

WHEN IS IT RESPONSIBLE TO GENERALIZE FROM A SINGLE INSTANCE?

*

DAVID BOTTING

Independent scholar
davidbotting33@yahoo.co.uk

Although large samples are always better than small samples and always confer greater justification on general claims than small samples, this paper argues for situations where even a small sample can justifiably be thought to be representative of the population and we are justified in believing, or having a pro-attitude towards, a general claim by generalizing from it. It is not fallacious to make inductive inferences in these situations from small samples. I will describe three such scenarios. On the other hand, when it cannot justifiably be thought to be representative of the population then it is always fallacious, irrespective of other considerations. I will describe one such scenario in which generalizing from a small sample has been claimed to be justified on the grounds of cognitive economy and will show that this claim is false unless the scenario reduces to one of the first three. Since generalizing from a single instance is a limiting case of generalizing from a small sample, I will focus on generalizing from a single instance. Whatever can be shown with regard to a single instance follows a fortiori for all small samples. As it turns out, it is very difficult for a reasoner reasoning in good conscience to commit to a fallacy of hasty generalization, and if a fallacy can only be committed by insincere reasoners and cannot be made by reasoners reasoning in good conscience, then it is not very interesting, as it is not really an error in reasoning. The reason for this is that for the reasoner to have reasoned fallaciously, they must have knowingly ignored evidence, and this is something that a sincere reasoner is not likely to have done. Equally, the charge that someone has committed this fallacy is not easy to substantiate and amounts to accusing the reasoner of insincerity.

Keywords: generalization, the fallacy of hasty generalization, induction, cognitive economy, disagreement.

1. Introduction

I am not offering any new analysis of the fallacy of hasty generalization in this paper. My aim is to show that it is difficult for a reasoner to commit this fallacy, and difficult likewise for a critic to justify an accusation that the reasoner has committed this fallacy. I will show this even for the cases where it seems like such an accusation is most justified, namely where the reasoner has generalized from a single instance. I will show through examples where generalizing from a single instance have been taken as paradigm cases of hasty generalization that the accusation that the reasoner has committed a fallacy of hasty generalization is not justified, provided only that some charity is given to the reasoner's inference and some attention shown to what attitude is being taken by the reasoner to be justified by the inference. I hope to show that these generalizations illustrate epistemically responsible behaviour, despite their being generalizations from a single instance. I will then show that when a critic accuses a reasoner of committing this fallacy, it often turns out that his reasons do not justify making this accusation but instead illustrate a substantive difference of opinion over the evidential basis of the generalization. When two reasoners have such a difference of opinion, neither is justified in accusing the other of generalizing hastily or fallaciously.

In the first section of the paper, I will describe generalization, and the inductive arguments that express the logical form of generalization, in general. I will be making two distinctions: between universal and statistical generalizations, and between probabilistic and proportional claims.

In the second section I will describe the fallacy of hasty generalization and will discuss three examples that have been given in the literature to illustrate this fallacy. These examples all involve generalizing from a single instance. I will argue that these examples only count as fallacious if one uncharitably interprets the conclusion as making a stronger claim than it is really making, and that in fact in all three examples the conclusion, when interpreted charitably, is one that it is epistemically responsible to draw. Although all three are generalizations from a single instance, I will show that none of them is fallacious. I will then give what I think *is* a good example.

Next, I will consider a case where, conversely, it is claimed that what appears to be a hasty generalization is not fallacious because the generalization is the outcome of the cognitively most efficient reasoning process available. I will show that this "argument from cognitive costs" is not sound. Unless the reasoner is in one of those conditions already identified where a small sample can justifiably be believed to be representative, generalizations are hasty and fallacious, irrespective of the cognitive economy. One cannot generalize from what one has no reason to believe is a representative sample and be justified in believing the general claim so inferred just because getting

a larger sample is cognitively expensive. Perhaps some kind of pro-attitude towards the general claim can be justified on these grounds, but not belief.

After this, I will describe how our beliefs about the evidence affect what it is epistemically responsible to infer and what needs to be the case for an accusation of inferring fallaciously to be justifiably made by an accuser. I will show that for a charge of fallacy of hasty generalization to be justified the reasoner must be drawing his inference from the wrong body of evidence and be culpable for taking this as his body of evidence; he must be *knowingly* disregarding something he himself takes to be evidence, or to have blinded himself to any such evidence. Knowingly disregarding evidence is not something that a sincere reasoner will ever do, and so this is not a fallacy that a sincere reasoner will ever commit; hence, an accusation against a reasoner that he has committed this fallacy is *ipso facto* an accusation that the reasoner is not sincere, and the burden of proof is on the accuser to substantiate such an accusation. Reasoning fallaciously cannot be attributed to the arguer on the basis of the form of the argument alone. To be interesting, a fallacy must be an error in reasoning that a reasoner could commit in all good conscience. It transpires that the fallacy of hasty generalization is not an error in reasoning that a reasoner could commit in all good conscience.

In short, the evidence must be evidence that the reasoner already has, or be in an epistemic situation where he ought to have it, so that it would be epistemically irresponsible for him not to have it and not to include it among the premises of the inductive argument. If it is not cognitively available to the reasoner and we charge him with fallacy, this is not really justified; rather, we are saying that we have a different body of evidence that leads to a different conclusion, in which case neither of us have reasoned badly, and the task is to convince the other to accept what we have taken as evidence. The matter is settled by a substantive debate on the evidence and making the evidence cognitively available to one another so that the other has it or ought to have it, and not by accusation and counter-accusation.

In the conclusion I will show that the burden of proof that the accuser needs to meet to justify his charge that another has committed the fallacy of hasty generalization is not an easy burden to meet, even in many cases where it seems easiest to meet, namely when the generalization is from a single instance.

2. Generalization

What is a generalization? When do we generalize?

In this section I want to argue that two different kinds of general claim, namely proportional claims and probabilistic claims, can be justified on the basis of instances.

(I will be using a frequentist interpretation of probability here.) Both kinds of claim make presuppositions that the reasoner must be justified in believing in order for their generalization to be justified: for proportional claims it is the claim that the sample is representative of the population, for probabilistic claims it is the claim that the sample has converged on a limit.

This section can be skipped by those readers familiar with the distinctions I am explaining here, or those who are only interested in the question whether proportional claims can be made justifiably from small samples. Its aim is merely to circumvent one apparently knockdown argument against generalizing from a single instance: for a probabilistic claim to be justified, there must be enough instances for the reasoner to believe that the sample has converged, and it is difficult to see how this can ever be justified if the sample consists of only one instance. A sample of one cannot be convergent. But a sample of one can be representative of the population, and there are situations where we can justifiably believe that a sample of one is so representative, and justify our generalizing from that single instance, taking “generalizing” here to mean inferring to a proportional claim.

