Chapter 15

Sufficient and Contrastive Explanation

15.1 Sufficient Causal Explanation

We may have found some causal factors for [1], namely [11] and [12], but how do we know that [11] and [12] suffice to provide a causal explanation of [1]? How do we know there isn't some essential [13] that must also be present? A causal explanation doesn't just search after *some* causal factors, it also searches after a *sufficient* explanation.

Suppose we have determined that $A_1 \Rightarrow\!\!\!\!> B, A_2 \Rightarrow\!\!\!\!> B, \ldots, A_n \Rightarrow\!\!\!\!> B$. We shall say that $A_1...A_n$ form a *sufficient set* of causal factors for B just in case it would be impossible for B not to happen (not to be the case) if $A_1...A_n$ had happened. In our situation, this is a counterfactual conditional, because we are performing a failure analysis of an actual incident, and so $A_1, \ldots, A_n, B$ did in fact all happen (were all the case).

15.1.1 The Causal Sufficiency Condition

These considerations lead to the *[Causal Sufficiency Criterion]*(CSC): $A_1, \ldots, A_n$ form a sufficient set of causal factors for B (equivalently, $A_1 \wedge A_2 \wedge \ldots \wedge A_n$ is a sufficient cause for B) if and only if

$$\neg B \;\Box\!\!\rightarrow\; \neg(A_1 \wedge A_2 \wedge \ldots \wedge A_n) \tag{15.1}$$

which leads to the

Causal Completeness Test (CCT): $A_1, \ldots, A_n$ form a *satisfactory set* of causal factors for B if and only if

$$\begin{aligned}&\wedge A_1 \Rrightarrow B\\&\wedge A_2 \Rrightarrow B\\&\wedge \ldots\\&\wedge A_n \Rrightarrow B\\&\wedge \neg B \;\Box\!\!\rightarrow\; \neg(A_1 \wedge A_2 \wedge \ldots \wedge A_n)\end{aligned} \tag{15.2}$$

This in turn yields another meta-axiom for causal histories:

Axiom 9 *MetaAxiom: For any B which has a causal factor (that is, for which there is an A such that $A \Rrightarrow B$ is established), the set $\{A|A \Rrightarrow B\}$ must be a satisfactory set of causal factors for B.*

15.1.2 The CCT for [11] and [12]

So is it the case that [11] and [12] form a sufficient set for [1]? We argue informally, because we have as yet not introduced a method for analysing the formal semantics of sentences containing intensional contexts, such as '*realizes that*' or '*believes that*' or '*judges that*'. It follows from [11] and the properties of the intensional context 'realize' that if one realizes that P, then P must be true (else one is, say, merely 'imagining that P'). So we imagine we should be able to prove

$$\vdash [11] \rightarrow \textit{AC is in the course of landing on a particular runway}$$

and this runway happens to be BRU RWY 25. Suppose that in the nearest worlds in which the AC had not landed on this runway (this is the $\neg B$ hypothesis, with $B = [1]$), then it could not have been the case both that the AC was landing and the CRW opted to continue: either the AC was not in landing configuration and process (e.g., someone had inadvertently triggered the TOGA, *take-off/go-around* autopilot switch) or the CRW had not opted to continue (i.e., they opted to go around, this being according to procedure the only available other option - let us presume that opting to do something does not necessarily mean a conscious explicit decision was taken: one can opt to continue simply by continuing). Since it could not have been the case that (AC was landing and the CRW opted to continue), it could not have been the case that [11] $\wedge$ [12] (by contraposition, since this implies the former). Thus the counterfactual is established and the CSC is fulfilled. Since we already know [11] $\Rrightarrow$ [1] and [12] $\Rrightarrow$ [1], the CCT is fulfilled.

