum damped informational value in that set. The existence of a maximum is insured as long as U_{LK} (and, hence, U^*_K) is finite.

Let $D(g/h, \underline{K})$ equal the damped informational value of a contraction of \underline{K} that minimizes loss of damped informational value among all contractions of \underline{K} removing h that also remove g . If it is type 1, we may write $D_1(g/h, \underline{K})$. If it is type 2, $D_2(g/h, \underline{K})$.

$D(h/h, \underline{K})$ is equal to the damped informational value of \underline{K}_h . This damped informational value will equal the cont-value of a maxichoice contraction removing h from \underline{K} that carries maximum cont-value in this set and, hence, minimum M^* -value whether we are thinking of damped informational value of type 1 or of type 2. $D_1(h/h, \underline{K}) = D_2(h/h, \underline{K})$.

Def. The degree of incorrigibility of $g = in(g/h, \underline{K}, \underline{LK}, M^*, U_{LK}) = D(h/h, \underline{K}) - D(g/h, \underline{K})$.²²

In the following discussion, I shall abbreviate this as $in(g/h, \underline{K})$ assuming that \underline{LK} , M^* and U_{LK} are given and held fixed.

$D_1(g/h, \underline{K}) = D_1(h/h, \underline{K})$ in two cases:

(a) At least one maxichoice contraction removing h from \underline{K} that carries minimum M^* -value and, hence, maximum undamped informational value among all maxichoice contractions removing h from \underline{K} also removes g from \underline{K} . This means that there is a cell in the dual ultimate partition U^*_K that entails both $\sim h$ and $\sim g$ and carries minimum M^* -value among cells entailing $\sim h$.

(b₁) At least one maxichoice contraction removing g from \underline{K} carries the same undamped informational value as \underline{K} itself. This means that the minimum M^* -value assigned a cell in U^*_K entailing $\sim g$ is 0. Consequently, the contraction removing h from \underline{K} that minimizes loss of

²²Levi, 1991, p.143 offers a characterization of the relation "at least as corrigible as" consonant with the definition in the text. In Levi, 1996, the penultimate paragraph on p.263 assumes the representability of incorrigibility by a real valued measure but imposes only ordinal constraints on the measure consonant with the 1991 formulation. The characterization at the bottom of the page contains a serious error and is inconsistent with the formulation one paragraph before. $in(g/h, \underline{K})$ is equated with $D(g/h, \underline{K})$ for cases where g is in \underline{K} but not in all contractions. Everything around this passage presupposes what I intended to say -- to wit, that $in(g/h, \underline{K})$ is a decreasing function of $D(g/h, \underline{K})$. As just indicated, I was not intending to specify a definite quantitative measure but to specify ordinal properties of any acceptable quantitative measure and to adopt the convention that incorrigibility be kept within the limits of 0 and 1. The specific measure proposed here satisfies these requirements as will any real valued measure that is a positive monotone transformation of the one given in the text.

type 1 damped informational value and satisfies the requirements of the rule for ties recognizes that cell as a serious possibility and, hence, removes g as well.

$D_1(g/h, \underline{K}) < D_1(h/h, \underline{K})$ if and only if both conditions (a) and (b₁) fail.

$D_2(g/h, \underline{K}) = D_2(h/h, \underline{K}) = D_1(h/h, \underline{K})$ if and only if either condition (a) or the following condition (b₂) holds:

(b₂) The minimum M -value assigned a cell in U^*_K entailing $\sim g$ is no greater than the minimum M -value assigned a cell in U^*_K entailing $\sim h$. That is to say, $D(g/g, \underline{K}) \geq D(h/h, \underline{K})$.

$D_2(g/h, \underline{K}) < D_2(h/h, \underline{K})$ if and only if both conditions (a) and (b₂) fail. In this case $D_2(g/h, \underline{K}) = D_2(g/g, \underline{K}) = D(g/g, \underline{K})$. By way of contrast, $D_1(g/h, \underline{K}) \leq D_1(g/g, \underline{K}) = D(g/g, \underline{K})$.

According to the definition of incorrigibility, all sentences not in \underline{K} as well as sentences in \underline{K} but not in \underline{K}_h carry minimum or 0 incorrigibility whether or not loss of undamped informational value is of type 1 or of type 2.

g is in \underline{K}_h if and only if $in(g/h, \underline{K}) > 0$. If instead of choosing a contraction removing h maximizing damped informational value when $in(g/h, \underline{K}) > 0$, we choose one where the damped informational value is "good enough" in the sense that it is at least as great as $D(g/h, \underline{K})$ and, satisfies the requirements of the rule for ties among such satisficing contractions, the contraction sought is \underline{K}_{hvg} . The difference in damped informational value between this contraction and the contraction \underline{K}_h that is the product of maximizing informational value and using the Rule for Ties is equal to $in(g/h, \underline{K})$.

Whether $D(g/h, \underline{K}) = D_1(g/h, \underline{K})$ or $D_2(g/h, \underline{K})$, the corresponding incorrigibility measure satisfies the postulates laid down on p.264 of Levi, 1996.²³

Moreover, the values of $D(h/h, \underline{K})$, $D_1(g/h, \underline{K})$ and $D_2(g/h, \underline{K})$ are all determined by the M -function given \underline{K}_h , and U^*_K .

Degrees of incorrigibility exhibit the formal properties of Shackle measures (Levi, 1991, p.143 and Levi, 1996, p.267). In particular, the degree of incorrigibility of a conjunction is the minimum degree of incorrigibility of a conjunct and for every consistent \underline{K} , either the degree of incorrigibility of g or of $\sim g$ is equal to the minimum or to 0 with respect to contractions removing h . Moreover, when a corpus \underline{K} is contracted by giving up h , items with the lowest degree of incorrigibility are given up.

In these respects, they are formally similar to measures of what Gärdenfors (1988) has called degrees of entrenchment. But there are some differences. According to Gärdenfors, all items in \underline{K} carry positive entrenchment. Only items not in \underline{K} carry 0 or minimal entrenchment. By way of contrast, 0 or minimal incorrigibility is assigned to h and all items not in \underline{K}_h .

This difference reflects another difference as well:

- 1) Incorrigibility is relative to M , \underline{K} , U^*_K and h . Incorrigibility can be determined uniquely by M (given \underline{K} , U^*_K , and h).
- 2) Entrenchment is relative to M , \underline{K} , U^*_K but not to h . That is to say, entrenchment can be determined uniquely by M (given \underline{K} , U^*_K).

Let $D(g/g, \underline{K})$ be the largest damped informational value assigned to a potential contraction of \underline{K} that removes g among all potential contractions removing g from \underline{K} . This is going to be the damped informational value of \underline{K}_g . The maxichoice contractions ingredient in the intersection of saturatable contractions removing g that constitutes \underline{K}_g

²³A minor emendation is necessary. "Logically equivalent" in (in1) should be replaced by "equivalent given \underline{LK} " and "logical truth" in (in3) should be replaced by "consequence of \underline{LK} ".

all carry the same damped informational value which in turn equals their common cont-value and this cont-value is, therefore, equal to the damped informational value of $D(g/g, \underline{K})$ regardless of whether it is type 1 or type 2. And this cont-value is determined by the M -function. We then have the following definition of degrees of entrenchment.

Def. $en(g/\underline{K}, \underline{LK}, U_{LK}, M) =$ the difference between the damped informational value of \underline{K} and $D(g/g, \underline{K})$ where $D(g/g, \underline{K})$ is equal to the maximum cont-value (minimum M -value) assigned a maxichoice contraction removing g from \underline{K} .

So entrenchment is dependent upon the M -function just as incorrigibility is. But it is independent of a sentence that by supposition is to be removed from \underline{K} . Gärdenfors's postulates for entrenchment are satisfied. In particular, only sentences not in \underline{K} carry minimum or 0 entrenchment. Moreover, the ordering with respect to entrenchment does not depend upon whether damped informational value is of type 1 or type 2.