Does this mean that only proportional claims are justified, and probabilistic claims are not? I do not think so, as the proportional claim can be taken to be evidence for a probabilistic claim: if I justifiably conclude that 80% of the B’s are A, I am also justified in believing that the probability of picking an A from the B’s is 0.8. However, I also wish to make the technical point that this inference from the proportional claim to the probabilistic claim (and *vice versa*) is inherently risky: the probability may not, in fact, be 0.8. This is not due to the fact that most inductive arguments are imperfect, because there is risk even in the case of perfect inductions.

Those willing to accept these points may proceed straightaway to section 3.

a. Two forms of inductive argument

At its simplest, a generalization occurs when we say that what is true for a sample of A’s is true of all A’s. To put it symbolically, it is the inferential step between the singular claims of particular A’s being B and all A’s being B, i.e.,

A_1 is B

A_2 is B

·

·

A_n is B

Therefore, $\forall x.A(x) \supset B(x)$

This is an inductive argument: its conclusion is a universal material conditional, denoted by the universal quantifier. If $A_1 \dots A_n$ is a complete enumeration of all the A's then we have a "perfect" induction and it is logically impossible for there to be a case where the premises are true but the conclusion false. In a perfect induction, the universal conditional follows conclusively, although not formally.¹

In the general case of induction, the enumeration is incomplete: the A's referred to in the premises are a sample taken from a wider population. In this case it is logically consistent for the premises to be true and the conclusion false, which is to say that the universal material conditional is not conclusively established to be true and that therefore inferring that it is true is ampliative, i.e., it goes beyond the evidence expressed in the premises.² These are "imperfect" inductions. As in all ampliative inferences, there is always a risk involved in imperfect inductions, because you know that what you infer may be false: in the case of inductive inferences this is called the "inductive risk." Despite this risk, we often rightly consider ourselves to be justified in believing that the conclusion is true when the evidence confers strong enough justification while conceding that further evidence could make us change our mind: if we assent to those premises and take those premises to express everything that we know to be relevant to the conclusion, then we are justified in provisionally believing the conclusion.³ Such arguments are neither semantically nor formally valid, but are often called "inductively valid."

This form of the inductive argument – where the conclusion is a universal material conditional – survived for a surprisingly long time. Probability theory had already been established (though still in its infancy) when DeMorgan noted that the inductive inference was just a special case of calculating the "inverse probability." This led to a widening of the conception of an inductive inference to any inverse probability. In this second, more sophisticated form of the argument we have statistical evidence

¹ The inductive argument is semantically valid but not formally valid, as there is no case where the premises are true and the conclusion is false.

² In a sense, though, what we infer is not that it is true but that it has a certain probability of being true given the premises, and this is not ampliative — the probability in question must, as a matter of necessity and deductive certainty, have the value that it has. Only if we miscalculate what this probability should be can there be an error in an inductive argument, as will be explained later.

³ May there be cases where the epistemically responsible thing to do is to have no beliefs either way? A conclusion that is made highly likely relative to one set of premises may be made highly unlikely relative to other premises, where either we do not know what those other premises are, or we know what they are but do not know whether they are true or false. Since we know of this possibility, we must be careful about drawing inductive inferences from premises, even when they do confer high justification on the conclusion, since we do not know that there is not another argument that confers high justification on the conclusion's falsity. Referring to the justification conferred as "strong enough" is to be understood as saying that it is strong enough despite these issues, strong enough to justify provisional belief and not just a waiting game of having no beliefs.

of the probability distribution of an attribute in a sample which is expressed in the premises, and we infer that the probability distribution of that attribute will be approximately the same in the population from which the sample was drawn as in the sample itself. It is not necessary that the attribute belong to *every* member of the sample or the population, or to put it another way, that the probability be one or zero; we conclude with a general, but not universal, claim.

Here is an example of the second form: suppose that in a sample of n A's there are m A's that are B and $n-m$ A's that are not B. We can express this new inductive argument this way:

A_1 is B

A_2 is B

.

.

A_m is B

A_{m+1} is not B

.

.

A_n is not B

Therefore, $p(A,B) = m/n$

The conclusion effectively states that the probability of an A being B in the population is the same as its probability in the sample.⁴

In a sense, the first form of inductive argument is just the special case of the more general form where $m=n$. However, we must be careful. Despite superficial appearances, $\forall x.A(x) \supset B(x)$ does not mean the same thing as $p(A,B)=1.0$; the universal quantifier makes a general claim about everything it quantifies over (i.e., A's), where this set may be finite or infinite, whereas the probability claim is a singular claim about the set of events in which a B is selected and the set of events in which an A is selected, where these are infinite sets. To be more specific, it claims that the frequency of selecting A's that are B is in a particular ratio to the frequency of selecting A's *simpliciter*. The claims do not mean the same thing and talk about different sets.

Nor can either one be inferred conclusively from the other. However, the truth of either claim can be used as evidence for the truth of the other. Let us go back again to the perfect induction. Here the universally quantified conclusion is established

⁴ This argument might describe sampling with replacement or sampling without replacement. Relatedly, the conclusion may be read either as a probability claim or as a proportional claim. Despite their similarity, and the fact that there appears to be no difference in terms of logical form, the differences are significant, as will be explained in detail shortly.

conclusively, but $p(A,B)=1.0$ is not established conclusively. The reason for this is that to establish $p(A,B)=1.0$ conclusively what we need to completely enumerate is not A's but selectings (with replacement) of A's, and this cannot be completely enumerated because it is by definition an infinite set. There is always a risk then in making a probability claim, since it can never be conclusively established or verified, even by a perfect induction (by selecting without replacement), whereas a claim that a particular proportion of A's are B is conclusively established by a perfect induction without any risk.⁵

Can we not say that, supposing that we have established conclusively the proportional claim that all A's are B, it is simply impossible to select an A that is not a B (since there are none), and this establishes the probability claim conclusively? This follows only if we can assume that the A's that are B will always be B and cannot become not-B.⁶ So, a complete enumeration is the strongest evidence we may have that $p(A,B)=1.0$, and certainly justifies making this probability claim, but it is still an ampliative inference that does not establish the probability claim conclusively. Similarly if, in the complete enumeration of n A's, m are found to be B's, this is strong evidence for the probability claim $p(A,B)=m/n$ but it does not imply it conclusively. If the enumeration is imperfect then the evidence is weaker, but otherwise it works in the same way.

