INFORMED DECISION MAKING

A Thesis

Submitted to the Graduate Faculty of
Louisiana State University and
Agricultural and Mechanical College
in partial fulfillment of the requirements
for the degree of Master of Arts

In

The Department of Philosophy and Religious Studies

By
Jonathan D. Tall
B.S., Northwestern State University, 2003
May, 2006

Acknowledgements

I would like to thank Dr. Kevin Elliott for all of his work with me both inside and outside of the classroom. Your dedication is unrivaled. Without your inspiration and grokking, none of this would have been possible.

Table of Contents

Abstract

This thesis will address the issue of appealing (or deferring) to the authority of expertise. The effects of the social character of knowledge highlight two points with regard to this issue. First, they leave the layman in an epistemically inferior position. Thus, the layman must appeal to the authority of experts. The second point, related to the first point, is the implicit role of trust in deferring to expert authority. Though I will pay attention to each of these points, the focus of this thesis will be on the former. If we accept that one must appeal to the authority of experts in order to be on epistemically firm ground, then as a consequence we leave ourselves open to the problem of adjudicating between conflicting expert testimonies. The goal of this thesis will be to examine a process by which one can arrive at the most epistemically justified position given the amount of information that one has at one's disposal.

Introduction

In today's world, it appears as though we are surrounded by controversy. To some degree, this is nothing new. After all, it is a rumpus that we can trace back to ancient times. Plato writes in the *Charmides*,

> [W]hen a person claims to know something, will our friend be able to find out whether he knows what he says he knows or does not know it. (Cooper 1997, 657)

Even then, people were trying to find a way around this problem of knowing when someone is an expert. Today, we can say the problem has evolved. Some authors, like David Bloor (1976), feel that knowledge is different today. Knowledge is no longer something that one man possesses, but rather a conglomeration of understanding from several specialized fields.

John Hardwig writes in "The Role of Trust in Knowledge,"

> Modern knowers cannot be independent and self-reliant, not even in their own fields of specialization. In most disciplines, those who do not trust cannot know; those who do not trust cannot have the best evidence for their beliefs. (1991, 693-694)

Hardwig is making a similar point to Bloor. There is an implicit trust that is built into the production of knowledge in the modern era. When one does research, even in their own field, they must trust that the backs of giants on which they stand were firmly based in sound scientific practice. That is to say, they must trust that the other experts, whose work they reference, are indeed in possession of the expertise that we attribute to them.

The question of the expertise of experts is a difficult one to assess, and is by no means easier today than it was 2 millennia ago. The increased specialization of fields contributes to the creation of more experts to work in those areas. As knowledge becomes more complex, more experts must come together from differing fields to unearth it. In two of his papers, Hardwig discusses an experiment that measured the lifespan of charm particles in physics. "The paper

reporting the results of this experiment has 99 authors, in part because it took about 280 person/years to do the experiment" (1991, 695). Not only was there a time factor, but the experiment itself required scientists from many different fields of particle physics.

Naturally, this is a lot of trust that is going around. The trust that we place in the knowledge this team has extracted of charm particles is based on a very complex and fragile system. It is safe to say that the likelihood of any individual scientist making an error is quite small, and we would expect that as we pool more scientists together, we would be decreasing the chance for error. After all, is this not how peer review operates in a scientific community?

Though I do not claim to deny that peer review can indeed weed out some forms of bias, the point I wish to make here is that we should not rule out what Kahneman and Tversky (1973) called the "conjunctive and disjunctive fallacies." Under these conditions, people often misinterpret the probability of an event occurring. One common example of such a misconception is in risk assessment. It is believed that since each component in a nuclear reactor has such a negligible chance of failure, that the system is reliable and safe. What often gets neglected is that since the system is quite large, a failure in any one piece can cause failure to the system. Thus, there is an *increasing* chance for failure as we add more parts to the system, despite each part having such a small probability of failure.

If we turn back to the example of charm particles, we can see how the failure of one scientist *could* have had a significant impact on the results of that experiment, and thus our knowledge of charm particles. What I am trying to highlight here is the importance of having a rigorous system for evaluating expert opinion. In cases like the charm particles, what is at stake is our understanding of a small area of particle physics. Other cases have much higher stakes since they involve international trade law, international environmental policy, and people's lives.

The effects of CFCs on the ozone and greenhouse gases on global warming are two examples of public policy issues that require, or required, well informed decisions (see Weiss 2003). On the one hand in these debates is the global economy. Unnecessary restrictions on industry based on speculative concerns can do severe damage to the global economy. On the other hand, failure to acknowledge the risks posed by non-regulation can lead to irreversible global environmental damage. Therefore, the effects of the decisions require that they be made with a strong degree of epistemic justification, i.e., a well informed basis.

The decision to regulate CFCs was an early decision which proved to be quite beneficial. As we later discovered, CFCs were destroying the stratosphere which in turn led to the creation of a hole in the ozone over Antarctica. We now face a similar problem with respect to global warming. Unfortunately, in this case an informed decision has not been swift in coming. Though global warming is an incredibly complex issue, 178 countries came together in 2001 and recognized the importance of this problem, and as with the CFCs, they have moved for stiffer regulations. The US was not one of them, and has only recently begun to take action beyond the Kyoto Protocol due to political pressure.

Another example is derived from international trade law, and to some degree human health. The use of genetically modified food is a matter of serious concern. Some countries are claiming that we do not know enough to simply eat these products carefree. For example, the European Communities tried to block the import of beef from the US, since many farmer use growth hormones. The other side of this issue concerns world hunger. Without genetically modified crops, many countries will simply fail to get enough food. Though the World Trade Organization has denied the right of the EC to block trade, this issue is far from resolved. The health and hunger of many people lie in the balance.

Finally, I will turn to two other issues involving public health and well being. The first is a hot issue in toxicology known as *hormesis* (see Kaiser 2003). Hormesis is a condition by which chemicals exhibit opposite properties at high doses and low doses. Many scientists working on hormesis wish to focus on chemicals that are beneficial at low doses, and toxic at higher doses. One popular example involves dioxin. Dioxin, when given to rats in small doses, gives them a 50/50 chance of dying from liver cancer. However, when given in even smaller doses, it has been shown to inhibit tumors.

There are many examples of hormesis in nature. Another popular one involves arsenic, which can cause immune cells to proliferate faster when given in small doses. But, not all cases of hormesis are good. Some scientists point to the fact that though dioxin suppresses breast tumors, it can also promote liver tumors. In other cases, toxicants called 'endocrine disruptors' were found to be more toxic at smaller doses than at high doses. Thus, the government decisions surrounding the regulation of these chemicals, and their effects on the environment and public health need to be well informed due to the complexity of the issue.

The final case that I will reference here involves AIDS in South Africa. As we will later see, there is a strong debate between whether or not more research is needed on the topic of HIV as *the* cause of AIDS. Ward Jones (2002) will argue that there is greater epistemic justification in the majority held position, while President Thabo Mbeki argues for consideration of the alternatives. The ultimate decision in this case determines whether or not a significant portion of the population is given anti-HIV medication as treatment for, and the prevention of, AIDS.

Each of these cases demonstrates the profound effects that follow from public policy decisions. In each of these cases, we would hope that the ultimate decision maker takes every precaution available to them to ensure a well informed decision. Thus, it is my intention in this paper to give a deeper examination of the problem of deference to expert authority. In the first

three chapters of the paper, I will explore some of the previous work in the area. My goal in these sections will be to critically evaluate each of these three positions. I intend to harvest the fruitful arguments, so that I may then, in the fourth chapter, synthesize these ideas into a process that will yield a higher degree of epistemic justification, and thus a more informed decision.

Each of the first three chapters constitutes a different author's perspective on this issue. The first chapter will provide a nice technical vocabulary derived from the concepts in social empiricism. This vocabulary will allow for a neutral approach to issues like "bias" in science. As such, it will be indispensable to discussing and classifying the motivations that one has for siding with a particular position. Furthermore, the discussion of social empiricism will also provide a framework for understanding what it means to have, to be in, and to lack consensus between scientists. The understanding of this will be crucial, since it will bring more clarity to the various positions, such as the majority position. In addition to this, it will also bring added focus to the relation between the positions with respect to productive science.