Suppose that $en(g/\underline{K}) < en(h/\underline{K})$ where both h and g are in \underline{K} . The following claims then hold:

$$(1) D(h/h, \underline{K}) < D(g/g, \underline{K}).$$

$$(2) D_1(g/h, \underline{K}) \leq D_1(h/h, \underline{K}).$$

(3) Hence, $in(h/h, \underline{K}) \leq in(g/h, \underline{K})$ if damped informational value is of type 1.

(4) On the other hand, if damped informational value is of type 2, $in(h/h, \underline{K}) = in(g/h, \underline{K})$

Thus, if the assessment of damped informational value is type 1, we cannot say that if g is less well entrenched than h in \underline{K} , we should remove g when removing h from \underline{K} (as Gärdenfors does on p.87 of 1988.) g may be more incorrigible and, as a consequence, less vulnerable to removal when removing h . On the other hand, if the assessment of damped informational value is of type 2, g must be as incorrigible as h so that removal of h will require the removal of the less well entrenched and equally as incorrigible g as well.

Under the suppositions made here, entrenchment is a better indicator of what is to be given up when it is stipulated that g is to be given up even when type 1 damped informational value is used.

$$(5) D_1(h/g, \underline{K}) < D(g/g, \underline{K}) \text{ because of (1).}$$

(6) $in(g/g, \underline{K}) < in(h/g, \underline{K})$ in agreement with the entrenchment ordering.

Gärdenfors, of course, thought of entrenchment as based on AGM partial meet contractions where the Recovery Condition is satisfied. Entrenchment so conceived cannot be derived from either type of assessment of damped informational value where Recovery can fail. And it cannot be derived from assessments of undamped informational value either as we have seen previously unless contraction is restricted to maxichoice contraction. If we wish to say as Gärdenfors does say on p.87 of 1988 that when \underline{K} is contracted, the sentences removed are those carrying minimal entrenchment while avoiding commitment to the Recovery condition, minimizing loss of damped informational value of type 2 satisfies our requirements.²⁴

²⁴In Levi, 1991, I thought of entrenchment as a rephrasal of what I had in mind by incorrigibility as did Gärdenfors in 1988. Because of this, I insisted in modifying Gärdenfors postulate (EE4) which fails for incorrigibility and, hence, so I thought, for entrenchment. In Levi, 1996, I acknowledged the differences between Gärdenfors's idea and mine. Still the definition of entrenchment given on p.265 is unsatisfactory. I prefer the definition furnished here. In addition, I falsely asserted that \underline{K}_h is the set of positively entrenched sentences in \underline{K} minus the set of sentences carrying entrenchment at least as small as that of h . This is, indeed, true if damped informational value is of type 2; but in Levi, 1996 I was still using damped informational value of type 1 or, more accurately, had not clearly separate the two.

Using damped informational value of type 2, it becomes apparent that whether one uses entrenchment or incorrigibility as a tool for representing appraisals of vulnerability to being given up is largely a stylistic matter. What is substantive, so I claim, is that neither mode of assessment acquires any significance for either belief change or suppositional reasoning except when derived from a representation of the agent's demands for information.

The interpretation of Shackle's formalism for potential surprise and belief in terms of inductive expansion rules supplies a way to understand the idea in terms of other ideas including subjective probability. In a similar spirit, deriving entrenchment or incorrigibility from the epistemic utility lost in contraction -- i.e., in terms of loss in informational value -- offers an understanding of entrenchment that merely taking it as a primitive idea constrained by postulates fails to provide. This is especially important since measures of entrenchment and incorrigibility are, subject to minor adjustments, subject to the same formal requirements as Shackle's measure of potential surprise and belief. And because of this formal similarity, it may be (and has been) tempting to think of entrenchment and degree of belief as usefully embeddable in a single ordering. A careful look at their interpretations, however, should disabuse us of such temptation.

I am suggesting, in the first instance, that such demands are representable by an informational value (or utility) determining probability over \underline{L} or, at least, over U_{LK} relative to \underline{LK} . Undamped informational value and damped informational value of types 1 and 2 are defined in terms of it. In the most general setting, I argue that demands for information are represented by a convex set of such M -functions. The ramifications of considering such sets in assessing contraction strategies will be addressed in section 21.

From this point of view, damped informational value is but a manifestation of the evaluation in terms of M -functions. This point is important to appreciate in addressing the question of iterated revision -- a topic I shall not consider here. Nonetheless, the utility of information function derived from an M -function is central to a decision theoretic account of contraction as seeking to minimize loss of informational value.

21. Indeterminacy in Informational Value

The main focus of the discussion to date has been on whether agents called upon to contract their state of belief seek to minimize loss of undamped informational value, damped informational value of type 1 or damped informational value of type 2. No matter which kind of loss is being minimized, the assessment of loss of informational value is itself a value judgment reflecting the inquirer's goals. Moreover, the assessment of loss of informational value is derivable from an assessment of informational value determining probability (as represented by an M^* -function relative to U_{LK} or by an M -function relative to U_K and \underline{K} in the context of expansion).

No assumption was made, however, as to whether there is a M^* - value that everyone ought to adopt. I do not believe that there is such a standard. We should allow that inquirers may have honest disagreements concerning the choice of M^* - function.

On the other hand, the choice of M^* -function can make a difference in the way in which inductive expansions and in the way contractions are evaluated. Insofar as it matters to us how disagreements as to what to add via induction or to remove in contraction, we should give some account of how inquirers who do disagree can move to an evaluation of informational value based on the aspects of their conflicting assessments that they share in common.

Consider a context where inquirers are debating which of rival theories to add to their corpora. Sometimes the disagreements between such inquirers reflect differences

in their research programs with respect to the explanatory virtues of the rival theories as much as it reflects a difference in the evidence. Demands for informational value may seem to overwhelm assessments of risk of error as contributors to the dispute. In such cases, it is desirable that inquirers who have good reason for resolving their disagreements be capable of doing so without begging controversial issues including differences with respect to their research programs and, hence, their demands for information. There is a way to formally represent such shared agreements and how they relate to the critical assessment of expansion strategies.

Demands for information that are ingredient in the goals of inductive expansion are representable by measures of content or undamped informational value and the M - functions that determine them. (Levi, 1967a, 1967b). In Levi, 1984, ch.7 and 1980, I suggested that the shared agreements between inquirers who assess undamped informational value by different M - functions could be represented by the set of weighted averages of the several M - functions (their "convex hull"). Each of the "permissible" M - functions in such a convex set could be combined with the utility of truth through weighted averaging to form a permissible epistemic utility function after the proposals of Levi, 1967a and 1967b. A permissible expected epistemic utility of any potential expansion strategy in a set determined by \underline{K} and U^*_K given a permissible credal probability function and a permissible epistemic utility function is then well defined. Each permissible expected epistemic utility function defines a set of potential expansion strategies (potential answers) that are optimal with respect to that expected utility.

According to the decision theory I have advocated since Levi, 1974 and elaborated in Levi, 1986, a rational agent should restrict his choices to *E-admissible* options. An E-admissible option is one that ranks best in expected utility according to at least one permissible expected epistemic utility function in the set of permissible expected utility functions defined over the range of available options. The set of expected utility functions could be larger than a singleton because the agent's judgments of probability are indeterminate and, hence, representable by a set of probability distributions (over states) or because of conflict or indeterminacy in the agent's values and goals that is representable by a set of utility functions. Problems of inductive expansion are situations where both kinds of indeterminacy can emerge. When the options are potential expansion strategies, these should be E-admissible. These are the expansion strategies that are optimal according to at least one permissible expected epistemic utility function.

In general, when more than one option is E-admissible, secondary value commitments may be invoked to recommend choice from subset of the E-admissible options. In some contexts for example, an agent may place a premium on security and choose an E-admissible option that maximizes his level of security among the E-admissible options. Some such idea seems to make sense out of the violations of the independence postulate that are held to be manifested in the choices favored by experimental subjects when confronted with Ellsberg and Allais predicaments without interpreting the behavior as violations of independence. (Levi, 1986.)

What is true in general applies to inductive expansion. When there is but one permissible expected epistemic utility function, there will always be a unique weakest optimal potential expansion strategy so that we can recommend breaking ties in expected epistemic utility by rejecting the smallest set of elements of the ultimate partition that maintains optimality. This is a special case where two or more options are E-admissible and a secondary criterion is invoked to reduce the set of E-admissible options to a smaller set. When two or more expansions are admissible, a wise person tries to suspend judgment between their conclusions. This injunction is coherently acceptable when there is a uniquely permissible expected epistemic utility function and all E-admissible options are optimal options.

Can the rule for ties be extended to apply when more than one expected epistemic utility function is permissible? Can we recommend choosing the inductive expansion strategy that is the intersection of all the E-admissible ones? We could if that intersection were itself E-admissible. In that event, we can recommend choosing the uniquely weakest E-admissible expansion strategy.

Unfortunately we cannot always guarantee a uniquely weakest E-admissible potential expansion strategy when more than one expected epistemic utility function is permissible. When there is one, it can be recommended as in the case where there is a uniquely permissible expected utility function. Otherwise, the best we can do is to recommend selecting an expansion strategy that is (1) E-admissible and is such that (2) no other E-admissible expansion strategy is weaker than it is. Call this the *weak rule for breaking ties*.

The upshot is that in the context of theory choice and other types of inductive expansion where there is conflict and, hence, indeterminacy in assessments of undamped informational value and credal probability, we cannot guarantee a unique recommendation of an potential expansion strategy. We cannot say that the intersection of the admissible potential expansion strategies should be adopted; for that intersection may not be E-admissible.²⁵

In inductive expansion, undamped informational value is aggregated with the utility of avoiding error to obtain an epistemic utility function characterizing the concern to obtain error free information.

In contraction, undamped informational value (or the M^* -function that defines it) is transformed into damped informational value (of type 1 or, as I now prefer, type 2). But just as the recognition of several M^* -functions as permissible is tantamount to recognizing several measures of undamped informational value and measures of epistemic utility as permissible, so too it requires recognition of several distinct measures of damped informational value as permissible as well.

In evaluating contraction strategies, however, we do not aggregate informational value with avoidance of error because in contraction the issue of avoiding error does not arise.

Because it is unnecessary to aggregate informational value with the utility of avoiding error and to worry about uncertainty, when attempting to identify optimal contractions relative to a single permissible evaluation of loss of damped informational value, any two evaluations that yield the same weak ordering of the potential contraction strategies will yield the same recommendations. One might argue that quantitative differences in assessment of informational value and, hence, in M^* -values are irrelevant to the assessment of contraction strategies.

Moreover, we are not really concerned with preserving the weak ordering induced by the M^* -function over its entire domain but only over elements of U^*_K ; for the evaluation of damped informational value of the second kind is uniquely determined by this assessment. That is to say, the weak ordering of all potential contraction strategies with respect to damped informational value of the second kind is uniquely determined by the weak ordering of all potential maxichoice contractions.

Quantitative considerations cannot be so readily ignored in the case of inductive expansion. For one thing, interest in (undamped) informational value is weighted against the concern to avoid error. For another, the concern to avoid error implies a concern to minimize risk of error. When facing risk, quantitative dimensions of value become relevant.

²⁵If, however, we consider the result of reiterating the effort at inductive expansion via bookkeeping to a stable conclusion (see Levi, 1967, 1980 and 1996), it turns out that the conclusion will be univocal if we follow the weak rule for ties as presented here and will coincide with the result of bookkeeping using the strong rule.

Consider then an agent with corpus \mathcal{K} concerned to remove h and facing a dual ultimate partition $U^*_\mathcal{K}$ with assessments of (undamped of type 2) informational based on a convex set of M -functions. In particular, consider three elements x, y and z of $U^*_\mathcal{K}$. Let the set of permissible M -functions be the convex hull of M_1 assigning values 0.1, 0.03 and 0.01 to x, y and z respectively and of M_2 assigning 0.01, 0.03 and 0.1 to the same elements. All other elements of the dual ultimate partition implying h receive higher M -values and, hence, could not qualify as optimal maxichoice contractions according to any permissible M -function. According to some permissible M -functions, meets of x and y and of z and y qualify as optimal contractions. But no permissible M -function recognizes the meet of all three as optimal. So it would appear that only the weak rule for breaking ties is operative in such cases. So the meet of three maxichoice contractions is not E-admissible and, hence, not admissible at all.

Keep in mind, however, that in contraction quantitative dimensions of informational value do not seem relevant. At least this is so when exactly one M -function is permissible. But in that case, we could just as well have adopted any other M -function preserving the same ordering over elements of $U^*_\mathcal{K}$. We need not preserve the order over other subsets of elements of $U_\mathcal{K}$.

Quantitative considerations do arise when we consider two M -functions such as our M_1 and M_2 to be permissible. We have supposed that the set of permissible M -functions is the convex hull of these two. But if we had considered the convex hull of M -functions that are positive monotone transformations of these two (i.e., preserve the weak ordering generated by these two M -functions when restricted to elements of $U^*_\mathcal{K}$, we would have a much larger set of M -functions. In this set, it will turn out that potential contractions that are meets of maxichoice or saturatable contractions that recognize x, y and z as possibilities are optimal according to some permissible M -functions.

This result is quite general. If we are given n M -functions M_i as permissible, there are two ways to proceed:

Method 1: Take the convex hull \mathbf{C} of the M_i 's. Construct the set of positive monotone transformations of members of \mathbf{C} . These give the permissible weak orderings of potential contractions with respect to assessments of damped informational value of type 2.

Method 2: Form the convex hull \mathbf{D} of the positive monotone transformations of the M_i 's. Construct the set of positive monotone transformations of \mathbf{D} . $\mathbf{C} \subseteq \mathbf{D}$. These give the permissible weak orderings of potential contractions with respect to assessments of damped informational value of type 2.

As illustrated above, method 1 does not guarantee that the recommendation of a strong rule for tie breaking will yield an E-admissible contraction. According to the decision theory I favor for cases where there multiple permissible rankings of the options, the use of the strong rule is then unacceptable as a general secondary criterion.

On the other hand, using method 2 does guarantee that the strong rule will always yield an E-admissible option.

Which method is acceptable given the supposition that loss of damped informational value of the second kind is to be minimized? Can we justify using method 2 rather than method 1 on the grounds that minimizing loss of damped informational value requires the use of ordinal information alone?

The answer to this question is not entirely clear to me. The pivotal difference between methods 1 and 2 concerns what count as potential resolutions of conflict between rival assessments according to each method. Method 1 takes the original quantitative assessments of undamped informational value or M -value, converts them into quantitative assessments of damped informational value of type 2, *considers potential resolutions of conflict between these rival assessments as weighted averages*

of the original quantitative assessments and then ordinalizes the results. Method 2 begins by ordinalizing the quantitative assessments of damped informational, taking potential resolutions as weighted averages of all quantitative measures that represent these orderings and then ordinalizing again.

Method 1 takes the quantitative assessments of undamped informational value and the corresponding quantitative assessments of damped informational value seriously. Method 2 considers them irrelevant to assessing what is to count as a potential resolution of a conflict between different demands for informational value. I have argued that the quantitative dimension is important in address inductive expansion. Should we therefore respect it in dealing with contraction even though it plays no role in determining either entrenchment or incorrigibility orderings relevant for identifying recommend contractions?

If we do, we preclude the use of the strong rule for ties in the context of contraction. This seems very unattractive. When two contraction strategies are admissible, we surely ought to be in a position to retain a posture of suspense or doubt regarding their conclusions.

