

Forthcoming in *The Information Society*.

## “Epistemic Values and Information Management”

Don Fallis  
School of Information Resources and Library Science  
University of Arizona

Dennis Whitcomb  
Department of Philosophy  
Western Washington University

**Abstract:** The philosophy of information is concerned with the nature, management, and use of information. Thus, it should be able to help us make better decisions about how to manage information (e.g., decisions about intellectual property laws, collection development policies, and Internet evaluation techniques). These decisions have *knowledge acquisition* as one of their principal goals. Thus, one way to improve these decisions is to clarify our *epistemic values*. In this paper, we combine *epistemology* and *decision analysis* in an attempt to assist people in this important task.

**Keywords:** Applied Epistemology, Decision Analysis, Epistemic Value, Knowledge Acquisition, Knowledge Management, Information Ethics, Information Management, Information Science, Philosophy of Information, Value Hierarchy.

### 1. The Philosophy of Information

The philosophy of information is concerned with “how information should be adequately created, processed, managed, and used” (Floridi 2002b, 138). Although the term *philosophy of information* was coined only a few years ago, these sorts of issues are not new to philosophy (cf. Floridi 2002a, 44).<sup>1</sup> But recent advances in information technology have made these issues increasingly critical.

In contemporary life, some of the most important decisions that people must make are about the management of information (e.g., about the collection, organization, distribution, and evaluation of information). Legislatures have to decide which privacy and intellectual property laws to adopt, libraries have to decide which information resources (books, journals, databases, etc.) to collect and how to organize them, and individuals have to decide whether to trust the information that they find on *Wikipedia* or

---

<sup>1</sup> Many eminent philosophers of the past have discussed how we ought to manage information. For example, in the *Republic*, Plato argues that it is a good idea to misinform citizens on certain subjects. Also, in the *Phaedrus*, he argues that a speaker who intends to misinform someone must know the truth about her subject.

on the Internet in general. In this paper, we combine *epistemology* and *decision analysis* in an attempt to better equip people to make such information management decisions.

### 1.1. Applied Epistemology

The philosophy of information connects to several branches of philosophy. For instance, many theorists have applied *ethics* to information management, addressing such issues as freedom of speech, online privacy, and digital rights management (Floridi 2001, Moore 2005). But while ethics is certainly important to information management, *epistemology* is central (Fallis 2004a, 102, Floridi 2004b). We create and disseminate information because it allows people to *acquire knowledge* (Fallis 2006b, 486-490). Thus, we should try to make information management decisions that promote knowledge acquisition. For example, intellectual property laws should encourage people to discover and disseminate new knowledge, and library acquisitions policies should provide patrons with materials that make them more knowledgeable. Since epistemology is the study of what knowledge is and of how people acquire knowledge, it can potentially help us to manage information in ways that lead to the acquisition of knowledge (Goldman 1999, 161-217).<sup>2</sup>

Most contemporary epistemology is divorced from practical decision making. It tends to focus instead on such questions as whether we know anything at all, given that we might be massively deceived like the characters in *The Matrix* (Steup and Sosa 2005). Even epistemological work in the philosophy of information tends to be highly abstract. This work often applies information theory to traditional epistemological topics like the nature of knowledge (Dretske 1981, Floridi 2004a).<sup>3</sup> But just as information theory can be applied to epistemology, epistemology can be applied to practical information management decisions. In the same way that philosophy of information needs applied ethics, it also needs *applied epistemology*.

Work in applied epistemology tends to focus on *individual* decisions about what to *believe* (Levi 1962, Battersby 1989, Goldman 1989, Maher 1993, Fallis 1997, Bishop and Trout 2005). But epistemology can be applied to all sorts of decisions where knowledge acquisition is at stake. R. W. K. Paterson (1979), for example, used epistemology to evaluate curriculum policies in education. Philip Kitcher (1993) looked at how the scientific community ought to be organized to promote knowledge acquisition. Alvin Goldman (1999) subsequently applied epistemology to social policies in several other areas, such as the legal system and the political system. Most important for our purposes here, epistemology can also be applied to *information management decisions*.<sup>4</sup>

---

<sup>2</sup> As Floridi (2002a) has pointed out, the philosophy of information is much broader than epistemology. But epistemology is still a central component of the philosophy of information (Fallis 2006b, 508). In survey articles on the philosophy of information, epistemology typically features prominently (Herold 2001, Floridi 2004b, Greco et al. 2005).

<sup>3</sup> Fred Dretske (1981, 86) claims that “*K* knows that *s* is *F* = *K*’s belief that *s* is *F* is caused (or causally sustained) by the information that *s* is *F*.” Floridi uses results about communication protocols to argue that the Gettier problem is unsolvable. See Maruyama 1961, Popper 1972, von Foerster 1980, and Harms 1998 for other applications of information theory to standard debates in epistemology.

<sup>4</sup> When epistemology has been applied to information science issues, it has usually gone under the label *social epistemology* (see, e.g., Shera 1970, Goldman 1999).

A few epistemologists have begun to apply their work to information management (Suppe 1985, Floridi 1996, Goldman 1999, 161-217, Fuller 2002, Fallis 2004a).<sup>5</sup> In addition, quite a lot of applied epistemology has been done by information scientists (Shera 1970, Neill 1982a, Wilson 1983, Nitecki 1993, Dick 1999, Alavi and Leidner 2001, Andersen 2002, Furner 2002, Hjørland 2002, Budd 2004, Fallis 2004b, Svenonius 2004, Thellefsen 2004, Day 2005, Fallis 2006b, Gault 2006, Zins 2006, Rowley 2007). Some of this work focuses on purely conceptual or foundational questions. For example, Jennifer Rowley (2007) surveys how the concept of information has been distinguished from the concept of knowledge within the discipline. But these information scientists always have an explicit or implicit application to concrete issues in information management in mind (Shera 1970, 88, Dick 1999, 305). For example, many of them (e.g., Shera 1970, Andersen 2002, Svenonius 2004, Thellefsen 2004) have looked at how information should be *organized* so as to facilitate knowledge acquisition. But epistemology has also been applied to many other activities in information management, such as collection development and reference service in libraries (Fallis 2006b, 497-498).

We think that epistemologists working in the philosophy of information should pursue such applications. Moreover, since people need to make many important *decisions* about the management of information, we argue that epistemology should be put explicitly into the framework of *decision analysis*. A critical component of this project, that has not received very much attention, is to clarify the *epistemic values* that should guide such decisions. In this paper, our central focus will be to fill this critical gap in the literature.<sup>6</sup>

## 1.2. Layout of the Paper

In section 2, we discuss why *epistemic values* are so important in the context of information management decisions and we briefly survey what is currently known about epistemic values and what more needs to be discovered. In section 3, we suggest that *epistemic value hierarchies* need to be constructed to make good information management decisions and we show how work in epistemology can help us to construct such value hierarchies. In sections 4 through 6, we focus on three specific areas where epistemology can help: viz., identifying epistemic values, identifying constraints on the structure of epistemic value hierarchies, and identifying some default structures for epistemic value hierarchies. Finally, in section 7, we discuss the possibility of constructing sample epistemic value hierarchies for particular types of information management decisions by considering epistemic value hierarchies for *collection development* decisions.

## 2. Decision Analysis

---

<sup>5</sup> In fact, several of the early philosophers of information (e.g., Plato 1961 and Nietzsche 1997) cited by Floridi were doing applied epistemology of just this sort.

<sup>6</sup> In this paper, we focus on how epistemology can help us make practical information management decisions. But just as work in applied ethics has influenced ethics more generally, work in applied epistemology can influence epistemology and the philosophy of information more generally (Floridi 2002a, 38).

There is a standard procedure for making good decisions: *decision analysis*. This procedure has been applied to all sorts of decisions, including locating nuclear power plants, developing environmental policies, and deciding on medical treatments (Keeney 1992, Kirkwood 1997, Gregory et al. 2005).

Decision analysis features many detailed techniques, but in essence, it boils down to the following steps.<sup>7</sup> First, we figure out what alternatives are available to us. Second, for each alternative, we figure out what the outcome of making that choice is likely to be.<sup>8</sup> Third, we figure out how *valuable* each of the possible outcomes is. Finally, we simply choose the alternative that leads to the most valuable outcome. While we do not want to subject all of our decisions to this sort of careful analysis, we do want to use this procedure for our *important* decisions. As suggested above, information management decisions often fall into this category.

The focus of empirical research has largely been on the second step in this process (Fallis 2007c, 30). Many researchers have tried to determine what the consequences of various information management policies are likely to be (see, e.g., Bishop and Trout 2005, 24-53). For example, researchers have looked at which techniques for evaluating information on the Internet are most likely to lead people to acquire true beliefs (see, e.g., Fallis and Frické 2002, Kunst et al. 2002).

There is, however, an important sense in which the third step in this process (i.e., clarifying our values) is the most critical (Keeney 1992, Saracevic and Kantor 1997, 562). You have to think clearly about what you value in order to produce it. For example, in order to pick a good job, you need to clarify your thoughts about what makes jobs good (e.g., money, excitement, flexible hours). In addition to allowing us to evaluate our alternatives, clarifying our values can often suggest *new alternatives* (Keeney 1992, 9). Thinking carefully about what makes jobs good, for example, might lead you to apply for different types of jobs or jobs in different parts of the country. In fact, clarifying our values can sometimes suggest *new decision opportunities* (Keeney 1992, 16-17). Thus, even if you were not on the job market, thinking carefully about what you value might lead you to decide that you should be on the job market. For these sorts of reasons, we focus here on the third step in decision analysis.

## 2.1. Focusing on Epistemic Values

Our typical reason for seeking information is to acquire knowledge, or at least beliefs with knowledge-making properties like truth and justification (Fallis 2006b, 497-498). For example, suppose that you are deciding how to evaluate information on the Internet. Presumably, you are doing so in order to acquire *true beliefs* (e.g., about the safety and effectiveness of the medical treatment that you are considering). Thus, when people

---

<sup>7</sup> The normative structure of decision analysis is the same as any other articulation of *decision theory*. The techniques have simply been tailored to facilitate practical decision making.

<sup>8</sup> Actually, this step is a bit more complicated when there is significant *uncertainty* about what the outcome is likely to be. See the section below on “How Epistemic Values are Distributed.”

make such information management decisions, the relevant values tend to be *epistemic* values. Work in epistemology can help us to determine exactly what these epistemic values are.

In focusing on epistemic values, we are focusing on a critical issue in making information management decisions. But it is important to note that we are bracketing several other issues that are also critical in making such decisions. First, we are focusing on only one of the two roles that knowledge plays in information management decisions. In order to make any type of decision, we have to know what the outcome of each alternative is likely to be and we have to know which outcomes are valuable (Kirkwood 1997). In other words, knowledge is always a necessary input to the decision making process. However, our focus here is only on knowledge (and other epistemic values) as a desired *output* of information management decisions.

Second, even in the context of information management decisions, epistemic values are not the only values that we care about. We also prefer to make decisions that are ethical, legal, have low financial costs, etc. However, our focus here as applied epistemologists is to simply get clear about the relevant *epistemic* values. Other specialists will have to help clarify the other values that we have. In addition, at some point, we must determine the relative weights of all these values.

Third, epistemic values may not always be the *ultimate* objective in information management decisions (Fallis 2006b, 487-488). For example, the ultimate objective of legislatures is hopefully to promote the well-being of citizens. The ultimate objective of individuals is presumably to live a good life. Knowledge is sometimes a constitutive part of such ultimate objectives. In other words, knowledge is sometimes valuable for its own sake (Goldman 1999, 75).<sup>9</sup> But whether or not knowledge is valuable for its own sake, it is typically valuable *as a means*. In other words, knowledge usually gets at least some of its value from being useful (Goldman 1999, 73-74).

Nevertheless, knowledge is a “fundamental objective” in the context of information management decisions. In other words, it is part of what we directly aim to achieve when we make such decisions. Thus, epistemic values should take center stage in information management decisions. Also, it is more efficient to focus on knowledge rather than all of the many valuable things that this knowledge leads to (Fallis 2004a, 103).

## **2.2. Objectives from Information Science and Knowledge Management**

Despite the growing importance of information management decisions, there has been little careful analysis of epistemic values in this context. Information science is the discipline that has spent the most time studying our objectives when we make decisions about the management of information (Saracevic and Kantor 1997). And information scientists are well aware that knowledge acquisition is a fundamental objective in such

---

<sup>9</sup> One of the authors (Whitcomb) believes that knowledge *always* has at least some value for its own sake. The other author (Fallis) is agnostic on this point.

decisions (Hamburg et al. 1972, 110, Neill 1982b, Budd 2004). However, there has not been much analysis of our *epistemic* values in the information science literature.

This gap in the literature is understandable. It is hard to measure the amount of knowledge people acquire.<sup>10</sup> It is also hard to isolate the effects of a single information management decision on knowledge acquisition.<sup>11</sup> As a result, information scientists have typically chosen to focus on correlated objectives, such as people being exposed to documents, that are easier to measure (Hamburg et al. 1972, 110-111).

While this gap in the literature is understandable, it is also lamentable. The objectives that information scientists have identified are often good means to knowledge acquisition. But things like exposure to documents are valuable *only* insofar as they lead to knowledge acquisition. We do not value for them their own sake. And, according to decision analysis, in order to make effective decisions, it is vital to focus on the objectives that are fundamental and not derivative (Kirkwood 1997, 22, Fallis 2006b, 488-490). As John Rawls (1971, 91) put it, “it is irrational to advance one end rather than another simply because it can be more accurately estimated.” In other words, we do not want to be like the proverbial drunk who looks under a lamppost for his keys, even though he lost them elsewhere, simply because the light is better under the lamppost.

Knowledge management (KM) is an area of information science that does focus explicitly on knowledge. However, KM researchers typically take knowledge to be a special kind of information rather than a special kind of belief (Alavi and Leidner 2001, 111, Rowley 2007, 172-174).<sup>12</sup> And, while the creation, preservation, and management of knowledge *qua* information is certainly a critical means to our goal, our goal is the acquisition of knowledge *qua* belief (Saracevic and Kantor 1997, 532, Budd 2004, 363, Fallis 2006b, 495-496). That is why, for example, there is no value to preserving materials in an archive unless it is possible that somebody at some time will use them to acquire knowledge.

Some KM researchers do define knowledge as something that is crucially related to the cognitive states of human beings (Alavi and Leidner 2001, 109). But these researchers do not think that the distinguishing features of knowledge have anything to do with the epistemic properties, such as truth and justification, typically discussed by epistemologists.<sup>13</sup> In fact, some KM researchers explicitly set aside work in

---

<sup>10</sup> Of course, it is possible to measure knowledge acquisition (Hunt 2003). In fact, we do so all the time. For example, we regularly give students tests to determine how much they know. But it is more difficult to measure the knowledge acquired by, for example, library patrons than the knowledge acquired by students in a classroom (Hamburg et al. 1972, 111).

<sup>11</sup> A number of other factors (e.g., educational institutions, family, community) have an effect on the amount of knowledge that people acquire (Hamburg et al. 1972, 110-111). But this is a concern for almost any important decision (e.g., many factors beyond emissions laws have an effect on air pollution).

<sup>12</sup> There are also a few philosophers (e.g., Popper 1972, 106-152) who discuss knowledge as something that exists independently of human beings.

<sup>13</sup> For example, Alavi and Leidner (2001, 109) claim that “what is key to effectively distinguishing between information and knowledge is not found in the content, structure, accuracy, or utility of the supposed information or knowledge.”

epistemology as being irrelevant to their project (Alavi and Leidner 2001, 109). But even KM researchers have to assume that knowledge does have good epistemic properties. If it did not, there would be little point in theorizing about how to manage it. Nonetheless, these researchers do not explicitly say what these properties are. In order to clarify what these good epistemic properties are and thus fill this gap in the KM literature, we need to appeal to epistemology.

### 2.3. Value-Theoretic Epistemology

Value-theoretic work has flourished in recent epistemology (Riggs 2006). Some theorists have examined our “epistemic goals” (Riggs 2003, Alston 2005). Others have asked why knowledge is better than mere true belief (Kvanvig 2003, Sosa 2007). In closely related work, virtue epistemologists have developed a theoretical focus on character traits like the love of truth (Zagzebski 1996). Additionally, several philosophers of science have tried to identify the *epistemic utilities* of scientists (see, e.g., Levi 1962, Maher 1993). And Alvin Goldman (1999) has proposed a theory of *veritistic value* that he has used to evaluate various social policies and institutions.

These projects are a good start, but they are not developed enough for use in information management decisions. Three avenues for further development are particularly salient. First, the work on epistemic goals typically only considers the value of true belief (and the disvalue of false belief). Many other epistemic values (e.g., knowledge and understanding) also matter in information management, and epistemologists have only begun to address these other values (Kvanvig 2005). Second, the question of why knowledge is better than true belief is only one among many questions about epistemic values. A number of other questions (e.g. about the relative weights of different epistemic values) arise when we make information management decisions. Third, the work on epistemic utilities only addresses scientists. However, a wide variety of people (e.g., library patrons, students, the public at large) are affected by information management decisions.

We need to fill these three gaps if we are to make value-theoretic epistemology useful in information management. The best way to fill them is by putting value-theoretic epistemology into the framework of decision analysis, a framework that allows decision makers to clarify their views about the values that their decisions should promote.<sup>14</sup> In the remainder of this paper, we begin developing value-theoretic epistemology into a tool for decision making using the framework of decision analysis.

### 3. Constructing Epistemic Value Hierarchies

If epistemology is to be put into the framework of decision analysis, exactly what needs to be done?

---

<sup>14</sup> At the moment, most work in epistemology is only comprehensible to, and only read by, other experts in epistemology.

According to decision analysis, in order to make a good decision, we first need to construct a *value hierarchy* (Keeney 1992, 55-98, Kirkwood 1997, 11-40). A value hierarchy is an organized structure that captures all of the values that are relevant to a particular decision.



Figure 1: A Simple Value Hierarchy for Information Management Decisions

Many different values can be important in information management decision making: *epistemic* ones, *financial* ones, *moral* ones like respecting privacy, *legal* ones like following copyright law, and possibly even *aesthetic* ones like spreading beautiful writing. And information management decisions should be made on the basis of all-things-considered value hierarchies that include all of the relevant values.

In many information management decisions, however, a simple list of values or objectives may not be sufficient to allow us to choose between our alternatives. As a result, each of the objectives might need to be further clarified. For example, the objective of “minimizing financial costs” is typically subdivided into the objectives of “minimizing fixed costs” and “minimizing variable costs.”

As noted above, our focus here will be exclusively on clarifying the *epistemic* objective of “maximizing knowledge acquisition.” Our claim is that applied epistemology can help people make better information management decisions by facilitating the construction of an *epistemic value hierarchy*. This hierarchy will represent one branch of the all-things-considered value hierarchy.

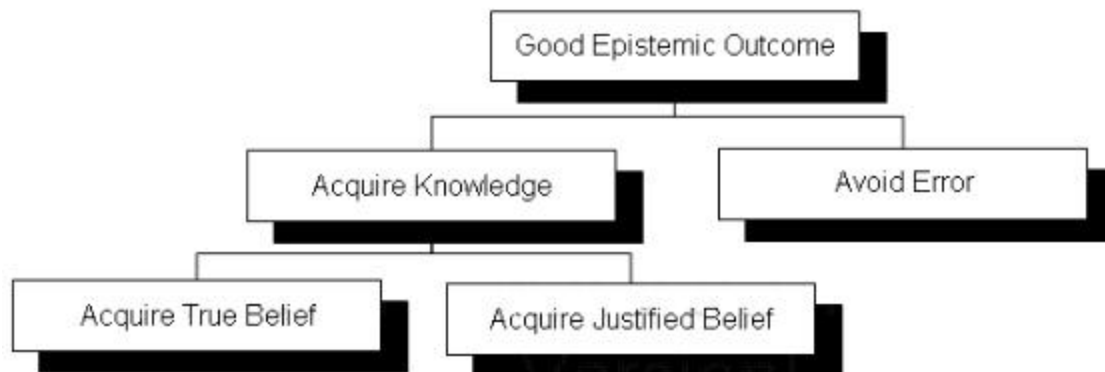


Figure 2: A Simple Epistemic Value Hierarchy

A common strategy for constructing value hierarchies in general is to take a value hierarchy constructed by someone with similar values and make modifications (Kirkwood 1997, 21). For example, most people looking for a job are concerned with whether the work will be interesting and with how much they will get paid (Kirkwood 1997, 34-35). Their values will typically just differ in the details (e.g., in the relative importance assigned to the different values). However, since no one has ever constructed an epistemic value hierarchy for making information management decisions, the strategy of using an existing hierarchy as a template is just not an option. Instead, we will have to begin by constructing epistemic value hierarchies from scratch.

There are two steps in constructing a value hierarchy from scratch. First, we have to determine what sorts of things are valuable in the context of a particular decision. As noted above, when making information management decisions, we are typically concerned with acquiring and retaining true belief, justified belief, wisdom, and understanding, and avoiding the opposites of these things, both in ourselves and in others. Second, we have to determine how these values are structured. That is, we have to determine the relative importance of these values, determine which values are only valuable as a means to other values, determine which values are only valuable in the presence of other values, etc.

A *complete* value hierarchy (i.e., one that includes all of the relevant values and that assigns definite relative weights to these values) is typically needed to make good decisions (Kirkwood 1997, 16). In particular, a complete epistemic value hierarchy will be needed to determine which alternatives do better in terms of “maximizing knowledge acquisition.” Ideally, leveraging work in epistemology would allow us to construct a complete epistemic value hierarchy that could be used for any information management decision.

However, epistemology may not be able to tell us *exactly* which values to include in an epistemic value hierarchy and *exactly* how these values are structured. First, there may be no epistemological fact of the matter about the relative importance of some epistemic values. In other words, the relative importance of different epistemic values may depend on non-epistemic considerations (Fallis 2006a, 182-183). Second, even if there is an epistemological fact of the matter, we may not yet know what it is. For example, philosophers (e.g., William Clifford and William James) often disagree about the relative importance of avoiding error and acquiring new true beliefs. Finally, exactly what it means to “maximize knowledge acquisition” may be relative to the context. In other words, which epistemic values are relevant and what their relative importance is might depend on the type of decision that is being made (Fallis 2006b, 500-503). For example, while avoiding error may often be more valuable than acquiring new true beliefs (e.g., when we are seeking medical information), acquiring new true beliefs may sometimes be more valuable than avoiding error (e.g., when we are seeking entertainment information).<sup>15</sup> In addition, how knowledge should be distributed *across people* arguably depends on the context (Goldman 1999, 96, Fallis 2007b).

---

<sup>15</sup> Similarly, the *accuracy* of the information may sometimes be more important than the *precision* of the information (cf. Floridi forthcoming). But sometimes it will be the other way around.

Even so, we can leverage work in epistemology to carry out the following important tasks. First, epistemology can help us to identify the things that are epistemically valuable. Second, it can help us to identify constraints on the structure of these epistemic values. More specifically, epistemology can constrain how these epistemic values fit together into a hierarchy, what their relative importance is, and how they should be distributed over topics, over people, and over time. Third, it can help us identify some default structures for epistemic value hierarchies. This will take us much of the way toward constructing complete epistemic value hierarchies for our information management decisions.

In the following sections, we will describe in more detail how these three tasks might be carried out by an applied epistemologist. We will also point out several places where additional research in applied epistemology needs to be done.<sup>16</sup> However, in order to construct a *complete* epistemic value hierarchy for a particular information management decision, it is important to note that we will probably also need to consult the values of the stakeholders in that decision. Admittedly, there are well-known difficulties with *accurately* eliciting the values of stakeholders (McNeil et al. 1982, Kahneman and Snell 1992). However, decision analysts continue to develop techniques for addressing these difficulties (Satterfield 2001). These techniques can easily be applied in the domain of epistemic values.

#### 4. Identifying Epistemic Values

Value hierarchies need to capture *all* of the values that are relevant to making particular decisions. If any relevant values are left out, we run the risk of making bad decisions. (We will end up shortchanging those alternatives that promote these missing values.) In order to make better information management decisions, we should identify as many epistemically valuable properties of beliefs as possible. The epistemological literature can help us to do so.

As noted above, epistemologists have largely focused on two epistemic values: true belief and error avoidance.<sup>17</sup> But there are further epistemic values as well. For instance, *knowledge* is something we should definitely not leave out of our epistemic value hierarchy. It is the *paradigmatic* epistemic good.

---

<sup>16</sup> Since work in epistemology has never been explicitly put into the framework of decision analysis, some questions about epistemic values are likely to arise that have not yet been addressed in epistemology (cf. Fallis 2004a). Thus, an applied epistemology that will help people make better information management decisions may require extending, as well as leveraging, work in epistemology.

<sup>17</sup> Information scientists (e.g., Wolkoff 1996, 87, Dick 1999, 315, Budd 2004, 364) have correctly noted that information professionals are not in a position to be the “arbiters of truth” for their clients and patrons. But that is a question about what the best *means* to achieving our goal is (i.e., it is part of the second step of the decision making process) and not about what our goal is. Admittedly, some information scientists (e.g., Shera 1970, 97) have claimed that information management decisions should not aim to promote true belief. But see Fallis 2006b, 490-495 for a detailed critique of this claim.

In addition, epistemologists have identified several components of knowledge that are valuable in their own right. For example, the standard account of knowledge tells us that a belief must be *justified* (i.e., adequately supported) as well as being true. And it is certainly epistemically valuable for a belief to be justified even if it does not count as knowledge.

There are many accounts of justification. *Evidentialists* have it that for a belief to be justified is for it to be based on good evidence. *Reliabilists*, by contrast, have it that for a belief to be justified is for it to be the product of a reliable process. Additional views abound. Pure epistemologists expend considerable effort disputing about which of these views are true. Without having to take sides in these disputes, the applied epistemologist can view each side as offering up potential epistemic goods (compare Alston 2005).

While justification and the other candidate conditions on knowledge are at the core of pure epistemology, they are not the only epistemic values (Riggs 2003, 347, Kvanvig 2005). It also matters that our beliefs are *explanatory* and *consistent*, and that we obtain higher-level epistemic goods like *understanding* and *wisdom*. Recent epistemology has begun to increasingly address these higher-level goods (Ryan 1999, Kvanvig 2003, 188-203). These values may be important to information management decisions, because different alternatives may be better or worse at promoting them (Neill 1982b).

Another category of epistemic values is also important in information management decisions. Traditional epistemologists typically talk about *beliefs* and their properties. But Bayesian epistemologists (e.g., Maher 1993, Kaplan 1996) talk about *degrees of belief* and their properties. For example, it is good epistemically to proportion your degree of belief to your evidence (Hume 1977, 73). It is also good epistemically to have a high degree of belief in a true proposition (Goldman 1999, 90). Such epistemic goods are critical because degrees of belief are particularly important to action. For example, you may believe that *p*, but not be sufficiently confident in your belief to act on it (see Hunt 2003, and for an alternative view, see Hawthorne and Stanley forthcoming).

Since an epistemic value hierarchy is only one branch of an all-things-considered value hierarchy, any or all of these epistemic values may in some contexts be outweighed by non-epistemic values.<sup>18</sup> For instance, sometimes it is better all things considered to have false beliefs than true beliefs (Stich 1990, 101-127, Zagzebski 2003, 22, Pietikäinen 2004).<sup>19</sup> In particular, there are information management decisions where it is bad to achieve certain epistemic goals. For example, libraries do not want anybody to acquire true beliefs about which books other patrons have checked out (McDowell 2002, 55-56).

But the aim of the applied epistemologist is to say what is *epistemically* good, not what is good all things considered (contrast Goldman 2002, 218-220). And even when false belief is good all things considered, it is bad epistemically. Of course, the epistemic good

---

<sup>18</sup> In addition, the relative weights of these epistemic values may well vary across decisions. We discuss this possibility in the section below on “How Epistemic Values are Weighted.”

<sup>19</sup> Studies have found that people who have an inflated view of their abilities tend to be happier than people who have accurate self-images (Taylor and Brown 1988).

should only be promoted when it coincides with the all things considered good (Bishop and Trout 2005, 20-21). But in information management, these two goods typically coincide (Fallis 2007c, 35-39). And even when they do not coincide, the epistemic good is still an ingredient of the all things considered good.

## 5. Identifying Constraints on Structure

In addition to determining which epistemic values can go into an epistemic value hierarchy, we also have to determine how epistemic value hierarchies can be structured. Decision analysis puts certain constraints on the structure of any value hierarchy. For example, the values at each level in a hierarchy must be mutually exclusive (Kirkwood 1997, 16-17). Epistemology provides us with additional constraints on *epistemic* value hierarchies.

As noted above, epistemology may not yield a complete epistemic value hierarchy for all information management decisions, but it can place limitations on what epistemic value hierarchies should look like. The applied epistemologist can help people to make better information management decisions by identifying such constraints. In subsection 5.1, we discuss constraints on how epistemic values *fit together* to form a hierarchy. In subsection 5.2, we discuss constraints on how *relative weights* can be assigned to epistemic values. In subsection 5.3, we discuss constraints on how epistemic values can be *distributed* (e.g., over time or over people).

### 5.1. How Epistemic Values Fit Together

In order to construct an epistemic value hierarchy, it is not sufficient to simply compile a list of epistemic values. We also have to determine how these values fit together to form a hierarchy. As we discuss in this subsection, epistemic values can be interrelated in a variety of ways. For example, some epistemic values are constitutive parts of other epistemic values. Also, some epistemic values are principally valuable as a means to the achievement of other epistemic values. Such interactions can have an important impact on how we construct epistemic value hierarchies.

#### 5.1.1. Constitutive Values

According to most epistemologists, knowledge is a special type of true belief.<sup>20</sup> In other words, truly believing that *p* is a constitutive part of knowing that *p*. This is important from the perspective of decision analysis because there is a danger of *double-counting* (Kirkwood 1997, 17). If we assign value to knowledge *and* to true belief, then in a sense we count the same value twice. This “double counting” can make an alternative look better than it really is. Thus, we cannot simply add knowledge as a value alongside true belief in an epistemic value hierarchy.

---

<sup>20</sup> A few epistemologists (e.g., Radford 1966) deny that knowledge requires belief, but they are a small minority.

We can steer clear of double counting by adding additional *levels* to the epistemic value hierarchy (see figure 2 above). In a value hierarchy, each level is a decomposition of the next higher level. That is, the values that fall under a particular value tell us what that value amounts to. For example, part of the value of knowledge is the value of true belief. Thus, even though the values *at each level* of the hierarchy must be mutually exclusive (Kirkwood 1997, 16-17), the value of knowledge and the value of true belief can both appear in an epistemic value hierarchy.

When we add an additional level to an epistemic value hierarchy, we have to be careful not to leave out any relevant values, however. When we decompose the value of knowledge, for example, we have to be able to say what *all* of its constituent values are. That is, we have to identify *all* of the things that are valuable about knowledge. While we do not have to identify all of the *constituents* of knowledge, the project in epistemology to give necessary and sufficient conditions for knowledge can help us to identify constituents of knowledge, such as justification, that are valuable.<sup>21</sup>

### 5.1.2. Means and Ends

Most epistemologists believe that justification is valuable as a *means* to the truth (BonJour 1985, 7-8, Sartwell 1992, Kvanvig 1998, 433-434, Zagzebski 2003, 14).<sup>22</sup> There are several reasons for thinking as much. First, beliefs that are justified are more likely to be true. That is, justification is *correlated* with truth. Second, true beliefs that are justified are more likely to be maintained (Mill 1978, 33-34). Third, we are better able to convince others to adopt our true beliefs if we have reasons for them (Jones 1997, 429).

Since justification is correlated with truth, gaining justified beliefs is a good way to gain true beliefs. For instance, since information professionals want their clients and patrons to gain true beliefs, one possibly strategy is to help them gain justified beliefs (e.g., by providing access to a wide variety of well-organized and authoritative information resources). Importantly, information professionals can engage in this project while refraining from being “arbiters of truth.”

But even though justification is highly valuable, it should not necessarily be assigned much weight when we construct an epistemic value hierarchy. When one objective is a means to another objective, there is the potential for double-counting (Kirkwood 1997, 22). Justification should appear in the hierarchy *only* if it has at least *some* value as an

---

<sup>21</sup> As Edmund Gettier (1963) showed, knowledge requires more than justified true belief. Some theorists claim that, whatever the additional conditions are, they have no effect on inquiry (Kaplan 1985, 355). Others deny this (Williamson 2000, 60-93). If the former theorists are right, then perhaps those additional conditions should be left out of our hierarchies. We take no stand on these difficult issues.

<sup>22</sup> Justification might also be valuable as a means to other epistemic goals. In fact, true belief and error avoidance might themselves be valuable as a means to further epistemic goals, such as understanding (Riggs 2003, 350).

end as opposed to a mere means. If its value is *wholly* instrumental, then we would in effect count its value twice over by including it in the hierarchy.<sup>23</sup>

### 5.1.3. Other Interrelations between Epistemic Values

Just as with knowledge, truth, and justification, other epistemic values can be interrelated in complicated ways. For example, understanding involves having certain sorts of knowledge (or at least true belief) and grasping explanatory connections (Grimm 2005). Also, wisdom involves having certain sorts of knowledge, such as knowledge of how to live well (Ryan 1999). And, just as with knowledge, truth, and justification, such interactions can have a profound effect on the construction of an epistemic value hierarchy.

In order to construct an epistemic value hierarchy that will allow us to distinguish between the available alternatives, we will need a detailed analysis of the various things (such as knowledge, wisdom, and understanding) that are epistemically valuable. This may not require the epistemologist to provide necessary and sufficient conditions for each of these things. But we do need to be able to answer questions of the following sort: Which things are constitutive parts of other epistemic values? Which things are primarily valuable as a means to other epistemic values? Which things are only valuable in the presence of other epistemic values?<sup>24</sup> Answering such questions is an important direction for future research in applied epistemology.

## 5.2. How Epistemic Values are Weighted

Once we determine what values are important in a particular decision and how they fit together in a value hierarchy, the next step is to determine their *relative weights* (Kirkwood 1997, 55-61).

Ideally, for any decision, we would be able to identify an alternative that does better with respect to all of our values than any other alternative. For example, you might hope to find a job that is more fulfilling *and* that pays more than any other job option that you have. However, whenever we value more than one thing, we will usually face tradeoffs. One alternative will do better with respect to one value, while another alternative does better with respect to another value. For example, one job may pay more while another job is more fulfilling.

The same thing happens with epistemic values. For example, believing the truth often conflicts with avoiding error. A particularly clear case of this problem arises when we evaluate information on the Internet. If we simply trust whatever information we find,

---

<sup>23</sup> Even if justification does not appear the epistemic value hierarchy for a particular information management decision, it should still appear in a *means-ends diagram* (Keeney 1992, 69-76). This diagram indicates exactly how those things that are valuable as means are connected to those things that are valuable as ends. The construction of this diagram is an important part of the second step of the decision making process.

<sup>24</sup> For example, it is often suggested that an explanatory belief is only valuable if it is also a true belief (Hempel 1965).

we will gain many true beliefs, but also many false beliefs. By contrast, if we only trust information on the Internet once it has been corroborated by a print resource, we will not make as many errors, but neither will we gain as many true beliefs.<sup>25</sup>

In fact, the conflict between gaining new true beliefs and avoiding error arises in many information management decisions. For example, there is the question of what to do about information that may be out-of-date. Consider the decision that a librarian must make when he purchases the new edition of an encyclopedia and has to decide whether to retain the old edition as well (Fallis 2006b, 501-502). The old edition probably contains a significant amount of information that is no longer accurate. However, the old edition probably also contains a significant amount of useful information that did not make its way into the new edition. At the very least, it is a good source of information about what people believed at the time that it was published. In order to decide what to do here, the librarian needs to determine the relative importance of patrons not acquiring false beliefs (from the out-of-date information in the old edition) and patrons acquiring true beliefs (from the useful information only found in the old edition). That is to say, the librarian needs to determine the relative weights of these different values (Edwards and Newman 1982, 52, Kirkwood 1997, 61).

As we suggested above, it may be impossible to provide a single answer about how to resolve these conflicts that applies to all decisions. However, there are some constraints on the matter. For example, since knowing that  $p$  involves truly believing that  $p$ , the value of knowing that  $p$  is always at least as great as the value of truly believing  $p$ .<sup>26</sup> Also, the value of acquiring new true beliefs can never completely outweigh the value of avoiding error. If it did, then people would do right epistemically to simply believe everything (Alston 2005, 32). And clearly, simply believing everything could never be right epistemically. Thus, even if epistemology cannot identify a specific tradeoff that is epistemically *required* for all decisions, it can rule out some tradeoffs as being epistemically *impermissible* for all decisions.

In addition to constraints on the relative importance of epistemic values, there are constraints on how epistemic value increases as degree of belief in the truth increases. For example, it cannot increase linearly (Maher 1993, 178-179, Fallis 2007a, 229). In fact, the function from degree of belief in the truth to epistemic value has to have a very specific shape (Fallis 2007a).

Of course, there seem to be limits to how much epistemology constrains epistemic value hierarchies. While many epistemologists (e.g., Hume 1977, 111) claim that avoiding error should always have much greater weight than acquiring new true beliefs, this would probably be too tight a constraint. There are many situations where it is reasonable for acquiring new true beliefs to have greater weight (Godfrey-Smith 1991). For example, the cost of failing to truly believe that there is a tiger in the area may be much greater

---

<sup>25</sup> This is analogous to the sort of conflict that can arise between *precision* and *recall* in information retrieval.

<sup>26</sup> Some philosophers (e.g., Jones 1997, 435, Kvanvig 1998) have argued that it is not always *strictly* greater.

than the cost of falsely believing that there is a tiger in the area.<sup>27</sup> Several epistemologists (e.g., Levi 1962, 57, Lehrer 1975, 71) have concluded that there are no constraints on the relative weights of these two epistemic values that apply in all contexts.

Some epistemologists (Maher 1993, 134-135, Kaplan 1996, 112-121) have also suggested that the value of having *consistent* beliefs always outweighs all other epistemic values. However, there are good reasons to think that this is not the case either (Lehrer 1975, Foley 1993, 162-173, Fallis 2005, 45-49).

In addition to identifying constraints, it is important to determine where epistemology does not constrain us in constructing our epistemic value hierarchies. At these points, non-epistemic values may play a role in determining the relative weights of epistemic values (Fallis 2006a). For example, the relative importance of acquiring new true beliefs and avoiding error may depend on exactly why these values are useful in the context of a particular information management decision.

### **5.3. How Epistemic Values are Distributed**

Epistemologists typically focus on identifying what knowledge is. There is much less discussion of how knowledge should be distributed over topics, over times, and even over people (Paterson 1979, Goldman 1999, Fallis 2007b). However, such questions about the distribution of epistemic goods are often the most critical when we are making information management decisions. Since different alternatives may lead to different distributions, we need to be able to say which of these distributions is epistemically better. But just as with the relative weights of epistemic values, it is not clear that epistemology provides a single answer that applies to all decisions (Fallis 2004a). Thus, we may need to consult the values of the stakeholders in particular information management decisions in order to construct complete epistemic value hierarchies.

#### **5.3.1. Over Topics**

Questions about the distribution of knowledge over topics arise, for example, in the grant allocation decisions of scientific funding agencies (Kitcher 2001), in the curriculum decisions of schools (Goldman 1999, 349-373), and in the collection development decisions of libraries (Fallis 2006b, 501-503). A library might have to decide between acquiring the latest text in neuroscience or the latest text in molecular genetics. The *subject-independent* epistemic values that epistemologists typically discuss, such as justification and truth, will not resolve this issue. A belief about neuroscience can be as true and as justified as a belief about molecular genetics. Thus, a complete epistemic value hierarchy would have to say what the relative significance of knowledge on these different topics is. There might also be cases where our choice is between more significant knowledge and better (e.g., more justified) knowledge. A complete epistemic value hierarchy would need to resolve such conflicts as well.

---

<sup>27</sup> Since having a true belief has positive value and having a false belief has negative value in this situation, these values seem to be epistemically permissible.

### 5.3.2. Over Time

Questions about the distribution of knowledge over time arise, for example, in decisions about how to evaluate information. It is plausible that, all other things being equal, acquiring knowledge quickly is better (Goldman 1999, 93).<sup>28</sup> However, a policy that leads to more knowledge will often require more time. For example, one is likely to end up with more true beliefs (and fewer false beliefs) on a topic if one spends more time researching it. And if a library spends money on interlibrary loan requests rather than on acquisitions, then patrons can acquire knowledge on a wider range of topics, but they will have to wait longer to do so (Fallis 2006b, 502). Again, a complete epistemic value hierarchy would need to resolve such conflicts.

### 5.3.3. Over People

Epistemologists usually focus on knowledge and true belief at the level of the individual. But, in information management decisions, we are often concerned that many individuals acquire knowledge (Kawaal 2002, Fallis 2004). For example, when legislators decide on intellectual property laws, they want many people to end up with more knowledge.<sup>29</sup> And when libraries make management decisions, they are concerned that many patrons end up with knowledge. This is suggested, for example, by the fact that libraries provide patrons with instruction on evaluating the reliability of information sources (Fallis and Frické 2002, 74). It is also suggested by the fact that information scientists do empirical studies on the accuracy of the information people receive at the reference desk (Hernon and McClure 1986).

Different alternatives in information management decisions will bring about different distributions of knowledge across people. For example, strong intellectual property laws may increase the amount of knowledge that people can acquire, but will almost certainly decrease the number of people that can acquire it (Fallis 2007d). And if a library spends money on outreach activities rather than acquisitions, then more people are likely to acquire knowledge, but there will be less knowledge that they can acquire (Fallis 2006b, 502). Yet again, a complete epistemic value hierarchy would need to resolve these conflicts.

### 5.3.4. Risk

Finally, there will invariably be some uncertainty about what the epistemic outcomes of a particular information management decision are going to be. As a result, a complete epistemic value hierarchy would need to say how to deal with *risk* (Kirkwood 1997, 129,

---

<sup>28</sup> Since it is possible to lose knowledge, there are actually further complications here. For example, one might prefer to acquire the knowledge right before one needs it so that there will be less chance of forgetting it.

<sup>29</sup> It might be suggested that legislators today are more interested in maximizing corporate profits (Vaidhanathan 2001). But at least according to the original intent of the United States Constitution, legislators should be trying to come up with intellectual property laws that “promote the progress of science and useful arts.”

Fallis 2007a). Decisions under risk can actually be treated as yet another kind of distribution problem (Broome 1991, 22-23). Essentially, a complete epistemic value hierarchy would have to say how epistemic goods should be distributed over possible states of the world that have different probabilities.

## 6. Identifying Default Structures

Even though epistemology may not always be able to provide tight constraints on epistemic value hierarchies that apply in all contexts, it can often suggest default structures for such hierarchies. For example, many epistemologists think that most of the epistemic value of justification consists in its value as a means to truth.<sup>30</sup> If these epistemologists are right, then we should take it as a default that justification does not have much weight. Also, most epistemologists think that avoiding error and maintaining consistency are more important than acquiring new true beliefs (Riggs 2003, 347). If these epistemologists are right, then we should take it as a default that error avoidance and consistency have a lot of weight. Having a list of such default structures is probably the next best thing to having an epistemic value hierarchy to use a template. As with the list of epistemic values themselves, the applied epistemologist can leverage existing work in epistemology to generate this list of default structures.

## 7. Constructing Epistemic Value Hierarchies for Specific Decisions

So far, we have been focusing on identifying values, constraints, and defaults that apply to epistemic value hierarchies for any information management decision. However, it may be possible to say something more substantive about the shape of epistemic value hierarchies for particular decisions. For example, different epistemic values may have greater importance in different information management decisions. In fact, sample epistemic value hierarchies for particular information management decisions might be constructed. Such sample hierarchies could be used by people facing such decisions to construct their own hierarchies. The construction of such sample hierarchies is an important direction for future research.

Sample epistemic value hierarchies might be constructed for decisions about which intellectual property policies to adopt (Fallis 2007d). They might also be constructed for decisions about how to evaluate information sources such as *Wikipedia* (Fallis 2008). In this section, we briefly consider what an epistemic value hierarchy might look like for a *collection development* decision about which information resources (books, journals, databases, etc.) to acquire.

When decisions are important, it is worth constructing a value hierarchy to base them on. And most collection development decisions *are* important. It has become extremely expensive in recent years for academic libraries to subscribe to a sufficient number of scientific journals (Branin and Case 1998). Difficult decisions must be made about which journals to subscribe to, and about which worthwhile subscriptions to cancel.

---

<sup>30</sup> One of the authors (Fallis) counts himself among these epistemologists. The other author (Whitcomb) does not.

Given the importance of these decisions, libraries often develop written *collection development* policies (Evans 2000, 69-90, Johnson 2004, 76-83). These policies typically include selection criteria such as the following (Symons and Harmon 1995, 27):

- Accuracy of information
- Appropriate writing style or presentation to suit the audience
- Authoritativeness and honesty
- Comprehensiveness and depth of treatment
- Diversity of viewpoint
- Importance of the subject matter
- Potential use
- Timeliness or permanence of topic

A simplified value hierarchy based on these sorts of selection criteria might look something like this:<sup>31</sup>




Figure 3: A Simple Value Hierarchy for Collection Development Decisions

The criteria in these policies are not explicitly epistemic. But it is clear that they are primarily designed to promote certain epistemic ends.<sup>32</sup> For example, the *accuracy* and *authoritativeness* of information resources are clearly important for promoting the acquisition of true beliefs. The *comprehensiveness* and *depth* of information resources are directly related to the amount of knowledge that can be acquired. John Stuart Mill (1978, ch. 2) famously argued that considering *diverse* viewpoints is the best way to acquire true beliefs. Finally, people will only acquire knowledge if they actually *use* and *comprehend* the information resources.<sup>33</sup> Since epistemic ends are implicit in these selection criteria, we can potentially construct epistemic value hierarchies based on the existing written policies.

---

<sup>31</sup> As noted above, each of these objectives might be further subdivided. For example, maximizing usage consists in maximizing circulation and maximizing in-house usage (Kraft and Boyce 1991, 68). Also, minimizing financial costs consists in minimizing subscription costs, minimizing processing costs, and minimizing storage costs.

<sup>32</sup> Many readers of *fiction* are primarily seeking entertainment rather than knowledge. But even when selecting fiction, epistemic objectives can matter. For example, readers of *historical fiction* typically expect to acquire true beliefs about the larger historical context of a novel. Moreover, fiction can be a rich source of knowledge about the human condition. By reading a novel like *Pride and Prejudice*, we can deepen our grasp of the moral and emotional dynamics of interpersonal relationships (Elgin 1996, Kivy 1997).

<sup>33</sup> Another common selection criterion is *currency*. Up-to-date information resources tend to bring about true belief because they are more likely to be based on our best evidence and because it is less likely that the facts have changed since they were published (Fallis 2006b, 501-502).

In addition to identifying the epistemic values implicit in these policies, we might also identify the *tradeoffs* implicit in them. For instance, consider the decision for an academic library of which scientific journals to subscribe to. Since the budgets are limited, the library might have to decide between acquiring a new *Geographic Information Systems* journal or a new *Oceanography* journal. In order to deal with such conflicts, collection development policies typically codify how “intensely” the library seeks resources within each subject area (Evans 2000, 78-79, Johnson 2004, 82). For example, *Geographic Information Systems* might have an intensity score of 5 while *Oceanography* might only have a score of 4. These scores represent tradeoffs between the different types of knowledge that patrons might acquire. In other words, collection development policies take a position on how knowledge should be distributed *across topics*.

Moreover, faculty and students often have very different information needs (Joswick and Stierman 1997). As a result, acquiring a journal to improve the collection for one group may mean not acquiring a journal to improve the collection for the other group. Thus, tradeoffs also have to be made between the different types of patrons that might acquire knowledge. In a similar vein, law libraries typically give greater weight to the information needs of judges and lawyers than they do to the legal information needs of the general public (Richmond 2003, 80).<sup>34</sup> In both of these cases, libraries implicitly take a position on how knowledge should be distributed *across people*.

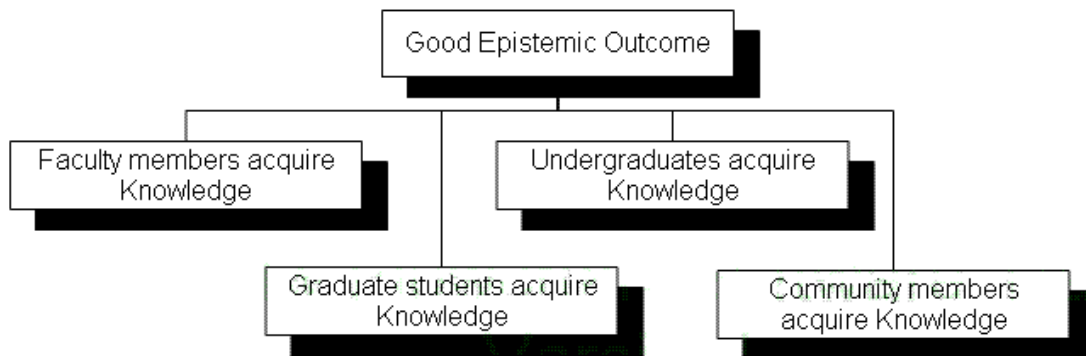


Figure 4: A Simple Epistemic Value Hierarchy for an Academic Library

Finally, it is important to note that, even if we do not manage to construct a complete epistemic value hierarchy for our information management decisions, diligent attempts to do so can still render us better off.<sup>35</sup> Simply thinking carefully about our values can often

<sup>34</sup> Public libraries frequently have to decide between acquiring information resources for which there is greater *demand* and acquiring information resources which have greater *quality* (Berry 1990). More knowledge can be acquired from quality information resources, but fewer people will acquire this knowledge if there is less demand for these particular resources. Thus, this is another collection development decision that requires an epistemic value tradeoff.

<sup>35</sup> For example, even if we cannot precisely determine the relative importance of some of our values, we still may be able to reach a definitive conclusion about what to do by using a *sensitivity analysis* (Edwards

lead us to make better decisions (Keene y 1992). And by trying to construct complete epistemic value hierarchies, we do just that.

## 8. Conclusion

The philosophy of information deals not only with what information *is*, but also with how we should manage it. In order to effectively address this question, the philosophy of information needs to appeal to work in epistemology. In particular, in order to make better information management decisions, people need to be able to construct epistemic value hierarchies. An applied epistemology that leverages work in epistemology and decision analysis can assist people in this important task.

## 9. Acknowledgements

We would like to thank audiences at the School of Library and Information Science, Indiana University, the Department of Library and Information Science, Rutgers University, and the School of Information Resources and Library Science, University of Arizona for their feedback. We would also like to thank Terry Connolly, Sandy Goldberg, Richard Healey, Pen Maddy, and especially Kay Mathiesen for helpful comments on earlier versions of this research. We would like to thank the anonymous reviewers for *The Information Society* for several helpful suggestions. This research was partially supported by a Professional Development Award from the Social and Behavioral Sciences Research Institute at the University of Arizona.

## 10. References

- Alavi, Maryam and Dorothy E. Leidner. 2001. "Knowledge Management and Knowledge Management Systems: Conceptual Foundations and Research Issues." *MIS Quarterly* 25:107-36.
- Alston, William P. 2005. *Beyond "Justification"*. Ithaca: Cornell.
- Andersen, Jack. 2002. "The Role of Subject Literature in Scholarly Communication: An Interpretation Based on Social Epistemology." *Journal of Documentation* 58:463-81.
- Battersby, Mark E. 1989. "Critical Thinking as Applied Epistemology: Relocating Critical Thinking in the Philosophical Landscape." *Informal Logic* 11:91-100.
- Berry, John. 1990. "Leaning Toward "Quality"." *Library Journal* 115(11):76.
- Bishop, Michael A. and J. D. Trout. 2005. *Epistemology and the Psychology of Human Judgment*. Oxford: Oxford.
- BonJour, Laurence. 1985. *The Structure of Empirical Knowledge*. Cambridge: Harvard.

---

and Newman 1982, 81, Kirkwood 1997, 82). That is, we may find that a particular alternative is best as long as our values fall within a certain range.

- Branin, Joseph J. and Mary Case. 1998. "Reforming Scholarly Publishing in the Sciences: A Librarian Perspective." *Notices of the American Mathematical Society* 45:475-86.
- Broome, John. 1991. *Weighing Goods*. Cambridge, Massachusetts: Basil Blackwell.
- Budd, John M. 2004. "Academic Libraries and Knowledge: A Social Epistemology Framework." *Journal of Academic Librarianship* 30:361-67.
- Clifford, William K. 1901. "The Ethics of Belief." Pp. 163-205 in *Lectures and Essays*, vol. 2, New York: MacMillan.
- Day, Ronald E. 2005. "Clearing Up "Implicit Knowledge": Implications for Knowledge Management, Information Science, Psychology, and Social Epistemology." *Journal of the American Society for Information Science and Technology* 56:630-635.
- Dick, Archie L. 1999. "Epistemological Positions and Library and Information Science." *Library Quarterly* 69:305-23.
- Dretske, Fred I. 1981. *Knowledge & the Flow of Information*. Cambridge: MIT.
- Edwards, Ward and J. R. Newman. 1982. *Multiattribute Evaluation*. Beverly Hills: Sage.
- Elgin, Catherine. 1996. *Considered Judgment*. Princeton: Princeton.
- Evans, G. E. 2000. *Developing Library and Information Center Collections*. 4th ed. Englewood, Colorado: Libraries Unlimited.
- Fallis, Don. 1997. "The Epistemic Status of Probabilistic Proof." *Journal of Philosophy* 94:165-86.
- Fallis, Don. 2004a. "Epistemic Value Theory and Information Ethics." *Minds and Machines* 14:101-17.
- Fallis, Don. 2004b. "On Verifying the Accuracy of Information: Philosophical Perspectives." *Library Trends* 52:463-87.
- Fallis, Don. 2005. "Epistemic Value Theory and Judgment Aggregation." *Episteme* 2:39-55.
- Fallis, Don. 2006a. "Epistemic Value Theory and Social Epistemology." *Episteme* 2:177-88.
- Fallis, Don. 2006b. "Social Epistemology and Information Science." Pp. 475-519 in the *Annual Review of Information Science and Technology*, vol. 40, ed. Blaise Cronin. Medford, New Jersey: Information Today.

- Fallis, Don. 2007a. "Attitudes Toward Epistemic Risk and the Value of Experiments." *Studia Logica* 86:215-46.
- Fallis, Don. 2007b. "Collective Epistemic Goals." *Social Epistemology* 21:267-80.
- Fallis, Don. 2007c. "Epistemic Value Theory and the Digital Divide." Pp. 29-46 in *Information Technology and Social Justice*, eds. Emma Rooksby and John Weckert. Hershey, Pennsylvania: Idea Group.
- Fallis, Don. 2007d. "Toward an Epistemology of Intellectual Property." *Journal of Information Ethics* 16:34-51.
- Fallis, Don. 2008. "Toward an Epistemology of Wikipedia." *Journal of the American Society for Information Science and Technology* 59.
- Fallis, Don and Martin Frické. 2002. "Indicators of Accuracy of Consumer Health Information on the Internet." *Journal of the American Medical Informatics Association* 9:73-79.
- Floridi, Luciano. 1996. "Brave.Net.World: The Internet as a Disinformation Superhighway?" *The Electronic Library* 14:509-14.
- Floridi, Luciano. 2001. "Information Ethics: An Environmental Approach to the Digital Divide." *Philosophy in the Contemporary World* 9:1-7.
- Floridi, Luciano. 2002a. "On Defining Library and Information Science as Applied Philosophy of Information." *Social Epistemology* 16:37-49.
- Floridi, Luciano. 2002b. "What is the Philosophy of Information?" *Metaphilosophy* 33:123-45.
- Floridi, Luciano. 2004a. "On the Logical Unsolvability of the Gettier Problem." *Synthese* 142:61-79.
- Floridi, Luciano. 2004b. "Open Problems in the Philosophy of Information." *Metaphilosophy* 35:554-82.
- Floridi, Luciano. forthcoming. "Understanding Epistemic Relevance." *Erkenntnis*.
- Foley, Richard. 1993. *Working Without a Net*. New York: Oxford.
- Fuller, Steve. 2002. *Knowledge Management Foundations*. Boston: Butterworth-Heinemann.
- Furner, Jonathan. 2002. "Shera's Social Epistemology Recast as Psychological Bibliology." *Social Epistemology* 16:5-22.

- Gault, Fred. 2006. "Measuring Knowledge and Its Economic Effects: The Role of Official Statistics." Pp. 27-41 in *Advancing Knowledge and the Knowledge Economy*, ed. Brian Kahin. Cambridge: MIT.
- Gettier, Edmund. 1963. "Is Justified True Belief Knowledge?" *Analysis* 23:121-23.
- Godfrey-Smith, Peter. 1991. "Signal, Decision, Action." *Journal of Philosophy* 88:709-22.
- Goldman, Alvin I. 1989. "Replies to Commentators." *Philosophia* 19:301-23.
- Goldman, Alvin I. 1999. *Knowledge in a Social World*. New York: Oxford.
- Goldman, Alvin I. 2002. "Reply to Commentators." *Philosophy and Phenomenological Research* 64:215-27.
- Greco, Gian M., Gianluca Paronitti, Matteo Turilli, and Luciano Floridi. 2005. "How to Do Philosophy Informationally." *Lecture Notes on Artificial Intelligence* 3782:623-34.
- Gregory, Robin, Baruch Fischhoff, and Tim McDaniels. 2005. "Acceptable Input: Using Decision Analysis to Guide Public Policy Decisions." *Decision Analysis* 2:4-16.
- Grimm, Stephen. 2005. *Understanding as an epistemic goal*. Ph.D. dissertation, University of Notre Dame.
- Hamburg, Morris, Leonard E. Ramist, and Michael R. W. Bommer. 1972. "Library Objectives and Performance Measures and Their Use in Decision Making." *Library Quarterly* 42:107-28.
- Harms, William F. 1998. "The Use of Information Theory in Epistemology." *Philosophy of Science* 65:472-501.
- Hawthorne, John and Jason Stanley. forthcoming. "Knowledge and Action." *Journal of Philosophy*.
- Hempel, Carl G. 1965. *Aspects of Scientific Explanation*. New York: Free Press.
- Heron, Peter and Charles R. McClure. 1986. "Unobtrusive Reference Testing: The 55 Percent Rule." *Library Journal* 111(7):37-41.
- Herold, Ken R. 2001. "Librarianship and the Philosophy of Information." *Library Philosophy and Practice*. <http://www.webpages.uidaho.edu/~mbolin/herold.html>
- Hjørland, Birger. 2002. "Epistemology and the Socio-Cognitive Perspective in Information Science." *Journal of the American Society for Information Science and Technology* 53:257-70.

- Hume, David. 1977. *An Enquiry Concerning Human Understanding*. Indianapolis: Hackett.
- Hunt, Darwin P. 2003. "The Concept of Knowledge and How to Measure It." *Journal of Intellectual Capital* 4:100-113.
- James, William. 1979. *The Will to Believe and Other Essays in Popular Philosophy*. Cambridge: Harvard.
- Johnson, Peggy. 2004. *Fundamentals of Collection Development and Management*. Chicago: American Library Association.
- Jones, Ward E. 1997. "Why Do We Value Knowledge?" *American Philosophical Quarterly* 34:423-39.
- Joswick, Kathleen E. and Jeanne K. Steirman. 1997. "The Core List Mirage: A Comparison of the Journals Frequently Consulted by Faculty and Students." *College and Research Libraries* 58:48-55.
- Kahneman, Daniel and Jackie Snell. 1992. "Predicting a Changing Taste: Do People Know What They Will Like?" *Journal of Behavioral Decision Making* 5:187-200.
- Kaplan, Mark. 1985. "It's Not What You Know That Counts." *Journal of Philosophy* 82:350-363.
- Kaplan, Mark. 1996. *Decision Theory as Philosophy*. Cambridge: Cambridge.
- Kawall, Jason. 2002. "Other-Regarding Epistemic Virtues." *Ratio* 15:257-75.
- Keeney, Ralph L. 1992. *Value-Focused Thinking*. Cambridge: Harvard.
- Kirkwood, Craig W. 1997. *Strategic Decision Making*. Belmont, California: Duxbury.
- Kitcher, Philip. 1993. *The Advancement of Science*. New York: Oxford.
- Kitcher, Philip. 2001. *Science, Truth, and Democracy*. New York: Oxford.
- Kitcher, Philip. 2002. "Veritistic Value and the Project of Social Epistemology." *Philosophy and Phenomenological Research* 64:191-98.
- Kivy, Peter. 1997. *Philosophies of Arts*. Cambridge: Cambridge.
- Kraft, Donald H. and Bert R. Boyce. 1991. *Operations Research for Libraries and Information Agencies*. San Diego: Academic Press.
- Kunst, Heinke, Diederik Groot, Pallavi M. Latthe, Manish Latthe, and Khalid S. Khan. 2002. "Accuracy of Information on Apparently Credible Websites: Survey of Five Common Health Topics." *British Medical Journal* 321:581-82.

- Kvanvig, Jonathan L. 2003. *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge.
- Kvanvig, Jonathan L. 2005. "Truth is Not the Primary Epistemic Goal." Pp. 285-95 in Steup and Sosa 2005.
- Lehrer, Keith. 1975. "Reason and Consistency." Pp. 57-74 in *Analysis and Metaphysics*, ed. Keith Lehrer. Dordrecht: Reidel.
- Levi, Isaac. 1962. "On the Seriousness of Mistakes." *Philosophy of Science* 29:47-65.
- Maher, Patrick. 1993. *Betting on Theories*. New York: Cambridge.
- Maruyama, Magoroh. 1961. "Communicational Epistemology." *British Journal for the Philosophy of Science* 11:319-27.
- McDowell, Ashley. 2002. "Trust and Information: The Role of Trust in the Social Epistemology of Information Science." *Social Epistemology* 16:51-63.
- McNeil, Barbara J., Stephen G. Pauker, Harold C. Sox, and Amos Tversky. 1982. "On the Elicitation of Preferences for Alternative Therapies." *New England Journal of Medicine* 306:1259-62.
- Mill, John S. 1978. *On Liberty*. ed. Elizabeth Rapaport. Indianapolis: Hackett.
- Moore, Adam D. 2005. *Information Ethics*. Seattle: University of Washington.
- Neill, S. D. 1982a. "Brookes, Popper, and Objective Knowledge." *Journal of Information Science* 4:33-39.
- Neill, S. D. 1982b. "Knowledge or Information: A Crisis of Purpose in Libraries." *Canadian Library Journal* 39:69-73.
- Nietzsche, Friedrich. 1997. *Untimely Meditations*. tr. R. J. Hollingdale. Cambridge: Cambridge.
- Nitecki, J. K. 1993. "Metalibrarianship."  
<http://www.twu.edu/library/Nitecki/metalibrarianship/>
- Paterson, R. W. K. 1979. "Towards an Axiology of Knowledge." *Journal of Philosophy of Education* 13:91-100.
- Pietikäinen, Petteri. 2004. "Truth Hurts: The Sociobiology Debate, Moral Reading and the Idea of "Dangerous Knowledge"." *Social Epistemology* 18:165-79.
- Plato. 1961. *The Collected Dialogues*. eds. Edith Hamilton and Huntington Cairns. Princeton: Princeton.
- Popper, Karl. 1972. *Objective Knowledge*. Oxford: Oxford.

- Radford, Colin. 1966. "Knowledge—by Examples." *Analysis* 27:1-11.
- Rawls, John. 1971. *A Theory of Justice*. Cambridge: Harvard.
- Richmond, Larry D. 2003. "The *Pro Se* Patron: An Ethical Rather Than Legal Dilemma." *Legal Reference Services Quarterly* 22:75-84.
- Riggs, Wayne D. 2003. "Balancing Our Epistemic Ends." *Nous* 37:342-52.
- Riggs, Wayne. 2006. "The Value Turn in Epistemology." in *New Waves in Epistemology*, eds. Vincent F. Hendricks and Duncan Pritchard. Aldershot: Ashgate.
- Rowley, Jennifer. 2007. "The Wisdom Hierarchy: Representations of the DIKW Hierarchy." *Journal of Information Science* 33:163-80.
- Ryan, Sharon. 1999. "What Is Wisdom?" *Philosophical Studies* 93:119-39.
- Saracevic, Tefko and Paul B. Kantor. 1997. "Studying the Value of Library and Information Services." *Journal of the American Society for Information Science* 48:527-63.
- Sartwell, Crispin. 1992. "Why Knowledge Is Merely True Belief." *Journal of Philosophy* 89:167-80.
- Satterfield, Theresa. 2001. "In Search of Value Literacy: Suggestions for the Elicitation of Environmental Values." *Environmental Values* 10:331-59.
- