

FOLLOWING CLUES: Do's and Don't's

Many people by profession have to follow clues. They include, among others, detectives (who follow clues left behind by criminals), scientists (who follow clues left behind by nature), physicians (who follow clues left behind by ailing bodies), and automotive mechanics (who follow clues left by malfunctioning motor-cars). By following clues these professionals find out things to which they do not have direct access. Detectives are not given advance notice informing them where and when a crime is to be committed, so that they can be there to witness it. Scientists have not been given a key to God's library where they can sit down comfortably and copy down the laws of nature. Physicians are not the microbes themselves crawling around in our bodies. And automotive mechanics—we would not let them work on our cars if they have no idea what is wrong with them before they start.

But not only professionals follow clues many ordinary people do as well. Judging by the popularity of detective stories there are probably more amateur sleuths than professional detectives. Besides taking our cars to commercial garages many of us work on them

ourselves on week-ends. In the early days of science many who contributed to its development were wealthy enough so that they did not have to work for a living. Even nowadays there are amateur scientists. For good reasons the law prevents us from being amateur physicians, but this has not stopped us from paying attention to our own symptoms and guessing what is wrong with us when we are not feeling well.

Needless to say, whether in following clues or in other kinds of activities amateurs usually (but not always) do not fare as well as professionals, but this is not to say that amateurs know nothing about following clues. Indeed, because so many people follow clues, what we should do when following clues has become part of popular culture. To cite just one example, who these days have not heard of the process of elimination? A murder has been committed. Clues tell us it is the butler or the maid or the master of the house. It is not the butler or the maid because of such and such. Therefore it must be the master of the house.

However, while much of what we should do when following clues has become part of popular culture, no one has collected together this body of knowledge and made a study of it. There are books aplenty written by knowledgeable people on how to collect or interpret particular kinds of clues—for example, how to record blood splatter

patterns and what to make of them, but there is no book on how we should follow clues in general. Blood splatter patterns are of interest to detectives investigating a murder, but not to a mechanic trying to troubleshoot a car. There are things of interest to specific kinds of investigators and there are things of interest to all investigators. Take as example the process of elimination again. This process is employed not just by detectives, but by all who follow clues. But the process of elimination is just one item in the investigator's tool box. What other items are there? And why? That is, why are all these items useful?

Do we need to ask why? Do we need to know the reasons for the many things we do when following clues? We know in practice what to do and when. Is that not sufficient?

Clearly, it is not sufficient. If we do not understand the reasons behind our practice we might achieve the wrong results. We expect clues to lead us to the truth, but if we make mistakes in following clues (because we do not understand what we are doing), we might reach not the truth but something very different. And this has been said of many of the things we *think* we should do when following clues. Take the process of elimination again. Is there any one who is not familiar with the argument that it is illegitimate? There is an assumption we make when employing the process of elimination. We assume that we have exhausted all possibilities. It is the butler or the maid or the

master of the house. *No one else*. Only when this is the case and we have eliminated the butler and the maid, can we then conclude that it must be the master of the house. But, critics of this process say, people who follow clues are working in the dark (otherwise why do they have to follow clues?). How can people who work in the dark know they have exhausted all possibilities? And if they do not know, how can they know it must be the master of the house? Let us suppose for argument's sake, they continue, that the master of the house confesses. Is this a vindication of the process of elimination? No, the critics say. If he confesses and we think we have caught the right person we could have become unwitting participants in a charade, one which the master of the house engineers: he is not guilty, but he wants to protect the person who is (and that person is neither the butler nor the maid).

Are these critics right? Is the process of elimination legitimate? Does this process depend on our exhausting all possibilities? Suppose we have failed to exhaust all possibilities; can we still find what we want to know?

It seems impossible that we can still find what we want to know if we fail to exhaust all possibilities. But how can we be sure that we

have not? Yet the process of elimination is in use! Sherlock Holmes is a champion of it.¹ Is Sherlock Holmes and a million others wrong?

Now Sherlock Holmes could be wrong, and even a million others, but before we say they are wrong we should look into the clue-following process. There is danger in giving too much weight to argument. Sometimes things that 'should not work' in fact does. Heavier-than-air machines 'should not be able to fly'; yet they do.

There are many things we do in practice the reason for which we do not know or are not clear. Sometimes there is not much we can do in this unhappy situation: looking for reasons is not always easy; sometimes we do not even know where to start. But in the present case; in the case of following clues; I think I know where to begin.

The present essay is a part of my study of how we should follow clues. I have given this study a name: Theseology, after Theseus (the name Theseology should not be confused with Theology). Theseus went into the Labyrinth to kill the Minotaur. To ensure that he could come out again he unwound a ball of thread on his way in. When it came time to leave all he had to do was to retrace his steps by following this thread. Following clues is a kind of retracing, and is often compared to Theseus following his thread.

¹ 'It is an old maxim of mine that when you have excluded the impossible, whatever remains, however improbable, must be the truth.' – *The Beryl Coronet*

From the word 'theseology' we can derive the verb 'to theseologise' and the noun 'theseologiser'. To theseologise is to follow clues. A theseologiser is a person who follows clues. Scientists are theseologisers; so is Sherlock Holmes; so are automotive mechanics. We can also make up the phrase 'the theseological process'; it means the same as 'the clue-following process'.

Concerning the clue-following or theseological process there is a common misunderstanding. People often think that the reason we can come to know by following clues lies with the way we reason. Things we wish to know by following clues are things to which we do not have direct access. Since these things cannot be known directly, the argument goes, they must be known indirectly, by inference. Now if this popular view is correct, Theseology should be a branch of Logic, since Logic teaches us how to reason. But Theseology is *not* a branch of logic. Theseology and logic have each its own distinctive subject matter. Logic is about how we should reason, *irrespective of what we are reasoning about*. Theseology is about how we should follow *clues*. In the process of following clues we have to reason, but this reasoning is directed towards helping us find clues and make sense of them. It is obvious that in following clues we have to reason, but that we have to reason when following clues does not make Theseology into a branch of logic. In photography too we have to reason, and photography

clearly is not a branch of logic. In photography we have to do other things besides reasoning, or no pictures will result. When following clues we also have to do other things besides reasoning, or we will not be able to find those things we want to know.

The knowledge we obtain when following clues is indirect. But that it is indirect does not mean that it is obtained through reasoning.² There is no doubt that reasoning plays a role, but it is a supporting role. When I recognise my good friend Sydney on the street no one would say I do so by reasoning. I know Sydney. When I see him I know it is him. Now to discover by following clues, we will soon see, is at heart a process of recognition. It is a process of recognition spread out through time, or, we may say, a process of recognition in slow motion.

Part I: Do's

Do look for as many clues as you can

When we follow clues we do so with a view to reconstructing some structure or structures that is hidden from us. A crime has been

² What is obtained indirectly is not necessarily obtained through reasoning. A microscope allows us to see the structure of a cell. We cannot see the structure of a cell directly, with our naked eyes. But a microscope is an instrument, not some form of reasoning.

committed. To find out what happened we follow clues. A crime has a structure: there is a motive for the crime, which leads to certain actions; these actions take place in a certain order; and *they* have consequences which are not arbitrary. For example, fingers not protected by gloves will leave behind fingerprints. By following clues we can find out what the motive for the crime was and also what happened. But this finding out is actually a reconstruction. Basing ourselves on clues we describe what is likely to have taken place; we do so *without having actually witnessed the crime* (since we were not present when the crime was committed). In this respect the crime is hidden from us. But that it is hidden does not mean we cannot find out. Even though we were not present we can reconstruct what happened by relying on clues. This reconstruction will not be perfect in every detail but in the better cases will have captured enough of the essentials.

Now not only crimes have structures, but other things as well—for example, the universe, or, to introduce something I have not mentioned before, ciphers (not all ciphers, but some of them). This is why scientists follow clues, and cryptanalysts (code-breakers). By following clues scientists can discover the structure of the universe, or rather the many layers of structures in the universe. They can do so without asking God. Similarly, by following clues cryptanalysts can

crack ciphers. They can do so without asking those who constructed these ciphers.

Now this last point is useful to keep in mind. *When we follow clues and things go our way, we do not have to be told.* In the old days people believed crimes could be solved only if the criminals confessed, so they employed all kinds of means to bring about this confession, including torture. But we do not share this view any more. We all know now that if we have enough clues we can reconstruct what happened even without confessions. Indeed, nowadays in most civilised nations people accused of crimes have the right to remain silent.

To solve a crime we do not need the criminal to confess. Similarly, to find out the many layers of structures present in the universe we do not have to ask God. And obviously, when we are trying to solve ciphers we cannot pick up the phone and ask our opponents how their ciphers work.

Why are clues so powerful? Why is it that by following them we can reconstruct structures hidden from us? Now it is not so hard to find an answer to this question. Examine the following:

SBR SBCTU DBCKERVS FCGG WTTXCR SFH FRRJD YTHE SHUWI

Here we have a series of letters that appear to be hiding a message (we do not have to be certain that it is, but it seems reasonable to suspect that it is). By examining it we will discover clues. For example, it might occur to us that the two letters SB at the beginning of the first two groups of letters might stand for TH (or some such). The groups of letters are words, we tell ourselves. The first two words both start with TH. SB is a clue. We would do well to treat it as such if we want to find out whether these words come together to form a message ...

So, what are clues?

From our example we can see without much difficulty that they are but the characteristics of structures, disguised. The English language has a structure. In English there are rules governing how letters are put together to form words and how words are put together to form sentences. As a structure the English language has certain characteristics. One of these is that many of its words start with TH. The first two words in this cryptogram could start with TH. We do not see TH only because they have been disguised as SB.