In order to make this argument, we have not only relied on particular properties of the intensional context '*realize that*', we have also implicitly relied on particular properties of referential expressions. For although the AC was landing on BRU RWY 25, the CRW may not have realised *this* fact in its entirety, namely that this runway was BRU RWY 25. It suffices that they realized that they were landing on a runway at the 'wrong' airport, and that it follows from the semantics of 'realize' that the AC was indeed landing on whichever runway it was landing on. It also is the case that the runway the aircraft was landing on is in fact identical with BRU RWY 25. But this identity has to be asserted outside the intensional context – intensional truths are not necessarily preserved under syntactical substitution of equireferential expressions (in fact, the definition of '*intensional context*' is roughly speaking that there are two equireferential expressions that yield different truth values when placed in this context). But these are subtleties of semantics that are well handled for our purposes by philosophical logic, and it strays from our purpose to do more than remark them here. The moral is, however, that arguments for counterfactual conditionals, whether for $\Rightarrow\!\!\!\!\succ$ or for the CCT, may depend on logical analyses of the node descriptions.

In general, it is not the case that satisfactory sets for a given B are unique. They depend on how one wants to analyse the situation; how much one wants to say, what depth and detail of analysis one wants to pursue, the limitations of the language chosen to express the nodes. All of this is very much the choice of the investigator. For example, in our incident, we regard it as sufficient that [11] and [12], whereas someone with great interest in the psychology and consequences of flight-deck decision making may wish to analyse these circumstances further into a causally-explanatory chain using, say a PARDIA (perception-attention-reasoning-decision-intention-action) model such as we have used.

A similar situation exists in pure axiomatic mathematics. One is provided with proof rules sufficient to get any job done, but what proofs are constructed and how are up to the individual wishes and skill of the user. Proofs may be more detailed or less detailed, easy-to-follow or cleverly slick, pro forma or creative. Yet the criteria for being a valid proof remain constant throughout the enterprise. So with WBA. We have no wish to regulate whether an analysis is most subtle, or how it indicates what future steps to take to prevent recurrences, or whether it must use the latest theory of human-machine interaction. We wish to lay out criteria and reasoning rules for providing a formally-complete causal explanation, according to assumptions that an analyst makes in a particular case. We thereby make the assumptions clear, explicit and precise, exhibit their role in the explanation, and make the reasoning clear. This is a worthy goal, just as a formal verification of an algorithm or a safety-critical design is worthy in so far as it gives us certainty that the system design fulfils its purpose. The issue of whether the algorithm or design is somehow the best or the most efficient is not addressed by the issue of verification. Likewise, we don't address such individual properties of causal explanation.

15.1.3 Inference Rules for the CCT

One can provide inference rules that ensure that the CCT is fulfilled. We introduce the relation $\Box\!\!\Rightarrow$ to stand for 'sufficient cause', and establishment of a number of instances of this relation will be the goal of a WBA. Thus one goal of the CCT is to show

$$A_1 \wedge A_2 \wedge \ldots A_n \ \Box\!\!\Rightarrow\ B \tag{15.3}$$

Were we to introduce $\Box\!\!\Rightarrow$ for the notion of 'necessary and sufficient cause', we could think about introducing the rule

$$\begin{array}{l} C \ \Box\!\!\rightarrow\ B \\ \neg C \ \Box\!\!\rightarrow\ \neg B \\ \neg B \ \Box\!\!\rightarrow\ \neg C \\ \hline\hline \quad C \ \Box\!\!\Rightarrow\ B \end{array} \tag{15.4}$$

