

A Preferred Treatment of Mill's Methods: Some Misinterpretations by Modern Textbooks

BRAM VAN HEUVELN *Binghamton University*

Abstract: A number of modern logic books give a misrepresentation of Mill's Methods as originally conceived by Mill. In this paper, I point out what I believe is a better presentation of Mill's Methods. This treatment is not only historically more accurate, but it also represents a better conceptual introduction to Mill's Methods in general.

Résumé: Un certain nombre de manuels courants de logique représentent mal les méthodes de Mill, telles que Mill lui-même les a originellement conçues. J'identifie ce qui me semble être une meilleure présentation des méthodes de Mill. Cette approche est plus exacte du point de vue historique et en général une meilleure introduction conceptuelle.

Keywords: J. S. Mill, Mill's methods, inductive reasoning, method of agreement, method of difference, joint method of agreement and difference, textbook treatments

1. Introduction

Modern logic textbooks refer to Mill's Methods as a paradigm example of inductive reasoning. Unfortunately, some of them misrepresent the methods as originally described by Mill himself. This particularly concerns what Mill called the Joint Method of Agreement and Difference. Understandably, many textbooks treat the Joint Method as some kind of combination of the Method of Agreement and the Method of Difference. However, a quick look at Mill's original writings reveals that Mill did not conceive of the Joint Method as such. Moreover, the treatment given by these textbooks is not only historically inaccurate, but often also conceptually confusing. Indeed, I argue that the Methods are best understood when closely following Mill's original presentation.

In Section 2 I give a general overview of Mill's Methods and the basic ideas behind them. Section 3 discusses the Joint Method in particular. In here, I refer back to Mill's own writings to show what Mill intended this method to be, and thus to demonstrate the general misinterpretation of the Joint Method by various modern textbooks. In Section 4 I show a number of specific ways in which the Joint Method has been misrepresented. Finally, in Section 5, I intend to describe what I believe to be a preferred treatment of Mill's Methods. This treatment is historically more accurate, while at the same time illuminating the important ideas behind Mill's Methods.

2. The Method of Agreement and The Method of Difference

Basically, there are five of Mill's Methods: Method of Agreement, Method of Difference, Joint Method of Agreement and Difference, Method of Residues, and Method of Concomitant Variation. In this paper, I will only be concerned with the first three methods, and with the third method in particular. This section therefore presents the Method of Agreement, the Method of Difference, and some related methods as presented by various modern textbooks. The next section will focus on the Joint Method.

2.1 *The Basic Idea behind Mill's Methods*

Let me start by giving the general idea behind Mill's Methods while at the same time also establishing some general terminology. When applying Mill's Methods, one is looking for the *cause* of a certain *phenomenon* (sometimes called *effect*). To do this, one considers a number of cases, called *instances*, in which the phenomenon is either present or absent. In those instances, certain *circumstances* will also be present or absent. Using Mill's Methods, one may be able to find reasons that a certain circumstance was the cause of the phenomenon.

As a running example to illustrate all the different methods and interpretations thereof, I will use a hypothetical case of food poisoning in a student dorm. In this case, a number of students became ill after eating in the dining hall. To find out what caused the food poisoning, a number of students were questioned about what they had eaten that night. The following table lists the results:

STUDENT	STEAK?	FRIES?	PASTA?	BEANS?	FOOD POISONING?
Carla	No	Yes	Yes	Yes	Yes
John	Yes	No	No	Yes	Yes
Tom	Yes	Yes	No	No	No
Mary	No	Yes	Yes	No	No

We see that there are two instances, Carla and John, where the phenomenon (food poisoning) is present. We call these *positive instances*. There are also two *negative instances*: Tom and Mary. Finally, there are four circumstances (steak, fries, pasta, and beans) that are assumed to have been the possible cause of the phenomenon. Mill's Methods embody a variety of techniques to find out what caused the food poisoning.

2.2 *Assumptions Behind Mill's Methods*

The application of Mill's Methods works under a number of assumptions. First of all, it is assumed that the listed circumstances do contain the cause of the phenomenon under investigation. Second, it is assumed that the cause was the same in all instances, i.e., that all instances are somehow comparable to each other. And

third, it is often assumed that a single circumstance was the cause of the phenomenon. However, all of these assumptions can be false. There could have been some other and unlisted circumstance that was the cause, maybe the phenomenon was caused by different circumstances in different positive instances, and maybe the phenomenon was the result of a combination of circumstances.

These assumptions reveal that Mill's Methods can never conclusively show that a certain circumstance is indeed the cause of a certain phenomenon, making them inherently inductive. Another implication of this observation is that one always has to decide whether Mill's Methods can even be reasonably applied to any particular case before applying them. This decision will have to be made from case to case, as there are no set rules for making this decision. In the food poisoning example, all of the assumptions are plausible. Given what we know about food poisoning, and given that all the students ate at the same dining hall on the same day, it is plausible that there was indeed a single food item that caused the food poisoning in all positive instances. And, we can safely assume that all the food items that potentially caused the food poisoning are listed. The food poisoning example is therefore very suitable for the application of Mill's Methods.

2.3 *The Method of Agreement*

John Stuart Mill wrote:

If two or more instances of the phenomenon under investigation have only one circumstance in common, the circumstance in which alone all the instances agree, is the cause (or the effect) of the given phenomenon. ([12], Book III, Chapter VIII, §4)

Mill's description leaves no room for misinterpretation: If there is a single circumstance that is present in all positive instances, then by the Method of Agreement we can induce that it was this circumstance that was the cause of the phenomenon. Indeed, I have seen no textbook that gave any other interpretation of the Method of Agreement.

In our food poisoning example, we have the following positive instances:

STUDENT	STEAK?	FRIES?	PASTA?	BEANS?	FOOD POISONING?
Carla	No	Yes	Yes	Yes	Yes
John	Yes	No	No	Yes	Yes

Since there is a single circumstance (beans) that is present in both positive instance, we can conclude that it was the beans that caused the food poisoning.

Some textbooks point out that it can be helpful to see the Method of Agreement as a process of *elimination* ([6], [8]). Thus, since Carla did not eat the steak, but did get sick, we have good reasons to believe that it was not the steak that caused the food poisoning. Similarly, the fries and the pasta can be elimi-

nated as the possible cause of the food poisoning as well. However, since the beans were eaten by both Carla and John, we still have reason to believe that the beans caused the food poisoning. In fact, since the beans are the *only* item that was eaten by both Carla and John, the Method of Agreement allows us to make this our tentative conclusion.