Clues are not the characteristics of structures plain and simple, but the characteristics of structures disguised. Even so, if we can remove their disguises we can discover what is behind them. And this is why we should look for as many clues as we can in an investigation—that is, when we are trying to reconstruct hidden structures. For the

more characteristics of a structure we know, the easier it is to reconstruct it, just as the more characteristics of a person we know the easier it is to locate that person and not confuse her with somebody else. If we know only the first name of a person we cannot find that person because it is likely that there are many people with that same first name. But if we know not just the first name, but also the last, and in addition who her parents are, when she was born, what she works at, and so on and so on ...; if we know all these many things our task will become easier. Now the same with reconstructing hidden structures. The more characteristics of the hidden structure we know, the easier it is to reconstruct it and not confuse it with some other structure. And this is why, as we have said, we should look for as many clues as we can. Take solving a crime for example. If there are many clues and we can decipher all of them (that is, remove their disguises), we should be able to find out many of the details of the crime. And it is obvious that the more details we have, the easier it is to determine who is responsible.

In solving a crime we will be interested in knowing who is responsible. But we should not expect that we can find out who is responsible just by one clue. We cannot say the butler is responsible because the butler was holding a smoking gun. The butler could have been shooting at the person who shot the victim. If we want to find out

who is responsible we will have to find out in some detail what took place. The more details we can uncover the easier it is to determine who is responsible.

In an investigation we should look for as many clues as we can. The more clues we have the easier it is to reconstruct the structures that have been hidden from us.

But while we should look for as many clues as we can in an investigation, in practice quite often we simply cannot find all that many, however hard we try. In crime investigation, very often the clues found are few in number. On these occasions we either do not know who is responsible, or we cannot be certain.

There is one kind of investigation in which on occasion we have managed to find a huge number of clues. These clues usually are not found in one day but over a long period of time. This has occurred in the investigation of nature (science) and is the reason why in the mature sciences we can have a high degree of certainty. In the younger sciences, due to the small number of clues, the results are often uncertain.

Do try to develop new clues from old

In an investigation we should look for as many clues as we can. The more clues we have, the easier it is to find the things we are

looking for. Clues are the characteristics of structures. The more characteristics of a structure we know the easier it is to reconstruct it. But it is not always easy to find clues, as we all know. Suppose in an investigation we have found only a few (say, just one or two); does this mean that our investigation is doomed? That we will never be able to find what we are looking for?

Those familiar with investigations will know that it does not. Most investigations start out with very few clues. Sometimes we think ourselves lucky if we can find just one. But that there are very few clues at the beginning does not necessarily mean that the investigation will never bear fruit. If it did, there would have been very few successful investigations, if any. One remarkable characteristic of the theseological process is that often, or often enough, new clues can be developed from old clues, so that even though the original number of clues is small eventually there will be enough to bring the investigation to an acceptable conclusion.

How do we develop new clues from old? We answer this question in the next section.

Do apply hypotheses to the evidence

In most investigations, if we are to succeed we will have to develop new clues from old, as we have explained in the last section.

Now in simple investigations developing new clues from old is something we often do automatically. As illustration, look to the cryptanalytic example we have been using.

SBR SBCTU DBCKERVS FCGG WTT CXR SFH FRRJD YTHE SHUWI

Here we have a cryptogram. Clues that we can easily detect in this cryptogram tell us that S probably stands for T, and F, for W (we leave out for the moment the other clues). Now if we did adopt these two hypotheses, most of us would very naturally go on to translate all the Ss in the message to Ts, and all the Fs, to Ws, as we have done below.

T T T W TW W T
SBR SBCTU DBCKERVS FCGG WTT CXR SFH FRRJD YTHE SHUWI

But once we have done this we will have developed a new clue, in the sixth word. The sixth word is SFH. Since S stands for T, and F, for W, the sixth word now appears as a three-letter word that starts with TW. What could the third letter be? It is easy to answer this question. But this is to say, TW? is a clue, and it is a new clue since it was not there originally: it only appears after we have made the translations.

Notice how easy it is to develop this new clue. We merely translate all the Ss to Ts, and Fs to Ws, and voila! a new clue appears. Developing this new clue is so easy that in practice it is not likely that we would have taken time to remark upon its appearance; instead, we are likely simply to take advantage of it and carry on with the decipherment.

In simple cases developing new clues requires no conscious decision on our part and takes next to no effort. It is something we do, as we have been saying, automatically.

But this will not always be the case in all investigations. In more complicated investigations developing new clues will require conscious and deliberate effort.

Suppose we meet with one of these more complicated cases. Suppose we find ourselves in a situation in which we have decided that we need new clues; what do we do? How do we consciously go about developing new clues from old?

To answer this question we look back to what we have done in our cryptanalytic example. How did we develop the new clue that suggested to us that SFH was in fact TWO?

If we retrace our steps and pay a little more attention this time, we will find that this new clue appears as a result of applying hypotheses suggested by old clues to the evidence (which in this case

is the cryptogram). Before we develop our new clue we have already proposed a number of hypotheses in answer to the clues already present. These include the two that say, respectively, that S stands for T and F, for W. When we applied these two hypotheses to the cryptogram, translating all the Ss to Ts and Fs to Ws, the new clue TW? appears.

So this is what we should do when we want to develop new clues from old. We apply the hypotheses suggested by old clues to the evidence. It is by applying hypotheses to evidence that we develop new clues.

We do not apply just any hypotheses to evidence. We apply the hypotheses suggested by old clues to the evidence. A hypothesis *not* suggested by any clues has little chance of being right. A hypothesis suggested by a clue has a much better chance. And of course, compared to clues yet to be developed, all the clues we already have are old.

However, we should notice that when we apply hypotheses to evidence we do not develop new clues every time, even when these hypotheses are suggested by old clues. When we apply hypotheses to evidence, we develop new clues only *sometimes*. In our example we apply the hypotheses suggested by old clues to most of the words in the message, but in only one of them—the sixth word—do we develop

a new clue. But since new clues are so important in an investigation, even though we do not develop new clues every time, we should keep applying hypotheses to evidence. The more often we do this the more likely we develop new clues.

Do try to gather a large body of evidence at the outset

One way to look for clues is to search for repeating patterns in the evidence. If we have a large body of evidence it is easier to detect these patterns. If the amount of evidence is too small we may not notice any pattern at all. In our example if our cryptogram had been shortened to just three letters, SBR, we would not have known that SB is a pattern.

In an investigation it is to our advantage to have from the outset a large body of evidence. The beginning of an investigation is usually the hardest, but if we have a large body of evidence it should be easier to find clues.

Of course evidence does not just fall into our laps; we have to look for it.³ So this is what we should do: we should look for a large body of evidence, and we should do this early in the investigation. Now we may not succeed, but at least we should try.

³ We consider how we should look for evidence in a later section.

When we have a large body of evidence early in an investigation it should be easier, we have said, to find those initial clues which will propel the investigation forward. In addition to this, if we have a large body of evidence right from the start, when it comes time to develop new clues our task should be easier too. For we develop new clues by applying the hypotheses suggested by old clues to the evidence. The larger the body of evidence to which we can apply hypotheses, the easier it is to develop new clues.

It is an advantage to have a large body of evidence early on in an investigation, but we should be aware it is possible sometimes to have more evidence than we need. In cracking the simple cipher in our example, if we are only interested in how the cipher works once we have enough ciphertext to work out how the twenty-six letters in the alphabet are translated, more will be superfluous. Now this is useful to know because in practice gathering evidence can at times be costly in terms of effort and resources. It is good to have a large body of evidence, but there is no point in having too much if having too much means wasting effort and resources.

How much evidence do we need? When do we have enough? How much is too much? These questions are not always easy to answer. For the amount of evidence we need depends on the structure or structures we want to uncover. The cipher in our example has a

simple structure. For this we need only a small amount of evidence (ciphertext). If the cipher had been more complicated we would have needed a much longer ciphertext. But in most investigations we would not have known in advance how complicated those structures are that we are trying to know. For example detectives may be thinking that they are solving a simple crime and that they have enough evidence. It may turn out that the crime is much more complicated than they think, so that they will have to look for more evidence. Since this kind of error is common in investigations, when gathering evidence it is better to err on the side of caution, better, that is to say, to collect more than is needed instead of not enough. So we often say, in an investigation we should collect a large amount of evidence. In this context the word 'large' is being used in a relative sense. We use our best judgement to determine how much evidence is enough and then we look for a little more to give ourselves a comfortable margin.

We have to keep in mind the following as well. In most investigations finding evidence is difficult. Left to ourselves we have a tendency to stop too soon when looking for it. If we do not constantly remind ourselves that we should gather a large amount, we might just succumb to this tendency and end up with not enough.

Do look for additional evidence if the evidence in your possession is not sufficient

In an investigation we need a large body of evidence in order that we can develop enough clues. It is preferable that we are in possession of this large body of evidence early in the investigation, so that we will have enough clues to set us on our way. Unfortunately this does not happen too often. In many investigations, despite our best effort, we would still find that the evidence we have is not enough. When this occurs we of course will have to look for additional evidence. Now no shame is involved in doing this. In an investigation we are looking for things that we do not yet know. Without knowing these things we cannot make an accurate judgement as to how much evidence we need, and what kind. We can only do our best but we should not be surprised if we discover later that even our best is not good enough.

How do we look for additional evidence? For this we go to the next section.