as a definition of $\Box\!\!\Rightarrow$. This would make the logic cleaner, since as we shall see in Chapter 20, the alethic and deontic operators are definable in terms of $\Box\!\!\rightarrow$, and this would make $\Box\!\!\Rightarrow$ also so definable. This would extend the completeness argument given there for a fragment of the logic to the whole logic. Technically very desirable, and very clean. However, one major consideration speaks against it, which we discuss in Section 15.1.4 below. Namely, if the facts 'on the ground', along with standard operating procedures, *logically entail* the occurrence of an event or a state, and the procedures were also followed, then the occurrence of the event or state can be intuitively completely explained simply by adducing the facts and the adherence to the procedures. However, not *all* of the standard operating procedures will be necessary conditions of the event - in general one could expect many of them not to influence the event at all. It would be a logical problem, interesting perhaps, but not always directly relevant to an incident investigation, to determine precisely which of the procedures might have been necessary to the occurrence of the event. We would not wish to enforce this investigation in all circumstances. Accordingly, we do not insist that $\Box\!\!\Rightarrow$ states necessary and sufficient causes, but only sufficient causes in these circumstances. Thus the first $\Box\!\!\Rightarrow$ rule is unidirectional:

$$\begin{array}{l} C \\ B \\ \neg C \ \Box\!\!\rightarrow\ \neg B \\ \neg B \ \Box\!\!\rightarrow\ \neg C \\ \hline \quad C \ \Box\!\!\Rightarrow\ B \end{array} \tag{15.5}$$

The question remains if the CCT, which has a variable number of premisses, is derivable from rules which have a fixed number of premisses. We claim the

CCT is indeed derivable from Rule 15.5 along with the rule

$$\frac{\begin{array}{c} A \;\Box\!\!\rightarrow\; C \\ B \;\Box\!\!\rightarrow\; C \end{array}}{(A \vee B) \;\Box\!\!\rightarrow\; C} \tag{15.6}$$

Theorem 1 *The CCT is derivable from Rules 15.5 and 15.6.*

The argument for this theorem, consisting of a hierarchical proof reducing a CCT-conclusion to the CCT-premisses using Rules 15.5 and 15.6 may be found in Appendix A.3.

15.1.4 Causal Sufficiency Through Procedural Necessity

There is also another way in which causal sufficiency can be explained. Suppose we are able to derive that

$$\Box \mathit{Procedures} \;\succ\; \Diamond \mathit{Event}$$

that is, that the occurrence of *Event* follows necessarily upon the implementation of the *Procedures*. Then it is a sufficient explanation why the *Event* occurred simply that the *Procedures* were followed.

For example, if we have implemented *Procedures* as a state machine, and this state machine is in state S, and through action A this state machine attains state S_1, then it is a causally sufficient explanation why the machine is in state S_1 that

- the specification of the state machine was followed; and
- the machine was shortly before in state S; and
- the action A then occurred.

This is to say the following, formally. Suppose we can demonstrate that from the *Procedures* and *Hypotheses* (former state, plus 'environmental' conditions being in order, plus action A being made) that the *Event* must occur:

$$\vdash_{TLA} \mathit{Hypotheses} \wedge \Box \mathit{Procedures} \;\Rightarrow\; \Diamond \mathit{Event}$$

then we can conclude that

$$(\mathit{Hypotheses} \wedge \Box \mathit{Procedures}) \;\succ\; \Diamond \mathit{Event}$$

by Rule 14.21. And we further suppose that the *Hypotheses* and *Procedures* were followed: then *Event* must occur, and *Hypotheses* and *Procedures* are a sufficient explanation for it:

$$\begin{array}{l} \mathit{Hypotheses} \\ \mathit{Procedures} \\ (\mathit{Hypotheses} \wedge \Box \mathit{Procedures}) \succ \Diamond \mathit{Event} \\ \Diamond \mathit{Event} \\ \hline (\mathit{Hypotheses} \wedge \mathit{Procedures}) \ \Box\!\!\Rightarrow \Diamond \mathit{Event} \end{array} \tag{15.7}$$

Contrast this rule with the rule we would have were we to define $\Box\!\!\Rightarrow$ as 'necessary and sufficient cause', i.e., as

$$A \Box\!\!\Rightarrow B \ \triangleq \ (A \Rightarrow\!\!\succ B \wedge \neg B \Box\!\!\rightarrow \neg A)$$

which in terms of $\Box\!\!\rightarrow$ alone is

$$A \Box\!\!\Rightarrow B \ \triangleq \ (A \Box\!\!\rightarrow B \wedge \neg A \Box\!\!\rightarrow \neg B) \wedge \neg B \Box\!\!\rightarrow \neg A)$$