In the second chapter, I will take a look at Ward Jones's (2002) approach to dealing with the problem of choosing between experts. Analysis of this argument will yield insights into how others have attempted to solve the problem at hand. This argument will provide an excellent starting position, since I will first need to address the holes left open by Jones' argument.

The third chapter will take a slightly different turn from chapter two. Whereas in chapter two I will examine an attempt to solve the problem, chapter three will look into a wider exploration of the problem. Though Alvin Goldman (2001) does not offer any solutions to the problem, his work will help to provide added insights into the issues that a viable solution would be required to address.

In the fourth chapter, I will bring together the constructive ideas from the previous chapters. I will begin by clarifying the relation between expert and layman, as well as some of

the issues that I see as facing each side. This groundwork, combined with the observations from social empiricism, will allow me to derive a process, or flowchart that one may use to arrive at an informed decision (qua a superior epistemically justified position).

This thesis will address the issue of appealing (or deferring) to the authority of expertise. The effects of the social character of knowledge highlight two points with regard to this issue. First, they leave the layman in an epistemically inferior position. Thus, the layman must appeal to the authority of experts. The second point, which can be derived from the first point, is the implicit role of trust in deferring to expert authority. Though I will pay attention to each of these points, the focus of this thesis will be on the former. If we accept that one must appeal to the authority of experts in order to be on epistemically firm ground, then as a consequence will leave ourselves open to the problem of mitigating between conflicting expert testimonies. The goal of this thesis will be to examine a process by which one can arrive at the most epistemically justified position given the amount of information that they have at their disposal.

Social Empiricism

Miriam Solomon argues for a new epistemology of science in her book *Social Empiricism* (2001). She introduces her position by first showing how others have focused on the individual scientist as the locus for reason and objectivity, and then claiming that she will instead try to show that the locus of reason and objectivity lies with the community. Solomon argues that in previous debates, although there was disagreement over the conclusions, an underlying assumption was accepted. She claims that this assumption, Enlightenment epistemology, became the framework for the debate. The goal of her book is to challenge this framework.

Whig Realism and Empirical Success

Solomon chooses to divide successes in science into two categories: empirical successes and theoretical successes. For example, predictive success would serve as an empirical success, while simplicity would serve as a theoretical success. The distinction between empirical and theoretical successes can be summed up quite simply: "*Empirical successes are contingent on the world outside the inquirers; theoretical successes are not*" (Solomon 2001, 17). According to Solomon, empirical success is a primary goal of science, while theoretical successes are not. This is not to say that theoretical successes are useless, but rather that theoretical successes are not a *necessary* goal of science.

Empirical success comes in a variety of ways. It can be predictive, retrodictive, technological, explanatory, or observational to name a few. This variation occurs because different types of success may be desired or required in various fields. The same variation applies over time, as new types of empirical success advance the field. These variations prevent a uniform definition of what constitutes empirical success in science.

There are, however, two features that are important to understanding empirical success: robustness and significance. Robustness refers to the ability to reproduce the empirical successes. Empirical success that cannot be reliably demonstrated over time is not noteworthy success. Significance, on the other hand, is a qualitative measure of empirical success.

> Empirical success is significant when it is mostly attributable to the theory, rather than to prior knowledge shaping the application of a theory. Such significance is clearly produced when a theory has *new* empirical success. (Solomon 2001, 29)

Solomon's view of significance is restrictive enough that it picks out precisely what makes significant (empirically) without overextending into cases that are purely descriptive. For instance, in psychology there have been many theories that can reliably explain any behavior. The problem is that these theories do not have falsifiable conditions, and as such these successes are not significant. These theories can only be applied via prior knowledge in a retroactive sense to new cases. Theories that are significant can be applied to new cases and generate new predictions prior to any new knowledge.

Solomon's new insights lead her to a new form of scientific realism that she terms 'whig realism.' "Whig Realism gives novel explanation of empirical success in science" (Solomon 2001, 33). Solomon creates this as a means of explaining how theories can generate significant empirical success, even when they are wrong. Since truth is as important to a theory as empirical success, Whig realism has methodological importance. It is whiggish in the sense that it involves looking back from the present and seeing the truth in past theories. To put it another way, Whig Realism looks back on past theories and claims that even though all of the statements of the theory are false, there is still 'some truth in the theory.'

To help demonstrate how there can be 'truth in a theory' without actual statements of the theory being true; we can look back at Priestley's phlogiston theory. Priestley described a theory in which there were invisible properties in the atmosphere that were conducive to combustion.

Phlogiston theory had very robust and significant empirical success, despite the fact that the theory was false. It was later discovered that oxygen was the characteristic responsible for empirical success of phlogiston theory. The discovery of oxygen allowed for a separation of the 'correct' portions of phlogiston theory, from the 'incorrect' theoretical structures.

Solomon makes further effort to clarify that Whig Realism does not mean that successful theories are 'approximately' true. It does not mean that successful theories are 'partly' true. And, Whig realism is not whig history. The 'Whig' in Whig realism refers solely to an explanation of the empirical success attained by theories in their own time. It is a way of understanding how theories that are mostly (or completely) false can achieve empirical success.

A Fresh Look

Several philosophers have argued over the appropriate incorporation (or removal) of values from science. For example, Ernan McMullin (1983) argues in favor of a 'pure' scientific practice that is free of non-epistemic values. Helen Longino (1990), by contrast, argues in favor of incorporating non-epistemic values into scientific practice as a necessary means of bridging the gap between data and hypothesis. Solomon brings a fresh look to this debate by introducing a new vocabulary that helps to focus the issue. Rather than referring to these factors as 'biases' or 'values,' she refers instead to 'decision vectors.'

Decision vectors allow us to describe the factors that affect scientific reasoning from an epistemically neutral position. The name 'decision vector' comes from the fact that they influence our decisions in a particular direction. It is epistemically neutral because it does not describe the influence as positive or negative. Sometimes, as Longino rightly points out, non-epistemic values can help push science in the right direction. Other times, as McMullin argues, these very same values can bias science in a negative manner. Because of the ultimate determination of whether or not a non-epistemic value is positive or negative is circumstantial,

Solomon argues that it is crucial to approach the issue from a neutral standpoint. Furthermore, the fact that we can't escape from psychological heuristics means that we need a way to talk about non-epistemic values neutrally, so that we can better understand how they operate in such as way as to positively or negatively affect us.

Solomon breaks decision vectors into two categories: empirical and non-empirical. Empirical decision vectors are preferences for theories with one or more kinds of empirical success. Non-empirical decision vectors are preferences for a theory for some other reason. So an empirical decision vector might be salience or availability of data, while a non-empirical decision vector might be elegance or simplicity. Salience and availability of data are empirical decision vectors because they demonstrate a preference for a theory with data. Elegance and simplicity are non-empirical decision vectors because they are not directly related to empirical success.

It is thinking in terms of decision vectors that allows us to put new perspective on scientific consensus and dissent. For example, they allow us to better specify adequate or inappropriate distribution of research effort.

> [T]here is a general need for a normative epistemology of science that applies at the social level. A descriptive epistemology is not enough. I have begun this…, suggesting that empirical decision vectors be equitably distributed, and non-empirical decision vectors equally distributed. (Solomon 2001, 95)

Since several theories may each have some empirical successes, research should be distributed accordingly. Thus, scientific dissent can be characterized as being such a case (one in which more than one theory has empirical success). Under such conditions, dissent is appropriate, and research should be divided equitably among the theories based on their empirical successes. Similarly, Solomon argues that non-empirical decision vectors should be distributed equally

among the theories. This equitable and equal distribution of decision vectors help to ensure that each theory gets a normatively appropriate amount of the research distribution.