Do make use of existing information to develop new information even if existing information is uncertain

Before we explain how we look for additional evidence we want first to put forward this reminder: Do make use of the information you already have to help you develop new information, *even if the information you already have is not all that certain.*

Let me first use an example to explain what I mean. Earlier we have used our cryptanalytic example to show that to develop new clues we should apply hypotheses suggested by old clues to the evidence. For example one of the hypotheses suggested by old clues is that S stands for T. Now this hypothesis is far from certain. Yet, despite its uncertainty, we apply it to the evidence. And this is proper. In this case, by applying this uncertain hypothesis to evidence we develop a new clue, TW?, which suggests that H stands for O. Now this is a new bit of information, itself also uncertain. However, while uncertain, it could lead to other new clues, which in turn could lead to other bits of new information. Now making use of uncertain information in this way—that is, in the process of following clues—is not only legitimate, but desirable, for if it goes on and on like this, after a while, when we look back at the earlier bits of information, we will find that these earlier bits are now much more certain than they were originally. In cracking our cipher the hypothesis that S stands for

T occurs early. When it first occurs it is uncertain, but later on when the whole message has been deciphered this same hypothesis is much more certain than it was at the beginning. In the clue-following or theseological process, not only is there no harm in making use of uncertain information in the development of new information, it is incumbent upon us to so make use of this kind of uncertain information.

But I should caution that we can make use of uncertain information this way only in the clue-following or theseological process. In such a process, uncertain information could lead to new clues, which in turn could lead to new information. When more and more information is obtained as a result of following clues, the earlier information becomes more and more certain even though it could have been very uncertain to begin with.

What happens if we are not following clues? Suppose I simply put forward one uncertain hypothesis after another, doing so without following clues; will the first hypothesis become more certain after I have put forward one hundred such hypotheses?

Clearly it will not. Here we have only a series of hypotheses, none of which is dependent on the others for clues. Each one of them being independent of the others, there is no reason why after one hundred of them the first one should become more certain. If we are

not developing new hypotheses from the older ones by following clues, the older hypotheses will not become more certain after the new hypotheses have been advanced.

Why should this be the case? Why is it so important that we should be following clues if we want the old hypotheses to become more certain?

There is a simple reason for this. If the hypothesis answering to a clue is too far from the truth, it will not lead to new clues. Without new clues, no new hypotheses or other kinds of information can be developed. The hypothesis answering to a clue could be uncertain to begin with, but if it leads to more and more new clues, it becomes less and less likely that it is too far from the truth; which is to say, the certainty that it is at least close to the truth increases.

Why should this be? Why can a wrong hypothesis, a hypothesis that is too far from the truth, not lead to new clues?

The reason for this is that clues are the characteristics of structures. When we misread a clue; when we interpret it in the wrong way; we are putting into the structure we are trying to reconstruct an element that in fact is not in that structure. This foreign element cannot then combine with the other characteristics of the structure to produce new clues. When SFH is translated into TW? we have a new clue. But when SFH is translated into KW? we cannot have a new clue.

S stands for T, not for K. When K is made into the first letter of this three-letter word by mistake, it cannot combine with W to produce a new clue. It cannot do this because it should not have been there in the first place. An English word has a structure. Just as we cannot put together at random any number of letters to form a word, we cannot put together at random any number of letters to produce a clue.

Since in the clue-following or theseological process we develop new information by following clues, if we are able to develop new information from old the old information is likely to be correct. If it were not, there would not have been those clues that led to the new information. It does not matter that the old information was originally uncertain. So long as it led to new clues it is likely to be true or close to the truth. The more new clues it leads to, the more certain that this is the case.

It is an important characteristic of the clue-following or theseological process that in this process we can rely on uncertain information to develop new information. Having illustrated this characteristic with an example, let us next see how this characteristic can help us find additional evidence.

In an investigation, when the evidence we have earlier found is not sufficient, we will have to look for additional evidence. But how do we look for additional evidence? Evidence for a crime, for example, is

not just anything we can lay our hands on. We cannot go into a gun shop, pick up a gun and say that it is part of the evidence. For the crime we are investigating all of the guns in the universe, with the exception of perhaps one or two, are not part of the evidence. How are we to look for that one or two guns that are relevant to our investigation?

To answer this question we follow the advice we have been advancing. In the clue-following or theseological process we should try to make the best use of the information already in our possession when trying to develop new information, *even though the information already in our possession may not be all that certain*. Let us suppose we are investigating a murder. At the crime scene we are able to retrieve one bullet, the one which, we assume but are not certain, killed the victim. We examine this bullet for score marks (when a gun is fired, the barrel of the gun, not being perfectly smooth, will leave score marks on the bullet). From these score marks we now know which gun to look for. We will be looking for not just any gun, but the one that leaves the same score marks. Now just as the score marks tell us which gun to look for, other clues at the crime scene could tell us where we are likely to find this gun. Now all this information could be very uncertain, but if we succeed in actually finding the gun, and then the owner, and then the murderer (the owner may not be the

murderer), together with other details of the crime because of additional *clues* that have been developed, the uncertain information we started with will become much less uncertain.

In an investigation we can make use of uncertain information to develop new information. That the information we rely upon is uncertain does not necessarily make it useless. All the information we make use of in an investigation will always have a certain amount of uncertainty associated with it. If we only allow ourselves to use information that is absolutely certain we will never find out anything. Indeed, if we hold to such an unreasonable standard we will be reduced to inaction.

It is important to point out however that although we would allow ourselves to use uncertain information in an investigation, this does not mean we can throw all caution to the wind. We would not use just any information. *We like our information to be as certain as possible, even though we do not require it to be absolutely certain.* For example, when we advance hypotheses we would want these hypotheses to be based on clues; we do not want them to be the result of wild guesses. Hypotheses based on clues are uncertain, but they are less uncertain than wild guesses.

Do pay attention to whether old clues lead to more and more new clues

This is an important task in an investigation. In an investigation we are trying to reconstruct structures that are hidden from us. Since the structures we are trying to reconstruct are hidden, we cannot determine whether we have succeeded by putting what we have constructed side by side with these hidden structures and comparing the two (sets). But while we cannot carry out this direct comparison, we can pay attention to whether old clues lead to more and more new clues. If we do see old clues leading to new clues, and these new clues leading to more new clues, and so on and on, we can derive from such an observation a progressively more accurate idea as to whether we have succeeded in our reconstruction. For (as we have explained in the last section) if our reconstruction had gone seriously amiss; if for example we had gravely misinterpreted some of the clues; we could not have developed more and more new clues.

On rare occasions a clue that has been incorrectly interpreted will lead to a false clue by accident. But false clues peter out; they do not lead to more and more new clues. One accident is rare. Two accidents in a row is even rarer. Since this is the case, when we manage to develop more and more new clues from the old ones we

will have good reason to believe that it is likely that the old clues have been correctly interpreted.

In an investigation we should try to develop new clues from old. If we succeed, not only will we find out more, but we will have reason to believe that what we have *earlier* found out is likely to be correct. We follow clues to reconstruct structures that are hidden from us. But that these structures are hidden does not mean we cannot tell whether our reconstructions are correct, that is, whether they correspond to the hidden structures. Direct comparison is not the only way to find out whether two structures correspond. We can also find out whether they correspond by seeing whether they leave behind the same clues.

Do recognise that there is a method

People not accustomed to following clues frequently have doubts that clues could ever lead us to the truth. We follow clues when the things we want to know are hidden. If these things were not hidden; if they were there for everyone to see; there would have been no need to follow clues. But if the things we want to know are hidden, whatever the clues tell us, how do we know what they tell us is true? If we can compare what the clues tell us to what is the case we can determine whether what they tell us is true. But when we are following clues we

cannot make this comparison. We cannot because what is the case is hidden.

Things hidden, it seems, should be unknowable because whatever we say of these things we can never compare what we say to the things themselves (since they are hidden). Yet in following clues this is exactly what we are trying to do. We are trying to know the hidden, which is unknowable.

Is it true that the hidden cannot be known? Of course the hidden cannot be known directly. If it could it would be a contradiction in terms, for then the hidden would be both hidden and not hidden at the same time. Since the hidden cannot be known directly, if it is to be known at all it will have to be known indirectly. Can the hidden be known indirectly?

Of course it can. This is what we are trying to do when following clues. When we are following clues we are trying to know *indirectly* things that are hidden from us. And it is possible, as experience has shown us over and over again. It is possible to know indirectly *if there is a method*. When we have a method we do not need to *compare*. Let me explain.

Suppose I have two huge piles of marbles. I know how many marbles there are in each pile. Now I put the two piles together to

form a even larger pile. How many marbles are there in this larger pile?

The *direct* way to find out is to count all the marbles that are now in the larger pile. But no one will want to do this because it will take too long, and be extremely boring to boot. And there is no need to waste our time in this fashion. As we all know there is a simpler and faster way. Since we know the number of marbles in each pile before we put them together, we can simply add up these two numbers following the *method* (algorithm) we commonly employ. This method is *indirect*: it will give us a number without counting. But although indirect, the number it gives us will be accurate provided we follow this method correctly. It is important that we follow this method correctly. If we do not we of course will not get the right answer (except by chance). But if we follow the method correctly we will always arrive at the right answer. We know the answer will be right *without checking it by counting the marbles*. When we have a method there is no need to *compare*.

Now the hidden cannot be known directly. But this does not mean it cannot be known at all. The hidden can be known *indirectly* if we have a method. And we have a method. We have been using this method whether we are aware of it or not. I call this method the theseological method, since we use it during the theseological process.

When we are following clues we are following the theseological method, or attempting to follow it. The theseological method tells us that if we want to uncover knowledge of things hidden we have to follow clues and develop new clues from old. Now this method is not easy to follow, but if we do manage to follow it we will arrive at right results. We know these results will be right without 'counting the marbles', that is, without having direct access to the things we want to know. Why should the theseological method be able to do this? We in effect have already answered this question in the last two sections. The method tells us to follow clues and develop new clues from old. Now when we succeed in following this method; that is, when we succeed in following clues and developing more and more new clues; the earlier clues are likely to have been correctly interpreted. The more generations of new clues they lead to, the more likely this is the case. A clue that has been seriously misinterpreted cannot lead to new clues.