In sum, there is a tension between two kinds of consideration: one arguing for method 1 on the grounds that the interest of expansion and contraction ought not to be divorced and the other for method 2 on the grounds that only ordinal considerations are needed in contraction and that the strong rule for ties should be deployed whenever one can with good conscience do so.²⁶

22. Fallbacks

When demands for information are conflicted in a manner that yields indeterminacy representable by a convex set of M -functions, each permissible M -function defines an assessment of damped informational value of type 2 and through that for each \underline{K} a permissible entrenchment ordering of the elements of \underline{K} . Each permissible entrenchment ordering yields a nested system of spheres in the sense of a Grove model. These spheres qualify as fallbacks in the sense of Lindström and Rabinowicz, 1990. If we consider all unions of the sets of fallbacks associated with each permissible ordering, we have a system of fallbacks of the sort considered by Lindström and Rabinowicz with an associated entrenchment ordering that allows for incomparabilities.²⁷

Lindström and Rabinowicz do not consider fallbacks to be contractions of \underline{K} . But they are contractions in the sense in which I have used the term since 1974. Maximal \sim -h permitting fallbacks in their sense are contractions of \underline{K} removing \sim -h that are E-admissible if permissible evaluations of losses of damped informational value are of type 2 so that an E-admissible contraction would have the properties of a mild contraction or a severe withdrawal. Lindström and Rabinowicz object to using mild contractions (and, by implication, damped informational value of type 2) on the grounds that contraction so conceived would be too drastic.

Part of their anxiety has been addressed previously by invoking coarse graining of potential contraction strategies with the aid of dual ultimate partitions and minimal corpora and another part has been addressed by pointing out that loss of damped informational value is not ordered by set inclusion.

²⁶If we follow method 2, the choice function that emerges will allow for failures of Sen's property β and intransitivity of revealed indifference in choosing contraction strategies.

²⁷Whether the entrenchment ordering is the categorical preference ranking of Levi, 1986 or the revealed preference ranking is unclear. And whether the permissible rankings are those allowed by method 1 or by method 2 is also unclear.

Whatever their reservations with fallbacks might be in discussing contraction, Lindström and Rabinowicz have no reservation with using it in discussing revisions. That is because they still are following the tradition according to which revision of \underline{K} by h when h is already in \underline{K} is taken to be \underline{K} itself. In that setting, it does not matter whether one uses a contraction function satisfying Recovery or not when considering revision. If it is advantageous to use mild contractions = severe withdrawals = fallbacks in discussing entrenchment -- especially when considering failure of weak order, one can use it without raising any problems for revision other than the loss of weak order. Given their reservations when contraction does not lead to revision and their legitimate objections to Recovery, Lindström and Rabinowicz in effect distinguish two kinds of contraction of interest to them: One they call "fallback" and the other "contraction".

I, on the other hand, think that their fallbacks will serve the purposes of contraction very well. This becomes important when one takes Ramsey revision seriously as I do.

My concern here is with the issue of noncomparability. Lindström and Rabinowicz take the position that when two or more $\sim h$ - permitting fallbacks are maximal (that is to say, are optimal according to some permissible complete entrenchment ranking) and, hence, are E-admissible in my sense, no further recommendation is to be made.

In section 17, I conceded that one might be driven to the conclusion that two or more contraction strategies should end up as admissible all things considered if one uses method 1 for determining the set of permissible entrenchment rankings; but I wish to emphasize that this conclusion is quite an uncomfortable one. The instincts of Gärdenfors and Makinson to favor skeptical responses where one suspends judgment in such cases seem quite sensible.

The problem, as I have insisted, is that in order to cater to these sensible instincts one must show that doing so is not merely to choose an available option and not merely the intersection of admissible options, but is to choose an admissible option. One can overcome this problem with the aid of method 2 and do so in a manner that is in keeping with a principled account of rational choice when utility goes indeterminate. The cost of doing so, however, is that the link between the assessments of informational value relevant to inductive expansion and the assessments relevant to contraction are somewhat weakened.

The weakening may or may not be harmless. That remains to be seen. My reservation with the attitude of Lindström and Rabinowicz concerns the complacency with which they confront the prospect of abandoning the skeptical view. All else being equal, even when entrenchment lacks a weak ordering, I would have thought the skeptical view an attractive one to endorse if one can honorably do so.

23. Infinite Ultimate Partitions: Finite Intervals:

Throughout this discussion of contraction, attention has been focused on basic partitions relative to \underline{LK} containing finitely many cells and the finite ultimate partitions and dual ultimate partitions relative to \underline{K} carved out of them. However, contexts arise where it is desirable to lift this restriction both for the purposes of inductive expansion and contraction. Krister Segerberg (1995, 1996) has devoted some attention to this topic in the development of his account of fallbacks. I shall summarize some partial suggestions regarding this topic by focusing on certain special kinds of cases where the need for thinking of infinity seems urgent.

My strategy will be to rehearse ideas I have presented in Levi, 1980 and 1996 concerning inductive expansion when the ultimate partition relative to \underline{K} is infinite and then draw some lessons from these suggestions regarding dual ultimate partitions

relevant to contraction and the properties of informational value determining M -functions relevant to defining undamped and damped informational value appropriate when dual ultimate partitions are infinite.

The first case I shall consider concerns the problem of estimating the value of a real valued parameter whose true value is believed to fall in some finite interval. That is to say, with the exception of at most a subset of Lebesgue measure 0, all point estimates in an interval from a to a^* are countenanced as serious possibilities relative to K and are elements of an ultimate partition U_K relative to K .²⁸

Suppose the expectation determining probability distribution Q , the informational value distribution M and the Lebesgue Measure are absolutely continuous in each other. That is to say, a subset of U_K has 0 Lebesgue measure if and only if it is assigned 0 Q -value and 0 M -value.

Consider some real point estimate h_x asserting that x is the precise true value of the parameter in question. Focus on the coarsened ultimate partition carved out of U_K by taking the hypothesis $h_{x,\epsilon}$ asserting that the true value is in the interval from $x - \epsilon$ to $x + \epsilon$ and the partition of the negation of this hypothesis into n alternatives each of which specify that the true value is in an interval of positive Lebesgue measure. All alternatives in this coarse grained partition carry positive Q - values and M - values (because of the absolute continuity conditions) so that we can apply the criteria for inductive expansion to determine which, if any, of the alternatives is rejected for any given value of q . Notice that the conditions for rejecting $h_{x,\epsilon}$ are the same regardless of how its negation is partitioned into finitely many alternatives. Hence, as long as it is assumed that $\sim h_{x,\epsilon}$ is partitioned into at most finitely many interval estimates of positive Lebesgue measure, we have a condition for rejecting $h_{x,\epsilon}$ that is the same regardless of the how partitioning is done.

Rejecting $h_{x,\epsilon}$, entails rejecting the truth of each point estimate in the interval. Let us suppose at a given level of boldness q there is an ϵ^* such that for every $\epsilon \leq \epsilon^*$, $h_{x,\epsilon}$ is rejected. We can then say that h_x (which is an exact point estimate in the noncountably infinite ultimate partition U_K) is rejected at the given value of q . That is to say, the hypothesis that x is the true value is rejected if and only if the hypothesis that the true value falls in a small interval around x is rejected relative to all to any finite partition of U_K of which such an interval estimate is an element no matter how small that interval around x might be.

If the cumulative distribution functions for the subjective probability Q and informational value determining probability M are continuous and differentiable so that they both define density functions $f(x)$ and $m(x)$ respectively, we can say that h_x in U_K is rejected at level q if and only if $f(x)dx < qm(x)dx$.

If this method of inductive expansion is reiterated, a result is reached that is stably accepted at the given degree of boldness q . If $q = 1$, the unrejected elements of U_K are those points for which $f(x)/m(x)$ is a maximum if there is one.

As already noted, U_K might, prior to inductive expansion of K , lack a set of points of measure 0. It might, indeed, lack the points for which $f(x)/m(x)$ is a maximum. In that case, at $q = 1$, iteration to stability will result in rejecting all elements of U_K . To avoid such contradiction, one should avoid being maximally bold. Given the continuity of m

²⁸Here the elements of U_K constitute a nondenumerably infinite set. In the case under consideration, we not only have the set but the elements of the set form an interval having a certain length and contain subsets that are subintervals of the interval from a to a^* and have lengths. (Elements of the set length 0.) Taking the class of all lengths, we can form other sets by taking countable intersections, unions and by complementation and extend the measure of length to apply to these sets to form the set function known as the Lebesgue measure. Not all subsets of U_K will have a well defined Lebesgue measure; but in this setting we will not need to rely on this point.

and f , at $q < 1$ a consistent stable conclusion will be reached. Adopting $q = 1$ is being excessively rash.