2.4 *The Inverse Method of Agreement*

Some textbooks list an additional kind of Method of Agreement, called the *Inverse Method of Agreement* ([8], [17]) or the *Negative Method of Agreement* ([11],[16]). Indeed, to differentiate the Method of Agreement from this alternative, the Method of Agreement is often renamed as the *Direct Method of Agreement* or the *Positive Method of Agreement*. The Inverse Method of Agreement is not listed by Mill as a separate Method, although we will see in Section 3 that Mill implicitly refers to this Method nevertheless. The Inverse Method of Agreement allows one to conclude that a certain circumstance is the cause of the phenomenon under investigation if this circumstance is the only circumstance (of those considered) that is absent in all negative instances.

As an example, let us consider all negative instances of the food poisoning case:

STUDENT	STEAK?	FRIES?	PASTA?	BEANS?	FOOD POISONING?
Tom	Yes	Yes	No	No	No
Mary	No	Yes	Yes	No	No

Since there is a single circumstance (beans) that is absent in all these negative instances, we can once again make the induction that it was the beans that caused the food poisoning.

Again, it helps to conceive of the Inverse Method of Agreement as a process of elimination. Thus, since Tom ate the steak but did not get food poisoning, we have reason to believe that it was not the steak that caused the food poisoning. Similarly, we can eliminate the fries and the pasta as possible causes, and hence we are left with the beans.

Some textbooks explicate the difference between the Direct Method of Agreement and the Inverse Method of Agreement as a difference between finding reasons for eliminating circumstances as a *necessary* condition or as a *sufficient* condition for the phenomenon respectively ([8], [17]). That is, if a circumstance is absent in a positive instance, then this circumstance cannot be a necessary condition for the phenomenon. Hence the Method of Agreement can be seen as ruling out circumstances as necessary conditions. Likewise, the Inverse Method eliminates circumstances as sufficient conditions for the phenomenon.

2.5 The Double Method of Agreement

The Direct Method of Agreement and the Inverse Method of Agreement can be combined into what some textbooks call the *Double Method of Agreement* ([8], [17]). Thus, the hope is that the Direct Method picks out a single circumstance as the cause of the phenomenon, while the Inverse Method picks out that very same circumstance. If that is the case, then the Double Method allows one to conclude that that circumstance is the cause of the phenomenon.

In the food poisoning example, the Double Method allows one to say that the beans were the cause of the food poisoning, since it was the only circumstance present in all positive instances, and the only circumstance absent in all negative instances.

With the above definition of the Double Method in place, some people may raise the following question. If the Double Method can be applied, then apparently one would have been able to find the answer already by merely applying either the Direct Method of Agreement or the Inverse Method of Agreement individually. So, doesn't that make the Double Method superfluous?

The rather obvious answer to this question is that both the Direct Method of Agreement and the Inverse Method of Agreement merely establish *reasons to believe* that a certain circumstance was the cause of the phenomenon. Hence, the more reasons to believe this, the stronger our tentative conclusion. Indeed, by combining both the Direct and the Inverse Method of Agreement, the Double Method gives us *more* reason to believe that it was indeed this circumstance that was the cause.

The Mill scholar Skorupski, however, gives an interesting alternative answer ([16], pp. 181-183). To understand his answer, let us change the food poisoning example a little bit:

STUDENT	STEAK?	FRIES?	PASTA?	BEANS?	FOOD POISONING?
Carla	No	Yes	Yes	Yes	Yes
Ann	Yes	Yes	No	Yes	Yes
Doug	Yes	No	No	No	No
Byron	No	Yes	No	No	No

In this case, both the Direct Method of Agreement and the Inverse Method of Agreement do not give us any unique answer, since the Method of Agreement points to either the fries or the beans, whereas the Inverse Method points to either the pasta or the beans. However, at this point we do of course want to combine the last two results and say that it is the beans that caused the food poisoning, since only the beans are pointed at in both cases.

We thus see that we can consider two forms of the Double Method:

1. *General Double Method*: If there is only one circumstance that is both present in all positive instances and absent in all negative instances, then we have reason to believe that that circumstance is the cause of the phenomenon.
2. *Specific Double Method*: If there is only one circumstance present in all positive instances, and that very same circumstance is also the only circumstance absent in all negative instances, then we have reason to believe that that circumstance is the cause of the phenomenon.

Obviously, the general Double Method can be applied whenever the specific Double Method can be applied. However, as the above example shows, the general Double Method can sometimes be applied where any of the other Methods fail. In fact, in the above example one cannot apply the Method of Difference either, as will become clear from the discussion of that Method in the following section.

2.6 *The Method of Difference*

John Stuart Mill wrote:

If an instance in which the phenomenon under investigation occurs and an instance in which it does not occur, have every circumstance in common save one, that one occurring only in the former, the circumstance in which alone the two instances differ, is the effect, or the cause, or an indispensable part of the cause, of the phenomenon. ([12], Book III, Chapter VIII, §4)

As with the (Direct) Method of Agreement, also this description leaves nothing to the imagination. Applying the Method of Difference, one looks for a positive instance and a negative instance that agree in the absence or presence of all considered circumstances except for one. If the latter circumstance is present in the positive instance, but absent in the negative instance, then that circumstance can be concluded to be the cause of the phenomenon. Many textbooks give this interpretation of the Method of Difference, although we will see some exceptions in Section 4.2.

In the food poisoning example, there are the following two relevant instances:

STUDENT	STEAK?	FRIES?	PASTA?	BEANS?	FOOD POISONING?
Carla	No	Yes	Yes	Yes	Yes
Mary	No	Yes	Yes	No	No

Since the one circumstance in which they differ (beans) is present in the positive instance, but absent in the negative instance, the Method of Difference allows us to conclude that it is the beans that caused the food poisoning.

Technically, all we can conclude is that the beans are likely to be a necessary condition of the food poisoning. That is, maybe it was some combination (e.g., beans *plus* fries) that caused the food poisoning. This is why Mill added the phrase “or an indispensable part of the cause.”

3. The Joint Method of Agreement and Difference

Although believed by Mill to be an additional and separate technique, this method is best understood as the combined use of the Method of Agreement and the Method of Difference in the same investigation. ([6], p.513)

This quote by Copi and Cohen regarding the Joint Method reveals exactly the kind of mistake that a number of textbooks are making when presenting the Joint Method. That is, although details vary from textbook to textbook, the Joint Method is presented as some combination of both the (Direct) Method of Agreement and the Method of Difference ([2], [6], [8], [10], [14], [15], [17]). Given the name of this method, this is quite understandable. However, a quick look at Mill's original writings reveals that this was never intended.