When we follow the theseological method our attention is on clues. The more clues we are able to develop from old clues the more certain we can be of our earlier results (that is, our interpretation of the old clues). To find out whether our earlier results are right we do not have to compare these results to the things that are hidden. That we do not have to make such a comparison may appear strange at

first, but should not any more now that we know when there is a method, there is no need to make this kind of comparison.

People not familiar with the theseological process frequently have doubts whether clues can lead us to the truth. But they are not the only ones. Even experienced theseologists will sometimes have the same doubts. The simple fact is, that we can get to know about things hidden is counterintuitive. It is therefore useful to keep in mind that although things hidden cannot be known directly, they can be known indirectly when there is a method. And there is a method: the theseological method, which tells us to follow clues and develop new clues from old.

Do look for patterns in the evidence

Looking for patterns in the evidence is often the first thing we try when looking for clues. Clues are the characteristics of structures. When we have found a pattern in the evidence it is likely that we have found one of the characteristics. Thus it is that in practice when we have found such a pattern, we will go on to ask what this pattern means. For clues are not just the characteristics of structures, but the characteristics of structures, disguised. In our example the pattern SB is a pattern; it suggests to us that it might stand for TH.

But we should be aware, patterns sometimes occur by accident. When we have found a pattern we should not be dogmatic that it must be significant. As we so often say, the reason for the pattern could be 'totally innocent'. However, if the pattern repeats itself, not just once but many times, the chance that it occurs by accident diminishes.⁴

Not all patterns are easy to find. Sometimes discovering a pattern requires a good bit of ingenuity. Take for example the motion of the planets. For a long time people in different parts of the world have noticed that the motion of the planets across the night sky is erratic: there is no noticeable pattern in these motions. But Kepler discovers that this is so only if these motions are viewed from the earth. If instead of using the earth as the reference point we use the sun, we will discover—as Kepler did after a large amount of calculations—that the planets all move in elliptical orbits. Kepler's discovery—that the planets all move in elliptical orbits *around the sun*—is one of the clues Newton uses in his discovery of the law of universal gravitation and his three laws of motion.

A pattern could be a clue because there could be a reason for the pattern. If a pattern *repeats* itself, it is likely that it is not accidental.

⁴ Even then we should not be dogmatic that the pattern must be significant for our investigation. There could be a good reason for the pattern, but it could be irrelevant to our investigation.

Now in detecting patterns as clues, we notice the pattern first and then have to figure out the reason for it. That is, we have to search among the structures we know for one that will produce the same pattern. Now in order that we can carry out this task it matters whether the pattern repeats itself indefinitely. Not all repeating patterns repeat themselves indefinitely; some will repeat themselves for a time and then, for a good reason, stop. When we are looking for a reason for the pattern, therefore, we cannot take it for granted that the pattern will repeat itself indefinitely. If we did, we could come up with the wrong explanation, or we may not be able to come up with any explanation at all.

If we notice a pattern repeating itself for a time and then stop, we know then that it is not repeating itself indefinitely. But suppose it has been repeating itself for a long time and has not stopped so far. Can we, just on this basis, take it for granted that it will never stop?

Obviously, we cannot, since we don't know yet the reason for the repetition. Sometimes a pattern could repeat itself for a long time before stopping. The earth has been circling around the sun for many years but this is not to say it will do so for ever and ever. A time will come when the earth will be no more because it has spiralled to its death by falling into the sun.

In an investigation it is important that we do not take for granted that a pattern that has been repeating itself will repeat itself indefinitely. If we did we could seriously hamper our investigation.

What can we do when we see a repeating pattern and we do not know whether it will stop?

This is what we could do. We could *assume* it will repeat itself indefinitely, but if this leads nowhere, we will have to try the other possibility, that is, that it will not repeat itself indefinitely.

Do look for unique occurrences

Looking for unique occurrences is another popular way of looking for clues. For, if the occurrence is unique there must be a reason why it is unique. If we can figure out this reason we will have found out something important about those structures we are trying to uncover. In one of the cases Sherlock Holmes investigated, a person staying at a hotel had *one* of his boots stolen even though the other was equally available. Now this is odd. Most of the time boots are worn in pairs. Why would anyone steal just one boot? Concerning this peculiar incidence Sherlock Holmes has this to say:

The more outré and grotesque an incident is the more carefully it deserves to be examined, and the very point which appears to complicate a case is, when duly considered and scientifically handled, the one which is most likely to elucidate it.

— *The Hound of the Baskervilles*

A unique occurrence could be a clue. It could even be a highly significant clue. But how do we determine whether an occurrence is unique? In the case of stealing boots, stealing one is unique, but this is only because we all know that boots are worn in pairs. Suppose we do not know that boots are worn in pairs; how do we determine uniqueness then?

Clearly, if we do not know, we would not be able to know that stealing one boot is unique. To determine uniqueness we have to know what the norm is. And this has important implications for investigators. Let me give a simple example.

A bank has been broken into overnight and a large sum of money has been lost. As part of the investigation all the employees are interviewed, including of course the manager. When asked what time she left work the previous day, the manager answers, '5:15.' The detective in charge of this case takes down this answer and after

comparing it with those given by the other employees thinks no more of it. Why? Because she finds nothing unusual about this answer. There is nothing unusual about the manager being the last person to leave. By 5:15 all the other employees had gone home.

In fact 5:15 is highly significant, and the detective would have known if she had asked a few more questions. For this particular manager has been working at this bank for a long time and without exception has always left punctually at 5:12. Without exception, that is to say, until the previous day!

Unique occurrences can be important clues, but uniqueness is not always apparent. Sometimes we fail to notice uniqueness for lack of sufficient information. This bespeaks again for having a large body of evidence. This is not to say that when we have a large body of evidence we are bound to notice uniqueness: the relevant information could be absent from this large body; but in an investigation we can only, and should try to, increase our chances.

Do try to increase your chances as much as possible

In following clues to uncover hidden structures by reconstructing them, success is not guaranteed in advance.⁵ However, there are

⁵ There can be many reason why success is sometimes not possible. For example, most of the clues could have been destroyed.

things we can do which can increase our chances. We have had occasion to mention some of these already. For example, why should we follow clues instead of making wild guesses? The reason (as everybody knows) is that following clues increases our chances. The hypotheses we advance in response to clues may not be right, but at least they have a better chance than mere guesses.

Also we have asked, why we should apply hypotheses to evidence? The answer we have given: it increases our likelihood of finding new clues. The more new clues we have the better our chance of finding those things we are looking for.

As we continue with this essay we will have more occasions to point out how we can increase our chances. At this time, let us simply draw attention to this useful rule. If we do not follow this rule; if we do not try to increase our chances even when it is possible; we could miss finding those things we want to know. If we follow this rule; if we do try to increase our chances; we could still fail, but then we can at least say we have done our best.

Do test when possible the hypotheses you have advanced

In the theseological process we advance hypotheses in response to clues. We find the clue SB. We ask, what does it mean? We answer by putting forward the two hypotheses: S stands for T, and B, for H.

Now when we put forward hypotheses in answer to clues it is possible that the hypotheses we put forward are wrong. A clue could mean more than one thing. S might not stand for T, but for something else. How do we know that the hypothesis that S stands for T is not wrong? One way to find out is to test the hypothesis. Suppose S occurs also in another word, ten letters in length. Suppose that all the other letters in this ten-letter word have been deciphered, so that S is the only remaining one. Now if this is the case and we substitute T for S, and find that T fits in, we can conclude that our hypothesis is likely to be right, that is, that S does stand for T. But if T does not fit in; if after translating S to T we do not produce a proper English word (we assume here that we are deciphering an English message); we will have reason to suspect that our hypothesis is wrong. Now the reasoning here seems to be simple enough; nevertheless, we would do well to make it clearer to ourselves.

Let us carry out the analysis in the following way and call it our Attempt A.

Attempt A

Let us first consider the case that T does *not* fit in.

If T does *not* fit in; if after we have translated S to T we do not obtain a proper English word; it seems clear that the hypothesis

cannot be right. The hypothesis, if true, leads to the conclusion that we should obtain an English word. The conclusion turns out to be false: we do not obtain an English word. Therefore the hypothesis must be false.

Now consider the other case: T fits in.

If T fits in; if after we have translated S to T we do obtain a proper English word; can we then conclude that the hypothesis is true? Here we meet with a major difficulty. The hypothesis leads to the conclusion that we should obtain a proper English word. This conclusion turns out to be true: we do obtain a proper English word. But we cannot conclude from this that the hypothesis must be true! It is well known in logic that false premises can sometimes lead to true conclusions. Take a simple case:

Premise 1: I am a millionaire.

Premise 2: People who have a million dollars will of course also have one dollar.

Conclusion: I have one dollar.

Now Premise 1 is false. But the conclusion is true! I do have one dollar. From the fact that I have one dollar I cannot conclude that I am a millionaire. Similarly from the fact that T fits in we cannot conclude that the hypothesis that S stands for T is true! But when we test a

hypothesis, this is what we do. When we test a hypothesis we derive from the hypothesis a prediction. If the prediction turns out to be right we conclude that the hypothesis is likely to be right.

Something is wrong with Attempt A. Attempt A does not help us understand actual practice.

What could be wrong with Attempt A?

Of course, actual practice could be wrong. Perhaps we should abandon actual practice. Perhaps we should say, even though the result of a test is as predicted, we are not a whit wiser as to the truth or falsity of the hypothesis. The hypothesis could be true or it could be false. The test has not told us whether it is one or the other.

However, before we reject actual practice, let us make a second attempt to understand the reasoning behind it. For Attempt A is a little strange. We make Attempt A to understand something we frequently do when following clues. In following clues, it is common that we test hypotheses when possible. But notice what we have just said. We have just said that we test hypotheses when possible *when following clues*. But Attempt A says nothing about clues!⁶

We now make a second attempt to understand the reasoning behind hypotheses testing when following clues. Let us call this **Attempt B.**

⁶ Attempt A focuses attention only on the relation between premise and conclusion

Attempt B

This time let us first consider the case when T fits in; that is, when the substitution of S by T produces a proper English word as is predicted.