Hence, it is possible for the universal material conditional to be true but the probability claim to be false. Equally, it is possible for the probability claim $p(A,B)=1.0$ to be true while the universal material conditional $\forall x.A(x) \supset B(x)$ is false. The reason for this is that the value of the probability is a point of convergence, that is to say, the frequency ratio in the long run, and this in itself does not rule out selecting an A that is *not* B. For example, take the prime numbers $\{\text{prime}_1, \text{prime}_2, \dots, \text{prime}_\infty\}$. There is an infinite number of prime numbers, but the frequency of their occurrence is a decreasing function of the natural numbers. Since the denominator in the frequency series is just an incrementation function of the natural numbers, from the facts that for some natural numbers, namely the prime numbers, in the corresponding selecting the attribute does not occur, but that for all non-prime numbers, in the corresponding selecting the attribute does occur, it follows that the probability will still be one. Consider the frequency series

$1/1, 1/2, 1/3, 2/4, 2/5, 3/6, 3/7, 4/8, 5/9, 6/10, 6/11, 7/12, 7/13, \dots$

⁵ This is the only time a conclusion is established non-provisionally in an inductive argument — whether or not the conclusion is universal does not make our belief any more or less provisional.

⁶ For essentialists, there are some A's and B's for which this *is* a safe assumption to make. Although we may disagree with them, it seems justified to accuse them of reasoning badly only if we can find some *internal* inconsistency in their position.

Whenever the denominator is non-prime (1, 4, 6, 8, 9, 10, 12) there is a positive instance, e.g., a B has been selected from the A's. Whenever it is prime (2, 3, 5, 7, 11, 13) there is a negative instance, e.g., a not-B has been selected from the A's. Early in the series there are roughly as many prime numbers as non-prime numbers, just as in the above the ratio after 13 selectings is near 0.5. After 1000 selections the ratio would be close to 1.0, due to the relative infrequency of primes among higher numbers. In the limit, then, the ratio will converge on 1.0. (It is worth noting, though, that however far we actually continue the frequency series, the final value will always be less than 1.0, that is to say, less than the value on which the series converges.)

For this reason a probability claim "All philosophers are bearded", meaning here $p(\text{Philosophers, Bearded})=1.0$, may be true while as a universally quantified claim it may be false. So, the inference from the probability claim to the universally quantified claim is also defeasible, although the circumstances under which the universally quantified claim does not follow are clearly bizarre and do not much undermine our justification for making such a claim; clearly, if we have already encountered a negative instance, we will not infer that the universally quantified claim "All philosophers are bearded" is true but rightly conclude that it is false.

When we sample *without* replacement, we tend to make a proportional claim. Suppose that there are ten philosophers in the room, and eight of them are found to be bearded. We can then conclusively infer the proportional claim that 80% of the philosophers *in the room* are bearded (having completely enumerated the philosophers in the room), but if we take this as a representative sample of philosophers *simpliciter* and infer the proportional claim that 80% of philosophers are bearded, this inference has an inductive risk. This is no different from the case of making a universal claim.

If necessary we can infer the probability claim from the proportional claim, where the probability claim is effectively a (counterfactual in this instance) claim about what we would have observed had we sampled *with* replacement. Just as I showed earlier that there is a difference between a probability of one and universal quantification, the probability claim does not mean the same as the proportional claim, though any proportional claim has a probability claim corresponding to it: on the basis of the proportional claim, a frequency claim is made that says something analogous to "Put all the philosophers into a room. Mix them up and then choose one at random. If it has a beard, make a mark "B". Replace the philosopher and repeat. The relative frequency of "B" in our pickings will converge around 0.8. It will not be exactly 0.8, but its variation from 0.8 can be made arbitrarily small by increasing

the sample size.” What is being made is a singular claim about two infinite classes (and about the whole of those classes and not about a proportion of them); it is not, strictly speaking, a general claim, though general claims concerning the members of the classes can be made on their basis.

Mutatis mutandi, sampling with replacement justifies a probability claim from which we infer a proportional claim that says what would have been observed had we sampled without replacement. Suppose that we make a probability claim that the probability of selecting B's from A's is 0.8. From this we can infer that, had we not replaced the A's that we selected and continued this until there were no A's left, there would have been found to be (roughly) 80% of B's.

This probability claim, made by sampling with replacement, is a frequency claim. The grounds of claims about infinite sets is evidence about finite sets, namely, the relative frequency in the sample so far observed—the so-called ‘practical limit’. The limit in the finite sample we have observed so far may not be the limit in the infinite series; the grounds support the hypothesis only on the assumption that the sample has converged on its final limit. The evidence itself says nothing about this assumption, so the probability claim presupposes something about the future for which evidence is lacking. Thus, although inductive generalizations have semantic content concerning the future, there are no grounds – inductive or otherwise – for that content; they are corollaries of the assumption that the world is predictable.⁷

We see from this that asking someone to make a probability claim is to ask a question that whoever answers the question can only answer with confidence either when they take themselves to be justified in believing that the point of convergence has been reached (and, that this is not likely to be due to bias, e.g., the selectings from the population are not genuinely random), or when they are justified in believing a proportional claim from which the probability claim can be inferred. If the answerer does not believe that the frequency series will converge, or that the point of convergence has not been reached in the sample, then he cannot answer, and a questioner who insists on an answer can hardly complain when the answer gotten in return involves hasty generalization. A charge by the questioner that the answerer has committed a fallacy would not be deserved — it would wrongly imply

⁷ This is often put as the claim that induction can be shown (mathematically) to work if anything does. What has not been shown, however, is that it is true that anything does. This is an assumption, made largely on the basis of pragmatic reasons, namely that we do best, or at least not worst, when we make this assumption. Also, the Straight Rule of Induction is not the only rule that can be shown to converge in the limit, and although in the long term they will all give the same value for the limit, in the short term (in which we must operate) different rules will give different values as practical limits. It is well outside the scope of this paper to discuss these complications.

3. Hasty Generalization

When we criticize the generalization as being “hasty,” we are saying that making a general claim about the whole population, whether in the form of a proportional claim or a probability claim, is not justified by the evidence given in the premises. Usually, our ground for saying this is that the sample size n is too small. This is especially so when we generalize from a single case — this is a hasty generalization if anything is, and so a perfect test case for discussion.

Here is an example from Engel (1976, p. 69): “I had a bad time with my former husband. From that experience I’ve learned that all men are no good.” Johnson and Blair (1994, p. 70) give an example of someone concluding that Calgary is not a friendly city from a bad experience at Calgary Zoo. It is easy to see how such experiences can prejudice the one who has a bad experience against a whole of which the offender is only a part. These are inductive arguments based on a single instance. Groarke and Tindale (2004, p. 287) note that one good experience, such as a carpet cleaning company doing a good job of cleaning a carpet, may convince us that that company always does a good job. In such cases the evidence is usually anecdotal.