3.1 *Mill and The Joint Method*

Here is what Mill writes about the Joint Method:

In cases ... in which it is not possible to obtain the precise pair of instances which [the Method of Difference] requires—instances agreeing in every antecedent except [circumstance] A . . . ; we may yet be able, by double employment of the Method of Agreement, to discover in what the instances which contain . . . [phenomenon] *a*, differ from those which do not.

. . . Suppose, then, that as we previously examined a variety of instances in which [phenomenon] *a* occurred, and found them to agree in containing [circumstance] A, so we now observe a variety of instances in which *a* does not occur, and find them agree in not containing A; which establishes, by the Method of Agreement, the same connexion between the absence of A and the absence of *a*, which was before established between their presence.

. . . This Method may be called the Indirect Method of Difference, or the Joint Method of Agreement and Difference; and consists in a double employment of the Method of Agreement, each proof being independent of the other, and corroborating it, ([12], Book III, Chapter VIII, §4)

Two things become clear from Mill's writing:

1. Mill thinks that it is possible for the Joint Method to be applicable in cases where the Method of Difference is not applicable. This *immediately* refutes the claim that the Joint Method would be the employment of both the method of Agreement and the Method of Difference.
2. In Mill's description of the Joint Method we recognize the Double Method. Indeed, with "a double employment of the Method of Agreement", Mill almost without any doubt meant the employment of both the Direct Method of Agreement and the Inverse Method of Agreement.

3.2 *The Literal Interpretation*

It is fairly certain that Mill intended the Joint Method of Agreement and Difference to be the Double Method. So why did Mill give the Joint Method such a misleading

name? The plausible answer is that the Joint Method, *through* a double employment of the Method of Agreement, *resembles* the Method of Difference. That is, like the Method of Difference, the Joint Method observes that there is a difference between positive and negative instances in their respective presence and absence of a certain circumstance. Of course, the difference between the Joint Method and the Method of Difference is that the Method of Difference all instances have to be alike with regard to these other circumstances, whereas for the Joint Method no such restriction applies. The Joint Method thus leads one to a weaker conclusion than the Method of Difference, which is why Mill sometimes referred to the Joint Method as the ‘Indirect Method of Difference’. It is certain, however, that the Joint Method does not contain the Method of Difference. Thus, recalling Copi and Cohen’s quote about the Joint Method, Mill regarded the Joint Method as a separate method, because it really *is* a separate method.

Still, when all is said and done, the ‘Joint Method of Agreement and Difference’ remains an unfortunate and confusing name for the Double Method. To add to the confusion, Mill made a failed attempt at a precise formulation of the Joint Method. Here is what Mill writes:

If two or more instances in which the phenomenon occurs have only one circumstance in common, while two or more instances in which it does not occur have nothing in common save the absence of that circumstance, the circumstance in which alone the two sets of instances differ is the effect, or the cause, or an indispensable part of the cause, of the phenomenon. ([12], Book III, Chapter VIII, §4)

First of all, notice that this definition does not describe any kind of combination of the Method of Agreement and the Method of Difference. However, it isn’t exactly the Double Method either. The problem lies with the second requirement that the negative instances “have nothing in common save the absence of” some circumstance. To illustrate this requirement, let us change our food poisoning example such as to make sure that all negative instances indeed “have nothing in common save the absence of” some circumstance:

STUDENT	STEAK?	FRIES?	PASTA?	BEANS?	FOOD POISONING?
Carla	No	Yes	Yes	Yes	Yes
John	Yes	No	No	Yes	Yes
Doug	Yes	No	No	No	No
Mary	No	Yes	Yes	No	No

The positive instances have only the beans as a common present circumstance, while the negative instances “have nothing in common save the absence of” the beans. By the literal interpretation of the Joint Method we can therefore conclude that it was the beans that caused the food poisoning

However, it does not seem to make any sense to require that the negative instances “have nothing in common save the absence of” some circumstance in

order to be able to draw the appropriate conclusion. That is, if we take Mill's definition from above literally, then in the original food poisoning example (see Section 2.1) the inference would be blocked by the mere fact that both Tom and Mary ate the fries. Had one of them not eaten the fries, then the inference could have been made, but now it cannot. I think that most readers will agree that such a separation is quite unintuitive. It is therefore most likely that Mill made a slip of the pen when writing the above definition, and that it should not be taken literally.

Interestingly, in the very first edition of *Introduction to Logic* published in 1953, Copi does offer the suggested literal interpretation as one of three alternative interpretations of the Joint Method ([4], p. 345). Even more interestingly, Copi then criticizes this interpretation in that it does not reflect any kind of combination of the Method of Agreement and the Method of Difference. Indeed, in the second edition, only the interpretation that takes the Joint Method to be the independent application of both the Method of Agreement and the Method of Difference is left ([5]). This interpretation is also found in the most recently published tenth edition and, I presume, in all editions in between ([6]). However, I already indicated that the latter interpretation is incorrect by itself. Copi therefore rejected the literal interpretation on the wrong grounds.

Similarly, Cohen and Nagel rightly reject the literal interpretation of Mill's definition of the Joint Method, but do so for the wrong reasons. According to Cohen and Nagel, it is 'absurd' to suppose that two instances would "have nothing in common save the absence of" some circumstance ([3], p. 260). That is, two instances will always have *something* in common (e.g. in both cases President Nixon was not present), so the rule can never be applied. However, following that line of reasoning, one would have to reject the Method of Difference as well, since it is equally absurd to think that two instances would differ in exactly one respect. And, for that matter, the Method of Agreement would become impractical too. The mistake Cohen and Nagel are making is that Mill's Methods are supposed to work on a predefined set of circumstances which, according to certain case-specific criteria, are thought to be relevant. Thus, the fact that all instances share the absence of President Nixon is, for most cases, completely irrelevant. Given the working constraints of Mill's Methods, it is in fact not all that hard to come up with a case where two negative instances do "have nothing in common save the absence of" some circumstance, as demonstrated in the example above. The literal interpretation should therefore not be rejected for this reason. The *right* reason to reject the literal interpretation is that there is simply no intuition that lies behind it.

3.3 Possible Causes of Misinterpretation

Obviously, the confusing name of the Joint Method may very well have caused the authors of modern textbooks to misrepresent the Joint Method as some kind of combination of the Method of Agreement and the Method of Difference. However, this does not explain the fact that most *older* textbooks present the Joint

Method *correctly* (as I argue throughout this paper) as the Double Method ([1], [7], [9], [13], [18], [20]). In fact, where some of these older textbooks explicitly point out that the Joint Method can and has been misunderstood ([1], [7], [18]), I have been unable to find actual instances of misrepresentation before 1950. Moreover, it seems that it was in particular during the 50's and 60's that a number of textbooks started to adopt some mistaken version of the Joint Method ([4], [5], [14], [17]). So, why this sudden shift?