Consensus, as Solomon puts it, is a special case of dissent. "Consensus can be viewed as a limiting case of dissent—when the amount of dissent approaches zero" (Solomon 2001, 118). Consensus occurs when one theory possesses all of the empirical successes[1]. When consent occurs, there is no longer a need to distribute resources to other theories. It is at this point that non-empirical decision vectors become the driving factor in research. Consensus comes in two types: shallow and deep. Since scientists can agree on a theory for different reasons, they may come to a 'shallow' consensus. 'Deep' consensus occurs when they all agree for the same reasons.

Solomon further adds that there is nothing intrinsically desirable about consensus. Consensus is not always conducive to science, and there are cases in which consensus has hindered or stagnated progress. Because of this, Solomon emphasizes that consensus is normatively appropriate in some instances and not others.

Social Empiricism

Given this framework, we can now fully realize what is meant by social empiricism. Social empiricism is an extended account of how dissent should be structured within the scientific community. As was stated earlier, consensus becomes an absolute limiting case of dissent. Solomon argues that one can stipulate the appropriate distribution of effort in science on a spectrum ranging across differing degrees of dissent.

Solomon argues that dissent in science should be organized as follows:

1. Theories on which there is dissent should always have empirical

[1] Solomon's claim that one theory can have *all* of the empirical successes is a controversial one, because of the underdetermination of theories by empirical evidence. Because this issue is far beyond the scope of this thesis, I will not debate Solomon's claim here.

successes.
2. Empirical decision vectors should be equitably distributed (in proportion to their empirical successes).
3. Non-empirical decision vectors should be equally distributed (the same number for each theory). (Solomon 2001, 117-118)

From this layout, we can see that as one theory comes to acquire more of the empirical successes, they will also garner more of the research distribution. Non-empirical decision vectors must then be divided between the remaining theories. It follows from this that once one theory acquires all of the empirical successes, we can formulate the appropriate conditions for consensus as:

1'. One theory comes to have all the empirical successes available in a domain of inquiry.
2'. This same theory comes to have all of the empirical decision vectors, since all scientists working productively (with empirical success) are working within one theory.
3'. Any distribution of non-empirical decision vectors is OK, but typically more will develop, over time, on the consensus theory, as the old theories fade away. (Solomon 2001, 119)

According to Solomon, consensus should occur only when one theory has all of the empirical successes. Because of this, consensus should not be seen as the endpoint in a domain of inquiry. Anomalies, or other circumstances, may contribute to new theories and potentially new empirical success. Under the following conditions, for example, Solomon believes that it is appropriate for consensus to dissolve back into dissent:

1". A new theory has empirical success that is not produced by the consensus theory. (So, the new theory deserves attention.)
2". Empirical decision vectors come to be equitably distributed.
3". Non-empirical decision vectors come to be equally distributed. (Solomon 2001, 119-120)

During this entire process, it is the empirical decision vectors that determine whether there should be dissent, consent, or dissolution of consent. Since these decision vectors are divided

among the members of the scientific community, they provide the 'social' aspect of social empiricism.

Solomon's account provides an excellent framework for my own account. I plan to build off her foundation and expand upon some of the issues that are still left open. Social empiricism is a good account of how the division of scientific labor should be treated from the perspective of the scientists. I am interested in what this means for the layman when it comes to expert testimony.

It is not clear from this account which theory is to be seen as more 'correct.' We often think of the majority opinion in a community as being a more epistemically justified position to which laypeople should defer. Yet, Solomon's account does not support this conclusion. She believes that each theory with empirical success should be considered to be a legitimate topic for scientific investigation. For example, Solomon argues that empirical success should be distributed equitably among the theories that the scientific community investigates. However, this approach does not work as well in the context of making policy decisions based on scientific information. A theory may have plenty of empirical success, yet have none of these successes in the policy area that is under consideration. Thus, the amount of empirical successes is not necessarily a measure of epistemic superiority. To put it another way, the amount of empirical success loosely demonstrates the amount of 'truth in the theory,' which under Whig realism does not mean that the theory is any more or less true than another theory (the theory can very well be entirely false!). Moreover, policy makers may need to choose a particular theory as a basis for action; it may not be possible to act on more than one theory at a time in the policy arena. For these reasons, we need a more adequate means of determining when it is appropriate to side with one or another theory in terms of appealing to experts.

Dissident vs. Loyalist

Ward Jones wrote a paper entitled "Dissident Versus Loyalist: Which Scientists Should We Trust," (2002) in which he addresses the very issue that Solomon left open. In other words, he considers under what circumstances it might be appropriate for policy makers to side with "dissident" scientists (i.e., those who accept scientific views that differ from those held by the majority of scientists in a particular field). Jones took up this issue after a speech by President Thabo Mbeki of South Africa, in which Mbeki compared the rejection of dissident scientific perspectives to the previous despotism there. In April of 2000, Mbeki claimed:

> We are now asked to do precisely the same thing that the racist apartheid tyranny we opposed did, because, it is said, there exists a scientific view that is supported by the majority, against which dissent is prohibited. The scientists we are supposed to put into scientific quarantine include Nobel Prize Winners, Members of Academies of Science, and Emeritus Professors of various disciplines of Medicine! (Jones 2002, 512)

These are strong words from Mbeki, but they do well to highlight the importance of this question of how to respond to "dissident" perspectives within the scientific community. In Mbeki's case, this policy decision determines whether thousands of people are given anti-HIV medication for the treatment or prevention of AIDS.

In order to further understand precisely what is at issue, let us look a little bit deeper at what exactly is being debated here. Mbeki's position is that one should examine the merit of the two available scientific positions, not the numbers of scientists who accept each position. Thus, for Mbeki, the two positions are epistemically equal. The other view, argued for by Jones, is that the numbers do matter. The dominant position is the community position, and thus it carries extra weight because it toes the party line. Jones thinks that the issues involved in evaluating the scientific positions can be quite complex and should therefore be left, as much as possible, to the scientists themselves to sort out.

Testimony and Bias

One way to put perspective on the dominant versus dissident debate is to treat them simply as testimonies. Under this view, Jones argues that we can examine the opposing scientific positions for bias. If one position demonstrates less bias, then we can say that the position is more likely to be accurate than the other position. However, Jones goes on to explain that the distinction is not that simple. Each side can be biased in an epistemic or non-epistemic way. Furthermore, the line between epistemic and non-epistemic bias is not always a clear one. Jones gives the following example:

> Jones is committed to a theory because Jones judges the theory to be more elegant than any of its rivals. (Jones 2002, 515)

Jones claims that this statement is ambiguous with respect to bias, since it ultimately depends on Jones' view of elegance. If Jones holds that elegance is a feature of a theory that is more likely to yield truth, then we can say that elegance is an epistemic bias for Jones. However, if by contrast, Jones does not believe that elegance correlates with truth, then Jones merely has a personal preference for elegance. This option would be a non-epistemic bias.

To help clarify this somewhat murky bias issue, Jones feels that we need to look for two characteristics, which if taken jointly, yield unbiased results. The first has already been mentioned. The decisions should be based on what the believer takes to be epistemic considerations. The second characteristic is that the decision should be 'readily affected by counter evidence' (Jones 2002, 516). Both of these characteristics are features of a communal scientific theory of acceptance.

In a communal scientific theory of acceptance, scientists publish their work. Through this act, they open their work to debate, and thus to counter-evidence. Arguments for and against the

published work cannot be advocated with the use of non-epistemic biases. The community works to sort through, and weed out biases, until it comes to agreement.

> It follows that the acceptance of a theory by a scientific community is evidence-determined and evidence-sensitive. Acceptance does not occur if members of the community successfully question the theory at hand. (Jones 2002, 519)

The arguments made on each side are matters of the evidence and inferences made from evidence to theory. Similarly, the challenges made to arguments are directed towards the evidence and inferences. This ensures epistemic discourse in the scientific community. Moreover, it points to the community acceptance of a theory as being truth-oriented.

If this is true, then naturally, dissidents (regardless of status) are still set in opposition to the community and the community accepted theories. From the vantage of the layperson, this means that the dissident view is not epistemically equal to the community view. Jones' point is simply that since the scientific community acts to filter out non-epistemic bias, then it less probable that the loyalist position will be affected by non-epistemic bias. The same is not true of the dissident position. Thus, in cases like Mbeki's, one should opt for the loyalist position.