Now we should remember that the hypothesis that S stands for T is suggested by a clue (SB). We test this hypothesis by applying it to a ten-letter word in which nine letters are already known (with S being the tenth letter). But what is this ten-letter word if not another clue? When we have a ten-letter word in which S is the only undeciphered letter, the other nine letters should give us a good idea as to what S stands for. So whereas previously we had only one clue (SB), we now have two (SB and the ten-letter word). They both point to S standing for T. This being the case, after the test we obviously should be more certain about the hypothesis we have advanced (that S stands for T). If we do not have a single clue, S for all we know could stand for anything. When we have one clue, the number of possibilities for S is reduced. When we have a second clue, especially a specific one (like the ten-letter word in our example), the number of possibilities is even smaller. For good reasons, therefore, when we test a hypothesis during the theseological process and obtain a positive result, we will

look upon the hypothesis much more favourably after the test than before.

Let us now consider the case when T does not fit in.

In this case the test produces a negative result: after translating S to T in the ten-letter word we do not obtain a proper English word. From this we infer that the hypothesis (that S stands for T) is wrong. And this is proper. Even when we had only one clue SB, we knew that S might not stand for T. Now we have a second clue and it tells us that S does not stand for T. It is always better to have two clues than just one. The more clues we have the clearer the message they send us.

We have offered two analyses—Attempt A and Attempt B—for the reasoning behind the practice of testing hypotheses in the theseological process. Attempt A can now be seen to be misdirected: it forgets that we are examining this practice as it occurs in the theseological process. Attempt A therefore should be rejected. Attempt B is better than Attempt A. Attempt B clearly acknowledges that we are examining this practice in the context of following clues.

In our example, to test the hypothesis that S stands for T we use a ten-letter word in which nine is already known. The remaining letter is the only uncertain letter, being the one to which the hypothesis is to

be applied. Now this is something we always try to do when testing a hypothesis. We always try to find a situation in which the only uncertainty is the hypothesis under test. When this is the case the result will tell us how good our hypothesis is. If there are other uncertainties besides the hypothesis under test, it may be difficult or even impossible to interpret the result. A test for a hypothesis is a second clue to the same hypothesis if the hypothesis is right. But if this second clue is too vague it will not help us pin down the hypothesis.

How important are tests? Are tests the only way by which we can find out how good a hypothesis is?

People often say yes to the last question. It is a common view that tests are the only means of finding out how good a hypothesis is. But they are not; there is at least one other way by which we can find out how good a hypothesis is. In fact we have already explain this other way. We have already said that a clue severely misinterpreted will not lead to new clues. Now hypotheses are proposed in answer to clues. If a hypothesis is seriously defective it will not lead to new clues. If a hypothesis leads to new clues which lead to even more new clues, such a hypothesis is likely to be true or close to being true. The more (generations of) new clues it leads to the more likely that this is the case. Tests are not the only means for finding out whether a

hypothesis is right (true or close to being true). Instead of testing a hypothesis we could wait and see whether it leads to more and more new clues.

In fact, this method of evaluating hypotheses, which tells us to pay attention to the development of new clues from old, is fundamental to the theseological process, and as such more important than tests. For tests depends on our knowing a great deal already, as we have seen. If we know already nine letters in a ten-letter world we can use this ten-letter word to test a hypothesis for the remaining letter. But how do we get to the stage where we know nine letters? When we start out we may not know a single letter at all. When little is known, how do we know that the hypotheses we advance in response to clues are correct? We can only know by seeing whether they lead to more and more new clues. If they do they are likely to be right even if we have not had a chance to test them. Tests are not the only way by which to find out how good a hypothesis is. In an investigation we have to follow clues and develop new clues from old. It is by following clues and developing new clues from old that we find out those things we want to know. This is to say, to find out things hidden from us we employ a method: we follow clues and develop new clues from old. I have come to call this the theseological method. The greater our success in following this method; that is, the greater our success in

following clues and developing new clues from old; the greater the chance that our hypotheses are right, especially the older ones.

To save time or resources, or because they are just impatient, experienced investigators would sometimes omit testing a hypothesis even when such a test is available. Instead of testing the hypothesis they would just 'plough on'. Now if the hypothesis leads to more and more new clues, they know it is likely to be right.

But these investigators will also warn you that avoiding tests is a risky business. In an investigation it is a good idea to test a hypothesis first, before moving on. For, most investigations are complex. In a complex investigation any mistake not caught in time can be difficult to unearth later on. When we avoid testing a hypothesis, no harm is done if our hypothesis happens to be right. But if it is wrong we could be creating trouble for ourselves.

In science a common way to test a hypothesis is to carry out experiments. By this we mean experiments in the laboratory or in the field. But sometimes such experiments are not actually possible for technological, financial, or other reasons. What can scientists do in such instances if they still want to find out how good the hypothesis is before they move on? This is what they sometimes do. They carry out the experiments 'in their head', relying on intuition or imagination to tell them what the results are likely to be. They call these thought

experiments. Needless to say, thought experiments are not as good as real experiments: intuition and imagination are not always reliable. Still, they are better than no experiment at all. And if the hypotheses they are supposed to test should later lead on to new clues, that they are thought experiments only, and not real, makes no difference.

Do double-check when possible

When things we do are important, to ensure that we have done them, or have done them properly, we double-check. In locking up for the night a bank manager has to make sure the alarms are on, so she checks and double-checks. Now in investigations we can correct mistakes, but correcting mistakes long after they have been made is often difficult. It is better that we make sure, or as sure as we can, that what we have done is right before we move on. So we check and double-check when we can. Testing hypotheses, we have seen, is a kind of double-check. A hypothesis has been suggested by a clue (SB). But clues can be vague. To make surer that the hypothesis suggested is right we look for a second clue (the ten-letter word), which as we have seen is the test.

But in investigations double-checks are not always possible. At the beginning of an investigation we often will not have known enough to be able to test a hypothesis (as we have already explained). In

these cases we simply plough on and hope for the best. If our hopes are dashed; if it should turn out later on that progress is impossible; we backtrack and try a different hypothesis.

In an investigation we should double-check when we can. The theseological process is not simple; success is not guaranteed; but the fewer mistakes we make the better our chances.

Do try to make explicit your assumptions

In an investigation we always have to make assumptions; without them the investigation cannot start. Look back to our cryptanalytic example. If we did not assume that the series of letters is hiding a message we would not even start looking for clues. If we are already looking for clues it must mean that we have made this assumption. Now when we are looking for clues we need to have some idea what we should focus our attention on. Should we for example pay attention to the size of the letters? Or perhaps we should forget about the letters altogether and focus only on the spaces between them. Now we might not actually have asked these questions before we start, but once we have started looking for clues it will mean that we have formed some assumptions as to how the cipher works and even what the language is in which the plaintext is written.

Now all this is very well if after making these assumptions we are able to meet with success after success, so that eventually the whole message is deciphered. At that point we know that the cryptogram is indeed hiding a message and that the other assumptions we have made are correct. But what if this does not happen? What if, after thousands of attempts, we still have not been able to decipher the message?

Those who have had any experience in following clues will know that in such an eventuality they will have to start thinking about changing some or all of their assumptions. The lack of success could be due to the presence of wrong assumptions. But before we can change assumptions we have to know what assumptions we have made. And this is one of the reasons why in an investigation we should try to make explicit assumptions. Assumptions in an investigation are often made unconsciously, without ourselves being aware that we have made them. Before any of these *hidden* assumptions can be corrected we first have to take them out from their hiding places, that is, make them explicit.

Making explicit assumptions is not necessarily easy. Assumptions are often so well hidden that we do not even know they are there. In practice we often unearth them one or two at a time, with long intervals in between.

Making explicit assumptions is especially important when progress in an investigation is difficult or impossible, but this is not the only time we should engage in this activity. Even when an investigation is going well (relatively speaking) it is a good idea to try to render explicit the assumptions we have made. Assumptions affect what we do in an investigation—for example, what kind of clues or evidence we should look for. If we are not clear about them, and we will not if they remain hidden, the lack of clarity could lead to unnecessary steps or to hesitation.

An investigation can reach a successful conclusion without our making explicit all the assumptions we have made. However, if we do take time out to clarify some of these assumptions we could as a result hasten success. For example, in some instances, because we know what we should do, instead of stumbling along, we could automate some of the steps.

Do try the simpler assumptions first

In forming assumptions sometimes we can form more than one set. In solving the cipher in our example, instead of assuming the plaintext to be in English we could assume it to be in Chinese. And if we do we will have to assume that the cipher works in a rather

complicated way. Now these two sets of assumptions cannot both be correct. How do we know which is correct?

The only way to find out is to see which set leads to results. If by following the 'English' set of assumptions we arrive at an English plaintext, then the English set is right. If the plaintext is in English it is not possible for a Chinese plaintext to turn up, so long as we are following clues. Clues are the characteristics of structures. If the plaintext is not in Chinese, it will not have Chinese characteristics.

Does this mean then that instead of trying the English set of assumptions first it is just as reasonable to try the Chinese set? We do not know which set is correct before we have obtained results. Before we have obtained results there is no reason why we should think that one set is more likely to be correct than the other.

But in practice, with an example like the one we have been using, most people will try the English set of assumptions first. Why?

The reason is not that they know from the outset that the English set of assumptions is more likely to be correct. They will try it out first because the English set is simpler. That is, it *appears* to them that it is easier to render this set of assumptions compatible with available evidence: the cryptogram. Of course, appearance could be wrong. They could find out, once they have started looking for clues and interpreting them, that it is not easy at all trying to decipher this

message on this set of assumptions. But if this should happen they could change assumptions. When we have a number of alternatives, it is better to try the easiest one first. If it is wrong, we will know soon enough, and then we can try the others.