One plausible answer to this question is that some authors have simply copied mistakes from others. Thus, once a few influential authors were actually thrown off by the name of the Joint Method, later textbooks simply propagated these mistakes. Indeed, although I can only speculate here, I can see how a widely used textbook such as Copi's *Introduction to Logic* may very well have been at the base of some of the current mistreatments of the Joint Method. Indeed, the first two editions of this book came out in 1953 and 1962, right at the time of the apparent shift ([4], [5]).

There is, however, an intriguing second explanation of the shift from correct to incorrect treatments. In Section 4, I will indicate that there are actually two subtly different ways in which the Joint Method has been viewed as some kind of combination of the Method of Agreement and the Method of Difference. One is where the Method is seen as the simple independent application of both the Method of Agreement and the Method of Difference. This is the interpretation given by Copi, and it is exactly the kind of treatment that the name of the Method would suggest (see Section 4.1). The second class of misinterpretations, however, combines the Method of Agreement and the Method of Difference in a much more complex way. Indeed, this second kind of misinterpretation seems to go beyond a mere confusion of the name of the Method (see Section 4.2). The question thus emerges where this type of misinterpretation originated.

My personal investigations suggest that this second class of misinterpretations was the result of a number of technical advancements on Mill's Methods. In particular, Von Wright's *Treatise on Induction and Probability*, which was first published in 1951, presents a highly technical account of a number of inductive tools that try to generalize the basic ideas behind Mill's original Methods ([19]). One particular technique described by Von Wright is to use positive instances to narrow down the list of possible causes, and to *subsequently* use negative instances to narrow down this list even more. It is indeed this technique that is quite typical of the second class of misinterpretations as described in more detail in Section 4.2. Moreover, at one point, Von Wright writes:

Thus there are two possibilities of a joint use of the methods. The first is to combine the Direct Method of Agreement with the Inverse Method of Agreement. The second is to combine the Direct Method of Agreement with the Method of Difference. We shall call the method originating from the first combination the Double Method of Agreement, and the method originating from the second the Joint Method of Agreement and Difference. ([19], p. 99)

Now, the distinction between the two possible combinations had been recognized before. For example, in his 1939 book *The Rhyme of Reason*, Holmes was quite explicit about this ([7]). However, the terminology given to the distinction by Von Wright was very unfortunate, for Mill used the 'Joint Method of Agreement and Difference' to refer to the first combination as indicated by Von Wright, and not the second. Indeed, together with the particular ways in which Von Wright combined the methods, this terminology was straightforwardly adopted by Skyrms in *Choice and Chance*, and subsequently by Hurley in the popular modern textbook, *A Concise Introduction to Logic* ([8], [17]). My hypothesis is that it is through this path that quite a few of the current misinterpretations of Mill's Joint Method have come around (e.g., [2], [15]).

Surprisingly, the scholarly literature seems to be quite silent on this whole subject. The only Mill scholar that I have found to take up Mill's Methods in substantial detail is Skorupski, who notes the awkwardness of Mill's formulation of the Joint Method, offers an alternative definition similar to the Double Method, but fails to observe the misinterpretations by modern textbooks ([16]).

4. Some Specific Misrepresentations

This section gives an account of some specific ways in which Mill's Methods have been misrepresented. All misrepresentations see the Joint Method as some combination of the Method of Agreement and the Method of Difference. However, as indicated earlier, there are roughly two different ways in which this has been worked out. These are:

1. The Method of Difference is presented correctly, and the Joint Method is presented as the application of both the Method of Agreement and the Method of Difference.
2. The Method of Difference is presented as something like an Inverse Method of Agreement, and is combined as such with the Method of Agreement to form the Joint Method.

This division is rather crude, as details vary widely from treatment to treatment. Not all misrepresentations fall under any of these two headings very well, while others seem to fall under both headings at once. Extra complications are introduced by explicit considerations of combinations of circumstances, or by considerations of the generality of the conclusion. Still, I think the classification is a useful one, and many misrepresentations can be seen as committing one of these two mistakes.

4.1 The Straightforward Combination

The first way in which the Joint Method is misinterpreted is by viewing it as a straightforward combination of the (Direct) Method of Agreement and the Method of Difference as presented in Sections 2.3 and 2.6. Thus, consider the original example:

STUDENT	STEAK?	FRIES?	PASTA?	BEANS?	FOOD POISONING?
Carla	No	Yes	Yes	Yes	Yes
John	Yes	No	No	Yes	Yes
Tom	Yes	Yes	No	No	No
Mary	No	Yes	Yes	No	No

The Method of Agreement can be applied to Carla and John, while the Method of Difference can be applied to Carla and Mary. Both Methods point to the beans as the culprit, and so by this interpretation of the Joint Method, we now have extra good reason to believe that it was indeed the beans that caused the food poisoning. That is, as with the Double Method, the plausibility of the conclusion obtained through this version of the Joint Method is greater than when applying merely the Method of Agreement or merely the Method of Difference.

A widely used textbook that gives this misinterpretation of the Joint Method is *Introduction to Logic* by Copi and Cohen ([6]). Another well-known textbook that gives this account of the Joint Method is *Logic & Philosophy: A Modern Introduction* by Kahane and Tidman ([3], pp. 341-348).

In the first edition of *Introduction to Logic*, Copi actually presents three ways to combine the Method of Agreement and the Method of Difference ([4]). The first one is the one we just saw, and the second is the literal interpretation as described in Section 3.2. Under the third interpretation, the Method of Difference is only applied *after* elimination of a number of necessary conditions according to the Method of Agreement. Thus, let us consider the revised food poisoning example:

STUDENT	STEAK?	FRIES?	PASTA?	BEANS?	FOOD POISONING?
Carla	No	Yes	Yes	Yes	Yes
Ann	Yes	Yes	No	Yes	Yes
Doug	Yes	No	No	No	No
Byron	No	Yes	No	No	No

Applying the Method of Agreement leaves us with the fries and the beans as possible causes. Thus, eliminating the steak and the pasta, we can apply the Method of Difference on Carla and Byron to obtain the conclusion that the beans were the cause of the food poisoning:

STUDENT	FRIES?	BEANS?	FOOD POISONING?
Carla	Yes	Yes	Yes
Byron	Yes	No	No

4.2 The Method of Difference as the Inverse Method of Agreement

The second group of misinterpretations is characterized by their use of the Method of Difference as a method that works very much like the Inverse Method of Agreement. That is, the Method of Difference is used to eliminate circumstances present in negative instances as possible causes of the phenomenon. When combining such a Method of Difference with the Method of Agreement, the resulting Joint Method is something that is, or is similar to, the Double Method. The problem, however, is that the Method of Difference is misrepresented, and that the Joint Method is still seen as a combination of the Method of Agreement and the Method of Difference. Again, details vary, but several specific treatments of the Joint Method as found in the literature can be classified as this kind of misrepresentation.