However, I am tempted to think that these examples make straw men of the reasoners accused. Granted: a *belief* that all men are no good is not justified on the basis of a single experience of a single man, or even many experiences of a single man, for that would show only that that particular man is no good and not that all men are no good. But is this what the argument is actually saying here? Rather, I think that what is being concluded is that, although there certainly are counter-instances out there, and therefore that the general claim *could* be false, this provides no particular reason to think that some instance will be one of them, and so it would not be wrong to draw the same inference for that particular person as one would if the universal generalization were literally true. Just because something might be true – the next man may be the source of a more positive experience – does not mean that one should behave under the assumption that it is. It follows that “Men are no good; therefore, Derek the man is no good” is a perfectly good inference when interpreted charitably and supposing that we have not suppressed any information that we know about Derek, and not, as it is often accused of being, a fallacy of sweeping generalization (i.e., that of drawing a conclusion about a particular from a generalization that has exceptions). Equally, the general claim licensing this inference is justified as an assumption and as a general rule to be applied.⁸

⁸ I deal with the fallacy of sweeping generalization in another paper (Botting, 2017).

It would be a fallacy if the woman concluded that all men are no good as an exceptionless generalization, which is not the case here, but it is not necessarily a fallacy to generalize from a single instance when the generalization concluded is one that has exceptions. Taking “all men are no good” as this kind of generalization, I think that it might be quite justified to draw this conclusion, and equally justified for the woman to draw the conclusion that some different man (other than the woman’s husband) is also no good on the basis of this generalization. My general point is that, assuming the exceptions are in the minority, there is good reason to think that a randomly selected individual is not an exception, and if there is no further reason to think that they are an exception (e.g., we do not know anything more about the individual other than his being male), the inference to and from the generalization are just as justified as they would be if, in fact, the generalization were exceptionless. The mere possibility of exceptions does not make us reason differently (from when we reason with exceptionless generalizations) and does not mean that we are unjustified when we reason that way. Moreover, if we *do* know that the individual is an exception, then we simply do not apply the general rule (i.e. the generalization). If we know that Derek is an exception to “All men are no good” then this amounts to knowing that Derek is good, so we would not then reason “(With exceptions) all men are no good; but, Derek is an exception; therefore, Derek is good”. This would be circular as the premise that Derek is an exception and the conclusion that Derek is good mean the same thing here. Hence, in the cases where we do reason using a generalization with exceptions, the generalization justifies the same inferences as it would if it were an exceptionless generalization, and this means that it is no fallacy either to generalize this way.

Now, it might be claimed that it is unjustified to make even a weak generalization with exceptions on the basis of a single bad experience—taking the generalization as having exceptions should not be taken as an excuse for a free-for-all in which there are no standards applied whatsoever, as I would not be the first to point out. This should not, though, lead us to dismiss the woman’s argument too quickly, especially when the attitude the woman takes is not belief but some kind of practical attitude. When it is a practical attitude, whether it is justified to have it depends in the end on a utility calculation: if the good experience is potentially good enough and the bad experience not too bad then one may be more prepared to give “the benefit of the doubt,” though without necessarily *expecting* a good experience or *believing* that it will be one; one is prepared to look for counter-instances, so to speak. Contrariwise, if the bad experience is bad enough then one will act under the assumption that the universal generalization is true, without believing that it is literally true, and one certainly will not go looking for instances either to confirm or falsify the generalization

but simply avoid those situations altogether as far as possible. I do not see how this is irrational; if it were, word-of-mouth recommendations would be next to worthless, and I would be no more justified in using the same carpet cleaning company next time in preference to any other I may choose at random from the phone book. Without some kind of idea of these utilities, these cases are under-described. To put it another way, “all men are no good” is a figure of speech that is not charitably interpreted as asserting literally that all men are no good or as justifying a belief that all men are no good, even taking the latter as a generalization with exceptions; it is a signal instead of her unwillingness to give the benefit of the doubt to some arbitrary instance of manhood that she knows nothing more about, and I think it presumptuous to suppose that she is wrong to do so or that doing so involves some kind of cognitive error. It is, then, justified for her to draw a general conclusion from a single instance, provided that she does not overstate it. A similar analysis applies to the case of the zoo visitor. Neither of these cases are fallacies, because their conclusions are not to be interpreted as expressing justified beliefs, although they are attitudes that it is rational to have in the circumstances and having taken the utilities into account.

I am not entirely convinced by the carpet cleaning example either, although for different reasons: with some things, any arbitrary instance can be considered to be average, and likewise whatever is true of the average will be true of the great majority of the population. At this point the phenomenon known as reversion to the mean might be mentioned. Consider the carpet cleaning company again. Suppose that one in ten of their cleanings is below average, one in ten is above average, leaving eight in ten as average. Suppose your carpet is cleaned and they do a good job. You do not and cannot know, on the basis of a single instance, whether their performance is above, below, or just average on this occasion, and it would certainly be wrong to believe that *all* of their carpet-cleanings are good. But again, this seems to interpret the conclusion of the argument very strongly.

Put it this way: does your knowledge that there are things that you do not know, and perhaps could know if you had a bigger sample, mean that you should not, in the sense of its being epistemically irresponsible, generalize now? That you should not generalize until you have been able to determine what an average performance is and what is below and above average for their standards? I do not think so. Most of the time performances are average because this is what it means to be an average performance, so any arbitrary instance is justifiably believed to be average. That we could be even more justified if we had a bigger sample does not mean that we are not justified now or that it would be fallacious to generalize from a smaller sample or even a single instance. Whether it is epistemically responsible to generalize now or

to wait depends mostly on how important it is, how quickly you need to act, etc. But as far as beliefs themselves are concerned (which do not depend on utility calculations),⁹ generally it is epistemically responsible to generalize now. If it was an above average performance and you generalize from this then obviously you will suppose a higher average than actually applies and your generalization will consequently be false, but not faulty. The next time you get the carpet cleaned it is not as good as the first time. You are disappointed with the carpet cleaning company and with your own inferential performance—“serves me right,” you might think, “for generalizing from a single instance.” Such disappointment, though understandable, is not really justified. The carpet company’s performance is simply reverting to its mean; similarly, your inferential performance was, in this particular instance, below average in so far as you inferred a generalization that is false, but this does not mean that generalization from a single instance is fallacious, since on its average performance generalization from a single instance (in these kinds of cases) will lead to true generalizations, generalizations that can then be justifiably applied to randomly selected instances. An above average clean (or one that is below average) will bring about a below average inference.