Then the following Rule 15.8 would be a derived rule:

$$\begin{array}{l} \mathit{Hypotheses} \\ \mathit{Procedures} \\ (\mathit{Hypotheses} \wedge \Box \mathit{Procedures}) \succ \Diamond \mathit{Event} \\ \neg(\mathit{Hypotheses} \wedge \Box \mathit{Procedures}) \ \Box\!\!\rightarrow \neg \Diamond \mathit{Event} \\ \Diamond \mathit{Event} \\ \hline (\mathit{Hypotheses} \wedge \mathit{Procedures}) \ \Box\!\!\Rightarrow \Diamond \mathit{Event} \end{array} \tag{15.8}$$

Theorem 2 *Rule 15.8 is derivable in EL from Rule 15.4.*

Proof Sketch: The proof proceeds as follows. The three hypotheses $\mathit{Hypotheses} \wedge \mathit{Procedures} \wedge \Diamond \mathit{Event}$ of Rule 15.8 entail $(\mathit{Hypotheses} \wedge \mathit{Procedures}) \Box\!\!\rightarrow \Diamond \mathit{Event}$, the first hypothesis of Rule 15.4. The hypothesis $(\mathit{Hypotheses} \wedge \Box \mathit{Procedures}) \succ \Diamond \mathit{Event}$ of Rule 15.8 entails $\neg \Diamond \mathit{Event} \succ \neg(\mathit{Hypotheses} \wedge \Box \mathit{Procedures})$ by the logic of $\succ$, which in turn entails $\neg \Diamond \mathit{Event} \Box\!\!\rightarrow \neg(\mathit{Hypotheses} \wedge \Box \mathit{Procedures})$, the second hypothesis of Rule 15.4. The third hypothesis of Rule 15.4 is $\neg(\mathit{Hypotheses} \wedge \Box \mathit{Procedures}) \ \Box\!\!\rightarrow \ \neg \Diamond \mathit{Event}$, which is also a hypothesis of Rule 15.8. Therefore the conclusion of Rule 15.4, which is also the conclusion of Rule 15.8, follows. □

We prefer Rule 15.7 to Rule 15.8. The following consideration speaks against acceptance of Rule 15.8. Suppose we're grilling sausages. If we leave the sausages under the grill, and the grill functions correctly, they will eventually cook. We put sausages under the grill, and turn the grill on. Unbeknownst to us, someone took the sausages from the grill and put them in the microwave for two minutes, and then back under the grill. Lo and behold, the sausages are cooked. It is a sufficient explanation for the fact that the sausages were cooked that they were under the grill and the grill was on and functioning correctly, but it is not *in this* *case* why the sausages cooked. In this case, they cooked partially under the grill and partially in the microwave. They cooked anyway, whether or not the grill was on or functioning. Therefore the condition, that had the grill not been on or functioning the sausages would not have cooked, is false. However, this condition is precisely the condition which appears as a hypothesis in Rule 15.8 which is absent from Rule 15.7.

Consider now the following. It *is* a sufficient explanation for the cooking of the sausages that they were put under the grill and the grill functioned. The world can be considered to split into two possibilities after this:

1. the sausages remained under the grill,and cooked
2. the sausages were removed, put in the microwave, cooked, and returned to the grill, and cooked more

Under both of these possibilities, the sausages cooked. From the event that the sausages were placed under the grill, it follows that they cooked, period. In this case, it seems that the placing of the sausages under the grill is therefore sufficient (under all possible circumstances, we would add) to result in the sausages being cooked. "*You can only get to there from here, and to no other place.*" This is, therefore, an appropriate *why...because...* explanation.