One clear instance can be found in *Invitation to Critical Thinking* by Rudinow and Barry ([15], pp. 218-221). In their book, Rudinow and Barry misinterpret the Method of Difference literally as the Inverse Method of Agreement as discussed in Section 2.4. A related misinterpretation of this category can be found in *The Elements of Logic* by Barker. In his book, Barker describes the Method of Difference as if it were the Inverse Method of Agreement, although the subsequent example seems to treat it as something like the Double Method all by itself ([2], p. 197).

The treatment of the Joint Method as found in Skyrms' highly detailed account of Mill's Methods belongs to this second class of misrepresentations as well ([17], pp. 84-105). Skyrms presents the Method of Difference as a *conditional* process of eliminating sufficient conditions used to find the likely cause of the phenomenon present in one particular instance. That is, one specific positive instance is used to obtain a list of possible causes, after which all circumstances from that list that are present in one or more negative instances are eliminated.

To illustrate Skyrms' Method of Difference, we can follow the original example, and ask what caused Carla's food poisoning. Thus we obtain the fries, pasta, and beans as the initial list of possible causes. We then use Tom and Mary to eliminate both the fries and the pasta from this list, leaving us with the beans as the likely cause of Carla's food poisoning. What caused John's food poisoning, however, is still an open question.

This Method of Difference can be combined with the Method of Agreement by allowing the use of positive instances in the conditional elimination process. Thus, consider the revised example:

STUDENT	STEAK?	FRIES?	PASTA?	BEANS?	FOOD POISONING?
Carla	No	Yes	Yes	Yes	Yes
Ann	Yes	Yes	No	Yes	Yes
Doug	Yes	No	No	No	No
Byron	No	Yes	No	No	No

When asked what caused Carla's food poisoning, we once again start with the initial list of fries, pasta, and beans as the possible causes. This time, however, the negative instances Doug and Byron can only be used to eliminate the fries. However, the other positive instance, Ann, can be used to eliminate the pasta, leaving us with the beans as the likely cause of Carla's food poisoning after all.

Skyrms calls this combination of the Method of Agreement and Method of Difference the Joint Method. Skyrms separates the Joint Method from the Double Method based on the generality of the obtained conclusion, meaning that the conclusion of the Joint Method only pertains to the one positive instance that this method starts with, whereas the conclusion of the Double Method pertains to all instances.

Hurley's popular modern textbook, *A Concise Introduction to Logic*, pretty much follows Skyrms' treatment ([8], pp.511-522). Like Skyrms, Hurley states that the obtained conclusion only applies to the specific positive instance to which the Method of Difference is applied. However, unlike Skyrms' Method of Difference, Hurley's Method of Difference requires a positive and a negative instance that differ in exactly one circumstance, just as the Method of Difference described in Section 2.6. Thus, following the original example, we can apply the Method of Difference to Carla and Mary, although the conclusion that it was the beans that caused the food poisoning only pertains to Carla, and not to John (unless a further inductive inference is made).

In Hurley's Joint Method, however, the requirements on the use of the Method of Difference are relaxed. In fact, the Method of Difference is used here as the Inverse Method of Agreement to eliminate circumstances. Thus, Hurley's Joint Method turns out to be exactly like Skyrms' Joint Method. Indeed, Hurley draws the same difference between the Joint Method and the Double Method in terms of generality of the conclusion.

In *Introduction to Logic*, Rescher misrepresents the Method of Difference as well. Rescher's Method of Difference looks for a circumstance present in one positive instance, but absent in all negative instances. This is very much like Skyrms' Method of Difference, but Rescher combines this Method of Difference with the Method of Agreement in a slightly different way. If there is one circumstance present in all positive instances, and if the Method of Difference points to that same circumstance whenever each individual positive instance is compared with all negative instances, then Rescher's Joint Method allows one to conclude that it was that circumstance that caused the phenomenon ([14], pp. 300-304). To illustrate Rescher's Joint Method, let us consider the original example:

STUDENT	STEAK?	FRIES?	PASTA?	BEANS?	FOOD POISONING?
Carla	No	Yes	Yes	Yes	Yes
John	Yes	No	No	Yes	Yes
Tom	Yes	Yes	No	No	No
Mary	No	Yes	Yes	No	No

The Method of Agreement points out that the beans are the only circumstance present in all positive instances. We can then compare each individual positive instance with all negative instances. Thus, comparing Carla with Tom and Mary, we find that of all circumstances present in Carla's case (fries, pasta, and beans), only the beans are absent for all negative instances. Similarly, when comparing John with Tom and Mary, we find that of all circumstances present in John's case (steak and beans), it is once again only the beans that are absent for both Tom and Mary. Hence, by Rescher's interpretation of the Joint Method, we now have extra good reason to believe that it was the beans that caused the food poisoning.

5. A Preferred Treatment of Mill's Methods

One may point out that the indicated misinterpretations are merely of historical interest, and that for educational purposes, it really doesn't matter exactly how Mill's Methods are presented. However, I argue that many interpretations result in a suboptimal understanding of Mill's Methods, and the lessons they teach us about inductive principles in general. Hence, even for educational purposes, it is important to straighten a number of things out.

In this section I will present my preferred treatment of the Method of Agreement, Method of Difference, and Joint Method. 'Preferred' is here to be understood as being consistent with Mill's writings as well as helpful to beginning logic students in providing some basic insights in Mill's Methods and the subject of induction in general.

5.1 *Mill's Methods as an Investigative Process*

As always, Mill's Methods can be applied to try and find the cause of a certain phenomenon. Now, when presenting Mill's Methods for the first time, it is probably best to understand 'cause' as a necessary and sufficient condition. That is, when looking for the cause of a phenomenon, we are looking for a circumstance in whose presence the phenomenon is present too, and in whose absence the phenomenon will be absent too. This is often the default understanding of 'cause', and for most educational purposes, there is absolutely no need to explicate 'cause' in any kind of technical way. Indeed, the food poisoning example would be a perfect first example to introduce Mill's Methods and to convey the basic ideas behind them.