In such kinds of distributions, it does not seem fallacious to generalize from a single instance. This also applies, with some caveats, to any attribute that should be normally distributed in a population. Height, for example, is an attribute that, relative to specific genders and races, is normally distributed, which is to say that the great majority of the population are within one standard deviation of the mean. Obviously, it would be unjustified to make a *universal* claim on the basis of what is true only in this region around the mean, but one is justified in making generalizations with exceptions of a kind such that, if one were to consider a single arbitrary instance of the population about which one knew nothing else about, one would be justified in believing that his height, for example, was in this region, and if one were to measure the height of this single instance, one would be justified in provisionally taking this to be the average, because knowing that the great probability is that this height will be near the average.¹⁰

What is common to both performances and normally distributed attributes is that the great majority of the population is ‘heaped’ around the average. However, caution

⁹ Our willingness to assert such beliefs may be different and require higher justification because of the utilities and risks involved, just as our willingness to make an assumption required lower justification. Here I am only concerned with the question of whether believing itself is justified.

¹⁰ Perhaps not every normal distribution is like this, as such a distribution with an extremely large standard deviation would not allow us to assume that a randomly selected individual is close to the average. Apart from made-up examples, though, most things normally distributed in the real world are not like this.

must be taken here. The generalization is not justified if we combine attributes. An “average man” will be described as one who has average height, weight, etc., yet there will be comparatively few average men as so described in the population; in other words, we have to take the attributes singly and be careful that they do not combine several things into one. One can also imagine scenarios where what we expect to be a normally distributed attribute is not. Suppose that the population of the world were wiped out except for giants and midgets: the average would still be somewhere in the middle, and yet there would be no men of average height at all. Unless we have reason to think that we are in such an outlandish, artificial situation, the inference from the single instance to the population is justified. Although we cannot always assume that the biggest proportion of the population is on or near the average, this assumption is generally safe for all simple normally distributed attributes, and consequently the general claim will be justified so long as we have not suppressed any evidence that we are in one of these kinds of abnormal situations. Again, our average inferential performance is better for making these inferences than if we did not, and we know this and can provide arguments that justify it. The fact that our inferential performance would be even better had we a bigger sample or more cognitive resources does not mean that we are not justified without them.

The bigger problem with generalization from a single instance is in knowing what exactly you should generalize with respect to. In the case of the disappointed zoo visitor, it should be noted that he does not conclude from his experience that all zoos are unfriendly, or that all Canadian cities are unfriendly, but that Calgary is unfriendly. We can only guess at his reasons for picking on Calgary,¹¹ but we may hypothesize that he had experiences that falsified the alternative generalizations. He may have had experiences that falsified a universal generalization with respect to Calgary also, but not to the extent of being prepared to give a Calgary tourist attraction the benefit of the doubt, especially when the cost of finding out is potentially another negative experience. Because of the costs, he is justified in making the assumption that Calgary is unfriendly and applying it as a general rule.

These, then, are two scenarios in which it is not fallacious to generalize from a single instance, that is to say, where it is justified to make a proportional claim about the population and taking this in turn to justify a probability claim.

¹¹ A larger sample would make it easier for us to discern the reasoner’s reference-class, but the reasoner himself does not need a larger sample in order to generalize, although a larger sample would obviously be preferable for the reasoner too and make his justification for the general claim stronger. The reasoner knows his own reference-class—it is not something he has to hypothesize from the sample.

Here is a third. The general claim that water boils at 100°C is not a claim that is made on the basis of a single sample of water (or so we may presume). But let us suppose that, after considering several different chemicals, we have found that all samples of that chemical have the same boiling point. We can inductively infer on this basis that a chemical that we have not hitherto tested will also have a constant boiling point, and therefore that one measurement of the boiling point will suffice to make the general claim. We therefore make a general claim on the basis of a single instance and treat any putative counter-instances by inferring instead that the sample is not of the same chemical.

Admittedly, this is only superficially induction from a single instance; rather, it is what Reichenbach calls a “concatenated induction” (and we should include the general claims about the tested chemical substances among the premises).¹² I am assuming here that the only justified reason one may have for believing something like chemicals’ having the same boiling point is through *a posteriori* means. Others may claim that this is justified on the basis of *a priori* metaphysical assumptions. Let us suppose that you generalize on the basis of a single instance because accepting such an assumption and that I accuse you of hasty generalization. Is this accusation justified? I have a substantial disagreement with you over the truth of your assumption and may not consider your reasons for holding this assumption to be good ones. But to justify the accusation the burden of proof is on me to show that these are unreasonable assumptions for *you* to hold—the fact that I would not consider them reasonable were *I* to hold them implies no fallaciousness in your reasoning. Were I to succeed in showing that you have reasoned badly in this instance, it would be by showing you that you had failed to notice some internal inconsistency, that is to say, that your assumption conflicts with other claims that you endorse. This amounts to claiming that you were not actually justified in the first place, but only thought that you were.

¹² It is true that the other cases I have mentioned, e.g., normally distributed variables, also depend on background conditions. I am not so sure that the reasoner needs to know those background conditions in order for their generalization to be a good one. Would reasoners be epistemically irresponsible if they generalized while in ignorance of this and/or did not include this fact about normal distributions among the premises? Or is the generalization justified anyway, and the truth about the distribution just explains why it is justified? I am not really committed on this issue either way, although I will suggest in a moment that we do need reasons for believing that the attribute in question is normally distributed. There is no doubt that it is better if we do.

What I certainly would deny is that including this knowledge about the distribution in the premises makes this generalization a concatenated induction too. It *would* be a concatenated induction if we had inferred something about this particular normal distribution from what is true about other normal distributions, but that is not what we have done here. In fact, since boiling point is plausibly a normally distributed attribute (albeit, as it turns out, with a standard deviation of zero, but it is assumed that we do not know this) we could generalize from a single instance even without the additional knowledge about the boiling points of other chemicals. A weaker general claim is justified by this generalization than by the concatenated induction, as we do not justifiably believe that the boiling point is *constant* (i.e. has a standard deviation of zero) but are assuming only that it is distributed normally and with a small standard deviation.

So far, I have been arguing that several cases that have been accused of being fallacies of hasty generalization, and even used as examples to illustrate the fallacy, are not justly accused except for reading the conclusion in uncharitably strong ways. I am now going to discuss an argument that claims to show, for different reasons, that we do not err when we generalize from just a few instances, but that this is what reasoners like us ought to do, and therefore no fallacy. I will argue that, as far as this argument says anything interesting, it is unsound. I will call this the “argument from cognitive costs.”