Counter-evidence

To close the issue at this point would be too hasty. Jones has put forth a compelling argument for the epistemic justification of deference to the majority. However, I do not think that this holds for anything but the most general of cases. Jones described Mbeki's position as follows:

> Policymakers should decide what to do on a case-by-case basis, examining the evidence for and against undertaking a particular course of action… Policymakers should hear both sides, but ultimately they should do what they judge to be the best course of action. They should follow their own assessment of the evidence. (Jones 2002, 513)

I do not think that Jones gives adequate respect to Mbeki's position. He does not think that policy makers should evaluate scientific positions on a case-by-case basis; rather, he thinks that

they should follow his general rule of respecting loyalist positions. Perhaps, he feels that laypeople are not capable of sorting through the evidence.

So how would a case-by-case analysis make a difference? Let's look back at some of the critical features of Jones's argument. First, Jones wants to separate epistemic and non-epistemic bias in science. The issue was tricky, since Jones argued that there were times when it is somewhat unclear whether a particular bias is epistemic or non-epistemic. Second, Jones listed two conditions that, taken jointly, would yield a result that is as free as possible from non-epistemic bias. Finally, he demonstrated that these conditions were met by a communal scientific theory of acceptance.

Solomon's social empiricism presented us with a different way of analyzing the same sorts of situations. Recall Solomon's discussion of epistemically neutral decision vectors and how they were employed. If we move to this neutral terminology, then we can suggest that Jones talk about empirical and non-empirical decision vectors rather than epistemic or non-epistemic biases. This, in turn, allows us to avoid the ambiguity of the term "bias." Factors such as elegance do not have to be categorized based on whether elegant theories are more likely to be true. Instead, elegance can be taken as a non-empirical decision vector and can be adequately analyzed by its relation to other decision vectors.

Furthermore, social empiricism meets Jones' required conditions for bias-free results, and also meets the conditions for a communal scientific theory of acceptance. In fact, social empiricism gives us specific conditions by which we are able to judge whether there should be scientific consensus or dissent.

In contrast to Solomon, consensus for Jones was a rather vague concept. Consensus occurred when the community accepted a particular theory. This acceptance is 'dependent upon the lack of challenges made against them" (Jones 2002, 519). Consensus for Jones is therefore

vague, since it is not clear what the conditions are for community acceptance. On the one hand, we might interpret this quite literally, and argue that dissident positions are challenges, and thus there is no community-accepted theory. In which case, Jones' argument for the communal acceptance of the dominant theory (as justification for its epistemic superiority) falls apart. If it is not communally accepted, it doesn't possess the special status he wants to grant it. However, if we interpret his definition of consensus loosely, then we might say that consensus occurs when challenges do not arise for a period of time. But how are we then to interpret 'do not arise for a period of time?'

In order to sort this out, we need to put some context on this situation. The case that Jones focuses on is AIDS, and whether or not HIV is the direct cause of it. Given the nature of this case, HIV is believed to be the cause of AIDS, and the issue was for many years closed. By social empiricist terms, we can say that the debate over HIV and AIDS reached consensus. For the social empiricist, this is because the dominant theory has all of the empirical successes. For Jones, it is because there are no challenges brought against it. The separation between these two views occurs in the next stage, dissolution of consensus.

For Jones, dissenters should be taken as those that oppose the majority position. This puts them at odds with the community. The community's acceptance of the theory seems to be a 'final word' in the sense that the challenges that are now raised by the dissenters are not seen as challenges to the community accepted theory. This conclusion is derived from the fact that the dominant position is still given the previous 'communally accepted' status (and the epistemic bonuses that goes with it), despite the fact that the community does not accept (wholly) the position.

From a social empiricist view, we can say that the dominant position can only hold such a status if it still retains all of the empirical successes. If that is not the case, then consensus is

dissolved back to dissent, and the new theory deserves attention, and distribution of research effort. Thus, it is assumed from Jones' argument that this is the interpretation he means. However, this is not from evidence. At no point does he state that the people to whom Mbeki refers do not have empirical success in their theory. One would hope that Mbeki would only side with those who had some empirical success. Similarly, one would hope that Jones would not make consensus the 'final word.' But, hopes are irrelevant. All that is certain is that Jones' argument is vague with respect to consensus and dissent, and since his argument hinges on a lucid account of consensus and dissent, we have to find his argument wanting. Therefore, we are still in search of a means of determining the experts to whom we should defer in cases of uncertainty.

Another Take

Alvin Goldman's paper "Experts: Which Ones Should You Trust," (2001) will serve to set the tone for the rest of this paper. In his paper, Goldman explores the issue from the viewpoint of the layperson. As such, he makes clear how he intends to define some of the key terms of the debate and then uses those definitions to delve deeper into some of the intricacies of our debate.

To begin, we should try to put some context on the relationship between layperson and expert. This understanding will help to elucidate the reasoning behind Goldman's definitions of his terms. The relationship between layperson and expert has been seen for the most part as a matter of testimony from the expert to the layperson. This has led some authors to argue that the relationship leaves the layperson as being 'blind.' The term 'blind' here is meant to note the difference between an expert appealing to an expert and a layman appealing to an expert. In the case of expert appealing to an expert, the person doing the appealing has some knowledge in the domain of the expert to whom they appeal. This allows the person to do what Phillip Kitcher calls "calibration" (Kitcher 1993, 314-322). The person can use direct calibration, in which he uses his knowledge about the subject domain to assess the 'weight' he should give to the other expert's testimony. Or, he can use an indirect calibration, in which he uses the opinion of another expert (of whom he has directly calibrated) to calibrate the expert's testimony. In either case, the person is acting with knowledge within the domain.

By contrast, the layperson appealing to an expert is "blind" in the sense that the layperson has no knowledge in the expert's domain. In this case, he has to take the testimony at face value. This means that when the layperson must decide between two expert opinions, he needs to find some means of justifying his choice of one expert over another. Unlike experts, he cannot fall back on his knowledge within the field to help classify one expert as being more credible. I don't

see the state of being "blind" as necessarily being a bad thing. The layperson begins the debate from a neutral stance. Thus, he might hear the testimony uncolored by calibration. Basically, my sentiment here is simply that calibration can give an expert a more proactive approach towards an expert's opinion. Whether or not calibration is *always* beneficial, I think is debatable.

Goldman uses a good example to describe the layperson's relationship to the expert. He describes the layperson as being in a similar position to that of a member of a jury.

> A listener—for example, a juror—did not himself witness the crime, and as no prior beliefs about who committed it or how it was committed… He wants to learn what transpired by listening to the testimonies of eyewitnesses. (Goldman 2001, 90)

The relationship is analogous because the eyewitnesses will claim to have knowledge of some event that the jury member knows nothing about. Furthermore, it is up to the jury member to sort out which eyewitnesses are more credible.

Having gotten a clearer picture of the relationship between laypeople and experts, we can turn to the terms. For Goldman, expertise can be divided into a few categories. First, there is an objective sense of expertise, "what it is to *be* an expert" (Goldman 2001, 91). This is better understood in contrast to a reputational sense, in which an expert is such because people "believe him to be an expert (in the objective sense), whether or not he really is one" (Goldman 2001, 91). The objective sense of expert relies on the individual having significant beliefs in true propositions, and fewer beliefs in false propositions than most people. However, this is not meant to imply that the expert is more accurate. It could well be the case that the few false propositions he possesses are critical.

Second, we divide experts into what Goldman calls the "strong" and "weak" senses. To understand this, we need to first make a distinction. We can say that a domain is made up of primary and secondary questions. Primary questions are questions that concern current research

in the field, while secondary questions are questions relating to evidence, or arguments related to primary questions. Thus, “the ‘strong’ sense of expert would be an expert who has extensive knowledge in a domain of both primary and secondary questions” (Goldman 2001, 92). The ‘weak’ sense would be someone who has extensive knowledge, but that “knowledge is more restricted to the secondary questions” (Goldman 2001, 92).