Although expressed in words, this reasoning has direct bearing on the description of procedures in a formal language such as TLA+, which we shall demonstrate later. A description in TLA+ of procedures can be 'higher-level' or 'lower-level', abstract or detailed. The level of detail of the procedures to be described will depend upon the description of the event whose occurrence these procedures explain. TLA+ has the ability to talk about timeliness of procedures to the level of detail required to do physics (so, a real-time component). But such details may as well be left out, if the consequence is (tense-logically) inevitable. If all that is required is to derive the consequence, it will not matter what path the world takes.

The following Rule 15.9 would be a derived rule of the logic with Rule 15.7 alone, since $\succ$ is transitive. However, since we also include the inference possibilities for $\Box\!\!\Rightarrow$ with Rule 15.5, strictly enlarging the set of provable formulae whose major connective is $\Box\!\!\Rightarrow$, we explicitly add it as a rule of EL:

$$\frac{\begin{array}{l} X \\ C \\ X \succ C \\ (C \wedge A) \Box\!\!\Rightarrow B \end{array}}{(X \wedge A) \Box\!\!\Rightarrow B} \qquad (15.9)$$

The justification for the rule is that if the truth of C along with the truth of A suffices as a causal explanation of B, and X entails C logically, then if X indeed is true, an alternative sufficient explanation for B is the truth of X along with A. (We remind the reader that sufficient causal explanations are not unique.) Again, such a rule would not be valid if we were seeking minimal sufficient causal explanations, since the formula $C \wedge Y$ would be such an X for *any* Y; and since Y can be anything, we can take it that some choice of true Y will be spurious to the causal explanation of B (although we will not establish this assumption rigorously here). Since Rule 15.9 is invalid for minimal causal explanations (which follows from considering that $\Box\!\!\rightarrow$ is not transitive), and Rule 15.5 is consistent with $\Box\!\!\Rightarrow$ providing minimal causal explanation, it follows that Rule 15.9 is not derivable from Rule 15.5, thereby necessitating its explicit inclusion in the set of rules of EL.

15.1.5 The Next Step

Our first step was to try to establish a series of causal factors for the incident itself, proceeding according to the use of the inference rules in a hierarchical proof. We were able to find a chain of causation from $\langle 2 \rangle$ to [1]. Thus we try to apply it on the next step, too. Why did the aircraft enter BATC area? - Unfortunately, searching for facts effecting $\langle 2 \rangle$ produces no satisfying answers because we cannot find any direct statement concerning causes for the transition from LATC area to BATC area.

Thus, we may reach a dead-end in trying to explain a relation $\hookrightarrow$ that occurs in the history, but for which no suitable causal-factor candidates appear in the history. To find the reasons leading to this node we need to expand our method by additional techniques.

15.2 Contrastive Explanation

Contrastive explanation concerns the explanation of facts of the form *why P rather than Q occurred*. Lewis [Lew86][pp229-230] suggests this may be accomplished by giving information about the causal history of P that would not have applied to the history of Q. Lipton [Lip91][p42] notes that this criterion allows for unexplanatory causes. J.S. Mill's *Method of Differences* [Mil73a][III.VIII.2] relies on the principle that a cause must lie among the antecedent differences between a case in which an effect occurs and a case in which it does not. Mill notes that this works best with *diachronic* (before/after) contrasts. Lipton [Lip91][p43] proposes the *Difference Condition*:

> *To explain why P rather than Q, we must cite a causal difference between P and not-Q, consisting of a cause of P and the absence of a corresponding event in the case of not-Q.*

(Lipton considers here that only events may be causes. We are considering causal factors to include nodes of all types, so relevant modifications must be made to this expression of the Difference Condition.) How may we apply such a principle of contrastive explanation, in some form, to our problem?