Do take time out to build up your background knowledge

Much of what we do in an investigation is based on intuition. This is especially true where the formation of assumptions is concerned. Now if our background knowledge is vast the chance is greater that intuition will be able to pick out from this body of knowledge those assumptions that would help us in the investigation. Use as example our SBR-cryptogram again. Those who know English will have an easier time forming assumptions after examining this cryptogram. Those who do not may never arrive at a useable set at all.

We need background knowledge when carrying out investigations. But we do not always know in advance what background knowledge we will need. A person may go through life without ever having to solve an English cryptogram. Now while this is the case it is still a good idea to have as large a body of background knowledge as we can. The larger the body the better the chance that the knowledge we need for a particular investigation will be available from this body.

How do we build up background knowledge? There is so much to know. Should we for example make a point of learning all the languages there are?

It is impossible for an individual to know everything. Even with languages the number one individual can become familiar with can only be a small fraction of all the languages there are.

What can we do then? Should we, for example, know a bit of everything? A few languages, a little chemistry, a smacking of psychology, and so on?

To most people the prospect of learning a bit about everything cannot be very attractive. Moreover, it may not be helpful either: the bits we need for an investigation could well be among those we have left out.

What other course of action can we take? How can we increase our background knowledge so that we will be better investigators?

One suggestion that has been made is that besides learning the basics investigators should be encouraged to cultivate each their own interests where the acquisition of background knowledge is concerned. In every kind of investigation there are usually certain basic things we need to know, things that will help us with the run-of-the-mill investigations in that particular area. Now if we do not want to be run-of-the-mill investigators; if we want to do better; we will have to

expand our background knowledge beyond these basics. How do we do it? We leave it to the individual. We leave room for the individual to develop each her own interests, and then we encourage these individuals to pursue these interests. Things we are interested in, we never know enough. Learning about them, finding out about them, doing research in connection with them—these we will do on our own and derive much satisfaction in doing so.

Of course even then we cannot guarantee that we will always have in our possession the background knowledge that we need for some of the investigations that will come up. There can be no guarantee of this kind. In preparing ourselves for future investigations we can only increase our chances.//

Do compare new cases with old cases even when they do not exactly correspond

Sherlock Holmes studies the annals of crime. Scientists take interest in investigations other scientists carry out. In general investigators—at least, the better ones—do not just carry out their own investigations, but pay attention also to what fellow investigators have been doing. This is one of their ways of building up background knowledge, background knowledge that will help them in their future

investigations.⁷ For this is what investigators often do: they look to past cases for inspiration—that is, for suggestions as to what clues to look for and how to interpret them. For this purpose they compare the case they are working on with past cases. They do not expect their current case will be exactly like some past cases; it is not exact correspondence they are looking for, just similarity. Provided there are similarities they can easily bridge the differences by themselves (for more on this, see the section, **Do look for analogies.**)

Do take an interest in the creation of structures and in studying their characteristics

To illustrate the importance of this point let us first use catching criminals as example. Some criminals are creative; they can invent new ways of committing crimes. If you are a detective and you only study old cases you are not likely to catch these creative criminals. If you want to be a good detective what will you do? Obviously, the thing to do is to anticipate the ways in which these creative criminals will be creative. This is to say, if you want to be a good detective you will have to think like a criminal—and not just any criminal but a creative

⁷ This is not to say they will not be interested in what their fellow investigators have been doing if this knowledge of what they have been doing will not help them in their future investigations.

one. You will have to think of new ways committing crime, and study these new ways to discover what clues they will leave behind. Now if one of the ways you have been thinking of happens to be the way one of these inventive criminals is following, that inventive criminal will be caught (because of the clues she leaves behind, which you are now able to detect). This is the reasoning behind the well-known saying, set a thief to catch a thief. Ordinary thieves can be caught the ordinary way; their *modus operandi* are well known. To catch an exceptional thief you have to think the way she does. If you manage to do so, she is caught. If you fail, she gets away. How to increase our chances of catching her? We imagine her planning not just one crime, but many, and hope that one of these imaginings corresponds to what she has actually done.

To solve crimes detectives practise by taking on the more common types first, the types that have been solved in the past. Scientists want to find out the structure of this universe, but they have no other universes to practise on. They have to start 'from scratch'; which is to say, they have to construct (theoretical) structures on their own and see if at least one of them leaves behind the same kind of clues that the actual universe leaves behind. For example, in what kind of universe will the planets move in elliptically orbits around the sun? Will they do so in a universe in which the force of gravitation obeys an

inverse cube law? Newton found out they will not, but they will if the force of gravitation follows an inverse square law. Did Newton find this out by studying old universes? Does any one before Newton know that planets will move in elliptical orbits if the force of gravitation obeys an inverse square law? No; there is no such past cases for Newton to rely upon. He finds out through his own initiative, by constructing theoretical models (more than one) and studying how they behave. From his models he can tell that in a universe in which the gravitational force obeys an inverse cube law the planets will not move in elliptical orbits, but in one in which it obeys an inverse square law, they will.

To figure out the meanings of clues we have to know what structures will leave behind these clues. Where does this knowledge come from? In some instances it comes from past experience: past cases help us deal with current cases; in others it derives from structures we ourselves have created and studied. Newton had to create on his own initiative new structures just as good detectives have to think of new ways of committing crimes.

Do observe the Small Steps Principle

Clues are the characteristics of structures. It is for this reason that by following them we can reconstruct those structures we want to

know. We can compare this to looking for a person, we have said (**Do look for as many clues as you can**). The more details/clues we have about this person the easier it is to find her and not confuse her with somebody else. But this is to say, each clue we have about this person tells us only a little. However, if we have enough clues we will be able to find this person even though we have never met her before. We therefore should not despise the amount of information each clue provides. Although this amount is small when compared to the total amount we need, if we have a large number of clues these small amounts will add up.

But not only should we *not* despise the small amount of information each clue provides, we should accustom ourselves to expecting only a small amount from each one of them. The amount of information any one of the clues provides is always small compared to what we eventually will know (if we are lucky). It is true that a clue can sometimes lead to a succession of other clues, and do so quickly (we call such a clue significant); still, even such a clue will only tell us a small fraction of what we eventually will find out.

In an investigation we find things out gradually, step by step. Since each clue can only tell us a little, the steps we take in an investigation are never huge. *In an investigation we can only take small steps, never big ones.* I call this the Small Steps Principle.

Most people familiar with following clues are aware of the Small Steps Principle even though they may not call it such. When they hear a person spinning a long tale as a result of just one single clue they will ask, 'How can you know so much from just one clue?'

We do not expect a single clue to tell us everything we want to know. We do not expect it to tell us even a large portion. We need many clues in an investigation. Which is to say, we need many steps, each small.

In an investigation we have to observe the Small Steps Principle. If we do not; if we take huge steps; we will not find the things we are looking for.

But it is not easy to follow the Small Steps Principle. Following clues is hard work and takes time. Often—or shall we say from time to time—even experienced investigators will become *impatient*, and after a certain point allow their imagination to *completely* take over.⁸ They have found a few clues, they have been patient with them so far, but suddenly one or two of these catches their fancy. From them they 'derive' huge amounts of information, enough they claim to enable them to declare victory. But this victory is a chimera. They have contravened the Small Steps Principle. The information they derive is

⁸ There is a place in the theseological process for imagination, but in this process imagination cannot be allowed to run amok.

not justified. And we can tell it is not justified simply by looking at the quantity.

Do be patient

Investigations are difficult. It is difficult to find clues, difficult to interpret them, and on top of this there is the Small Steps Principle: there are so many steps to take and they all, of necessity, have to be small. In an investigation if we are not patient we will never find the things we are looking for.⁹

Do permit yourself to make use of approximations

Most of us are brought up to think that in the search for knowledge we must work in an exact and precise manner. We must be careful with every step we take, and not allow ourselves to make any mistakes. Things have to be just so, not a whit more, not a whit less. Unless this is the case, we are taught, truth will escape our grasp.

But this is exactly *not* what an investigation is. Real investigations are messy and untidy. In investigations mistakes are common and exactitude and precision are seldom as high as we would

⁹ Even when we are patient we still may not find the things we are looking for, but nothing ventured nothing gained.

like them to be. Yet, despite the chaos, despite the inexactitude and imprecision, truth gradually emerges.

Why is this the case? How can order emerge from chaos? How can truth come from untruth?

We will have more to say in answer to this question later on. For the moment it is useful to notice that in an investigation we are permitted to employ approximations. Indeed, the clever use of approximations can sometimes make possible discoveries that otherwise may not be possible.

Approximations, we have to remember, are not the truth. They are departures from the truth (if only by a small amount). Yet sometimes they can lead to the truth. Why?

To answer this question it is useful to bring to mind the Small Steps Principle: in investigations we have to take small steps, never huge ones. We have the Small Steps Principle because in an investigation we have to follow clues. Each clue tells us only a little. When we have figured out a clue we have only taken one small step. Now an approximation is a small step away from the truth. When we inject an approximation into an investigation; when for example we answer a clue not by the truth but by an approximation; we have divided up a small step into even smaller steps (we have taken a small step forward within the small step, leaving other small steps within this

small step to be taken in the future). But there is no harm in this. In an investigation there is prohibition against huge steps, but no prohibition against dividing a step that is already small into many smaller steps. Indeed, if by doing so we can make it easier to answer some of the other clues, we are duty bound to do so. In an investigation the more we have found out, the easier the rest of the investigation tends to become. Put another way, the more small steps we take the easier the rest of the small steps. In an investigation the important thing is that we should be able to 'push forward'. It does not matter if we do so by taking fifty small steps or by expanding this fifty into one hundred or even a thousand smaller. So long as we are able to push forward, the rest of the investigation should become easier and easier (all things being equal).

