

Chapter 1

The Problem, the Claim, and the Plan

1.1 On Solving Problems

For more than forty years, everyone seems to agree that problem solving is central for the understanding of reasoning in the sciences. In the philosophy of science the idea starts at least with [Kuh62]; its consequences are most methodically studied in [Lau77]. Later, with the renewed interest for scientific discovery and scientific creativity, discovery was directly linked to problem solving—see especially [Nic80c] and [Nic80d], but also, for example [Sim77], [Sim87] and [LSBZ87].

There is something worrying about the situation. Simon and his followers—see for example the last three references—have done tremendously useful work in devising a set of computer programs that, once fed with the right data, come up with most impressive results. However, these programs turn out to have a rather restricted range of application, viz. to cases where the problem solving process can be reduced to finding an equation that fulfils certain criteria and is compatible with the data—the equation may refer to concepts that were not present in the data, but are constructed by the programs. The fact that Simon and followers underestimated the restrictive applicability of their programs led them to make some awfully anachronistic claims—see [Hal93] and already [Nic90, p. 9] for some tough criticism.

Let us now turn to work from philosophy of science in the area of problem solving. Kuhn, Laudan, and many other philosophers of science may say something about the quality of a proposed solution—often even this lacks precision—but hardly ever consider the process that led to a solution. The major exception in this respect is Thomas Nickles, who carefully discussed (from the 1980's on) the role of *constraints* in problem solving processes, including the change of constraints, and their (rational) violation—see [Nic80a]. Nevertheless, even Nickles does not offer a clue on how one should proceed in order to solve a problem.¹

That computational approaches appear too specific suggests that one should look for a more general approach. In the following chapters, I spell out a formal

¹Nickles has an excuse here: he does not believe in the existence of a general method for solving problems. I shall discuss this later.

approach to problem solving, which incidentally is extremely general. Some indications that such an approach is problematic or even impossible are discussed in Section 1.2. In the present section, I offer some further comments on problem solving in general.

From the start a terminological point should be considered. ‘Problem’ is a notoriously ambiguous term, referring on the one hand to difficulties that one tries to surmount, and on the other hand to questions that one tries to answer. The kind of problem solving I have in mind clearly concerns questions.

The (rather confusing) relation between difficulties and questions (both in their sense that is synonymous to a use of “problems”) requires some comments. Let me state from the outset that any question that is raised for a good reason, derives directly or indirectly from a difficulty. That I hurt my foot may be a difficulty, especially if I have to walk a long way home. The difficulty may give rise to some questions that I try to answer from my knowledge system—in more complex cases by extending my knowledge system. Remark, however, that one’s knowledge system may itself involve difficulties or run into difficulties. Thus my inability to find out whether a given logical system is decidable may form a very serious difficulty. Incidentally, whether it is difficult or not to answer a question is independent of whether the question derives from a difficulty. For example the question “What is the number of hairs on my head?” is certainly difficult to answer, but I can hardly see any serious difficulty that would form a good reason to raise the question.

Apart from what was said in the previous paragraph, the word “problem” in the title of these notes should be taken in the broad sense. No restriction on kinds of problems or on fields is intended. Moreover, solving scientific problems does not appear extremely different from solving so-called everyday problems. The kind of reasoning behind both processes seems to be similar, even if there may be differences in technicality, systematicity, etc.

That problems should be understood as questions requires a technical comment. Actually, I shall handle problems as *sets* of questions. Some sets may be singletons, but others are not, as will appear clearly in later chapters. The main reason for proceeding thus is the following. Suppose one starts from the (singleton) problem “Is A the case or not?”—formally $\{?\{A, \sim A\}\}$, a set containing a single yes–no question. Suppose next that one finds out that the question $?\{A, \sim A\}$ is answered (in one way or the other) if one is able to establish all of B , C and D . It seems natural then to introduce the questions $?\{B, \sim B\}$, $?\{C, \sim C\}$ and $?\{D, \sim D\}$. But in a sense these questions are connected. Only specific answers to all three questions are useful for solving the original problem. Hence, as soon as one question is answered in the wrong way, for example if the answer to the second question is $\sim C$, answering the other questions is (for all that has been said here) completely useless with respect to the original problem. Hence, it seems natural to consider the three questions as constituting a single derived problem, viz. $\{?\{B, \sim B\}, ?\{C, \sim C\}, ?\{D, \sim D\}\}$, and to remember that the problem was introduced in the hope to obtain the answers B , C and D . If one of the questions proves to be a dead end (that is, if an answer is obtained that is show that the original problem cannot be solved *along this road*), then the entire derived problem has to be discarded.

In view of what precedes, some readers will remark that a problem should not be handled as a set of questions, but rather as a set of questions *together*

with a set of pursued answers.² This is quite correct (and I started the approach with this view). However, it will turn out that the set of pursued answers will always be transparent. To be more precise, in the format that I shall use for problem solving processes, the set of pursued answers will always be apparent from a line that precedes the line at which a derived problem is written.

We now come to an important point: unless a problem has a readily available solution, solving it will require a reasoning process that has two important features. First, the process will involve *subsidiary problems* and/or *derived problems*—intuitively, subsidiary problems are inferred from the previous problem by purely logical means, while the inference of derived problems also relies on some declarative premises. Next, it is a *goal-directed* process in that it is guided by the problems that were raised. So, in contradistinction to the standard account of a proof, problem solving processes do not permit one to add steps because of the mere fact that they are correct, but require that the steps are useful, or at least possibly useful, with respect to the goal, which is to solve the original problem. From now on I shall say that the steps should be *sensible* in order to express that they are possibly useful for solving the original problem. Further discussion of the kind of restrictions that are involved follows in 1.6.

Let us, by way of an example, consider an old and well-known scientific problem solving process, which actually is at once a discovery process. Galileo Galilei wanted to find the “law of the free fall”. Actually, he derived them from a set of principles. But of course he also needed to perform some experiments that supported the law—see [DFMR92] for a fascinating study of the matter.




Figure 1.1: Inclined plane

Given the absence of an adequate instrument for measuring time, and given that objects fall rather fast, he had to go a long way to obtain an answer. First, he had himself constructed an inclined plane. He reasoned that the power that caused a ball to roll down from the plane was the very same power that acted on the ball if it was let to fall freely—he modelled both in Archimedean statics. Compare now the ball rolling down the inclined plane with a similar ball falling down the same distance (the distance between the top and the bottom of the plane). Invoking a postulate, Galilei derived that the proportion between the time needed for a ball to fall half of the distance to the time needed to fall the whole distance is identical to the proportion between the time needed by the ball to roll down half of the distance of the inclined plane and the time needed to roll down the whole plane. And similarly for other portions of the whole distance and of the whole plane. Figure 1.1 should clarify what I have in mind.

²If the original problem comprises a yes–no question, say because it is $\{A, \sim A\}$, both A and $\sim A$ may belong to the set of pursued answers.

The gain is that the movement of the ball on the inclined plane is relatively slow, and hence rather easy to measure, but nevertheless provides information about a relatively fast movement, the free fall of the ball—measuring this was far beyond the technology of the day.

This did not provide a ready solution. Galilei still needed to measure the times needed by the ball to roll down all of the plane, half of the plane, or another portion of the plane. Given the absence of chronometers and the like, an ingenious solution was necessary even for this. The solution Galilei came up with was the following. A leather bag filled with water was equipped with a small leather tube that warranted a constant flow of water. Starting with the tube closed by the pressure of his fingers, Galilei let the water flow at the same moment that he let go the ball at the top of the plane; when he heard the ball hit a wooden block (see Figure 1.2), say at one half of the plane, he closed the tube again. The amounts of water obtained for different positions of the block could be weighed and compared to each other. Needless to say, the experiment had to be repeated a couple of times for each position of the block in order to eliminate measuring mistakes. Apart from that, however, the problem was solved: Galilei had found a way to perform a set of measurements that happened to confirm the law of the free fall.




Figure 1.2: Inclined plane with wooden block

On the one hand, this example is a simple one: the solution does not require any conceptual changes. On the other hand, the example is sophisticated in several respects. First of all, the solution of the problem is a generalization rather than a singular observational statement. Next, the available knowledge was insufficient to solve the problem: new experimental data had to be collected. Moreover, the experiments had to be devised, and this was by no means a trivial task. It is not exaggerated to say that, given the availability of an induction logic, devising the experiments was the heart of the problem-solving process.

Maybe the example, however nice, is not completely realistic. According to [DFMR92], Galileo considered the observational data insufficient by themselves to warrant the correctness of the result (because his way of measuring time was not sufficiently precise). However, he was able to perform precise experiments that enabled him to verify the consequences of three different ‘principles’ taken together (the law of the fall, an inertia principle, and the squared times law). Even if all this be so, it does not follow that the example is misleading. Even if the outcomes of the experiments were merely empirical indications, it was important to design and perform them. Moreover, the design of the experiments that Galileo considered as reliable did not require ingredients that are fundamentally different from those of the process that led to the design of

the inclined plane experiments.

Further examples are easily imagined, but my view on problem solving, which agrees with the aforementioned literature, should be clear by now. Problems are solved by ‘deriving’ their solution from available knowledge, possibly extended with new empirical data, new generalizations, new hypotheses, etc. The means of derivation may vary from one problem solving process to the other. Deciding which data, generalizations, hypotheses, etc. are required, need not be straightforward. Solving a problem may require a long search process, involving many subsidiary problems and derived problems. The search processes required by the subsidiary and derived problems may involve different means and the involved kinds of ‘derivation’ may be quite diverse—in one search process the derivation may proceed merely by Classical Logic—henceforth **CL**—whereas another search process requires inductive generalization. Moreover, as I remarked earlier, the available premises and the logic may change during the search process. Nevertheless, the basic idea is goal-directed derivation.

1.2 Worries from the Philosophy of Science and from Erotetic Logic

The essential aim of the present notes is to devise a formal procedure (see Section 1.6) that can function as an explication (in the sense of [Car50]) of a problem-solving process. In claiming to present such a procedure, I may raise the reader’s suspicion. First, the enterprise might be suspected to be retrograde. Did the Vienna Circle philosophers not make great claims on the import of logic for understanding the sciences, and more generally any sensible knowledge, and have their views not been outdated for a long time? No doubt, the answer is twice “yes”. Next, the enterprise might be suspected to be utterly wrong. Did not even Thomas Nickles, presumably the most central philosopher of science who equated scientific discovery with problem solving, claim that there is no “logic of discovery”, for example in [Nic90], and does it not follow that there is no logic of problem solving? The correct answer is twice “yes” again.

I am not the type of person that likes to build up uncertainty in the readers’ minds, and hence comment at once on the suspicions. That the Vienna Circle’s views are outdated has nothing to do with its stress on logic. Moreover, the Vienna Circle identified logic with Classical Logic—henceforth **CL**. Of course, this was largely due to the absence of alternatives for **CL**. The Vienna Circle moreover adhered to a view on logic that, by present-day standards is very traditional, and hence made them blind for approaches as the one that I present in these notes. There is an excuse again: the lack of alternatives. Still, that the Vienna Circle’s formal enterprise failed cannot be taken as an indication that a formal approach to problem solving is bound to fail.

Tom Nickles is not particularly an adherent of **CL**, but, at least in the papers referred to, adheres to the same traditional view on logic as the Vienna Circle. I shall briefly describe the traditional view on logic and its alternatives in Section 1.5. If this view is broadened, Nickles arguments have to be revised as well. Moreover, some of Nickles’s claims are somewhat suspect. Thus he claims that there is no logic of discovery, but that there are discovery logics.³ He considers

³This is in line with the general disappointment that originated when people started to

it as typical for the latter that they are local, viz. content-specific. He goes on to discuss the question whether the existence of local discovery logics does not warrant the existence of a global discovery logic, for example in [Nic90, §6], but in doing so takes some steps that might very well be mistaken.

The matter is of utmost importance for the enterprise described in these notes. I said that the computational approaches were too specific, and hence that one has to look for a general approach to problem solving. If Nickles is right that there are only local logics of discovery,⁴ it makes no sense to aim at a general approach. So let me spell this out a bit.

First of all, there is no reason why the generalization of first-order local discovery logics would itself be second-order.⁵ If one tries to understand what Nickles means by local, viz. content-specific, discovery logics, one finds, for example, “as when natural law claims are turned into inference rules” ([Nic90, note 4]). This, seems to rely on a confusion. A logic, even **CL**, extended with such rules will lead to exactly the same consequences as the logic alone, provided the natural law statements are added to the premise set.

But even if a general logic of discovery existed, Nickles goes on to argue “[f]or historians and historically-oriented philosophers of science, the question [‘Does a logic of discovery exist?’] asks whether successful historical scientific practice has been based on such a logic; and the answer is clearly ‘no’”—[Nic90, p. 10] But why is the answer *clearly* “no”? I fear that Nickles takes this to be clear because no such logic was available to those practitioners. But this simply comes to denying the use of explication in methodological quarters.

Needless to say, no discovery logic can warrant correct results in the long run. Any logic or method can only be applied to available knowledge and within the available conceptual frameworks. If a general logic of discovery had been available to the Babylonians, it obviously could not have provided them with the sciences of our days. Although this seems obvious, it involves an important point that should be spelled out carefully.

Obviously, more data are available to us than to the Babylonians. Also, our conceptual frameworks are very different from theirs. The question then is: where do the latter come from? One might argue that, as new data become available, a prevailing conceptual framework may turn out to be problematic. This seems quite all right, but it imposes a requirement on problem solving processes, viz. that they may lead to conceptual changes. Once this is established—I shall have to return on this in subsequent chapters—we obtain a rather marvellous result. The reasoning process behind a problem solving process is clearly a dynamic one, often it is even non-monotonic—I extensively discuss dynamic reasoning processes in Chapter 4. The marvellous result is this: given that problem-solving processes are dynamic, later problem-solving processes may cause the results of earlier problem-solving processes to be overruled. To put it simply: solving a first problem may lead to a conceptual change, solving a second problem may lead to a further conceptual change, and solving

realize that the idea of a general problem solver, as proposed in [NS72], did not work out.

⁴Recall that Nickles sides with the mainstream in taking it that discovery processes “can be subsumed as special cases of the general mechanisms of problem solving” ([LSBZ87, p. 5])—actually he is largely responsible for making this the mainstream view in philosophy of science.

⁵This is a minor point in itself, except that the second-order story seems to support Nickles’s claim that a truly general logic of discovery exists at best abstractly, just as mathematical proofs exist even if no one discovered them until this day.

the first problem again after the second conceptual change may lead to a different solution, possibly involving another conceptual change. So at least in this sense, even a general logic of problem solving and discovery need not presuppose “ahistorical and asocial agents.” In sum, the dynamic character of problem solving undermines Nickles’s arguments against a general logic of discovery.

Yet, the relation with the relevant philosophy of science literature will be a touchy topic. Let me first state that I shall rely on the approach developed by Thomas Nickles ([Nic81], [Nic85], [Nic88], [Nic80c]), on the model developed by myself ([Bat85], [Bat92a] and [Bat92b]) and on the work of Joke Meheus ([Meh97], [Meh99a], [Meh99b], [Meh00a], [MB96]).

The remaining touchy point concerns the function of constraints in a formal approach to problem solving. Nickles uses ‘constraint’ in a very general sense: constraints restrict the problem-solving space in some or other way. Thus constraints impose requirements that a solution of the problem should fulfil—for example it should be compatible with a set of generalizations and theories. Constraints also present methodological instructions (does and don’ts) for solving the problem. Moreover, constraints define the problem: by changing the constraints, the problem becomes a different one. Remark in this connection that any ‘difficult’ problem will require that the constraints are changed during the problem solving process—this is nearly a matter of definition: if the constraints need not be changed, the problem is not a difficult one.

It seems unlikely that a formal approach to problem solving could take changing constraints into account. And yet, I shall show that it can. I shall return to the topic later, but let me state already now that a formal approach need not rule out that the premise set changes, provided it is clear at every stage of the formal *explicatum* which is the relevant premise set. I can add more already now. Suppose that the problem is to find an explanation for some or other fact, and that it turns out that the available theories and generalizations are unable to provide an explanation. In such case, the explanation can only be provided by a newly produced generalization or theory. In the approach that I shall spell out, this reduces to a change in logic: from the original set-up, which is traditionally labelled as deductive, one has to move up to a set-up in which (at least) an induction logic is applied, whence it is traditionally labelled as inductive. In other words, not only the premises may change, but also the available logics. All that is required is that such a change is made possible by the original set-up.

I now move to a different source of possible objections that are at least as difficult to handle as those stemming from the philosophy of science. This concerns the relation with standard erotetic logic, which I shall identify with Andrzej Wiśniewski’s work, for example [Wiś95], [Wiś96], [Wiśnt]. Previous work, for example [BM01], [Bat03b], and [Bat03a], suggests that the underlying erotetic logic of problem-solving processes differs in central respects from the standard approach, which comprises a definition of problem evocation (by a set of declarative premises) as well as a definition of problem implication (by a previous problem as well as a set of declarative statements). Actually the erotetic logic I shall rely on differs in two respects from the standard account. First, the standard account is obviously not sufficiently goal-directed. Thus, any question $?{A, \sim A}$ (Is A the case or is not- A the case?), erotetically implies $?{B, \sim B}$ for any B ,⁶ which is clearly unpalatable. Next, the standard account

⁶Actually, the implication requires to steps. $?{A, \sim A}$ erotetically implies the question

requires that every answer to a (directly) erotetically implied question at least partially answers the question by which it is implied. This means that any answer to the derived question together with the declarative premises entails a subset of the direct answers⁷ of the question by which it is implied. This requirement is clearly too restrictive for the purpose I have in mind. Thus if one tries to obtain an answer to the question $\{A, \sim A\}$ and it is given that $B \supset A$ (If B then A .), then it clearly is sensible to derive the question $\{B, \sim B\}$. If one obtains the answer B , the original question is at once solved. If one obtains the answer $\sim B$, one knows that trying to answer $\{A, \sim A\}$ by establishing B is a dead end.⁸

Incidentally, whenever the implied question is a yes–no question (a question of the form $\{A, \sim A\}$), the standard account enables one to justify the implication in two steps (see footnote 6). However, it does so only because it is not goal-directed. For implied questions of another form, for example whether-questions of the form $\{A_1, \dots, A_n\}$, the departure of the standard account seems unavoidable.

Given that the underlying erotetic logic appears to be not only systematic, but also attractive with respect to the intended application, I shall at present neglect the divergence from the standard account, and return to it later on. Meanwhile I shall use the terminology of the standard account.

1.3 Mastering Proof Heuristics

There is a particular set of problems for which logicians have a satisfactory problem-solving practice. The problems in the set may be called derivability problems because they have the form “Is statement A a consequence of the premise set Γ ?”, alternatively phrased: “Is there a proof of A from Γ or is there no such proof?”; in the formal metalanguage: “does $\Gamma \vdash A$ hold, or rather $\Gamma \not\vdash A$?” The notion of consequence (and proof) is ambiguous here in that it should be related to a specific logic. So the right question reads: does $\Gamma \vdash_{\mathbf{L}} A$ hold, or rather $\Gamma \not\vdash_{\mathbf{L}} A$? But actually, logicians have a satisfactory problem-solving practice for derivability problems in a vast collection of logics.

For logics from that collection—for the time being, I shall suppress the reference to the specific logic—the practice is indeed quite good. If there is a proof of A from Γ , logicians will be able to find it most of the time. If there is no such proof, they will be able to find out that there is none, most of the time. Most of the time they are even able to demonstrate that there is no proof. Needless to say, such a demonstration will not be itself a usual (object-level) proof. It will proceed by drawing a tableau, or performing another kind of procedure. Typically, if there is no proof of A from Γ , the procedure will stop—no instruction allows one to add a further step. The outcome of the tableau-method or other procedure, whether it leads to the answer “yes” or to the answer “no,” may then be shown in the metalanguage to be correct. For example, it can be proved in the metalanguage that (i) if the tableau-construction closes, then it

$\{A \wedge B, A \wedge \sim B, \sim A \wedge B, \sim A \wedge \sim B\}$ (Is A -and- B the case, or A -and-not- B , or not- A -and- B , or not- A -and-not- B ?) and this erotetically implies the question $\{B, \sim B\}$.

⁷But briefly, the direct answers of a question are the statements enumerated in the question (in the notation I am using). Thus the direct answers of $\{A, B, C\}$ are A , B and C .

⁸I am obviously presupposing **CL** or another non-paraconsistent logic in the present context. Paraconsistent contexts require careful attention that is postponed to a later chapter.

can be transformed (by a simple algorithm) into a proof of A from Γ , whereas (ii) if it remains open (comes to a stable stage that is not closed), then there is a model of Γ that falsifies A , and hence there is no proof of A from Γ in view of the soundness of the logic with respect to its semantics.

The previous paragraph contains three times the expression “most of the time,” and this expression does not always have the same meaning. The reason for this is that **CL** is not decidable, but that there is only a *positive test*⁹ for **CL**-derivability. This means that there is a (mechanical) procedure that, if there is a proof of A from Γ will stop after finitely many (possibly very many) steps with a positive answer, but might never stop if there is no such proof. So if there is a proof, it is in principle possible to find it. Of course, one still might need a lot of patience; if the procedure did not lead to an answer after ten billion steps, one might fear that it might never answer and erroneously decide that there is no proof.

The propositional fragment of **CL** is decidable, and so are a number of predicative fragments. Other logics are usually in a similar position, or are off a bit worse in a computational respect (less decidable fragments). As we shall see in subsequent sections, the situation of adaptive reasoning forms, and of adaptive logics as well, is much worse.

If A is not derivable from Γ , it still makes sense to start a procedure for checking whether A is derivable from Γ . It is indeed very much possible, even if one is in undecidable waters, that the procedure stops with the answer “no.” *If* it stops (and is correct in this respect), this warrants that there is no proof. So it makes sense to run a procedure that does not always stop, but nevertheless provides a criterion, if it stops, for non-derivability.

Most procedures that form a positive test for **CL** (or for some other logic) deliver results (sequences of formulas, trees of formulas, etc.) that are rather different from proofs, even if they may be turned into a proof if they provide a positive answer. This is a rather unfortunate situation. I believe indeed that proofs are excellent means to explicate our most natural forms of discursive reasoning. For example, just like proofs, those forms of reasoning have a sequential character.

I know that many people have claimed the opposite, more specifically have claimed that human reasoning proceeds essentially along semantic lines. I think, however, that this is a mistake. I readily admit that most model-based reasoning is not sequential. But I claim that most discursive reasoning is. As arguments I can only refer to my own introspection; I can add that I know of no experiment that shows me wrong; much more I cannot do.

The previous claim led to a good deal of research in the Ghent Centre for Logic and Philosophy of Science, some of which is already published, for example [BP01], [Pro02a] and [Pro02b].¹⁰ The underlying idea is to devise a form of proofs that display a prospective dynamics, which enables one to push most of the proof heuristics into the proof.

The idea underlying prospective proofs is that one does not proceed in terms

⁹See [BJ89] for such matter. In [BBJ02], which is the fourth edition of the previous book, the term “positive test” has been removed. It is now said that **CL** is partially recursive. This terminological change is not an improvement: it does not specify which part (derivability or undervability) causes the undecidability.

¹⁰Actually, this research started as a side-effect of adaptive logics. The corrective conditions of the latter were paralleled by prospective conditions.

of the premises, but rather in terms of a target. When the proof starts, the formula that one tries to derive is the target. At a later stage of the proof, other formulas may become the target in view of previous steps in the proof.

The main difference with usual proofs is that certain lines of the proof does not contain a simple formula, but a formula preceded by a set of (prospective) conditions. The general form of such a condition plus formula is

$$[B_1, \dots, B_n] A,$$

which indicates that the formula A is obtained if all of B_1, \dots, B_n are obtained.

Suppose that the target is A and that $C \supset (B \vee A)$ is a premise. The rules will first allow one to introduce this premise in the proof, and next to apply analysing steps to it. The premise will obviously be introduced on an empty condition:

$$C \supset (B \vee A).$$

The subsequent steps will be, first

$$[C] B \vee A,$$

because $B \vee A$ is derivable from the premise if C is available, and next

$$[C, \sim B] A,$$

because A is derivable from $A \vee B$ if $\sim B$ is available. The essential point is that the introduced conditions are subformulas or negations of subformulas of the complex formula occurring in the (local) premise.

Apart from such formula-analysing rules, prospective proofs have condition-analysing rules. These analyse a condition that cannot be obtained directly. Thus if it turns out that $C \wedge D$ cannot directly be obtained, the condition plus formula

$$[C \wedge D, B_1, \dots, B_n] A$$

will be analysed to

$$[C, D, B_1, \dots, B_n] A$$

because $C \wedge D$ can be obtained if both C and D can be obtained.

Finally, there are certain rules for eliminating formulas from a condition. Thus if C is obtained unconditionally, the last condition plus formula will be simplified to

$$[D, B_1, \dots, B_n] A.$$

It can be shown that the resulting procedure leads to a proof of the conclusion A from the premise set Γ iff (if and only if) A is derivable from Γ . If the procedure stops, A is not provable from Γ . Obviously, there cannot be a warrant that the procedure stops for cases where there is no positive test. Thus no procedure can possibly stop for all A and Γ for which $\Gamma \not\vdash_{\text{CL}} A$.

All this will become fully clear in Chapter 3. However, it should already be clear now that prospective proofs warrant a form of goal-directedness, which is what we were after. The central point of prospective proofs is that they are obtained directly from a natural proof-search method.¹¹ Thus if one tries to derive A and $C \supset (B \vee A)$ is a premise, it is natural to try to derive C and

¹¹That other proof-search methods are possible does not form an objection.

$\sim B$ from the other premises. I am not, at this point, very interested in the computational efficiency of prospective proofs. The main worry is that they should warrant that all steps in the proof are sensible with respect to deriving the conclusion from the premises. A different point is this: the more complex the calculations made outside of a proof, the more elegant will be the resulting proof. However, my aim is not to obtain elegant proofs, but to describe a sensible proof-search method. Incidentally, succeeding prospective proofs can easily be turned into standard proofs, if that is desired.

The results on prospective dynamics for **CL** readily reveal that it is easy enough to extend the kind of proofs for other logics ‘of the same type.’ It is not difficult to turn a set of inference rules into a set of rules for prospective proofs. The special case of adaptive logics will be discussed in Chapter 5.

1.4 Unusual Logics Needed

For all that was said in the previous section, a formal approach to problem solving seems straightforward: extend the prospective dynamics to the required logics and combine this with an erotetic logic that handles subsidiary and derived problems. There is, however, an important complication. The forms of reasoning that occur in problem-solving processes cannot be explicated by logics ‘of the same type’ as **CL**.

Let us start with a very common reasoning form: inductive generalization. To be more precise, I mean the reasoning by which one ‘derives’ from the data formulas of the form $\forall A$ (the universal closure of A) as well as **CL**-consequences of the data and the generalizations.¹² It is well-known that the involved consequence relation is non-monotonic: a generalization that is derivable from a set of data need not be derivable after the set of data is extended. This is one of the reasons why a logic explicating inductive generalization is not ‘of the same type’ as **CL**. There is, however, a deeper reason.

Non-monotonic consequence relations display an *external dynamics*. Suppose that $\Gamma \vdash_{\mathbf{LI}} A$ states that A is an inductive consequence (defined by the logic of induction **LI**) of Γ .¹³ If Γ is the set of data available at some point in time, then $\Gamma \vdash_{\mathbf{LI}} A$ enables one to accept A . At a later point in time, the set of available data might be $\Gamma \cup \Delta$ and if $\Gamma \cup \Delta \not\vdash_{\mathbf{LI}} A$, one has to give up the conclusion A . This dynamics is external in that it does not derive from the reasoning process itself. Both $\Gamma \vdash_{\mathbf{LI}} A$ and $\Gamma \cup \Delta \not\vdash_{\mathbf{LI}} A$ always were true and always will be true. If one knows them to be true, then one justly accepted A at the point in time where Γ was the set of available data, and one justly rejected A after Γ was extended with Δ .

Suppose that one is only interested in generalizations $\forall A$ in which A does not contain quantifiers or individual constants. The basic mechanism behind (thus restricted) inductive reasoning is joint compatibility. A generalization G is inductively derivable from a set of data iff it holds, for all sets Δ of generalizations, that $\Delta \cup \{G\}$ is compatible with the data whenever Δ is compatible with the data. Given that the data are singular formulas and given the form of

¹²Nearly always, background knowledge plays a role—this is discussed later on.

¹³Actually, the adaptive logic **LI** is described in [Batntb] and several further results are forthcoming in papers by Lieven Haesaert and myself, for example [BH01]. However, one needs not to know those systems in order to follow the argument in the text.

the generalizations, the matter is (effectively) decidable. Now, let us make the picture slightly more realistic and suppose that background theories are available. Suppose moreover, to keep things simple, that the data do not contradict the background theories. The available knowledge now consists of the data, the background theories, and the **CL**-consequences of the union of these sets. Which set of inductive generalizations is compatible with this knowledge is not in general a decidable matter. Worse, there is no positive test for it.

Given the absence of a positive test, how is it possible that people ever arrive at inductive generalizations? The answer is quite obvious: by reasoning. In specific cases, the reasoning may enable one to arrive at a final judgement. However, the reasoning cannot in general lead to a final judgement, even if the premises (data and background generalizations) remain stable during the reasoning process. It can, however, lead to a good estimate. Some people will arrive at a better estimate than others, and the efficiency of such reasoning processes may be studied.

Even if the reasoning does not result in a final judgement, one may consider its outcome sufficiently reliable for making a decision—one may know that a final judgement is impossible, one may consider it too expensive or time consuming to obtain a better judgement, etc.

The absence of a positive test makes the reasoning process necessarily dynamic. Even when reasoning from a stable set of premises, one will have to consider certain formulas as derived provisionally. In other words, it cannot be avoided that, at some point in the reasoning process, one considers as derived certain formulas that later have to be considered as not derived. This I shall call the *internal dynamics*. It is not caused by the introduction of new premises, but is a property of the reasoning process itself, even if it proceeds from a stable premise set. For example, if background knowledge is present, one cannot avoid deriving certain generalizations that later turn out to be incompatible with the (stable) available knowledge. This internal dynamics, caused by the absence of a positive test, forms the specific and problematic feature of the forms of reasoning that “cannot be explicated by logics of the same type as **CL**.”

Before proceeding to further examples of such reasoning forms, it is worth pointing out that the considered logic of inductive generalization is only a special case of a broader phenomenon. Whenever a new theory is adduced, it is supposed to be compatible¹⁴ with available knowledge (the data and formerly accepted theories, or at least part of them). As there is no positive test for compatibility, any reasoning that leads to accepting the new theory necessarily displays the internal dynamics and, except in the specific cases in which a final judgement can be reached, the decision taken as a result of this reasoning is necessarily defeasible and hence provisional.

A recent version of the theory of the process of explanation is presented by Ilpo Halonen and Jaakko Hintikka in [HH]. In Section 6, they discuss the conditions on (nonstatistical) explanations (with a number of restrictions). The conditions (I slightly change their notation) concern an explanandum Pb , a background theory T (in which the predicate P occurs) and an initial condition (antecedent condition) I (in which b occurs). Among the six conditions are the following:

- (iii) I is not inconsistent ($\not\vdash_{\mathbf{CL}} \sim I$).

¹⁴See [BM00] for the adaptive logic of compatibility (in the framework of **CL**).

- (iv) The explanandum is not implied by T alone ($T \not\vdash_{\mathbf{CL}} Pb$).
- (vi) I is compatible with T , i.e. the initial condition does not falsify the background theory ($T \not\vdash_{\mathbf{CL}} \sim I$).

Obviously, there is no positive test for any of the three conditions. In other words, no finite reasoning process can (in general) lead to the conclusion that Pb is explained by I and T .

So although the underlying ‘logic’ is \mathbf{CL} (see the formal conditions above), the reasoning process that leads to the conclusion that I and T together explain Pb cannot possibly be explicated in terms of \mathbf{CL} (at the object level). The reasoning is *about* \mathbf{CL} -derivability, and necessarily displays the internal dynamics. This is why it cannot be explicated by a \mathbf{CL} -proof but only by a dynamic proof.¹⁵

The logic of questions forms a further example. According to [Wiś96] and [Wiś95], where this problem is studied and solved, a question Q is *evoked* by a set of declarative statements Γ iff the (prospective) presupposition of Q is derivable from Γ but no direct answer to Q is derivable from Γ .¹⁶ Note that there is no positive test for non-derivability¹⁷ from Γ . Hence, although the definition is itself unobjectionable, only a dynamic proof may (in general) lead to the conclusion that Q is evoked by Γ .

Another example concerns handling inconsistency. Consider the case in which a scientific (empirical or mathematical) theory T was meant to be consistent and was formulated with \mathbf{CL} as its underlying logic, but turned out to be inconsistent. As we know from the literature,¹⁸ scientists do not just throw away such a theory. They *reason* from T in search for a consistent replacement. Obviously, they do not reason in terms of \mathbf{CL} , because this is known to lead to triviality. Nor do they reason in terms of some monotonic paraconsistent¹⁹ logic \mathbf{PL} —by invalidating certain rules of inference of \mathbf{CL} , paraconsistent logics deliver too poor a consequence set. In their reasoning, they want to interpret T *as consistently as possible*. After all, T was meant as a consistent theory.

Let us consider an utterly simplistic but instructive example. Let the ‘theory’ consist of $\sim p$, $p \vee r$, $\sim q$, $q \vee s$, and p . One obviously should not derive r from $\sim p$ and $p \vee r$ by Disjunctive Syllogism. This application of Disjunctive Syllogism presupposes that the falsehood of p , which gives us r in view of $p \vee r$, is entailed by the truth of $\sim p$; however, the theory asserts that p is true (see the last premise). However, as the theory was meant to be consistent, one will apply Disjunctive Syllogism to derive s from $\sim q$ and $q \vee s$. Indeed, q behaves consistently on the theory. To be more precise, q is consistently false on the theory, for $\sim q$ is obviously derivable whereas q is not, except of course by explicitly or implicitly applying Ex Falso Quodlibet to p and $\sim p$.²⁰

So the reasoning from T should proceed in such a way that one obtains T in

¹⁵See [BM01] for adaptive logics that explicate several forms of reasoning that underly the search for explanations.

¹⁶See [Meh01] for an adaptive approach to question evocation.

¹⁷I mean non- \mathbf{CL} -derivability, in agreement with the cited papers, but the matter is the same for any other sensible logic.

¹⁸See for example [Nor87], [Nor93], [Smi88], [Bro90], [Ner02], [Meh93] and [Meh02]. Not all these authors side with me on the required approach, but that is immaterial.

¹⁹A logic \mathbf{L} is paraconsistent iff $A, \sim A \vdash_{\mathbf{L}} B$ does not hold in general; if there is no A for which every formula is \mathbf{L} -derivable from $\{A, \sim A\}$, \mathbf{L} is said to be strictly paraconsistent.

²⁰One way to implicitly apply Ex Falso Quodlibet proceeds by first applying Addition to obtain $p \vee q$ from p and next applying Disjunctive Syllogism to obtain q from $\sim p$ and $p \vee q$.

its full richness, except for the pernicious consequence of its inconsistency. Precisely for this reason, the reasoning cannot proceed in terms of some monotonic paraconsistent logic **PL**. **PL** would invalidate certain **PC**-rules, for example Disjunctive Syllogism.²¹ However, as we saw from the previous example, the requested reasoning should not invalidate certain rules of inference of **CL**, but only certain *applications* of these rules. Let me express this more precisely. For certain rules, an application should be valid if specific involved formulas behave consistently on the theory, and invalid otherwise. Precisely this proviso causes the reasoning to be internally dynamic: there is no positive test for the consistent behaviour of some formula on a set of premises.²²

The forms of reasoning I have considered up to now display an internal dynamics and all of them concerned a single unstructured set of premises. In many cases, however, the premises are structured, usually as a n -tuple of sets. I shall now consider some examples of this type.

Let us return for a moment to inductive generalization. Usually background knowledge is available in addition to the data. Let us restrict the discussion to the case where the background knowledge consists of generalizations in the sense meant before. An obvious complication is that the data may falsify some of the background generalizations. So two forms of dynamics have to be combined. First, we retain the background generalizations in as far as they are not falsified by the data. Next, from the data and the retained background generalizations we obtain new generalizations by the aforementioned logic **LI**. Remark that it is in general impossible to perform the first selection (of background generalizations) before proceeding to the second selection (of new inductive generalizations). This means that both forms of dynamics are necessarily combined in the reasoning process. After deriving some new inductive generalizations, one may be forced to change one's judgement on the compatibility of some background generalization with the data, and this will affect the derivability of the new generalizations.²³

Not all background generalizations need be considered equally trustworthy. So instead of a set of background generalizations, one confronts a sequence of such sets, each having a different priority. In this case one has to combine a multiplicity of dynamics concerning the background generalizations with the dynamics that pertains to the new generalizations. Moreover, even certain falsified background generalizations may be considered as applying 'normally'. This means that an *instance* of the generalization is considered to hold unless and until proven incompatible with the data. Such 'pragmatic generalizations' may also be ordered by some priority relation. All this leads to more forms of dynamics (which, however, are all of three kinds)—see again [BH01] for the corresponding logics.

I now move to a very different example. Often problems are solved by a group of people. Each of the participants may have different relevant positions—see for example the personal constraints discussed in [MB96]. A participant may

²¹See [Meh00b] for an exception: the paraconsistent logic **AN** validates Disjunctive Syllogism (and all 'analysing' rules of **CL**) but invalidates Addition (and Irrelevance and similar rules).

²²Explicating this kind of reasoning was at the origin of the adaptive logic programme—see [Bat89], [Bat99], and many other papers.

²³The complication of falsifiable background generalizations is dealt with in [Batntb]. Most complications discussed in the subsequent paragraph of the text are handled in [BH01].

change his or her position in view of arguments adduced by other participants or in view of solutions of subsidiary or derived problems. As a result, the subsequent interventions of the participant will be mutually incompatible, even if the participant's position is consistent during each intervention. However, the participant will not state his or her full new position whenever there is a change. So, after an intervention, the participant's position has to be reconstructed from the sequence of his or her interventions. In order to do so, one starts with (the consistent part of) the last intervention, to this one adds that part of the previous intervention that is compatible with it, and so on.²⁴ Remark that, while doing so, one does not select statements that are made during an intervention, but rather their consequences.²⁵

Diagnostic reasoning forms a further example in which the premises are prioritized and hence require a multiplicity of dynamics. Reasoning proceeds from data on the one hand and expectancies (that may have varying degrees of trustworthiness) on the other hand. The expectancies, or rather their consequences, are retained (in their order of priority) until and unless proven inconsistent with the data. (See [WP99], [PW02] and [BMPV03] for the adaptive logics.)

In all examples mentioned before, the flat ones as well as the prioritized ones, the reasoning displays both the internal dynamics and the external dynamics. It is worth mentioning that, whenever the external dynamics (non-monotonicity) is present, the reasoning necessarily displays the internal dynamics (even if the premises are stable). The converse, however, does not hold. The *Weak* consequence relation, of Rescher and Manor—see, for example, [RM70] and [BDP97]—is monotonic. Nevertheless, it may be shown that the reasoning from premises to weak consequences requires an internal dynamics.²⁶ Some consequence relations that are monotonic as well as decidable may even be characterized (in an enlightening and attractive way) by a form of reasoning that displays an internal dynamics—see [Bat01] for an example.

The preceding paragraphs do by no means contain an exhaustive list of the reasoning mechanisms (or even of the types of reasoning mechanisms) that display an internal dynamics. Nevertheless, the problem should be clear by now. Essential forms of human reasoning, that play a central role in problem-solving processes, display an internal dynamics. So we need logics for explicating such forms of reasoning. Such logics are available; they are called adaptive logics. These unavoidably have some non-standard properties. More important, however, is that they are characterized in a formally stringent way, and that their properties are studied in agreement to the professional standards.

Let us now return to the plot spelled out in the first paragraph of this section. The prospective dynamics cannot in a simple way be generalized to adaptive logics. As we shall see in Chapter 5, the prospective dynamics for adaptive logics requires a procedure that goes through more phases than the procedure for **CL** and similar logics. We shall see that this complication can be handled. Moreover, the complication is directly relevant for arriving at a final judgement—remember the discussion of inductive generalization. Indeed, the prospective dynamics for adaptive logics provides us with *criteria* for deciding

²⁴So the interventions are naturally handled as having a certain priority, which is lower as the intervention is older.

²⁵Adaptive logics for this reconstruction are spelled out in [Ver03] and [Batnta].

²⁶ A is a Weak consequence of Γ iff it is a **CL**-consequence of some consistent subset of Γ —remember that there is no positive test for consistency.

whether the judgement reached is final. So, although there is no positive test for derivability,²⁷ the prospective dynamics often provides a warrant that a certain statement is finally derivable from a given set of premises.

1.5 The Traditional View on Logic

Sometimes people complain that adaptive logics are not really logics. They try to adduce arguments for this complaint by referring to the traditional view on logic, which they consider as deeply philosophically grounded and the outcome of a long tradition.

One should not worry too much about such objections. The point is not whether adaptive logics are called logics, rather than, say, giclos. The point is that they explicate common and important forms of reasoning.

Nevertheless, the claim that adaptive logics (and non-monotonic logics in general) fail to fulfil certain ‘essential properties’ of logics is not too impressive a claim. For one thing, logicians disagree about the nature of logics. Next, many kinds of logics that are today recognized as belonging to the mainstream of the logical tradition, were the object to sometimes nasty fights in the past: intuitionistic logic, modal logics, relevant logics, paraconsistent logics, . . . Moreover, the so-called traditional view on logic is by no means terribly old; or if it is identified as being old, then it was not very stable: it underwent several deep and dramatic changes during the last century. This can be illustrated by describing the traditional view that prevailed around 1900. Remark that this was not only a couple of millennia after Aristotle, with whom the tradition started, but after the turn to formal languages, after **CL** was spelled out, and after the axiomatization of the major mathematical theories was well on its way.

By that time, that is around 1900, the traditional view implicitly or explicitly held that a logic has a number of properties, of which I enumerate the most important ones.

L.1 A logic **L** is a *function* $\mathbf{L} : \wp(\mathcal{W}) \mapsto \wp(\mathcal{W})$ that assigns to every set of (closed) formulas a unique consequence set.

That the logic **L** allows one to derive the formula A from the premise set Γ is standardly written as $\Gamma \vdash_{\mathbf{L}} A$ (or $\Gamma \vdash A$ where no confusion arises). The consequence set of Γ , $\{A \mid \Gamma \vdash_{\mathbf{L}} A\}$, is written as $Cn_{\mathbf{L}}(\Gamma)$ (or $Cn(\Gamma)$ where no confusion arises).

L.2 The *deductive correctness* of an inference is judged by Aristotle’s criterion: $\Gamma \vdash A$ iff it is impossible that A is false when all members of Γ are true.

In order to make this precise, formal semantics was invented. It associated with every logic a set of models, which verify certain formulas and falsify others. The criterion then becomes: every model of the premises²⁸ verifies the conclusion; in other words, no model verifies the premises and falsifies the conclusion.

L.3 Reflexivity of the inference relation: from a set of premises every premise is derivable: if $A \in \Gamma$, then $\Gamma \vdash A$; in other words $\Gamma \subseteq Cn(\Gamma)$ (from L.2).

²⁷This form of derivability will be called “final derivability.” It is a stable form of derivability, in contradistinction to ‘derivability at a stage,’ which may vary from one stage to the next as a result of the internal dynamics.

²⁸A model of Γ is a model that verifies all members of Γ .

- L.4 Transitivity of the inference relation: what follows from the premises together with their consequences, follows from the premises: if $Cn(\Gamma) \vdash A$, then $\Gamma \vdash A$; in other words, $Cn(Cn(\Gamma)) = Cn(\Gamma)$.
- L.5 Monotonicity of the inference relation: what follows from the premisses, follows from an extension of the premises (if $\Gamma \vdash A$, then $\Gamma \cup \Delta \vdash A$).
- L.6 *Formal correctness*: the correctness of an inference is determined solely by its form (in the formal language).
- L.7 *Static* proofs: every initial fragment of a proof²⁹ is itself a proof; and a proof remains correct if it is extended.
- L.8 *Consistency*: Γ is false if an inconsistency is derivable from Γ (a formula A together with its negation $\sim A$). As all inconsistent premise sets are false, they have no models.
- L.9 *Ex falso quodlibet*: $A, \sim A \vdash B$.

Here is a well-known proof of L.9:

- | | | |
|---|------------|-----------------------------|
| 1 | A | Premise |
| 2 | $\sim A$ | Premise |
| 3 | $A \vee B$ | 1; Addition |
| 4 | B | 2, 3; Disjunctive Syllogism |

Semantically *Ex Falso Quodlibet* is demonstrated as follows: there is no model that verifies both A and $\sim A$ and falsifies B , which is true because no model verifies both A and $\sim A$ (by L.8). Remark that B plays no role here.

- L.10 Compactness: what is derivable from a premise set, is derivable from a finite subset of the premise set: if $\Gamma \vdash A$, then $\Gamma' \vdash A$ for some finite $\Gamma' \subseteq \Gamma$.
- L.11 Decidability (or effective decidability): there is an algorithm to decide (for every Γ and A) whether $\Gamma \vdash A$ holds.
- L.12 Finite axiomatizability: every definite knowledge domain (arithmetic, mechanics, ...) can be described by a finite set of axioms or by a finite set of (metalinguistic) axiom schemas.³⁰
- L.13 Demonstrability of consistency: if a set of premises (or a theory) is consistent, then it is possible to demonstrate this by ‘absolute’ means.

Property L.13 is much stronger than property L.11. It indeed requires not only that, for all A thus far considered, either $\Gamma \not\vdash A$ or $\Gamma \not\vdash \sim A$ was established; it also requires a metatheoretic proof that there is no A for which $\Gamma \vdash A$ and $\Gamma \vdash \sim A$ hold. This is demonstrated by absolute means if the demonstration does not depend on the consistency of another set of premises or theory, for example arithmetic or geometry—see also [Murar] for some interesting background materials.

²⁹A proof of A from Γ is a list of formulas, that ends with A and every member of which is either a premise (a member of Γ) or follows from previous members of the list by a rule of the logic. An annotated proof is a list of lines, every member of which consists of a number, a formula and a justification indicating by application of which rule the line was added. The concept of proof was made precise around 1900 (by Hilbert). Later one sometimes considered infinite proofs as well as rules drawing a consequence from an infinite number of previous members of the proof.

³⁰More correctly, such a domain can be described by a finite set of axioms or axiom schemata and by a logic. The resulting theory is the set of formulas that is derivable from the axioms by the logic.

During the first half of the twentieth century, it was proved that some of the aforementioned properties do not hold. Thus (full predicative) **CL** is not decidable, contrary to what L.11 states. Moreover, elementary arithmetic (the set of arithmetic truths) cannot be axiomatized in the sense of L.12, unless in case it is inconsistent (first Gödel theorem), nor can its consistency be demonstrated by absolute means in the sense of L.13 (second Gödel theorem).

The desirability of certain other properties seem a matter of taste, rather than the result of a deep insight. Hilbert already considered rules that give up L.10—see, for example, [Murar]. Non-monotonic logics obviously do not fulfil L.5 and fulfil only a restricted form of L.4; paraconsistent logics fulfil neither L.8 nor L.9. The essential point, however, is that properties L.11-13 *had* to be given up, after they had been accepted ever since anyone became aware of them—for some a couple of thousand years.

1.6 Logical Systems *vs.* Logical Procedures

On several occasions, I have referred to procedures. I better clarify what I mean by this. As we shall consider procedures for proofs and for (proof-like) problem-solving processes, I also better spell out how a procedure relates to a set of (inference) rules, which provides the standard definition for a proof.

The simplest way to characterize a procedure is as a set of rules to which are permissions and obligations attached—I shall sometimes use “instruction” to denote a rule with permissions and obligations attached to it. Applying a procedure leads to a sequence or tree of entities (often formulas or lines). The permissions and obligations depend of the stage of the sequence or tree. This means that one needs a rule for starting the application, and next rules that may or must be applied in view of the stage of the application.

Consider for example a tableau method for checking expressions of the form $A_1, \dots, A_n \vDash B$. The starting rule will instruct one, for example, to write a list of formulas consisting of A_1, \dots, A_n and of $\sim B$. This list is the initial stage of the tableau, viz. of its only branch. The other rules will be of the form: if formulas of specific forms occur in the list, then it is permitted to do this-and-that, for example to extend the branch with certain other formulas, or to split the branch in two, extending the left subbranch with certain formulas, and the right one with other formulas. A tableau procedure may forbid one to extend a branch as soon as it is closed, that is, as soon as a formula C as well as its negation $\sim C$ occur in the branch. The application of a procedure stops if it arrives in a stage to which it is not permitted to apply any rule.

In Chapters 3 and 5, I shall consider several proof procedures and problem-solving procedures. These will also consist of a starting rule, which leads to the first line of the proof or problem-solving process, and of rules that may be applied on the proviso that lines of a certain form occur in the proof. Remark that the usual definition of a logic in terms of a set of rules is a special case of a procedure: there is a set of rules and a universal permission is attached to them—every rule may be applied whenever one wants.

While some procedures are deterministic (at every stage only one rule can be applied), the proof procedures and problem-solving procedures that I shall describe in later chapters leave quite some room for choices. This will not prevent them from stopping, for specific premises and a conclusion, after a

certain number of steps. Given that choices are possible, some proofs may be more efficient than others (for example in being shorter). Actually, I shall not worry about efficiency. The main aim is to articulate procedures warranting that all steps are sensible. If this is realized, the resulting proofs or the resulting problem-solving processes are goal directed. Indeed, such procedures do not focus on abstract matters, such as models, semantic consequence, derivability, or the ‘existence’ of a proof, but on a very concrete issue: the next step to be taken in an attempt to find out whether a conclusion A is derivable from a premise set Γ .

1.7 The Plan

In Chapter 2, I shall first introduce and study a proof-search procedure for **CL**. Here the prospective dynamics will be introduced at the most basic level: searching for a proof of a conclusion from a set of premises. A remarkable fact will surface at this point: the *Ex Falso Quodlibet* rule (to derive B from A and $\sim A$) turns out to be an unnatural and isolated addition to an otherwise systematic procedure. Dropping *Ex Falso Quodlibet* will lead to a prospective dynamics for a logic which I shall call **CL**⁻. I shall show that it is a very natural logic, even if it can only be described in terms of a procedure, and that sensible real life proofs ‘in terms of classical logic’ are obtained by the procedure for **CL**⁻ and not by that for **CL**. In Chapter 3 a variant of the procedure for **CL**⁻ will be combined with problems (sets of questions) and this will result in a simple problem-solving procedure.

In both procedures, lines may be marked in one of several kinds of marks. These marks will be governed by definitions. Their function is to control the dynamics of the proofs and problem-solving processes.

The backbone of the problem-solving processes will concern a problem of the form $\{?A, \sim A\}$, solved by deriving one of the direct answers, viz. A or $\sim A$, from a premise set. I shall at once discuss some extensions of the backbone. A simple example concerns the availability of a set \mathbb{A} of questions that can be answered by observational or experimental means. We shall see that it is not difficult to build this into the problem-solving process. If it turns out, for example, that a problem cannot be solved by deriving the relevant answers from the premises, and if these answers may be obtained by answering a question in \mathbb{A} , the question will be ‘put to nature’ and its answer will be introduced as a new premise.

Adaptive logics will be introduced in Chapter 4. I will restrict the discussion to adaptive logics in so-called standard format. This format characterizes the logics by a lower limit logic, a set of abnormalities and an adaptive strategy. Some adaptive logics are weaker than **CL**—these are sometimes called corrective. An obvious example are inconsistency-adaptive logics, which are invoked if a theory that was meant to be consistent, turns out to be inconsistent. Other adaptive logics extend **CL**—these are called ampliative. Logics of inductive generalizations are of this kind. They do not only allow one to derive **CL**-consequences of the premise set (data and possibly background knowledge), but moreover allow one to derive inductive generalizations as well as the consequences of these generalizations and of the premises.

Devising a prospective dynamics for adaptive logics will be the main topic

of Chapter 5. As the considered adaptive logics are in standard format, a single approach will be sufficient for obtaining the prospective dynamics for all the logics. Finally, in Chapter 6, I shall consider some further extensions and some open problems.

It is one thing to show that the prospective dynamics has the required meta-theoretical properties. A different, and philosophically more important point is whether the problem-solving processes enable one to explicate their *explicanda*. As always, the proof of the pudding will be in the eating: the reader will have to be convinced by the results. I shall not attempt to show that the proposed approach is adequate for solving ‘all’ kind of problems. This is related to the way in which the approach originated. The first results on the prospective dynamics showed that the proof heuristics may be largely pushed into the proofs themselves. This could be easily generalized to answering questions by deriving one of their direct answers from a set of premises. That, in turn, led to the idea to invoke other logics, and to extend the problem-solving processes with, for example, the set \mathbb{A} and the means to handle it. The approach presented in the next chapters is the outcome of a growing enthusiasm: simple and specific problems proved to be solvable by formal means, and later the range of problems that could be handled started growing. So it seems advisable to continue in a rather opportunistic way, trying to further extend the range of problems that can be handled. Maybe it will turn out eventually that certain problems are beyond the reach of the present approach. But even this would obviously not form a good reason for not applying the present approach to problems that can be handled by it.

There is a set of problem-solving means that is not at present integrated in the approach, viz. all those means that rely on model-based reasoning. I hope to study the matter closer in the future, and to find a suitable general way for conceptualizing model-based reasoning and for applying both adaptive mechanisms and prospective dynamics to it. The fact that the prospective dynamics from Chapter 2 may be represented diagrammatically, and such results as mentioned in [Meh99b], might very well contain clues for the desired conceptualization.