So, once again, the goal is to find the cause of a certain phenomenon that occurred in a number of instances. This being the goal, I argue that Mill's Methods are best seen as a reflection of the investigative process that follows. More precisely, *rather than being a set of separate methods, Mill's Methods represent the successive stages that one goes through when trying to find the cause of a certain phenomenon.* Below I will illustrate this with the food poisoning example.

Stage 1: What is the first thing that one looks at when presented with a number of students that came down with food poisoning? The obvious first thing to do is

to see what all of these students ate. Thus, if all these students ate a particular food item, then that is reason to believe that that food item was the cause of the phenomenon. In fact, if there is only one food item that all these students ate, then we have even better reason to believe that this food item was the cause of the food poisoning. For example, let us apply this to the food poisoning case:

STUDENT	STEAK?	FRIES?	PASTA?	BEANS?
Carla	No	Yes	Yes	Yes
John	Yes	No	No	Yes

Since both Carla and John ate the beans and, moreover, since the beans are the only food item eaten by both Carla and John, we form the hypothesis that it was the beans that caused the food poisoning. Notice, however, that no mention is made of the fact that Carla and John got food poisoning, since the only students that we look at in the beginning are exactly the ones that did come down with food poisoning. Indeed, the human psychology is such that initially, it simply goes against our intuition to look at students that did not get sick. Step 1 therefore consists of merely looking at the positive instances to set up an initial hypothesis. However, this hypothesis can be strengthened by negative instances in a number of ways, as represented by Stage 2.

Stage 2: At this point, we have a tentative explanation for the food poisoning: it was the beans. However, could this explanation be incorrect? Could there be any alternative explanations? In this stage, we use negative instances to confirm the original hypothesis, and rule out alternative ones.

The obvious way to disconfirm the hypothesis that the beans caused the food poisoning is to find a case in which someone ate the beans, but did not get food poisoning. It is therefore at this point that we realize that we may want to take a look at students that did not get food poisoning. In particular, suppose that we had the following student:

STUDENT	STEAK?	FRIES?	PASTA?	BEANS?	FOOD POISONING?
Tom	Yes	Yes	No	No	No

Since Tom did not eat the beans, the original hypothesis does not get disconfirmed, and thereby actually gets strengthened.

It is important to realize that negative instances can strengthen the hypothesis in another way too. For example, since Tom did eat both the steak and the fries, we have reason to eliminate those food items from our list of possible causes, leaving us with just the pasta and the beans. Indeed, would the beans have been the only item that was not eaten by any of the negative instances, then we would have extra good reason to believe that it was the beans that caused the food poisoning.

A very particular way in which negative instances can strengthen the original hypothesis is as follows. Remembering from Section 2.2 the ways in which Mill's

Methods are not deductive, we realize when looking at Carla and John, that maybe Carla's food poisoning was caused by the fries, and John's by the steak. Thus, maybe there are multiple causes at work. However, the information we have on Tom would disconfirm such an alternative explanation, and hence strengthen the original one. In general, then, Stage 2 strengthens the hypothesis from Stage 1 by finding negative cases where (a) the suspected circumstance is not present, and (b) other potential circumstances are present.

Stage 3: The extra information given on the negative instances helped us confirm our original hypothesis. However, it is still possible maybe the different cases of food poisoning were caused by different food items, or maybe some strange kind of combinatory effect took place that either caused or suppressed the food poisoning. In fact, realizing those possibilities, there is one thing in particular that should make us somewhat nervous, and that is that the eating patterns of Carla, John, and Tom were quite different from each other, as the reader can confirm for his or herself:

STUDENT	STEAK?	FRIES?	PASTA?	BEANS?
Carla	No	Yes	Yes	Yes
John	Yes	No	No	Yes
Tom	Yes	Yes	No	No

In general, the more differences between the different instances, the more possibilities are opened up for alternative explanations. However, this works the other way too. That is, the more comparable the students' diets are to each other, the more confidence we can have in our hypothesis. In particular, consider the following student:

STUDENT	STEAK?	FRIES?	PASTA?	BEANS?	FOOD POISONING?
Mary	No	Yes	Yes	No	No

Since, except for the beans, Mary ate everything Carla ate, we now have excellent reason to conclude that it was the beans that caused the food poisoning.

5.2 The Hierarchy between the Methods

The account of the previous suggests to apply, in order, the following three operations:

Step 1: Agreement in Presence

See whether a certain circumstance is one of the few circumstances present whenever the phenomenon is present.

Step 2: Agreement in Absence

See whether the circumstance that was tentatively marked as the cause of the phenomenon in Step 1 is one of the few circumstances absent in all negative instances.

Step 3: Agreement on other Circumstances

See whether the circumstance that was pointed out as a possible cause by Steps 1 and 2 is one of the few respects in which the positive instances differ from the negative instances.

In the above steps, we can recognize the basic ideas of Mill's Methods. Step 1 reflects the Method of Agreement. Step 2 reflects the Inverse Method of Agreement which, combined with Step 1, gives the Joint or Double Method (same thing!). Finally, Step 3 reflects the Method of Difference.

The reader will undoubtedly notice that all three of the above steps are actually a bit weaker than the original Methods. This is intentionally so. Since we are dealing with induction, the data will often be messy, and the stringent requirements of the original Methods are often not fulfilled. However, we can still draw some useful inferences even when the original requirements are not fulfilled. For example, if there are three circumstances present for all positive instances, then we can at least narrow down our search to those three circumstances, and temporarily disregard the others in our investigative process. If there had only been two, then we would have been more confident, and we would have been maximally confident (at least as the Method of Agreement is concerned) if there was only one. Something similar goes on with the Inverse Method of Agreement: the fewer circumstances there are which are absent for all negative instances, the more confidence we have for each of those that they are causally connected to the phenomenon under investigation. And for the Method of Difference it holds that the more similar the instances are, the more confidence we have in our conclusions when comparing them. Thus, as with all of induction, our hypotheses come with a certain degree of confidence, and the requirements on Mill's original Methods merely represent an interesting limit case, although certainty is never obtained. It is therefore important to reflect those other gradations and possibilities too, and that is what the above three steps represent.

The different degrees of confidence are not contained to each individual step. The cumulative nature of the evidence gathered by the three steps suggests that the Method of Agreement gives the weakest support, the Method of Difference the strongest support, and that the strength of the support of the Joint Method lies somewhere in between these two. This is a healthy way to think about these Methods, since each additional step reveals more similarities between the positive instances on the one side, and the negative instances on the other side. And obviously, the more consistent the instances reveal a certain pattern, the more confident we can be about our conclusion. Maximal confidence is obtained in the case of total agreement except for one circumstance.