In practice investigations are frequently processes of gradual approximation. If we want to be good investigators, instead of succumbing to a phobia against approximations on the ground that they depart from the truth, we should cultivate a habit of using them to our advantage. One of the clues that enabled Newton to discover the theory of universal gravitation was based on an approximation. He and other scientists at the time were asking, why do the planets move around the sun in elliptical orbits? Now these orbits are not exactly elliptical; they cannot be because there is not only gravitational

attraction between each planet and the sun but also gravitational attraction between the planets themselves. The orbits are only *approximately* elliptical. But this approximation has not hindered Newton's investigation. On the contrary, it has made Newton's investigation easier. Imagine what would have happened if instead of asking why the orbits are elliptical, Newton had spent his time trying to determine the exact orbits of the planets.¹⁰

Do look for analogies

In an investigation, to reconstruct those structures we are interested in we have to look for clues and find out what they mean. Now clues are the characteristics of structures. In trying to reconstruct a structure, if the structure we are trying to reconstruct resemble some structure we already know, it should be easier to find clues and interpret them. For since the two structures resemble each other they should have similar characteristics; knowing the characteristics of one should therefore help us find the characteristics of the other. Put another way, since we know what clues the structure we know will leave behind it should be easier to find the clues for the structure we want to know. Here we see the reason why in investigations we often

¹⁰ Will we ever know the exact orbit of the planets?

look for analogies. Does this crime I am investigating remind me of some crime I know? Is this part of nature similar to some other part?

It is useful to notice that when we rely on analogies during an investigation the analogies do *not* have to be exact. In deciphering dead languages a common trick is to compare the dead language to some known language. Now no two languages are exactly alike, but this does not matter. If the two languages we are comparing are sufficiently close we should be able to decipher large parts of the dead language based on the analogy. Once we have done this; once we know enough of the dead language; we can go native, that is, we depend on context to help us decipher the rest (just as we depend on context to learn new words even with our mothertongue).

It is important to notice that in an investigation when we make use of an analogy to help us detect clues and decipher them we are *not* engaging in an analogical argument. We are *not* saying that because two things are similar in some respects they are similar in *all* respects. Analogical arguments are invalid because most of the time two things similar in some respects are *not* similar in all respects. Logicians frequently warn us against analogical arguments and they are right in doing so. But we should be clear that the warning is against analogical arguments, not against the use of analogies in investigations. Analogies can be used in different ways, for different

purposes. In investigations we use analogies to help us find clues and interpret them, not for drawing the conclusion that two things are similar in all respects because they are similar in some respects. All analogies will break down after a certain point, but this does not mean they cannot be useful in an investigation. It is true that we cannot always predict where or when an analogy will break down and as a result of this ignorance make mistakes, but in an investigation we can correct mistakes.

In an investigation we should not expect an analogy to hold indefinitely even when it has been holding up well so far. Rather, we should expect it to break down at some point. Our hope is, by the time we reach that point we will have found out enough about those structures we want to know that we do not have to depend on the analogy any more.

In an investigation, when an analogy breaks down after it has served its purpose exciting discoveries sometimes occur. For why should the analogy break down? Why if it is not because we have entered new territory? If we were still in old, well-traversed terrain the analogy should still be able to guide us. But it could do so no longer; it has broken down; we are now on our own: we have to depend on context for clues. And context could sometimes reveal to us surprising things, things we might never have dreamt of before. This occurred in

the discovery of the positron, which came about after scientists had discovered the limitations of the miniature solar system model of the atom. The atom is not exactly like a miniature solar system, but by the time scientists discovered this they had enough clues to fashion new theories in their attempt to understand the atom. These new theories led them to the conclusion that a particle with the same mass as the electron but with a positive charge should exist. No one expected such a particle but, lo and behold, it is there—as they discovered by examining *old* photographic plates. These plates were made to record the tracks left by electrons, but on a few of them they discovered tracks exactly like those left by electrons but curved in the opposite direction, indicating a particle similar to the electron but with a positive charge. Positrons were unexpected before they were predicted by theory. They were so unexpected that scientists did not notice them even though they had photographed their tracks.

In an investigation, that an analogy should break down after a while is not necessarily a sad event. If by that time it has helped us find out enough of those things we want to know we do not mind the breakdown at all, since we can now go native (that is, depend on context for clues). Indeed, now that we are behaving as the natives behave we could be open to new experience and new discoveries, experience and discoveries not possible before we turn native.

Do review often

In the course of an investigation we should take time out now and then to review what we have done. This is important not only when the investigation is not going well; it is important even when the investigation is going well. As we have pointed out, in an investigation we can make use of approximations. This is to say even when we make little mistakes here and there we can still advance. In this respect an investigation is very different from setting out a proof in geometry. In the latter every line we put down has to be true; we cannot make even a single mistake. If one should occur, all that follows will be worthless. But an investigation is very different. In an investigation we follow clues. Clues can be looked upon as little gaps. We close these gaps one or a few at a time; there is no requirement that we should close them all at once. Since some of the closings have to be left to the future anyway, there is no harm that in closing some of them we, intentionally or unintentionally, create a few more, smaller ones, in the process.

But at some point these smaller gaps will have to be closed, or they will pile up and lead us astray. But closing them is easy, so long as we would take time out to do it. For in closing the bigger gaps one by one we will have found out more and more of the structures we are interested in knowing. Now that we know more it should be a simple

matter closing these smaller gaps. And this we do in reviews, which we say we should conduct from time to time.

Correcting minor mistakes in an investigation is called fine-tuning. In investigations it is common that we arrive at rough results first and then fine-tune them during reviews. The fine-tuning may have to be done not just once, but many times. After each fine-tuning we acquire a clearer idea of where we are in the investigation and where we should be heading. This is another reason for reviewing often.

Do retrace your steps if you think a major mistake has been made

There are very few investigations (if any) in which we sail smoothly from the beginning to the very end without encountering any difficulties. Difficulties are common in an investigation. Sometimes they are so great that forward movement is no longer possible. Now there can be many reasons why this is the case, but one possibility is that one or more major mistakes have been made. When major mistakes are made in an investigation clues will dry up.¹¹ Without clues progress is impossible.¹²

¹¹ By this we mean we cannot detect them.

¹² Minor mistakes in an investigation will not cause impasse in an investigation. They can be corrected during reviews.

To correct major mistakes the common practice is to retrace our steps. Since when a major mistake is made progress will soon become impossible, we check our most recent steps first. If the mistake is there and we can correct it, we should be able to resume progress. If we cannot find the mistake there we will have to go back farther and farther, until eventually it is found.

How do we know which step is a mistake? This is usually not an easy task and may include a fair amount of trial-and-error. For, if the mistake were easy to detect we would not have made it in the first place. But since it is causing the impasse if we can replace the wrong step with the right one we should be able to progress again. But how do we know what step is the right one? This again is difficult. If we knew, we would not have taken the wrong step. So at this point we engage in trial-and-error. Suppose we suspect Step A is wrong. We replace it with Step A1. If after doing this the investigation is able to move forward, A is the wrong step and A1 the right one. If after replacing A with A1 progress is still impossible, we replace A with A2 and see what happens. If whatever we put in place of A progress is still impossible, we may want to suspect a different step instead of A.

Now in the above we are assuming that we have made only one major error. If we have made only one major mistake, by replacing the wrong step with the right one we should be able to progress again.

But suppose we have made not just one major mistake, but more than one. Suppose also that Step A is one of the wrong steps. Now if this is the case replacing A by the right step will not rejuvenate the investigation (since there are other major mistakes besides Step A). What do we do then?

Correcting a single major error is difficult enough. Correcting multiple major errors is even more difficult. But there is a way for doing it, for which we now go to the next section.

Do consider, if necessary, re-doing the whole investigation or a large part of it

In an investigation, when we suspect major errors we usually suspect first that there is only one. When based on this suspicion we still cannot rejuvenate the investigation after many trials, we will have to start suspecting that there could be more than one major error.

Multiple major errors cannot be corrected the same way as a single major error. To correct a single major error all we need do is replace the one wrong step with the right one through trial and error. Once this is done the investigation should be able to advance again. But to correct multiple major errors we have to replace all the wrong steps at the same time. Even if we miss one the investigation still will not advance.

The common way to correct multiple major errors is to re-do the whole investigation, or the part of it in which the major errors are suspected to have occurred. The hope here is, the same mistakes will not be made a second time. How do we know that our hope is realised? We will know if in the new attempt we are able to venture beyond the point at which the investigation was stalled. If this happens we have corrected the multiple major errors.

Do consider taking in new blood

To correct multiple major errors we re-do the whole investigation or a large part of it, hoping that this time we will not make the same mistakes. But what is the likelihood that we will not? If the mistakes were not easy to make we would not have made them in the first place. If we have made them the first time the likelihood is high that we will make them a second time.

What can we do then? How can we increase our chances of *not* making the same mistakes a second time?

One thing we sometimes do is to call in new blood. Instead of having the same people working on the same investigation a second time, we hand over the investigation to a new group of people. The hope here is that this new group will look at the situation in a different way and as a result not make the same mistakes.

When we call in new blood, to better improve our chances we should prevent the new group from knowing what the old group has done. If we do not; if we let the old group brief the new group, for example; the new group might, consciously or unconsciously, approach the investigation more or less in the same way and thus make more or less the same mistakes. Ignorance is not always bliss, but in this case it is.

Do carry out independent, parallel investigations when called for

When major mistakes have been made in an investigation progress (in the investigation) will be impossible. One way to correct the situation is to hand over the investigation to new people, in the hope that they will not make the same mistakes. Instead of waiting for an investigation to bog down before we call in new people we sometimes carry out independent, parallel investigations right from the start. We do this especially with investigations in which major mistakes are easy to make. By keeping these parallel investigations independent; that is, by preventing the different groups in charge of these investigations from communicating with each other as much as possible; we increase our chance of success. This way it is less likely

that they all will make the same mistakes, so that even if some of them are bogged down, not all of them will.