4. The Argument from Cognitive Costs

This issue is usually raised in relation to a particular model of our cognitive lives. There are two cognitive systems: System 1 and System 2. System 2 is “quick and dirty” — it is highly fallible but cognitively cheap. It is our default system, the one in general use, and guides us well enough most of the time. System 1 is stricter and holds our cognitive behaviour to a higher standard, but is cognitively expensive. For reasoners like us of bounded rationality and limited cognitive resources, System 2 is usually the better one to use, and this allows generalizations to be drawn from smaller samples than would be considered sufficient for a System 1 process. The argument, then, is that the accusation of “hasty generalization” is (at least sometimes) actually holding our inductive inference to a higher standard than is appropriate, viz., a System 1 standard, and that generalizations that would be hasty were they inferred by a System 1 process and would imply some kind of malfunction in a System 1 process, are inductively valid when inferred by a System 2 process.¹³

In the case where we make an *assumption* on the basis of a small sample, I agree that we may do so on the basis that further sampling would be cognitively costly, just as we did when we thought that further sampling would be emotionally costly. If this is what the objection amounted to then it would be true but not very interesting; it would be the charitable (re-)interpretation of the general claim itself that is doing all the work, and cognitive costs would be just one kind of cost among others.

I take it, then, that it is not just “making an assumption” that is being taken to be justified, but belief as such. Now, the argument runs something like this: you know that by generalizing from a small sample you might draw as a conclusion a general claim that is false. But enough ‘hasty’ generalizations of this System 2 type succeed to make belief justified, so unless there is a particular reason to justify the additional

¹³ Gabbay and Woods (2006) give arguments like this. The details are unimportant as I only want to make clear here the general structure that such an argument would take.

cognitive costs of System 1 and invoke the stronger standards of System 2, it is not fallacious to generalize, even from small samples, and in extreme cases from a single instance. This might seem to be of the same pattern as the argument I made earlier that our inferential performance will be right more than not when we make inferences from a single instance of a carpet-cleaning company's performance, and that our inferences revert to a mean just as the cleaning does, and this mean will generally be sufficient to justify belief in a general claim, even though we know that we will be wrong a certain proportion of the times that we make this inference. But there is a significant difference between that argument and what is argued here: in the former, the justifiability of this generalization depends on the nature of what the instance is an instance of, while in the latter it makes no difference, but a simple counting up of the number of generalizations made by the System 2 process and the proportion of those that are inductively valid, which is to say, it is a second-order induction.

Now, clearly there are cases, already discussed, where 'hasty' generalizations can be shown to succeed in a large enough majority of cases to justify outright belief. Consider again the case of a normally distributed attribute. Clearly, it is better to have a greater sample size, and when we do, we can calculate the inverse probability with greater confidence. Even so, if the attribute is normally distributed, then, within a certain margin of error depending on the standard deviation, our conclusion is likely to be true to the extent that our believing it (and not just our making an assumption) is justified. Provided that we have reasons for thinking that the attribute is normally distributed in the first place, we are justified. We have a good reason for thinking that the sample, though small, is representative within a certain margin of error. We may or may not consider the additional cognitive expense of further sampling worthwhile, but it is not the prohibitive cognitive cost of further sampling, but simply the nature of the case, that justifies our generalizing from a small sample.

In order to be claiming anything interesting, then, the claim must be that, *even when we have no reason for thinking that we are in one of these situations*—that is to say, when we have no reason to think that the sample is representative (and even, arguably, when we have reason to think that the sample is *not* representative)—we should follow the general policy of hasty generalization simply on the basis of an unfavourable comparison of the relative cognitive costs and perhaps the importance of the question we are trying to answer. While I am quite prepared to say that we may be justified in making an assumption on this kind of basis, I am unwilling to say that we may be justified in *believing* a general claim on this basis. In so far as an argument from cognitive costs says anything interesting, then, it must say that *belief* itself is justified, and this, I think, is false.

I think the argument from cognitive costs falters on a distinction between a particular reasoning process being the right process to go through and the outcome of the process being the right thing to believe. It is not inconsistent to suppose that some general claim is the right thing to believe but that the process by which we would have that belief is not the right process to go through, or that a process that is the right process to go through leads to a belief that is not the right belief to have. Suppose that a System 2 process of generalization is the best process, in terms of cognitive economy, to go through, and the result of going through it is the general claim that all A's are B. Yet we also know that this generalization was made on the basis of a non-representative sample. Should we believe that all A's are B? I do not think so. In fact, I think that in this situation we should not generalize at all, and surely doing nothing is cognitively cheapest of all. If for some reason we are forced to reason one way or the other (through System 1 or System 2), then all we can offer is a best guess, and not a belief. A belief that all A's are B is not justified and not epistemically responsible, even on the hypothesis that a process of generalization is epistemically responsible (which I doubt it is).

Is the matter different if we suppose only that we do not believe that the sample is representative, without necessarily believing that it is non-representative? I do not think that this matters much. The fact remains that we know that the outcome of the process is one that we have no good reason to think is actually true, and this conflicts with truth being the aim of belief. As before, this follows even if we concede that the process is one that we ought to – or that it is epistemically responsible to – go through.

We can effectively run the same argument the other way too. Suppose that we know, never mind how, that 9 A's out of 10 are B's, and that this is what the outcome of a System 1 generalization would be, whereas the outcome of a System 2 generalization would be as before that all A's are B. Which generalization are we justified in believing? It depends on how you evaluate the counterfactual. If we think that we are justified in believing that 9 A's out of 10 are B's, where the only way we could have come to have that belief is through a process in System 1, then evaluating this claim as a 'backtracking' counterfactual seems to suggest that we ought to go through System 1, irrespective of its cognitive cost. In short, we ought to go through whatever process leads to the belief that we are justified in believing and ought to believe. However, if we treat believing that 9 A's out of 10 are B's as a 'local miracle', then it does not follow that we ought to go through System 1; in fact, it seems that if we go through any system at all we should go through the cognitively cheaper System 2 while still remaining the case that believing that 9 A's out of 10 are B's is justified and

that believing that all A's are B—the outcome of the System 2 process—is unjustified. The second 'local miracle' approach seems to me the correct one when considering what we ought to believe; just as it does not follow from the fact that I ought to go through process X that I ought to believe whatever the outcome of X is, so also it does not follow from the fact that I ought to believe something that would be the outcome of X that I ought to go through process X.

It could be objected that this is a cheat because I built into the example that we know, and are therefore justified in believing, that 9 A's out of 10 are B's. Let's choose a more modest example, then, and suppose only that we believe that we have no good reason to believe that the sample is representative. In such a case what we ought to believe about the ratio of A's to B is: nothing. In this counterfactual circumstance we should have no belief. Again, if we evaluate this as a 'backtracking' counterfactual then whatever it is that we ought to do, one thing that we ought *not* to do is generalize, either through System 1 or System 2, since the outcome of both is a generalization. If we treat failure to have the belief as a 'local miracle' then nothing follows about which process we ought to go through. But, once again, it is strange to see failure to have a belief as any kind of miracle—to see it as a miracle implies that doing nothing is not an option, and doing nothing is cognitively cheapest of all. Moreover, the conclusion that we ought to do nothing is one that we have reached completely rationally by reasoning.