This particular treatment thus places the Method of Agreement, the Joint Method, and the Method of Difference in a helpful conceptual hierarchy. The Method of Difference is reflected as a stronger method (i.e., a method giving a stronger conclusion) than the Joint Method, while the Joint Method on its turn

subsumes the Method of Agreement. Notice that this actually makes the Method of Agreement part of the Method of Difference, which is just one more reason why any combination of the Method of Agreement and the Method of Difference does not make much sense in the first place. Another reason is that the Method of Difference is already seen as giving the strongest possible evidence for the conclusion, so combining it with any of the other methods seems quite useless. The best way to think about the Method of Difference is as a controlled experiment, and controlled experiments gives one about as good as support as one can get. This is an important insight, and to obtain that insight, the above account should prove very helpful. Moreover, such an account seems precisely what Mill had in mind:

It thus appears to be by the Method of Difference alone that we can ever . . . arrive with certainty at causes.

. . . The Method of Agreement is chiefly to be resorted to as a means of suggesting applications of the Method of Difference, . . . or as an inferior resource in case the Method of Difference is impracticable.

. . . This indirect method, therefore, can only be regarded as a great extension and improvement of the Method of Agreement, but not as participating in the more cogent nature of the Method of Difference. ([12], Book III, Chapter VIII, §4)

Also notice the following comment made by Stebbing in her book *A Modern Introduction to Logic*:

It is a mistake to regard the Joint Method as a more satisfactory form of the Method of Difference, although logicians not infrequently make this mistake. The Joint Method is to be employed only when the more stringent requirements of the Method of Difference cannot be secured. Mill himself clearly recognized this. ([18], p.338)

5.3 A Note on 'Agreement'

The Method of Difference is usually applied to one positive instance and one negative instance. However, the Joint Method seems to require at least two positive instances to apply the Method of Agreement on, and at least two negative instances to apply the Inverse Method of Agreement on. So how is it possible that the Joint Method is subsumed by the Method of Difference, as claimed in the previous section? The answer to this question is that there are two ways in which we can interpret 'agreement' when talking about Mill's Methods:

1. First of all, with 'agreement' can be meant an agreement *between* instances. Thus, for example, we can apply the Method of Agreement when all positive instances 'agree' on having a certain circumstance present.
2. Second, 'agreement' can also be understood as an agreement *within* instances. Thus, for example, we can apply the Method of Agreement when it holds that there is an 'agreement' between the presence of the phenomenon and the presence of the circumstance.

The Method of Agreement is normally understood in the first sense. Indeed, the Method of Agreement is usually illustrated with two or more instances that are similar in a certain respect. This is an important way to think about agreement since, as argued in the previous section, it is the similarities between the instances that do a lot of work. Therefore, although one technically could apply this method to a single instance, that would be quite counterintuitive. In fact, some textbooks explicitly require two or more instances, and so did Mill in his original definition (see Section 2.3).

However, understanding ‘agreement’ in the second, alternative sense, is important too. Thinking about agreement in that way, the Method of Agreement can actually be applied to a single instance in a very natural way. For example, suppose that we had the following student:

STUDENT	STEAK?	FRIES?	PASTA?	BEANS?	FOOD POISONING?
Heidi	No	No	No	Yes	Yes

Heidi got food poisoning, but the only thing she ate was the beans. Hence, it makes perfect sense to conclude that it was the beans that caused the food poisoning, even though there is only one instance on which we base this conclusion. Indeed, in this case the Method of Agreement works by noting the ‘agreement’ between the presence of the beans and the presence of the food poisoning.

Of course, the conclusion we draw in relation to Heidi is very much dependent on whether or not all suspicious food items were indeed listed. This is even more apparent when we consider the following case:

STUDENT	STEAK?	FRIES?	PASTA?	BEANS?	FOOD POISONING?
Skip	Yes	Yes	Yes	No	No

In this case, we want to conclude that it was the beans that caused the food poisoning since, apparently, it was not the steak, fries, or pasta. We can do so using the Inverse Method of Agreement, with ‘agreement’ once again understood as an agreement within, namely as an agreement between the absence of the beans and the absence of the food poisoning.

These last two examples are very educational. They make some of the background assumptions of Mill’s Methods quite explicit, and they reveal an alternative way of looking at the notion of ‘agreement’ which, on its turn, explains why the Method of Difference does subsume the Joint Method after all. For example, consider the following:

STUDENT	STEAK?	FRIES?	PASTA?	BEANS?	FOOD POISONING?
John	Yes	No	No	Yes	Yes
Tom	Yes	Yes	No	No	No

Since we can use the Methods of Agreement on a single instance, we realize that we can use the Joint Method to conclude that it was the beans that caused the food poisoning. However, we would have been even more confident of our conclusion had John eaten the fries, or Tom not eaten the fries. Indeed, in that case, we could have applied the Method of Difference.

5.4 Advancements

In the previous sections I have indicated what I believe to be a preferred treatment of Mill's Methods, as it illuminates the basic intuitions that ground these Methods. The most important intuition is that a causal connection can be found between circumstances and phenomena by finding certain regularities and similarities between the different instances. And, the higher the degree of similarity and regularity, the greater the confidence in the conclusions that we draw. It is exactly these kinds of intuitions that are reflected by the indicated hierarchy between the different methods. Hence, for educational purposes alone, the Methods should be presented as such.

Within the context of this hierarchy, however, details may vary from presentation to presentation. That is, as long as the basic hierarchy is maintained, one may want to play with the exact restrictions under which one can apply a certain Method, or extend the Methods in various ways. This section presents some ways in which such variations and advancements can be made. Most of these variations have been encountered at some place or other in this paper, but the instructor of Mill's Methods may find it helpful to have them all in one place.

First of all, one could use the strict requirement of having exactly one circumstance present in all positive instances in order for the Method of Agreement to apply, or use the weaker restrictions as I argued for in Section 5.2. In fact, one could very well loosen these restrictions even more. For example, one may want to apply the Method of Agreement even when not all positive instances have a certain circumstance present. Indeed, it is perfectly possible for some student to get food poisoning without having eaten the beans or, conversely, maybe some student ate the beans but did not get sick. However, notice that even in these cases the confidence of our conclusion comes with a certain degree of confidence, depending on the percentage of instances that reveals a certain pattern, and thus ultimately depending on the similarity between the different instances as usual. Hence, for educational as well as practical purposes, I would actually recommend such a move. In fact, one possible strategy that may help to convey these general notions of induction is to present Mill's Methods under the most stringent restrictions at first, and to gradually weaken these restrictions in the suggested ways.