Carrying out independent, parallel investigations saves time, since we do not have to wait for one investigation to bog down first before starting another. Also, if more than one of these parallel investigations produces the same results, the additional sets of result can serve as double-checks.

Do allow yourself to make use of third party results

Third party results are results obtained in investigations carried out by people other than ourselves. In investigations we cannot avoid using this kind of results. In carrying out an investigation many things are assumed 'known'. I look at the world much of the time through glasses. I assume known that glasses most of the time do not materially distort things seen through them. Did I ever conduct an investigation to determine this? No, I never did.

Should I? We have to make use of third party results in investigations. Before we use these third party results should we not make sure they are correct? And how can we make sure if not by an investigation?

Carrying out an investigation is certainly one way to determine if a third party result is correct. Scientists, for example, would

sometimes check on a third party result by re-doing the investigation that produces it. But checking on a third party result this way can be time-consuming, and costly in terms of resources, human and otherwise.

But re-doing an investigation is not the only way to check on a third party result. We can also check on a third party result by examining the history of the investigation that produces it. If from such a history we find out that the investigation has been properly carried out, we know then that we can use the third party result in our own investigation. And this scientists do too, to some extent, when they read scientific papers.

Why do I say 'to some extent'? I have added this phrase because written accounts of investigations do not always give us a full and detailed history of what happened. Moreover, it is not always easy to determine whether some steps in an investigation have been properly carried out. For example, an investigator might have written down that such and such a clue opens up possibilities A, B, and C. Is she correct? Could there be a D besides A, B, and C?

When we read the history of an investigation that produces a third party result we do not usually obtain from what we read a complete history which will allow us to determine unerringly that the result is correct. This is not to say that these histories are useless, but

we should be careful as to how much they can help us determine whether third party results are correct.

So far we have looked at two ways of evaluating third party results: re-doing the investigation and examining the history. Neither of them is ideal from a practical point of view. Is there anything else that we can do which will help us evaluate third party results?

There is a third way. To determine how good a third party result is we can simply use it in our investigation. If it leads to more and more new clues it is likely to be right. If this third party result were seriously wrong, it would have led to impasse.

Some will ask, if we have this third way, why trouble with the other two? Why, for example, should we examine the history of the investigation that produces the third party result? Why not just use the third party result in our investigation and see what happens?

We should know it is risky to accept third party results without questioning them. It is true that if a third party result leads to more and more new clues it is likely to be right, but if it is wrong, seriously wrong, it could cause major difficulties. In an investigation we usually make use of not just one third party result but many. If they are all accepted on faith and more than one of them is seriously wrong, our investigation will reach an impasse from which it will be difficult to extricate ourselves. When we use third party results we

have to be careful how likely to be correct they are before we use them. On those occasions on which we cannot re-do those investigations that produce these third party results, which is most of the time, we will have to pay some attention to the histories of these third party results, even if these histories are sketchy. In investigations success is not promised in advance, but we should try our best to increase our chances.

Do make use of the process of elimination when possible

To find out things hidden from us we follow clues. But clues do not always pinpoint; most of the time a clue will suggest a number or possibilities. In investigating a murder, for example, a clue we have found could have suggested to us that it is the butler, or the maid, or the master of the house. Now in such a situation one thing we could do is to see which of these possibilities leads to more and more new clues. The one that does is likely to be the truth. But often instead of doing this we could engage in a process of elimination and by doing so arrive at an answer faster. Trying out different possibilities to see which leads to more and more new clues takes time, but if in the meantime there are other clues that will allow us to eliminate all of the possibilities except one, that remaining one is likely to be the true answer. And this we should explore during an investigation. That is,

after some clue has suggested to us a number of possibilities we should try to see if there are some quick means of eliminating some of them. The ideal is to have only one possibility remaining after the elimination. If there is only one it is likely to be the truth. (If there is more than one after the majority of the possibilities has been eliminated, it is still a step forward, since we will now have much fewer possibilities to examine.)

What if all the possibilities have been eliminated, as it frequently happens?!

When this occurs one possible reason is that we have not exhausted all the possibilities before we carry out the eliminations. If we think this is the case we will have to backtrack and see what we have left out.

Suppose that we meet with the ideal situation, that is, after carrying out the eliminations we have one possibility left. Is this remaining possibility *necessarily* the truth?

We have to be careful here. Clues do not lead to absolute certainties. The process of elimination is a process we carry out when following clues. Thus even when there is only one possibility left we should *not* conclude that this one possibility absolutely must be the truth. The process of elimination depends on our exhausting all possibilities. Being human, we might not have exhausted all

possibilities. For this reason, when after a process of elimination we have only one possibility left we should test this possibility to see whether it leads to more and more new clues. If it does it is likely to be the true answer. If it does not we will have to backtrack to see what possibilities we have left out. The process of elimination is a short cut in the theseological process (instead of following a thousand and one leads we follow one, or a few); it is not its replacement.

In employing the process of elimination we assume that we can exhaust all the possibilities. Is it possible that in theory, and therefore in practice too, we can never exhaust all possibilities?

If both in theory and in practice we can never exhaust all possibilities, we will not be able to make use of the process of elimination. For however many of the possibilities we eliminate there will always be an infinite number left. *But this is not the kind of situation we meet with when following clues.* In the theseological process the possibilities are suggested by clues. Clues always narrow down the field. A clue cannot suggest just any thing and still be a (useful) clue. A clue always rules out as well as point to. A clue sometimes points to a certain range in which the true answer can be found. Now in this range there could be an infinite number of possibilities, but in this kind of cases, when circumstances are appropriate we can use the process of elimination to narrow down the

range. For example, we could say, the true answer has to be in one of three regions within this range: regions A, B, and C (each of which also contains an infinite number of possibilities). The true answer is not in A and B. Therefore it must be in C. Now it is true that even after the process of elimination we do not know what the true answer is, but we know that it is in C. And if this process can be repeated we should be able to come closer and closer to the true answer even though we will never know exactly what it is.

To repeat, it is true that if we cannot exhaust all possibilities we cannot employ the process of elimination. But when following clues we in theory should be able to exhaust all possibilities. This is because clues always narrow down the field. They rule out as well as allow in.

Do ask questions of the unknown through experiments

In one of his cases Sherlock Holmes was asked by a client who was being blackmailed to retrieve some photographs in the possession of a certain 'lady'. It was known that the photographs were in her house, but where in her house nobody knew but her. In order that he could take possession of these photographs—that is, steal them—Sherlock Holmes devised a simple stratagem. He started a fire in the house (which fire was in fact a false alarm) and watched the first thing

the lady did. And what was the first thing the lady did? She rushed to the spot where the photographs were hidden.

Sherlock Holmes wants to know the location of the photographs. He wants to know where they are. Obviously, there is no point in asking the lady in question; she will never tell. But this does not mean that Sherlock Holmes cannot find out. He can still ask the same question, in the form of an experiment: he starts a fire. The result of the experiment tells him where the photographs are.

Without a confession from the blackmailer Sherlock Holmes succeeds in finding the answer to his question. His question is not addressed to the blackmailer but to the situation in which the blackmailer finds herself. A blackmailer needs an instrument of blackmail—in this case, the photographs. As an instrument of blackmail the photographs are of great value to this lady. When there is a fire they will be among the first things she will want to save.

In theseologising there are many questions for which we want answers. Sometimes some of these questions can be answered by experiments. In these cases these questions are not addressed to those who already know (for example, the criminals themselves, or God), but to the *structures* we are investigating. Since structures have characteristics these experiments can provide the answers we are

looking for. In escaping from a fire most people will want to take with them things they regard as valuable.

As we all know experiments are common in science. In them scientists ask questions of that part of nature they are investigating. How many helixes are there in the DNA molecule? How much electric charge is there on *each* electron? These questions were answered by experiments.

We do not usually associate experiments with cryptanalysis, but in fact in cryptanalysis experiments are sometimes carried out. Just as Sherlock Holmes can induce a blackmailer to reveal the location of photographs being used in the blackmail, cryptanalysts can induce their opponents to reveal to them the meaning of certain code words, for example. Take a simple case. Suppose you want to find out what code name your opponents have assigned to General Eisenhower. To this end you send out to your own people an important, secret message, faked, in which the name General Eisenhower will be mentioned. Now if this message is picked up by your opponents and repeated among themselves using their own cipher (which you have broken), you will be able to find out what code name they have assigned to Eisenhower. For example, suppose you say in your original message that 'Eisenhower is visiting the island of Crete on March 15'. Now if the same message when repeated among your opponents

comes out as 'Shepherd is visiting the island of Crete on March 15', you know that Shepherd is the code name your opponents are using for General Eisenhower.

In following clues we can ask questions of the as yet unknown and find answers. This usually requires a certain amount of ingenuity (as in Sherlock Holmes's case) but is worth attempting. In uncovering the hidden the more we know the more we can find out.

Do be tenacious

It is difficult to follow clues. In following clues success is not promised in advance. If we want to find out anything by following clues we will have to be tenacious: we will have to keep trying even though we meet with disappointment after disappointment. In following clues we should expect many failures before we succeed, if we succeed at all.

Do maintain a certain degree of optimism

Failure is hard on our psychology. In tasks we carry out, if we keep failing, in time we will feel depressed and want to give up. Now failure is common when following clues. If we want to be successful when following clues we will have to instil in ourselves a certain degree of optimism, and maintain it even in difficult times.

Do be skeptical—in a healthy way

A certain degree of skepticism is always called for in investigations. When following clues the chance that we have made mistakes, major or minor, singly or in plural, is always there. There is never absolute certainty, however many new clues we may have reaped.