Setting aside the worries of the previous paragraph, can we insure the existence of a maximum value for $f(x)/m(x)$? Both $f(x)$ and $m(x)$ have finite and positive maxima and finite minima given the assumptions under which we are operating. Consequently, the ratio $f(x)/m(x)$ cannot be greater than the ratio of the maximum f -value and minimum m -value. If the minimum m -value is positive, the maximum must exist. If it is 0, there is no finite maximum. So for $q = 1$, stable inductive expansion into inconsistency would be recommended. Again, however, if $q < 1$, stable expansion into inconsistency is avoided.

Suppose the corpus is \underline{LK} with U_{LK} structured in the manner just discussed. Inductive expansion has taken place so that the expanded corpus \underline{K} states that the true value falls in some subinterval from b^- to b^+ or, perhaps, in the union of the set of subintervals of the interval from a^- to a^+ . The dual ultimate partition U_K consists of all the rejected point estimates.

It is important to keep in mind that losses of informational value by contraction using elements in U_K^* are not ordered by the density function $m(x)$. Each such point carries 0 M -value. If we used the ordering determined by the M -function and employed assessments of losses of type 2 damped informational value any contraction of \underline{K} removing some sentence in \underline{K} would automatically yield \underline{LK} . That contraction should yield \underline{LK} no matter what is contracted seems unreasonable.

Fortunately, type 2 assessments are not based on the ordering of the cells of U_K^* with respect to M -value in cases where the dual ultimate partition is carved out of a basic ultimate partition consisting of the points in a set of finite Lebesgue measure. Instead, attention is focused on finite coarse grained partitions of U_K^* into finitely many cells each of which carries a fixed positive but small Lebesgue measure and proceeding as we have done in the case of finite ultimate partitions. By doing this, we are treating all points in a cell as incurring equal loss of damped informational value. But comparisons of points in different cells are determined by the M -values of the cells.

Consider now the variant situation where the interval from a^- to a^+ is divided into finitely many intervals. Each interval carries positive Q -value and M -value and conditional Q -distributions and M -distributions are defined for each subinterval satisfying the conditions of the first case. But neither of the corresponding unconditional distributions over the total interval is continuous

As before, for any point estimate, we can adopt the rejection rule formulated for the case where the Q and M distributions are continuous over the entire interval with the following modification: h_x is rejected if and only if $f_i(x)Q(i)/m_i(x)M(i) < q$ where f_i and m_i are the densities conditional on the true value being in the interval I . Iteration to stable conclusions can proceed as previously.

The informational value determining M -value orders the points in the full interval in n different equivalence classes corresponding to the n intervals. Dual ultimate partitions will be treated in the same way as before.

24. Countably Infinite Ultimate Partitions:

Let U_K or U_{LK} contain countably many cells. According to the criteria for inductive expansion proposed in Levi, 1980 and 1996, rejection rules are formulated as follows:

Cell d of U_K is *n-rejected* if and only if there is at least one subset of n elements A of U_K containing d , such that if the corpus specified that exactly one element of A is true, d would be *stably* rejected at the given level of boldness q .

Cell d of U_K is rejected at the given level of boldness q if and only if there is an n^* such that d is *n-rejected* for every $n \geq n^*$.

The idea behind this rule is that d should be rejected provided that for all sufficiently large n , d is stably rejected relative to the n -fold subset A of U_K of which d is a member most favorable to rejecting d at the given level q . If d avoids rejection according to this test, it goes unrejected.

In the countably infinite case, the Q and M functions can be either countably additive or finitely but not countably additive. Moreover, the finitely but not countably additive distributions can “purely finitely additive” distributions that assign equal value to all cells and, hence, assign 0 value to all cells or they can be weighted averages of purely finitely additive and countably infinite distributions. (Schervish, Seidenfeld and Kadane, 1984.)

In order to apply the rejection rule just given, conditional probabilities on finite subsets of U_K must be well defined even though the unconditional probabilities for the finite subset are equal to 0. As de Finetti argued, there should be no objection to such definability for subjective probabilities and the same, so I contend, is true for informational value determining probabilities.

Thus, even though two finitely additive probabilities assign 0 value to all elements of U_K , they may, nonetheless differ in the conditional probabilities they assign to a given finite subset A of U_K . One might assign equal probability $1/n$ to all elements of A whereas another might assign unequal probabilities.

The difference may be captured by means of a σ -finite measure defined over U_K -- i.e., a countably additive measure F assigning positive values to the elements of U_K but where the countable sum need not be finite. One can use the F -function to define conditional probabilities in the usual way. $Q(d/A) = QF(d)/\sum_{d' \in A} QF(d')$ where A is a finite subset of U_K and $M(d/A) = MF(d)/\sum_{d' \in A} MF(d')$. When Q (or M) are countably additive, $QF = Q$ (or $MF = F$).

This means that it is possible to weakly order the elements of U_K with respect to the given probability. If d and d' are cells in U_K , let A be the join of the two cells. Given any finite set B that is a superset for A , the probability of d is at least as great as the probability of d' conditional on A if and only if the same holds conditional on B . Moreover, the comparison remains the same when the conditional probability is defined conditional on the true value falling in an arbitrary finite superset of A . This holds whether or not the probability measure is countably additive or merely finitely additive.

The inductive expansion rule or rejection rule for the countably infinite case has the following defect. There are Q and M distributions relative to which all elements of U_K are rejected for all positive values of q . To avoid this result, we may either abandon the inductive expansion rule or impose restrictions on the allowed Q and M distributions. In my judgment, the inductive expansion rule may be retained by imposing relatively modest constraints on the range of M -functions.

It would be undesirable to rule out countably additive Q -functions playing the role of expectation determining probability distributions because they are commonly employed in applications. Moreover, it seems desirable to be in a position to assign equal probability to each element of a countably infinite U_K both when probability is expectation determining and informational value determining. Indeed, one of the motivations for introducing finitely but not countably additive Q - or M -functions has always been to allow for probability judgments of this kind.

If an informational value determining M -function is countably additive over a countable domain and all elements of U_K carry positive M -value, there will be no smallest positive value for M . The greatest lower bound of such values will be 0. Otherwise the total probability for U_K would be infinite. Take any element d of U_K and let $M(d) = M^*$. Take the ratios $M(d'')/M^*$ for d'' such that $M(d'') < M^*$. These are also bounded by 0.

Suppose the inquirer's expectation determining probability function Q is finitely additive and is represented by a QF -function assigning equal positive value to each element of U_K . For any element x of subset A_n of n elements of U_K , the probability of d conditional on the true value falling in A_n , is $1/n$. I shall now argue that the use of a countably additive M -function must lead to inductive expansion into inconsistency for every positive value of q relative to such a Q -function. I conclude that countably additive M -functions should not be used.

Every element of U_K is 2-rejectable at every level of boldness $q > 1/2$. Just take any element d , identify a d' such that $M(d')/M(d)$ is very close to 0. This is always possible under the assumption that the M -function is countably additive. d is rejected on the assumption that d or d' is true for values of $1 \geq q > 0.5$. So d is 2-rejectable for values of q in that range. Add a d'' to the set of alternatives that carries lower M -value than d' . d is rejected for values of q ranging from near 0.33^+ to 1 and, hence, is 3-rejectable. d remains n rejectable for increasing n and for values of q whose lower bound approaches 0 as n increases. So d is rejectable for every positive value is q . No matter what positive degree of boldness is exercised, it is possible for circumstances to arise where the rejection rule leads to inductive expansion into inconsistency. Since we are allowing Q -functions represented by constant QF -functions, we must remove countably additive probability functions as informational value determining.