The instructor may also want to pause at the number of instances on which the Methods are being applied, for it is a general inductive principle that the more instances are being used, the more confident we become in our conclusion. The case of using a single instance is instructive in its own respect as illustrated in

Section 5.3. Conversely, the question can also be taken up to which instances the conclusion is supposed to pertain. Does it make a claim about one particular instance (see Section 4.2), some of the given instances (e.g. those that reveal a certain similarity), all of the given instances, or does it even pertain to instances not listed? Here, of course, the confidence goes down as we increase the number of instances that we want our conclusion to be about, which is another important general principle of induction.

When discussing Mill's Method, some textbooks make an explicit distinction between necessary and sufficient conditions (see Sections 2.3 and 2.4). Now this is certainly an important distinction that can be nicely demonstrated by applying Mill's Methods to a few well-chosen examples. However, I think that it is important to realize that one can easily apply Mill's Methods without at any point being explicitly concerned about this distinction. To see this, observe the following. One *could* say that eating the beans was merely a necessary condition of the food poisoning, in that the students also would have had to have the right digestive and other bodily properties in order for the food poisoning to take place. However, such a remark seems rather strange in that we would simply assume such properties to be there unless we were told otherwise. In fact, we regard eating the beans alone as a sufficient condition for the food poisoning. In general, then, there are almost always certain conditions that need to be satisfied in order for the phenomenon to take place, but that, under normal circumstances, are in fact satisfied. Indeed, the circumstances that we do list when applying Mill's Methods are ones that are thought to possibly effect the phenomenon in a more interesting way. It follows, then, that any kind of interesting distinction between necessary and sufficient conditions will have to pertain to those circumstances listed. In other words, the distinction only comes up when we start to consider the possibility of a combination of listed circumstances causing the phenomenon. Such combinations, however, are a rather special case. Indeed, for most applications of Mill's Methods, it is a single circumstance that can be marked as the cause. Hence, one may want to wait contemplating the distinction between necessary and sufficient conditions until one has a good grasp of the more basic concepts of Mill's Methods.

Still, one may want to consider the possibility of having combinations of circumstances being the cause of the phenomenon. At the same time, one could also consider multiple causes. Or one can do what some textbooks already do, which is to consider the possibility of the absence of a certain circumstance being the cause of the phenomenon ([8], [10]). Now, in these three cases, one may recognize the three basic logical operations of conjunction, disjunction, and negation, respectively. Indeed, Mackie gives a detailed account of the possibility of any kind of logical combination of circumstances being the cause of the phenomenon ([11], Appendix). This was also anticipated by Skyrms ([17], pp.106-110). It should be noted, though, that such a technical advancement only works under very stringent assumptions. Such advancements that allow one to deal with the aforementioned possibilities are therefore probably better left to the highly advanced student.

6. Conclusion

I have argued that a number of the most widely used logic textbooks have misinterpreted Mill's Joint Method of Agreement and Difference. Instead of viewing the Joint Method as some kind of combination of the Method of Agreement and Method of Difference, Mill most likely regarded the Joint Method as the Double Method. We have seen the following evidence for this:

1. Mill's writings about the Joint Method are, aside from the definition, in perfect agreement with the Double Method (Section 3.1).
2. Mill's symmetrical phrasing of the definition was probably intended to capture the highly symmetrical nature of the Double Method (Section 3.2).
3. The Joint Method as described by Mill was generally understood as the Double Method by logic textbooks published in the first half of this century (Section 3.3).
4. The Double Method is stronger than the Method of Agreement, but weaker than the Method of Difference. This is exactly how Mill talked about the Joint Method (Section 5.2).

The Joint Method, when interpreted as the Double Method, forms a conceptually helpful hierarchy with the Methods of Agreement and the Method of Difference. Therefore, such an interpretation is not only historically more accurate, but also represents a better conceptual introduction to Mill's Methods and the principles of induction in general. It is therefore preferred in two ways.

Acknowledgments

I would like to thank the two referees of an earlier version of this paper, for their helpful comments and suggestions. I particularly want to thank Eric Dietrich for the many ways in which he made this paper possible.

References

- [1] R.B. Angell, *Reasoning and Logic*. Meredith Publishing Company, 1964.
- [2] S.F. Barker, *Logic*. McGraw-Hill, 5th ed., 1989.
- [3] M. Cohen and E. Nagel, *An Introduction to Logic and Scientific Method*, Harcourt, Brace and Company, 1934.
- [4] I.M. Copi, *Introduction to Logic*. MacMillan, 1st ed., 1953.
- [5] I.M. Copi, *Introduction to Logic*. MacMillan, 2nd ed., 1961.
- [6] I.M. Copi and C. Cohen, *Introduction to Logic*. Prentice-Hall, 10th ed., 1998.
- [7] R. Holmes, *The Rhyme of Reason*. D. Appleton-Century Company, 1939.
- [8] P. J. Hurley, *A Concise Introduction to Logic*. Wadsworth Publishing Company, 6th ed., 1997.
- [9] S. Jevons, *Lessons in Logic*. MacMillan, 1919.
- [10] H. Kahane and P. Tidman, *Logic & Philosophy: A Modern Introduction*. Wadsworth Publishing Company, 7th ed., 1995.

- [11] J.L. Mackie, *The Cement of the Universe*. Oxford University Press, 1974.
- [12] J.S. Mill, *A System of Logic*. Longmans, Green and Co., 1941 (reprint).
- [13] W. Minto, *Logic: Inductive and Deductive*. Charles Scribner's Sons, 1904.
- [14] N. Rescher. *Introduction to Logic*. St. Martin's Press, 1964.
- [15] J. Rudinow and V. E. Barry, *Invitation to Critical Thinking*. Harcourt Brace, 1999.
- [16] J. Skorupski, *John Stuart Mill*. Routledge, 1989.
- [17] B. Skyrms, *Choice and Chance*. Dickenson Publishing Company, 1966.
- [18] L.S. Stebbing, *A Modern Introduction to Logic*. Methuen & Co., 1961 (reprint).
- [19] G.H. Von Wright, *Treatise on Induction and Probability*. Littlefield, Adams & Co., 1960.
- [20] P. Wheelwright, *Valid Thinking*. The Odyssey Press, 1962.

*Bram van Heuveln, Program in Philosophy, Computers and Cognitive Science
Department of Philosophy, Binghamton University
Binghamton, New York 13902 U.S.A.
bram@turing.paccs.binghamton.edu*