Evaluation

With all of this knowledge of experts and the relationship between layperson and experts, what are the ways in which laypeople can evaluate experts? The first of these would naturally be to examine the arguments made by experts (Goldman 2001, 93). We can divide the claims made by experts into two categories: esoteric and exoteric. Esoteric claims are domain-specific claims of which the layperson cannot determine the truth values (Goldman 2001, 94). These are claims that may come from deep understanding of the domain. Exoteric claims, on the other hand, are claims that are outside the layperson’s typical domain of experience, but have a truth value that the layperson may have access to (either now, or at some point in the future) (Goldman 2001, 94). This distinction is rather important since exoteric claims are claims that the layperson can reason with relatively well. Esoteric claims, however, are not only difficult to assess truth values for, but also limit the layperson’s ability to see the relationship between premises, or between premises and conclusion.

These claims become more important when we understand that laypeople can accept arguments based on direct or indirect justification. Direct justification of an argument is when the layperson understands the premises, as well as how the premises relate to each other and the conclusion (Goldman 2001, 94). Thus, the layperson is justified in accepting the argument because the argument makes sense. By contrast, indirect justification occurs when the layperson accepts the argument based on the superiority of argument given by an expert (Goldman 2001,

95). Here we are referring to the expert's skill at arguing, whether it come in the form of debate or something else. Goldman uses the example of one expert presenting a counter-claim to another expert's argument. When given the chance for rebuttal, the other expert is not able to counter each of the claims. Thus, the layperson may decide to go with the expert whose testimony was not defeated, not because he is convinced of the argument, but rather indirectly, as his challenger could not offer similar counter-claims (Goldman 2001, 95).

The second way that laypeople can evaluate expert testimony is to assess the quantity of experts that back the opinion in question (Goldman 2001, 97). This way is commonly used and has even become a precedent in the courtroom. However, the question of numbers can be very misleading. As such, Goldman divides the issue into two related questions. The first is still the quantity of experts, but the second is an appraisal of the experts' credentials. The second question is aimed at sorting out the credentials of the people who agree with the expert.

The two examples that Goldman uses to highlight why numbers don't *necessarily* impart more credence to an opinion are the issue of rumors, and the issue of a guru (Goldman 2001, 98). Rumors are usually widely believed. They are often well-circulated and accepted despite the fact that no-one has access to a "credible source" of the information. As is the case with rumors known as 'urban myths,' most rumors are false regardless of the number of people that believe them.

The other example is a bit more important. This example involves a guru and his students who slavishly believe what he believes. The reason that this issue is more important is because it highlights the second part of the numbers issue--that of the credentials of the people giving credit to the expert. One should look at the *conditional independence* of those that back the position. In the case of the guru, his students will believe whatever it is that he believes. Therefore, they do not have independence from his position. Perhaps, several scientists are getting funding from the

same company, and so they hold the same viewpoint on an issue, say the effectiveness of that company's drug.

This latter point brings us to the next condition for evaluation, interest and bias (Goldman 2001, 104). Here again, the issue of conditional independence is important. However, in this case, the conditional independence isn't primarily between scientists, so much as between the scientist and his position. We can reformulate this issue by asking, 'Is there some outside reason that is motivating the scientist to want to advocate this position?' Scientists may exaggerate their findings (or the importance of their findings) in order to acquire (or maintain) funding. Goldman cites an article by Friedberg et al. (1999) in the *Journal of the American Medical Association* which examined the relationship between research reports made by company sponsors and independent researchers. Goldman says of this report, "Unfavorable conclusions were reached by 38% of nonprofit-sponsored studies, but by only 5% of pharmaceutical company-sponsored studies" (Goldman 2001, 105). These are incredible findings. It helps to point out that the need for more in-depth evaluation by the layperson.

Our final consideration for evaluation by laypeople is the past track record of the expert in question. Goldman quickly points out that this can appear to be a tricky issue. After all, if the layperson has no knowledge in the domain of inquiry, how can he evaluate the expert's past record of success? Is it not like the 'calibration' issue we discussed earlier?

To resolve this issue, we need to go back to the distinction between esoteric and exoteric claims. Some claims that may be exoteric now, may have been esoteric for a period of time previously. Goldman uses the following example: "There will be an eclipse of the sun on April 22, 2130, in Santa Fe, New Mexico" (Goldman 2001, 106). Today, the statement would be considered esoteric. However, on that very day, any layperson in Santa Fe will be able to walk out into the street and evaluate the correctness of the statement. It will then be an exoteric

statement. This example is a bit extreme, in the sense that the people in 2130 won't be evaluating the past track record of the expert (qua the claim here in 2006) in order to assess a claim the expert would currently be making in 2130. However, the point is still moving. Claims that were once esoteric can become exoteric overtime. It is precisely these types of claims that would serve as one means for the layperson to evaluate the past track record of experts.

None of these methods of evaluation are meant to be taken as the only means through which the layperson can appraise the credibility of experts. Goldman has merely explored the domain of this debate by focusing on the terms and the relationships that emerge. As such, Goldman's work will serve as the perfect launching point for our synthesis.

Synthesis

I want to move in a slightly different direction than my predecessors. To begin, I want to use a notion of 'expert' that is more general. Under this account, 'expert' is taken to refer to an individual who has significant experience in a domain of inquiry. Naturally, this experience will also have to be current. This definition will envelop both Goldman's strong sense of expert and his weaker sense. I somewhat dissolve his distinction because I wish to advocate an approach that can be taken more generally. I want to be able to put all experts on the "same page" so to speak, since the relationship between expert and layperson is alike in many ways regardless of expert level. In appealing to experts, it is not always clear whether we are dealing with a strong or weak form of expertise. And, as far as our discussion goes, the experts, that are being referred to as experts, will fall into both categories. However, I do not wish to completely ditch Goldman's distinction. I find the 'strong' and 'weak' notions to be useful, and will even apply them in a similar fashion to laypeople. For the remainder of the discussion, however, I will use expert to cover both senses of the term unless explicitly stated.

My next change will occur with the understanding of layperson. I will consider a layperson to be an individual who is appealing to the authority of experts. As with experts, this will be a general grouping that will have both a strong and a weak sense. In the strong sense of layperson, we are referring to individuals that have little or no prior knowledge in the domain of inquiry. Likewise, the weak sense of layperson is meant to refer to an individual who might have some degree of knowledge in the domain of inquiry (such as another expert). Since my account is meant to be taken generally, many of the factors that will be discussed influence both experts appealing to experts as well as laypeople appealing to experts.

Again, I do not mean to downplay the distinction between strong and weak laypeople. Just as with strong and weak experts, I find the distinction to be meaningful. However, I find

these sub-categories to be useful at a later point in the discussion. The present analysis will apply to anyone appealing to the testimony or authority of another regardless of the status of each. To further expand on this, let me describe the relationship in a bit more detail.

The Relationship Revisited

When we take the terms ‘expert’ and ‘layperson’ in a general sense, we are describing positions with specific guidelines and obstacles. The differences in weak or strong laypeople (and similarly experts) influence how these ends are met and how difficult the obstacles are to overcome. Nevertheless, though the degree of effort and specific approaches needed to meet guidelines and overcome obstacles may be different, the overall structure is uniform.

We can begin with the expert. What are the expert’s guidelines? And what are the obstacles that she faces? The expert is in a unique position, since she has knowledge at her disposal that is not necessarily available to everyone. Her ethical objective is to present this information in an understandable and unbiased manner. This is by no means an easy task. In order to do this, she must be motivated to try to understand and overcome the barriers that prevent effective communication with the layperson. The authority granted to her position by the expert-layperson relationship confers a duty on her (qua expert) to the layperson. In the case of science, this can be seen as a duty to the public good by scientists who receive public funding (through public or political means)(see Shrader-Frechette 1994). One finds the same duty placed on the head of a local church and even on the coach of a youth soccer team. Anytime that one finds an expert-layperson relationship, these duties and guidelines apply.