Given the same Q -function, we should now consider finitely additive M -distributions. There are three cases to consider:

- (a) The MF -function has a minimum in all subsets of U_K .
- (b) The MF -function has no minimum but has a positive glb.
- (c) The MF -function has no minimum but has a 0 glb.

Case (c) can be shown to lead to inductive expansion into inconsistency for all values of q by the same reasoning used in the countably additive case. Case (b) consists of cases where inductive expansion into inconsistency occurs for values of q less than 1 but greater than some specific positive value depending on the glb. Even so, any specific level of boldness can lead to inductive inconsistency provided the glb is low enough. We may be prepared to prohibit $q = 1$; but it is far from obvious that we want to prohibit any other positive value. So it may be concluded that cases (b) and (c) should be ruled out.²⁹

Let us now turn to Q -functions. Consider first countably additive Q -functions. In this case, such Q -functions always have a least upper bound. Provided that the M -functions are finitely additive and have MF -functions with minima, rejection rules avoid expansion into inconsistency as long as $q < 1$ and avoid expansion into inconsistency in this case as well as long as the least upper bound is also a maximum. The same is true for finitely additive Q -functions. If the QF -function is bounded from above, no expansion into inconsistency is allowed for the case where $q < 1$ and if the upper bound is a

²⁹In Levi, 1996, p.197, I contended that finitely but not countably additive M -functions should not assign elements of U_K more than finitely many different probability values. That is true enough; but irrelevant. The question of relevance here is whether the σ -finite function MF should take on more than finitely many values. Every subset of U_K should have a cell with a minimum MF -value in the subset. But there is no requirement needed that it have a maximum.

maximum, in the case where $q = 1$. However, if the QF -function is unbounded, inconsistency is possible.

The chief reason for running through these considerations in connection with inductive expansion is to attend to the problem of contraction. Here we are concerned with the MF -function defined for U^*_K . This function is a restriction of the MF -function for U_{LK} and this by reasoning already deployed for inductive expansion must have a minimum. The upshot is that the ordering determined by the MF -function must have a minimum in U^*_K and in every subset of U^*_K so that the smoothness (Kraus, Lehmann and Magidor, 1990) or stopperedness (Makinson, 1989) or limit (Lewis, 1973, Segerberg, 1996) condition must be satisfied. Instead of imposing the condition directly, I have proposed here and in Levi, 1996 to derive it from properties of the informational value determining M -function and the idea that inductive expansion rules should not allow inductive expansion into inconsistency.

25. Estimating a real valued parameter lying on the real line:

If U_K consists of all real point estimates on the real line, we may think of proceeding as recommended in section 23. But it will not do the partition the real line into two sets: a small interval around x and the rest; for the rest receives Q an M -value 1. Nor does it hold to divide the real line into intervals of equal Lebesgue measure L . We now have a countably infinite partition.

But, when facing a countably infinite partition, we can deploy the techniques of section 24.

If the problem is to determine which of the interval estimates so constructed is true, apply the inductive expansion rule specified in section 24 at a given level q . Those intervals that are rejected are L -rejected. If such an interval is L -rejected so are all point estimates contained in it.

Having introduced the notion of a point estimate being L -rejected, we can then state conditions under which it is rejected mimicking the ideas of section 23.

A point estimate is rejected if and only if there is an L^* such that for all $L < L^*$, the point estimate is L -rejected. For $q < 1$, some point estimate must go unrejected.

It is clear from this that the MF -function used must have a minimum. The same must be true, therefore, in the finite interval case discussed in section 23. Observe, however, that in the finite interval case, even countably additive MF -functions can have minimum values.

These considerations are aimed, as before, at the case of contraction. As in section 23, we may suppose that U^*_K consists of cells incompatible with K of some fixed small Lebesgue measure. According to the restrictions imposed here, there should be a minimum - M^* -value assigned to cells in $U^*_K = U_{LK}/U_K$. Smoothness, stopperedness or a limit assumption is secured by linking the assessments of losses of damped informational value of type 2 to assessments of undamped informational value in inductive expansion and worries about permitting expansion into inconsistency that arise of constraints implying such an assumption are not in place.

These reflections also furnish another more general reason to be suspicious of undamped informational value of type 1 in the context of contraction. The primary motivation for assigning 0 M^* -value to elements of U_{LK} concerns cases where the basic ultimate partition is infinite as it is in examples such as those considered here. Damped informational value of type 2 yields violations of the Recovery postulate only if the assignment of M^* -value of 0 is to be understood as the assignment of 0 to $m(x)dx$ in the

finite interval estimation task of 2.3 or 0 *MF*-value to an interval of positive Lebesgue measure in the setting of 2.5 or 0 *MF*-value to a cell in the setting of 2.4. In these settings 0 means “impossible” in the sense of “incompatible with LK”.

But as long as LK is the minimal corpus, contraction from the expansion K of LK forbids contractions allowing such impossibilities to be countenanced as serious possibilities. That calls for altering LK which had not been my intention. This consideration may be added to the others that support the use of damped informational value of type 2 rather than of type 1.

Appendix A

Base Contraction and the Filtering Condition

Throughout this discussion, contraction has been understood to be a transformation of the current state of full belief to a weaker potential state of full belief relative to which consistent expansion to a potential state of full belief incompatible with the current state is available. I have focused attention on potential states of full belief only insofar as they are representable by deductively closed sets of sentences (*corpora or theories*) in some suitably specified language \underline{L} . A contraction can be represented by a transformation of initial corpus \underline{K} representing belief state \mathbf{K} to another potential corpus \underline{K}_h representing belief state \mathbf{K}_h . h is a sentence in \underline{L} representing the potential state of full belief also represented linguistically by the $Cn(\{h\})$.

Several authors (most notably A. Fuhrmann, 1991 and S.O. Hansson, 1989) have objected to the use of deductively closed theories or corpora to represent potential states of full belief because they lack surveyability by means of finite resources. They suggest instead that attention be focused on contraction of finite sets of sentences whose closures represent potential states of full belief. Such a finite set \underline{A} is called a base, basis or belief base for its closure \underline{K} . $Cn(\underline{A}) = \underline{K}$.

There can be no objection to alternative equivalent representations of potential states of full belief. But those who explore the properties of base contraction do not seem to understand contraction of the base \underline{A} of a theory \underline{K} by sentence h in a manner equivalent to contraction of \underline{K} by h .

In both cases, h is taken to be a sentence in \underline{K} . Otherwise, contraction is just an identity transformation. However, in contraction from a base, the *options* for contraction are restricted by the requirement that removal of h from \underline{K} be representable as the deductive closure of some subset of \underline{A} . For example, if \underline{K} is the deductive closure of $h \wedge f$ and $\underline{A} = \{h \wedge f\}$, the base contraction of \underline{K} removing h is restricted to contraction that removes $h \wedge f$ and yields as corpus the set of logical truths. If, however, $\underline{A} = \{h, f\}$ or $\{h \wedge f, f\}$, $Cn(\{f\})$ is also a potential contraction removing h that may be evaluated along with the set of logical truths with respect to their merits as potential contractions removing h from \underline{K} with respect to the goals of contraction -- such as minimizing loss of informational value in one of the senses discussed previously. (Notice that contraction to the logical truths violates the Recovery Condition. Contraction to the logical consequences of f does not.)

There are other bases available for the same \underline{K} that impose different restrictions on contraction removing h . Suppose for the sake of simplicity that the minimal corpus \underline{LK} is the set of logical truths and the dual ultimate partition consists of $h \wedge \sim f$, $\sim h \wedge f$, and $\sim h \wedge \sim f$. Every potential contraction in the sense of a meet of saturatable contractions is representable as the deductive closure of $h \wedge f$ with a disjunction of some subset (including the empty set) of the elements of this dual ultimate partition. If every such disjunction is a member of a basis \underline{A}^s , the set of available contraction strategies for contracting from \underline{K} relative to \underline{A}^s is identical with the set of meets of saturatable contractions. Relative to that basis (which is a finite basis) there is no restriction at all on the potential contractions relative to the dual ultimate partition that are optional for the inquirer. Relative to some such contractions, Recovery will be violated.