In any incident there is a contrast to be explained, namely why things happened the unfortunate way they did rather than the way *they should have done*, or the way *we expected them to have done*. There are (at least) two *possible worlds* here: one is the actual world, the way things actually happened, and the other is the world in which the things happened which should have (there may be more than one of these). Let us call this latter the *deontically-correct* world(s). In the example, in the actual world the aircraft landed at BRU, and in the deontically-correct world the aircraft landed at FRA. According to any of the three conditions proposed, we may look for a relevant difference in the causal history (let's just say history). Looking at 14.1 again, we find it in the transition from $\langle 3 \rangle$ to $\langle 2 \rangle$, because in the deontically-correct world, this should have been a transition $\langle 3 \rangle$ to $\langle 2b \rangle$, where

$\langle 2b \rangle$ AC in Maastricht ATC area

Note that $\langle 2b \rangle$ and $\langle 2 \rangle$ are contraries: they cannot both hold in a given history. We are trying to look back here to the first general difference. Recall from Axiom 1 that causal priority mirrors temporal priority, so that any relevant causally-prior difference will also be temporally-prior: we thus look for the lastest difference in events, represented by the truth of $\langle 3 \rangle \hookrightarrow \langle 2 \rangle$ rather than that of $\langle 3 \rangle \hookrightarrow \langle 2b \rangle$ and look for causes of that difference.

15.2.1 Proceeding by determining earlier contrast

The sources we used up to now do not provide much information on the procedures and laws of civil aviation. Having this *expert knowledge* at some point of incident and accident analysis is essential, as we have already remarked. What we need to know to explain $\langle 2 \rangle$ is, that there are

$\langle 21 \rangle$ LATC procedures

to handle flight traffic within the area it controls as well as from and to this area. Such general procedures are stated explicitly in documents such as the U.S. *Air Traffic Control Handbook*, *Federal Aviation Regulations* (FAR) and *International Civil Aviation Organization* (ICAO) standards, which describe general operating and flight rules as well as standard instrument procedures and communication. Procedures are also complicated, there are a lot of them, and they may not be precisely, let alone formally, stated. And the procedures themselves might be inoptimally designed, or contain loopholes. All this makes our analysis more difficult. So be it. The WBA analysis principles remain the same. One advantageous outcome of a formal approach is to be able to demonstrate clearly why, how and where more precise description of procedures is required.

A description of procedures is similar to a specification of a system. The system is said to fulfil a specification in the same manner that events and states are said to conform to procedures: if the procedures hold (if the specification is fulfilled), the events and states must necessarily follow. Our purpose here is to 'debug' procedures and specifications: that is, we do not know in advance whether they were in fact followed and we are attempting to find out.

We can use deontic reasoning for specifications similarly to the manner in which we used this reasoning in Section 14.6. Suppose *Spec* is a specification of a system or subsystem that is relied on by the ATC system. Suppose E is a state or event whose occurrence follows logically from *Spec*, that is:

$$\vdash_{TLA} Spec \rightarrow \Diamond E \tag{15.10}$$

Since we're dealing with a failure situation, we don't know whether *Spec* is in fact fulfilled, but we *do* know it ought to be: We can use the axiom

Axiom 10 $\vdash O(Spec)$

and thus handle system reasoning similar to the manner in which we handled reasoning concerning air traffic procedures: we may conclude by Rule 14.22 that, given these circumstances, $O(\Diamond E)$. That is: the values of the state predicates should be changed by events exactly as specified.

In our examples, a part of the procedures is to get

⟨22⟩ ATC data

– especially flight plans – from adjacent ATCs for all aircraft entering an ATCs area and passing this information according to each

⟨5⟩ aircraft's flightplan

to aircraft and other ATCs concerned. This information helps us to explain the circumstances which led the aircraft from LATC to BATC. It is important to mention that there is a presumption that the data ought to be processed as described above, rigorously conformant to ATC procedures. However, intuitively if everything *did* occur as determined in the procedures, something must be wrong with the processed data, since we know they landed in BRU instead of FRA. Although this information is given directly in our primary source (it is pretty clearly the result of inference rather than direct knowledge), we can also derive it from the facts we already collected by logic:
Comparing the part of the model we have analysed so far of what *did* happen (figure 15.1 solid lines) to what *should* have happened – namely