The expert is not without obstacles in this relationship. There is a great deal that she has to overcome. First, the expert must be able to translate from “special-ese” into, pardon the expression, “layman-ese.” This can be one of the most difficult tasks that the expert can face in this relationship. If an explanation is too technical, she risks losing the attention of her audience

and thus making the communication ineffective. On the other hand, by simplifying the message, she may further alienate her audience, who may feel talked down to or that perhaps the advice was so obvious that they failed to see the point.

Here, we can note one instance in which the degree of expert and layperson may be shaded. A “weak” expert may not use as strong of a technical vocabulary and thus might not experience a language barrier that is as distant. Similarly, a “weak” layperson may already have acquired a decent technical vocabulary and thus not need the same treatment as a “strong” layperson. What matters here is not the degree of the gap between them, but simply that there is this language barrier to overcome.

Second, the expert must be sure to avoid errors of framing and other biases (see Beauchamp & Childress 2001). Framing stems from presenting information in a way that encourages a particular response. For instance, many doctors will tell patients that they have a 90% chance of survival rather than telling them that they have a 10% chance of dying. Even though the information is the same, the effect of the framing commonly produces a more favorable outlook when framed by survival than when framed as a chance of death. This puts the expert in the difficult position of having to give an accurate assessment based on her knowledge, while keeping the point clear for her audience. She is in the precarious position of stating the fact of the matter, without up- or down-playing it. This is difficult since the layperson may not always understand why certain results are significant, so the critical point lies in communicating the significance, without biasing the information with interpretation.

Finally, the expert needs to understand her audience. This is to some degree a corollary of the previous two obstacles. Though it expands on them, this obstacle also addresses new conditions. For instance, there are times when the expert is asked to give an opinion that is outside the scope of her particular specialty, given that her field might be closely related to the

one at hand, or perhaps because her field is the closest applicable field. This is a different situation from the expert's removal of bias. In bias removal, the expert weeds out decision vectors that may cause *her* to overstep her bounds. Here, the expert needs to understand what parts of her background are applicable to the task at hand, so that she can adequately address the issue. If she does not understand the motivation behind seeking her opinion, then no amount of bias removal or framing will help yield an accurate opinion.

The lay perspective carries its own guidelines and obstacles. The layperson can have numerous reasons for seeking expert advice. Perhaps she seeks an expert opinion on some issue for which she has no knowledge. Or, she may just as well be seeking a second opinion to help her make an informed decision. Regardless of what drives her motivations, she needs to pay special attention to the hardships that accompany her position. For example, one of the most important points that she needs to understand is that experts are fallible. Expert testimony will not grant certainty. One need only take a brief look throughout history to see how concepts have changed over time. The notion of a 'caloric'[2] immediately comes to mind. Because of this, the layperson should not hold expert opinion as absolute, undisputable fact. Likewise, she should not be seeking an absolute fact of the matter. Such expectations will lead only to disappointment.

The layperson should not seek answers from an expert that is outside the scope of her specialty. Though the layperson may not always be in a position to make such a distinction, she should be aware that just because a question can be asked does not mean that it will have an answer. This is especially important when it comes to matters of political importance. One should always be prepared to accept that there is no definitive answer for the time being, and that

[2] Caloric was once a proposed 'fluid' whose movement constituted the gain or loss of heat. It has since been replaced with the term 'calorie,' which represents the amount of heat required to raise 1 gram of water 1 degree Celsius at 15 degrees Celsius.

more research is needed. Making hasty decisions may end up costing more in resources than taking the time for more careful research.

The layperson should also not seek to merely support her own intuitions. Seeking an expert to back up one's opinion is not an adequate means of seeking expert testimony. The point of deferring to experts in the first place is to try to get a more informed viewpoint on the matter. Though we may often be biased towards our gut reasoning, seeking expert testimony to back up such gut feelings is no more objective than making the decision alone. One is seeking only to validate prior assumptions and is not in effect open to other points of view, thus rendering the matter essentially closed from the beginning.

Here, it should be noted that 'calibration' can have an effect. A "weak" layperson may be able to calibrate an expert's testimony. However, they should be aware that even such calibrations can serve to heighten this effect. Just as we should not seek to merely back up our intuitions, we should also not seek out information to simply "validate" our calibrations. We should note that this only applies to situations in which the information concurs with previous calibrations.

Another hang-up that she should be aware of is that not all expert opinions are good expert opinions. What I mean here is that some experts use their knowledge to advance particular views rather than to support objectivity. This may be one of the most difficult snags with which the layperson must deal. She must assess as much as she can about the person presenting her with information. What are this person's affiliations and motivations? Are they advocating this view because it is the case, or because they have some bias towards promoting it? Recall the Goldman example in which a study from an independent pharmaceutical testing agency yielded a more critical result than one from the company's own testing team. What is the content of their argument? As in, are they deliberately trying to go over your head to sway you with technical

esoteric jargon, or are they sidetracking you from the technical side too much with simple analogies?

The layperson is not always in a position to distinguish between those who are experts and those that are trying to con them, but this does not absolve the layperson of the need to question the expert's motivations. Accepting an opinion on blind-faith is no better than making an uneducated guess. She should remember that the goal in deferring to an expert is an effort to educate herself more efficiently before making a decision. What I have discussed thus far by no means covers the wide range of impediments that encumber each side in this relationship.

Informed Decision Making

Thus far, we have explored some of the different ideas surrounding the problem of expertise. Now it is time to bring everything together in a more uniform manner. Having explored the relationship of the expert and layman, we are now in a position to design a schema, or flowchart, that one can follow in order to make an informed decision.

If we recall, Jones made a strong argument for the use of numbers in our decision process. His argument was aimed at demonstrating the epistemic superiority of the majority position, and I believe it is safe to say that his advice is good advice, if we have no time or resources available to look deeper into the issue. Thus, we can say that our initial starting position within the flowchart to be:

1- *If no further information is available, go with the majority held position.*

Since this is the starting position, I feel we need to further clarify what exactly this entails. If we again recall Jones' argument, we can see that the majority held position, all else being equal, has a degree of epistemic superiority. Since all experts are fallible, the majority may still be on the 'wrong' side of this debate. However, if we cannot get any further insights into the problem, then we have to side with the position that is more likely to be on superior ground. It is only in the

event that we have more time, or more resources available to us, that our previous critiques of Jones hold up.

For instance, Goldman produced examples that demonstrated how going by numbers could yield a problematic result. But, the steps that we can take to avoid, or even assess, such a situation depends on our capabilities of having access to more information. It will rarely be the case that we do not have such time or resources. In the cases of public policy, one would hope that more extensive steps are taken to fully understand the problem. As in the AIDS case, the ultimate decision from the case affects whether or not a significant portion of the population receives HIV medication. Decisions like these should always merit some time devoted to exploring each side of the issue.

In the event that we are able to access more information on the issue, then we can move from our starting position to step 2.

2- Why is the minority a minority?

This is an important first step to take in the assessment of the two conflicting opinions. In this step, we are trying to understand the minority position in its relation to the community. Here, we are not looking at the issue of the numbers, but at the justification of their position. We can translate this question using Solomon's terms as: 'Does the minority position have any empirical success?'

We might find this question to have a simple answer, 'no'. In such an event, we can effectively rule out the minority opinion and move on to the end of our flowchart. By contrast, the answer may in fact be 'yes,' in which case, we can then proceed to step 3.

3- Are there any empirical successes on either side in the domain of inquiry?

It is important for us to pause here for some added clarification. We might be tempted to ask ourselves, 'Why not go directly from step 1 to step 3?' It would appear as though the only

thing of interest to us is the particular domain in question (and hence the empirical success specific to that domain). Following from this, it may not immediately be clear why we should look at all at the empirical successes in general (as per step 2), since they may indeed not be in or related to the domain in question.