At best, an argument from cognitive costs can tell us that we ought to go through one system or another, and not that we ought to believe their outcomes. In those cases where they do produce justified beliefs, that is to say, beliefs that we ought to believe, it is not their cognitive costs that are relevant. Obviously, there is always a sense in which we are more justified when the belief has been produced by a System 1 process than by a System 2 process—this is a truism. When this extra boost in justification is not worth the extra cognitive expense, we are content to have a more weakly justified belief. However, it would be wrong to call such a generalization hasty or fallacious in circumstances where there is some reason to think that the sample is representative, and this could even be true, as we have seen, when the sample is a single instance. In circumstances where there is no reason at all to think that the sample is representative, no appeal to cognitive savings is going to justify belief in the general claim. The standards for a belief's being justified are not context-dependent in the sense that the standards for assuming, or accepting, or asserting, clearly are; if the outcome of a System 2 process is not a justified belief (e.g., because the sample is not believed to be representative), arguing that it was the correct process to go through given the cognitive resources available will not

make the belief any more justified. Cognitive costs simply cannot have this kind of effect on what we ought to believe.

A better example of the fallacy of hasty generalization, then, is one not involving anecdotal evidence, given that we can usually make justified general claims on the basis of such evidence provided that we have no evidence of being in one of the exceptional situations, and if we did have such evidence, that evidence ought to be included in the premises. Therefore, I think that the examples of arguments given by Engel, Blair and Johnson, and Groarke and Tindale, which do not have such premises, describe perfectly reasonable examples of reasoning, provided their conclusions are interpreted charitably as not literally making universal claims but stating a rule that, when applied universally, will produce better inferential behaviour on average than if we refused to draw the inference because of a mere possibility of being wrong, where a bad inferential outcome is identified either with drawing an inference that leads to a false belief (when belief is the issue, as in the cleaning example) or with leading to an assumption that incurs negative experiences or other costs (as in the “All men are no good” and “Calgary is an unfriendly city” examples). By following the rule, we should have on average fewer false beliefs and fewer negative experiences than if we did not.

We have already encountered a better example of the fallacy of hasty generalization:

All the philosophers in the room have beards.

Therefore, all/most philosophers have beards.

This seems fallacious, and we can safely assume that there are no utility considerations. As before, the problem seems to be that, even if we have no direct knowledge of any philosopher that is not bearded, the sample is not large enough to justify the generalization, or any proportional or probabilistic claim about the population. In such cases, is it justified to charge a reasoner with reasoning fallaciously?

Occasionally, other grounds than sample size are appealed to. For instance, the sample might be biased. This could be quite innocent and due to the fact of what evidence is available to us, or it could be that we have the sample we do because we have gone looking for confirming instances and avoided areas where falsifying instances are more likely to be found.¹⁴ “A sample,” Groarke and Tindale (2004, p. 290) say, “must be sufficiently large to give us confidence that its characteristics are not due to chance . . . [and] must also avoid bias.” Is it justified to charge a reasoner with reasoning fallaciously in this case?

¹⁴ In the testing of a scientific hypothesis, for example, experimental results must be replicated, not only by different scientists in different labs, but the experimental conditions must be varied.

I do not think so. When we criticize the generalization, often we are criticizing the methods of data collection used rather than the inductive inference itself; since the inference draws a conclusion relative to a certain body of evidence, the inference itself will be valid provided only that there is no mathematical mistake in performing this calculation,¹⁵ which is not a fallacy.

The point I am making is that accusations that a reasoner has actually committed a fallacy of hasty generalization are very hard to substantiate; when we look at the details, the criticism aimed at the generalization do not justify an accusation of having inferred fallaciously but of inferring at all from inadequate evidence or having inadequate methods of gathering evidence. Charges that a fallacy of hasty generalization has been committed is usually motivated by beliefs about the evidence. For example, I would suggest that what makes Johnson and Blair criticize the generalization that Calgary is an unfriendly city is evidence that they have that it is not an unfriendly city—evidence, however, which may not be available to the reasoner they are criticizing. It is to this that I now turn.

5. The Evidential Basis of Generalization

Let us recap. Assume that anything that a reasoner takes as evidence is included in the premises of the argument. Since the generalization always makes a claim relative to the evidence, an argument like

A_1 is B

A_2 is B

·

·

A_m is B

A_{m+1} is not B

·

·

A_n is not B

Therefore, $p(A,B) = m/n$

will always be inductively valid. If we were to add further evidence, we would quite simply have another argument.

If we charge the reasoner with hasty generalization, then, we effectively accuse of him of suppressing a premise. Such a suppressed premise might be simply that

¹⁵ For example, if from the earlier example the reasoner inferred $p(A,B)=0.9$, where this equals $m/n+1$ instead of m/n , then he has simply miscalculated and not committed a fallacy.

the sample size is too small. If the reasoner knows this then he ought to add it to his premises, and if he generalizes despite knowing this then this amounts to generalizing from what he does not believe to be a representative sample. It is difficult to see how a reasoner, when reasoning sincerely, could actually commit this fallacy; one would have to suppose that the reasoner chooses to ignore, or momentarily forgets, something that he himself takes to be evidence or knows about the evidence, that is to say, something that he takes to undermine the inference which he notwithstanding proceeds to make.

Now, in most cases when the sample size is small, believing that the sample is *not* representative is the safest default assumption to make and the one that is epistemically responsible. I have explained cases where this is *not* the case and where even a single instance may be taken to be representative, viz., attributes that belong to all of a natural kind (or is constant for all instances of such a kind, like the boiling point), normally distributed attributes and attributes like performance where reversion to the mean applies. If the reasoner justifiably thinks that in this particular case one of those situations obtains then he can equally justifiably overturn this assumption.

When, then, is a charge of hasty generalization justified? We cannot justifiably accuse the reasoner of an inductively invalid inference, since, as we have already said, it is not invalid. We cannot accuse the reasoner of ignoring evidence when the reasoner is not aware of the evidence. We can criticize the data-collection techniques and experimental design, and perhaps find evidence of bias. We can argue that the reasoner, having found this evidence, should have known better and that his performance was below acceptable standards for making general claims of the kind the reasoner was trying to make.

However, as accusers we must beware of making an accusation of fallacy when what we really have is a substantive disagreement about evidence. Suppose that the accuser, having different evidence (and thus different premises) draws a conclusion contrary to that of the accused. Both arguments will be inductively valid, and neither accuser nor accused can justifiably blame the other for bad reasoning. In effect, the accuser is claiming that her evidence is the better evidence, that it is a more representative sample. If the accused accepts this as evidence and yet refuses to take account of it in his argument, then he is ignoring evidence that is available to him. This is what a fallacy of hasty generalization amounts to, when justified, and since it involves knowingly ignoring evidence, it does not seem very plausible to attribute such a mistake to a minimally competent reasoner who is reasoning sincerely and who is not deliberately putting forward an argument that he himself does not believe to be a good one in a conscious effort to deceive.

What, however, if the accused does not accept the counter-evidence? Then what we have is a substantive disagreement about the evidence and not a fallacy. A reasoner may have a number of principled reasons for rejecting a putative piece of evidence. In fact, there are principled reasons for which we do this as a matter of course in certain situations (as we saw in the previous section): if someone suggested that a particular sample of pure water did not boil at 100°C or a particular piece of iron was not magnetic then we would not accept this as counter-evidence to the general claims involved but as evidence that the samples are not of pure water and iron respectively.¹⁶

6. Conclusion

This paper has described a number of cases in which it is epistemically responsible to generalize from a single instance and to believe the general claim inferred, that is to say, where a single instance can be taken to be representative of the population sufficiently to justify a generalization with exceptions: a) performances and other cases where one justifiably believes that “reversion to the mean” applies, as it is reasonable to suppose in such cases that the single instance is representative of the mean; b) cases of attributes that one justifiably believes are normally distributed in a population (and have small standard deviations); c) cases of attributes that one justifiably believes belong always (contingently) or essentially (necessarily) to the population. When there is no reason to believe that a single instance can be taken to be representative of the population, it is fallacious to generalize from it, and issues of cognitive costs do not suffice to overturn the default assumption that the sample is not representative. In the cases of (a) and (b) the conclusion is only probable and the conclusion should be stated accordingly, since in these cases we are only saying that the attribute in question is ‘heaped’ around the average, and so there is always a mathematical possibility, even in large samples, that the sample is not representative, though clearly the larger the sample the more likely it is of being representative. Arguably, all generalization is justified in this way: it is because the ratio of representative samples to non-representative samples increases with the size of the sample that we are justified in thinking that the probability distribution in the population is the same as what we have found in a large enough sample. Thus, denying the justifiability of the kind of inferences in (a) and (b) would amount to casting the whole process of inductive generalization into doubt. In the case of (c) a stronger, universal

¹⁶ The case is the same as in analogies: what one person considers to be a relevant difference (and thus takes the instance as being outside of the reference class and so as not part of the statistical evidence), another may not.

claim seems to be justified too, and the sample size is irrelevant, a single instance being as good a sample, and as representative of the population, as many instances.

These beliefs may be false. Suppose, for example, that you believe that height is normally distributed among the population of all men. This is false, because it ignores racial and other relevant differences between men — it is the wrong reference class, and as long as you have encountered other races this is something that you ought to know. It is difficult to think of an excuse for your failing to know this, so were you to believe on the basis of a sample of men that their height should be near such and such a number, an accusation that this commits the fallacy of hasty generalization seems justified in these circumstances.

But let us suppose further that you belong to a bygone age where encounters outside of one's immediate community are rare, and that you encounter an explorer who tells you of pygmies and what you would consider to be giants over six feet tall. You may or may not believe the explorer's fantastic tales. Before your encounter, you were justified in believing that what is the average height of all the men in your own community is the average height for all men, period. Depending on how credible you find the explorer, you may still be justified in this belief after the encounter.

Now, the explorer knows better — he knows that this belief is false by the evidence of his own senses. But would the explorer be justified in accusing you of bad reasoning? Have you committed a hasty generalization?

I do not think so. Instead, there is a substantive disagreement over a premise or assumption — the explorer has evidence that you do not. If the explorer is to convince you that you have actually committed a fallacy, he must do so by appealing to evidence that you already have but have somehow ignored. But if there is such evidence, this amounts to claiming that, in fact, you were not justified in the first place, but only thought that you were by failing to take some piece of evidence that you already had into account. To justify an accusation of fallacy the accuser must attribute to the reasoner knowledge of undermining evidence and disregard of that undermining evidence simultaneously. This is possible: you may not know your undermining evidence under that description, just as you do not know all the deductive consequences of your beliefs, that is to say, although you know it you have not appreciated that it undermines your inference. If you simply have no knowledge of the undermining evidence, though, there is no fallacy, but only an inference drawn from a false, but justifiably believed, premise.

Further, if one is presented with a new piece of evidence, one may have reasons that justify rejecting that evidence; e.g., a piece of iron that is not magnetic, a piece of metal that does not conduct electricity, may be justifiably rejected by the essentialist

or even by the non-essentialist (since even *a posteriori* laws may become so well-entrenched as to be virtually unfalsifiable), it being both preferable and epistemically responsible to believe that the sample provided is not iron or metallic. To justify the accusation of fallacy, the piece of evidence must be such that the reasoner has no excuse for not knowing it or for failing to appreciate its relevance to the inference. If the reasoner does know that a piece of evidence undermines his inference and yet simply chooses to ignore it and draws his conclusion anyway, then he is not reasoning sincerely anymore.

The result? The fallacy of hasty generalization, as knowingly ignoring evidence, is not one that is easy to make, and the charge that someone has made it is not easy to substantiate. To substantiate it amounts to accusing the reasoner of insincerity. What is normally found, instead of bad reasoning, is substantive disagreement over the premises.

Lastly, I would restate my disagreement with the examples of Engel (1976), Johnson and Blair (1994), and Groarke and Tindale (2004) — these are only hasty generalizations if we interpret their conclusions in an unnecessarily strong way; otherwise, the inferences they describe, although obviously fallible and known to be so, are not epistemically irresponsible, even though all draw general conclusions from a single instance. Even in these supposedly paradigm cases, an accusation of committing a fallacy of hasty generalization is not justified.

References

- Botting, D. (2017). The fallacy of sweeping generalization [Paper presentation]. In C. Winkler. *Networking Arguments: Selected Papers from the 20th Biennial Conference on Argumentation*, Alta, Utah, U.S.A.: Routledge
- Engel, S. M. (1976). *With good reason*. New York: St. Martin's Press.
- Groarke, L. and Tindale, C. (2004). *Good reasoning matters!* 3rd edition. Oxford and New York: Oxford University Press.
- Harman, G. (1965). Inference to the best explanation. *The Philosophical Review*, 74(1), 88-95.
- Johnson, R.H. and Blair, J.A. (1994). *Logical self-defense*. New York: McGraw-Hill.
- Reichenbach, H. (1968). *Experience and prediction*. Phoenix Books. Chicago: The University of Chicago Press.
- Gabbay, D. M., & Woods, J. (2006). Seductions and shortcuts: fallacies in the cognitive economy.