One can also consider as a basis all disjunctions with $h \wedge f$ of elements of the dual ultimate partition that entail $\sim h$ and no others. The set of contraction options will be restricted to partial meet contractions. Recovery must then be satisfied.

It is easy to see that contraction from a consistent, finitely axiomatisable corpus where the set of options is the set of all contractions of \underline{K} removing h relative to finite \underline{U}_K

is a special case of contraction from a basis as conceived by both Fuhrmann and Hansson.

Let \underline{U}_K be the disjunction of all elements of the ultimate partition U_K relative to the current corpus \underline{K} . \underline{U}_K implies \underline{LK} . So does every element of the dual ultimate partition \underline{U}_K^* . Every potential contraction of \underline{K} is representable as the logical consequences of the disjunction of \underline{U}_K with a subset S of elements of \underline{U}_K^* . Assuming that \underline{K} is finitely axiomatisable and that the basic partition relative to \underline{LK} is finite, we have the finite surveyability sought by Fuhrmann and Hansson.

If \underline{K} entails h , each such potential contraction will be a potential contraction of \underline{K} removing h if and only if at least one of the elements of S entails $\sim h$. If members of a subset of such potential contractions are counted as the ones to be considered optional, the basis is constituted by a set of axioms for \underline{K} consisting of \underline{U}_K and each of the disjunctions of \underline{U}_K with a subset S that represents an option in the subset. If the axiom system contains \underline{U}_K or some sentence equivalent given \underline{LK} with it as the sole axiom, we regard contraction to \underline{LK} as the sole available option. The basis used could include all disjunctions of \underline{U}_K with subsets S of \underline{U}_K^* that contain at least one element entailing $\sim h$. That is to say, it could recognize all potential contractions as defined above to be available. It could be more restrictive and recognize only disjunctions of \underline{U}_K with sets of members of U_K all of which entail $\sim h$. The available contractions would then be partial meet contractions. There are, of course, many other kinds of bases as well.

What I have been arguing here is that contraction from a basis is equivalent to contraction from a theory or corpus relative to a restriction of the set of contractions considered optional in some specific way. The case where no restriction is imposed except that imposed by relativisation to the dual ultimate partition (or the basic partition) is a special case associated with a special kind of basis. So is the case where the restriction is to partial meet contractions.

I have urged that once the minimum corpus \underline{LK} and the dual ultimate partition \underline{U}_K^* are given, the available options for contracting \underline{K} by removing h should be the meets of all subsets of saturable contractions removing h from \underline{K} . According to the understanding of a basis as specifying what the set of available options for contraction are to be, the basis should consist of the set of disjunctions of \underline{U}_K with some subset S of members of \underline{U}_K^* . No restriction should be imposed on contraction options except those introduced by reference to the minimum corpus \underline{LK} and the dual ultimate partition \underline{U}_K^* . Elimination of such options can be justified but the justification should show such options to be suboptimal. One should not eliminate such options by stipulation. Base contraction allows for eliminating contraction strategies by stipulation.

Insisting on base contraction because of a desire for surveyability misses the mark. Surveyability can be achieved by requiring finite axiomatisability and finite dual ultimate partitions. (I do not insist on such finitisation as a general rule. But I see no reason to prohibit consideration of it.) There is no need to appeal to base contraction except in the degenerate sense in which contraction of a finitely axiomatisable theory relative to a dual ultimate partition is a base contraction.

In sum, if non degenerate base contraction is to be taken seriously, the case for doing so will ignore the demand that restrictions on the options for contraction be justified by appealing to the goals of the inquiry.

Illustrations offered by Fuhrmann and Hansson suggest that they are committed to restricting the options for contraction for reasons having little to do with the aims of inquiry. Recently Makinson (1997) has been quite explicit about it.

Fuhrmann and Hansson do invoke another consideration when illustrating the import of bases by means of examples. There is Hansson's famous example in (1989, 117-8) concerning someone looking for a hamburger in town on a public holiday. He

knows there are two hamburger joints A and B in the town but at the beginning he is in suspense as to whether either one of them is open. He meets a friend eating a hamburger and concludes that one of the restaurants is open. He then has a basis $\{A \vee B\}$. He then sees the light on restaurant A and concludes that it is open. His basis is now $\{A, A \vee B\}$. Upon reaching restaurant A, he finds the restaurant closed. He needs to replace his corpus by one containing $\neg A$. Should it contain $A \vee B$? Hansson and Fuhrmann both think it should.

Change the story a bit. The person looking for a hamburger does not meet anyone munching a hamburger. But he sees the lights of restaurant A. So his basis is $\{A\}$. If he then sees the restaurant to be closed, he should not retain $A \vee B$.

The point to emphasize here is that what seems to determine the constitution of the basis is that in some unexplained sense, certain beliefs are acquired directly. In the first scenario, $A \vee B$ is acquired directly (by seeing the person eating the hamburger). It is also deducible from A.

In the second scenario, $A \vee B$ is not acquired directly but only via deduction from A. A, on the other hand, is acquired directly.

I do not know what I am attributing to Fuhrmann and Hansson when I allege that they draw a distinction between beliefs acquired directly and those that are mere consequences of what is acquired directly. Neither Fuhrmann nor Hansson use the terminology I have used. They let their examples speak for themselves. What is clear is that the constitution of the base \underline{A} is determined by how a belief is acquired or, perhaps, by how it is justified.

Given the base \underline{A} , Fuhrmann goes further and introduces a “filtering condition” that stipulates that $A \vee B$ should be removed in the second scenario when A is removed because the only reason it is in the corpus \underline{K} is the presence of A. Remove that reason and $A \vee B$ should be removed also. (Fuhrmann, 1991, 184-6.)

On the other hand, in the first scenario, $A \vee B$ should be retained because there is an independent reason additional to the reason for having A in \underline{K} in the first place for having $A \vee B$ in \underline{K} .

There is more than a whiff of epistemological foundationalism ingredient in the appeal to the filtering condition. It appears that when a consequence of the basis \underline{A} is challenged (i.e., when a sentence h is to be removed from $\underline{K} = Cn(\underline{A})$), we need to attend to the way in which an item gained entry into \underline{K} in order to determine what is to be given up. This is so whether we are concerned with the causal factors that contributed to the entrance into \underline{K} or with the reasons that justify such entry.

I follow C.S. Peirce in rejecting any such pedigree epistemology. (See Levi, 1980, ch.1.1.) Deciding what to give up, according to my proposals, depends upon determining which contraction strategy minimizes loss of informational value. The approach to fixing belief I have advocated for three decades is neither foundationalist nor coherentist. I am a pragmatist. As Peirce insisted, in inquiry we seek to remove doubt. We should emulate scientific inquirers in seeking to remove such doubts in a manner that avoids error as well. But in giving up beliefs, avoidance of error cannot be an issue. From the point of view of the inquirer, all of his current full beliefs are true. In giving up beliefs, one cannot be importing error. So there is no risk of error. But doubts are introduced. Information is lost.

In Hansson's first scenario, one might indeed not give up $A \vee B$. The reason would be that in doing so one would have to call into question background assumptions like the circumstance that a person was seen eating a hamburger or that the presence of someone eating a hamburger is a reliable indicator of the fact that some hamburger joint is open. The loss of information incurred might seem too great to bear. This

consideration has nothing to do with the way one acquired the belief that $A \vee B$. The point to consider is that the agent also believes that someone is eating a hamburger. In giving up A, one will either have to give up this claim (one form of loss of information) or give up the claim that $A \vee B$ so that the presence of such an agent is left unexplained or retain $A \vee B$. The latter option incurs less loss of informational value than the other alternatives regardless of how the agent acquired his convictions in the first place.

In the second scenario, there is no need to explain the presence of a hamburger eating agent or to call into question the reliability of the presence of such a person as a reliable indicator of the fact that some hamburger joint is open. The loss of type 2 damped informational value incurred by giving up $A \vee B$ together with A could very well be no greater than giving up A alone so that giving up both is to be recommended.

Hansson's example is thus accommodated without appealing to any reasons for coming to believe at all. Only losses in informational value are taken into account.

Both Fuhrmann and Makinson (1997) seem to think that if one does not appeal to what Makinson takes to be a justificatory structure that is associated with a basis, one is not contracting from a basis but from a theory and in this setting the Recovery Condition should obtain. That is because the filtering condition cannot be applied. Why the fact that the filtering condition cannot be applied should justify insisting on Recovery is not clear to me. Neither Fuhrmann nor Makinson offer any considerations establishing this result. And I have sketched here procedures whereby Recovery can be violated even though the filtering condition does not apply.

The point I wish to belabor is that justification is needed, when it is needed, for *changes* in belief and not for the beliefs the inquirer currently has. As a consequence, once it is settled that h should be removed from \underline{K} , whether some item f in \underline{K} is to given up as well or not depends entirely on the inquirer's *demands for information*. These demands are expressed by the dual ultimate partition U^*_K and the assessments of these corpora for undamped informational value. These two modes of evaluation indicate which kinds of losses of information matter to the inquirer and how different losses that do matter to the inquirer compare with one another.

Why should one ever give up information? This and related questions are discussed in Levi, 1980 and 1991. In this essay, only one issue is under discussion. Given that h is to be removed from \underline{K} , which of the many ways of contracting should be adopted? The salient difference between alternative contractions removing h is that different types of information are lost. Doubt is injected in different ways. The choice between alternative contraction strategies should be based on a consideration of which strategies minimize loss of valuable information or inject doubt in the least disturbing manner. The choice should not be based on the grounds invoked to add the information being considered for removal. Information is valuable if it contributes to explanatory intelligibility, to prediction, to the understanding of the subject matter under study. Such informational value is poorly correlated with the basis on which the information was acquired in the first place

In this essay, I have sought to remove some obscurities from the technical development of my views on contraction--views I have endorsed in broad outline since the early 1970's.³⁰ The important work of AGM prompted me to elaborate my ideas in more technical detail than I had done hithertofore. (Levi, 1991, 1996.) Several errors and obscurities crept into these accounts calling for amendments. But the account of

³⁰My first elaboration of these ideas appeared in "Truth Fallibility and the Growth of Knowledge". I read this paper at several Universities from 1971 onward and finally at the Boston Colloquium for Philosophy of Science in 1975. The paper was not published by the Boston Colloquium until 1983. It was reissued as chapter 8 of Levi, 1984. Another essays of the period that emphasize that in contraction loss of informational value is to be minimized is "Subjunctives, Dispositions and Chances" that appeared in *Synthese* 34 in 1977 and is reprinted as chapter 12 of Levi, 1984

CONTRACTION AND INFORMATIONAL VALUE

contraction I offer remains faithful to the philosophical perspective I have adopted for at least three decades. Contraction is not justified in virtue of the need to justify beliefs already held. And when it comes to evaluating contraction strategies, the concern is to minimize loss of valuable information. The pedigree of beliefs is irrelevant.

REFERENCES

- Alchourrón, C., Gärdenfors, P., and Makinson, D. (1985), "On the Theory of Logic Change: Partial Meet Functions for Contraction and Revision," *Journal of Symbolic Logic* 50: 510-30.
- Fuhrmann, A. (1991), "Theory Contraction through Base Contraction," *Journal of Philosophical Logic* 20: 175-203.
- Fuhrmann, A. (1997), *An Essay on Contraction*, Stanford: CSLI publications and Folli.
- Gärdenfors, P. (1988), *Knowledge in Flux*, Cambridge, Mass.: MIT Press.
- Hansson, S.O. (1989), "New Operators for Belief Change," *Theoria* 55, 113-32.
- Hansson, S.O. (1991), "Belief Base Dynamics," Ph.D. dissertation at Uppsala.
- Hansson, S.O. and Olsson, E.J. (1995), "Levi Contractions and AGM Contractions: a Comparison," *Notre Dame Journal of Formal Logic* 36, pp.103-119.
- Harman, G. (1986), *Change in View: Principles of Reasoning*, Cambridge, Mass.: MIT Press.
- Kraus, S., Lehmann, D. and Magidor, M. (1990), "Nonmonotonic Reasoning, Preferential Models and Cumulative Logics," *Artificial Intelligence* 44,167-201.
- Levi, I. (1967), *Gambling with Truth*, New York: Knopf. reissued in 1973 in paper Cambridge, Mass: MIT Press.
- Levi, I. (1974), "On Indeterminate Probabilities," *Journal of Philosophy* 71: 391-418.
- Levi, I. (1980), *The Enterprise of Knowledge*, Cambridge, Mass.: MIT Press. reissued in 1983 in paper.
- Levi, I. (1984), *Decisions and Revisions*, Cambridge: Cambridge University Press.
- Levi, I. (1986), *Hard Choices*, Cambridge: Cambridge University Press.
- Levi, I. (1991), *The Fixation of Belief and Its Undoing*, Cambridge: Cambridge University Press.
- Levi, I. (1996), *For the Sake of the Argument*, Cambridge: Cambridge University Press.
- Lewis, D. (1973), *Counterfactuals*, Cambridge, Mass.: Harvard University Press.
- Lindström, S., (1990) "A Semantic Approach to Nonmonotonic Reasoning: Inference Operations and Choice," unpublished manuscript.
- Lindström, S. and Rabinowicz, W. (1990), "Epistemic Entrenchment with Incomparabilities and Relational Belief Revision," *Uppsala Philosophy Reports*, No.2, Department of Philosophy, Uppsala, Sweden.
- Makinson, D. (1987), "On the Status of the Postulate of Recovery in the Logic of Belief Change," *Journal of Philosophical Logic*, 16: 383-94.

CONTRACTION AND INFORMATIONAL VALUE

Makinson, D. (1989), "General Theory of Cumulative Inference," in *Non-Monotonic Reasoning* edited by J. de Kleer, M.L. Ginsberg, and E.Sandewall, pp.1-18. Lecture Notes on Artificial Intelligence no.346. Berlin: Springer-Verlag.

Makinson, D. (1997), "On the Force of Some Apparent Counterexamples to Recovery," *Normative Systems in Legal and Moral Theory* ed. by E.G.Valdés, W. Krawietz, G.H.vonWright and R. Zimmerling, Berlin: Dunker and Humboldt.

Meyer, T.A., Labuschagne, W.A., Heidema, J. (1998), "A semantic weakening of the recovery postulate," unpublished manuscript.

Pagnucco, M. and Rott, H. (1997), "Severe Withdrawal (and Recovery)", unpublished manuscript.

Peirce, C.S., (C.P.) *The Collected Papers of Charles Sanders Peirce*, 8 vols., Belknap Press. References in the text take the form (CP.n.m) where n is the volume and m the paragraph number.

Quine, W.V. (1970), *The Web of Belief*, New York: Random House.

Rott, H. (1993), "Belief Contraction in the Context of the General Theory of Rational Choice," *Journal of Symbolic Logic* 58: 1426-50.

Schervish, M.J., Seidenfeld, T. and Kadane, J.G. (1984), "The Extent of Nonconglomerability of Finitely Additive Probabilities," *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete* 66, 205-66.

Seegerberg, K. (1995), "Belief revision from the point of view of doxastic logic," *Bulletin of the JPL*, v.3, pp.535-53.

Seegerberg, K. (1995), "Three Recipes for Revision," in *The Parikh Project*, Uppsala: Uppsala Prints and Preprints in Philosophy, no.18, 102-12.

Spohn, W. (1988) "Ordinal Conditional Functions: A dynamic theory of epistemic states," ed. by W. Harper and B. Skyrms, *Causation in Decision, Belief Change and Statistics, II* Kluwer Academic Publishers, 105-34.