Though this is a valid point, it must be stressed that step 2 becomes critical for two reasons. First, we are moving in steps based on the amount of information that we have available to us. As we noted in the initial step, it may be the case that we have no other information accessible. It follows from this, that just because we are able to have access to *some* information, it does not imply that we have access to *much more* information. Step 2 requires very little in the way of extra information on the problem at hand, yet it still provides us with a quick solution. If we recall the problems that arise by simply going by numbers, then we know that it may not always be the best position to take. However, if we can, by using very limited information, determine a powerful factor such as empirical success, then we can make a decision that is not purely a matter of numbers. There will be an important justification for siding with the majority. Keep in mind, empirical success does not ensure truth. However, empirical success shows that there is at least 'some truth in the theory.' Consequently, a lack of empirical success is just hypothesizing. And, we are certainly better to side with a position that can back itself up with data.

The second reason that step 2 is important lies in the fact that it gives us a means to calibrate step 3. Just because a theory possesses some degree of empirical success does not mean that it has empirical success in the domain in question. Therefore, step 2 becomes the launching point for step 3. In step 2, we are to assess the empirical successes of the theory, while in step 3, we are to evaluate them. Though we might find a simple answer to step 3, such as only one side having empirical success in the domain of inquiry, and thus we can simply opt to side with that

position, it can just as easily be the case that neither side has any empirical success in the domain of inquiry. Through the evaluation of step 2, we know that there are empirical successes outside the domain of inquiry, and we can turn to those for further examination.

Let us explore such a case. If we determine that as a result of step 3, neither side has an empirical success in the domain of inquiry, we are left evaluating the other empirical successes of each side. It should be stressed that at this point, we are essentially going to be making a decision with a lot of risk. Since neither side has any empirical success regarding the specific phenomena in question, we have drastically reduced our chances of getting a correct assessment of the problem. This does not mean, however, that we cannot still make the best of what we have. Should such a case occur, we should look at two things. First, are there empirical successes in related domains? And second, how unified is the theory based on its empirical success?

Thus far, we have explored both what to do in the case that only one side has an empirical success in the domain of inquiry, and what we should examine in the event that neither side possesses any. Now, let us examine the third option. What should we do if both sides have empirical successes in the domain? For this, we need to turn to step 4.

4- Can we reduce the community to a group of specialists?

This step is crucial in many respects, so it is important that we take the time to understand the many facets of this step.

First, we might be tempted to ask ourselves, ‘Why do we not do this earlier?’ Or, we might ask, ‘Is this not what we have done initially by deferring to expert opinion?’ I think that these are valid concerns, but as with earlier steps, I think further intuitions will help to resolve these questions. It is not going to be the case that the expert to whom we defer is indeed the specialist himself.

Let us suppose that we have a committee of members from various fields. One of them may be an expert in physics, generally. This means that he has a significant understanding of the questions that drive research in many of the sub-fields of the physics community. In this respect, he may not have specialization in the area of inquiry, but may be a member of the board because he would have better insight into which sub-field (and thus which specialists) the problem might default to. Furthermore, if we examine the series of steps that we have taken thus far, we can see how someone in such a position might be able to adequately resolve the above steps.

First, he can directly inform the committee of what the majority held position is, as well as, what some of the minority positions might entail. He would have access to this by being a 'generalist.' Second, he would also be able to account for any empirical successes that the minority might have. Again, this would come from his knowledge as a generalist, and from publications in respected peer-review journals. Finally, he would also have access to knowledge of whether or not each side has empirical success in the domain of inquiry. He could derive such knowledge again from his expertise in the current issues of the various sub-fields. Nevertheless, even though this expert might be a good choice for membership on some committees, it might not be wise to trust his judgment on whether the majority or minority position is most likely to be correct (because he does not have detailed expertise in the domain under consideration).

Now if we grant the feasibility of this case, and I certainly don't see why not, then we have established a situation in which one might need to defer to a group of specialists themselves. In this case, the expert to whom we have been deferring simply doesn't have the expertise in the area of inquiry, despite his knowledge of the major issues in that area.

This is not the only reason that we might want to reduce the community to a group of specialists. By reducing the community from the community at large to a group of specialists, we are essentially removing some of the problematic factors that come along with taking the word of

the community on the whole. Remember, that earlier we established in our critique of Jones, as well as portions of Goldman's work, that the numbers are not necessarily accurate. The inaccuracies can be attributed to a number of factors.

One way is to derive them from an error of heuristics. Much of the community may be aware of the major questions and to some degree the success that a particular domain may have. These articles may receive more attention than some of the minor articles that explore the more technical details that are only important to the specialists themselves. Because of this, the availability heuristic may affect the community-at-large. In psychology, the availability heuristic states that we make decisions based on the information that is available to us via recall. For example, Kahneman and Tversky (1973) asked people questions such as, "Are there more words that start with the letter 'k' or words with 'k' as the third letter?" Despite the fact that there are more of the latter, it is easier for us to recall words in the former category. Thus, we misattribute the former as being more correct.

It is entirely possible, though by far not necessarily the case, that one might base their opinion off of the amount of articles that are available to them in recall. The number of articles available to them does not necessarily reflect the content of those articles. However in the group of specialists, it is more likely that they will have access to all of the information available, and beyond that possess a significant understanding of each article's importance. Under much the same format as Jones' argument for the superiority (epistemically) of the majority opinion over the minority opinion, we can say that the opinion of the majority of specialists would be epistemically superior to the majority of the community itself. It may very well be the case that the majority view of the specialists is the same as the majority view of the community at large. But, this does not follow necessarily, and thus merits the extra examination.

By reducing the community-of-experts-at-large to a community of specialists, we are essentially returning ourselves back to the initial position in our flowchart. Barring further investigation, we are in a position to follow Jones' advice and simply go with the numbers. The steps we have taken up to this point have increased the likelihood that going with the numbers will ensure the higher degree of accuracy. Essentially, the reduction can be seen to simply make the numbers "mean more" than they did previously.

Though we have improved the significance of the numbers, we have not yet removed all doubt. It can be said that in reducing the community, we tried to remove those who are siding with an opinion because perhaps it advocates the same theory that they operate under. One potential obstacle that appears from this reduction would be the possibility that the specialists simply exaggerate their case. They may give the impression that certain results in their field are more profound than they actually are.

To accommodate this possibility, let us look at our endpoint.

5- *Evaluate for factors that have a negative influence.*

This is the step that we will ultimately find ourselves at, and it is by far the most difficult of all. In this step, we are looking specifically for factors that Solomon referred to as "'non-empirical decision vectors." It is in this step that we will do our last refining of the issue. It is here that we will be able to accommodate issues like the one of specialists overstating their case. By analyzing the non-empirical decision vectors, we may be able to determine to what extent extraneous factors are influencing the importance granted to a particular result.

This is not the only method we can use. In this last step, we are trying to weed out as many biasing factors as possible. One area of importance that Solomon's method misses is the issue of numbers. Remember that even when we progress fully through the chart, we ultimately arrive back at the issue of numbers. Some argument was given for the numbers being more

accurate here, but regardless of how improved the case is, we still have not completely handled some of the factors that arose in Goldman's exploratory work. Recall that one point that Goldman stressed was the degree of separation between scientists. If one scientist is the crony of another, and he holds the view that he does because of his non-independence from the other scientist, then his view is not adding anything to the debate. Recall the example that Goldman gives of the guru and his followers; they all believe what they believe because the guru believes such.

Therefore, we can use Goldman's condition of independence to determine the degree to which the numbers mean what they do. This would be especially important to the issue of numbers that is left, when we complete a cycle through the chart. It is indeed possible that in reducing to the community of specialists that the matter of independence becomes increasingly important. After all, not only do the numbers mean more, but they are also derived from a smaller and in all probability 'tighter' group of scientists.

Finally, as part of this last step, we need to focus on the incredibly difficult point of assessing not just the community, but ourselves. Though this is a difficult evaluation to make, it is nevertheless important. When evaluating ourselves, we should pay special attention to some of the psychological factors that can influence decision-making. Though these factors affect both the experts and laypeople equally, it is more important to the evaluation of our self. We should recall that this effort has been made so that *we* can make better decisions. Though we may choose to defer to the expert opinion, the choice we make is ultimately our own. As such, we should pay special attention to the factors that can affect this process.