⟨2*b*⟩ AC in Maastricht ATC area

succeeded by

⟨2*b*1⟩ AC in Frankfurt ATC area
[2*b*11] AC lands at FRA

(figure 15.1, dashed symbols) following the proper direction to Frankfurt – we can easily determine the difference between these possible "*worlds*", with the split lying at node *langle*3⟩. The situation we're considering here is determinate: we know that one of these worlds occurred and one didn't. Mill's criterion tells us that the difference must lie in factors present at node *langle*3⟩. The aircraft was handed off by LATC to a different ATC than the one it should have been handed off to. It is therefore time to consider ATC procedures in more detail.

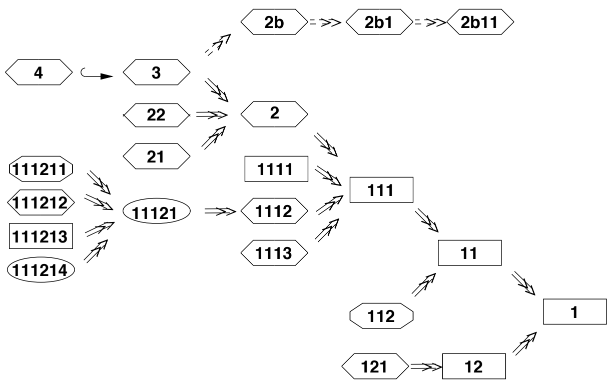

Figure 15.1: Possible and realized causalities

Before we consider ATC procedures in detail, we should remark on the situation in which we may wish to use a Difference Criterion, but we do not have enough information (or have not inferred enough information) to be able to tell precisely which state pertained or event happened. There are two ways of handling such a case. Either we can abstract from the details, and find one more general way of expressing what was the case and contrast it with what should have been the case, or we can consider all the possibilities individually. The first possibility involves ignoring some information, and intuition may tell us that if we're having trouble searching for a causal factor with more detail, we're hardly going to be more successful with less. However, the second possibility involves considering all the possible alternatives, the 'possible worlds'. Having more than one possible world to consider creates additional work for an analyst, because every fact newly introduced into one world's history often needs to be introduced in all other worlds, too, namely when the fact is independent of the uncertainty which creates the possible worlds. In such a case, considering all alternatives would lead to exponential growth of the number of worlds to be worked on and thus an exponential growth in the amount of work we have to do. To avoid this ultimately unmanageable situation, one must find a way to eliminate some of these worlds. Lewis suggests weighting each world by assigning possibilities based on expert knowledge [Lew86]. By defining a threshold, worlds with lower possibilities can be taken out of focus to some extent.

Such techniques represent a level of sophistication beyond that with which we currently feel confident of, and on which we can demonstrate practical application. Rather than trying to determine such probabilities, we shall simply try to represent the possible alternatives, and we shall find this approach tolerable for the present example.

15.2.2 Mill's Other Methods

Mill elaborated five *Methods of Induction* [Mil73a]. We have chosen to use the Method of Difference alone because, in Mackie's words [Mac74, p. 148],

> The Method of Difference is .... the principle method for discovering and establishing neolithic causal regularites, while variants of it, such as Mill's Method of Residues, and related devices like the Method of Agreement and some of the things covered by the name of the Joint Method, can also play some part.

(A *neolithic causal regularity* is a singular causal occurrence, as opposed to a causal law.) That the Method of Difference is most fundamental amongst Mill's methods suggests to us that it may not need supplementing. Mill's other methods address lawlike, that is, repeated, regularities and thus have less obvious application to our need. However, we should not hesitate to use any of Mill's other methods as a means of bootstrapping a causal inquiry should one prove to be needed, towards the main goal of explaining incidents satisfactorily.