One of these factors that has yet to be mentioned is called 'hindsight bias.' Hindsight bias is akin to the "knew it all along" feeling. It may be the case that we had the correct intuition before ever evaluating the expert testimony. We should be careful, however, not to discount the

effort that was made. Simply because we had the correct intuition does not mean that we were justified in that intuition. One major point of all of this effort was to arrive at the most epistemically justified position, and not the best hunch.

Hindsight bias can be a contributing factor to another condition called the 'confirmatory bias.' This bias causes us to seek out information that validates previously held intuitions. Because of this, we are more likely to discount an opinion that does not coincide with our own. Naturally, we can see how this plays off of the hindsight bias. If something confirms our previously held intuitions, we are more likely to feel we 'knew it all along.' Similarly, the feeling of having already known the answer in turn contributes to our confidence is our previously held intuition, and thus we discount other opinions.

One approach that may aid us in this situation would be to try our best to analyze the non-empirical vectors (including our initial intuitions) that might be influencing our own judgments. It may serve us well in our evaluation of the theories with respect to ourselves to create a new list of decision vectors. Rather than focusing purely on the decision vectors that apply to the scientists (and the reasons for their positions) is to examine ourselves from a similar perspective. Under such an approach, we could derive what are the empirical reasons that we have for accepting a certain position, as well as the non-empirical ones.

By composing such a list, we would put ourselves in a better position to examine some of Goldman's categories, such as direct versus indirect justification. Our list of decision vectors may help us to see that we were appealing to one theory over another directly, because we arrived at it purely as a matter of empirical decision vectors (in the new list), or indirectly, because we are basing the decision on non-empirical vectors.

All in all, the evaluation of interests and bias serves to filter the final result of our inquiry. Though perhaps applicable to each stage of the process, this sorting would merely muddle the

previous steps. Each of the steps is intended to remedy a particular obstruction in our process of becoming more informed. I think that the process functions best if we leave each step as concise as possible, and thus more apt at achieving its desired effect.

Potential Criticism

Before I go further, I would like to respond to what I see are some potential objections to the process. The first criticism that one might offer would be to object to the return to numbers. It seems as though part of this process was designed to move away from the numbers. There were arguments made to highlight the problem of numbers, yet in the end, we are turning back to a majority versus minority issue, which should be subject to the same hazards as the larger community.

I think this criticism is a bit misleading. The problems that were raised concerning the issue of numbers were not meant to 'rule out' numbers as a means of evaluation. They were meant to stress some of the complexities that come along with choosing by quantity. It would be a fruitless endeavor to keep reducing the community down to two or so disputers. Instead, we want to make an evaluation of the numbers. In some cases, as Jones might argue, numbers are all we have, and that's fine. Given the circumstances, we did the best we could, and numbers do confer *some* degree of epistemic justification. However, if possible, we want to reduce the size of the community that we are deriving our number from. In this way, we restrict the community in order to get a more specific answer. The important point that needs to be stressed here is that the restricted group *may not be divided like the community*. In such a case, the majority opinion of the community may only be the minority opinion among specialists.

To help make the numbers matter, we also have incorporated some of Goldman's suggestions. Once we reach the final step, part of weeding out our biases lies in a further filtering

of the numbers. It is at that stage that we take into account the conditional independence of the scientists, as well as, any indirect justifications.

Another potential criticism revolves around the degree of persistence. Does the layperson need to get so specialized? It seems that we are trying to turn the layperson into a specialist himself with the degree of information we expect him to understand and sort through. The question I must retort with is why must the layperson feel obligated to become a specialist? Is it not the layperson's judgment that assesses their problem? Is it not also the time and resources that the layperson has available that do not also dictate the extent to which they must explore the issue? If the lay person thinks that they are content to cut it off their investigation early, then that is their choice. Obviously, the depth to which it should be pursued is contingent on the situation at hand. Some situations, like in a courtroom, will require significantly more attention, while other situations may require less. These are criteria for increasing the reliability of the informed decision process by the layperson, not absolute guidelines.

One final potential criticism that I would like to address concerns the abilities of the layperson. Can it not be said that we are assuming too much on the part of the layperson? Can the layperson actually accomplish this feat? Yes, I believe they can. Remember, this is a two-way street. Part of the interplay between expert and layperson is the responsibility of the expert to present his case in a manner that can be understood. The whole weight does not bear down upon the layperson alone. This is a means to improve the capacity for the layperson to evaluate the claims of experts. Communication is to some degree, out of his hands. There are plenty of cases where the layperson may not be able to get all of his questions answered. For example, we can imagine time constraints serving this role. Or in other cases, the layperson may receive the information indirectly, perhaps from a media source.

Conclusion

The goal here has been to explore the previous literature in hopes of deriving a process that can account for the many complexities that come along with the issue of expertise. With this in mind, I tried to take a more general approach that would be as applicable to everyday decisions, as it is to important matters, like public policy. It was also my goal to bring together many important aspects of the literature together. The work of Solomon and Goldman explored different aspects of the territory, but in their exploration left us with lingering questions. Jones, instead, attempted to answer a question, but failed to account for the scope of the domain. It was through a synthesis of these views that this more comprehensive account was developed. I think the work suffices to fill in the gaps left behind, and bring some unity to the different theories.

Bibliography

Beauchamp, T. and J. Childress (2001), *Principles of Biomedical Ethics*, 5th ed. Oxford: Oxford University Press.

Bloor, D. (1976). *Knowledge and Social Imagery.* London: Routledge & Kegan Paul.

Cooper, J. (1997). *The Complete Works of Plato.* Indianapolis: Hackett Publishing Company, Inc.

Fischhoff, B., Slovic, P., & Lichtenstein, S. (1981), Lay Foibles and Expert Fables in Judgments About Risk. *The American Statician,* Vol. 36, No. 3:2 (August): 240-253.

Goldman, A. (2001). Experts: Which Ones Should You Trust? *Philosophy and Phenomenological Research,* Vol. 63, No.1 (July): 85-110.

Hardwig, J. (1985). Epistemic Dependence. *The Journal of Philosophy,* Vol. 82, No. 7 (July): 335-349.

- (1991). The Role of Trust in Knowledge. *The Journal of Philosophy,* Vol. 88, No. 12 (December): 693-708.
- (1994). Towards an Ethics of Expertise. In D. Wveste (Eds.), *Professional Ethics and Social Responsibility* (pp. 83-101). Lanham, MD: Rowman and Littlefield.

Jones, W. (2002). Dissident Versus Loyalist: Which Scientists Should We Trust? *Journal of Value Inquiry,* Vol. 36: 511-520.

Kahneman, D., & Tversky, A. (1973). On The Psychology of Prediction. *Psychological Review,* Vol. 80: 237-251.

Kaiser, J. (2003). Sipping From a Poisoned Chalice. *Science,* Vol. 302: 376-379.

Kitcher, P. (1993). *The Advancement of Science.* New York: Oxford University Press.

Longino, H. (1990). *Science as Social Knowledge.* Princeton: Princeton University Press.

Pierson, R. (1994). The Epistemic Authority of Expertise. *PSA:Proceedings of the Biennial Meeting of the Philosophy of Science Association,* Vol. 1994, Volume One: Contributed Papers: 398-405.

Shrader-Frechette, K. (1994), *The Ethics of Scientific Research.* Lanham, MD: Rowman and Littlefield.

Solomon, M. (2001). *Social Empiricism.* Cambridge: MIT Press.

Weiss, C. (2003). Scientific Uncertainty and Science-Based Precaution. *International Environmental Agreements: Politics, Law, and Economics,* Vol. 3: 137-166.

Vita

Jonathan Tall was born in Jefferson Davis Parish in 1979. He spent part of his high school years attending Louisiana School for Math, Science, and the Arts and the rest of his time at Jennings High School. He began his undergraduate work with a scholarship to the Louisiana Scholar's College at Northwestern State University in 1997. He later changed to the General College at Northwestern to lighten the coursework from his double major in theatre and psychology. He took a brief period away from school, and upon return he would change to a single major in psychology with a minor in philosophy. He graduated with a B.S. in psychology in 2003.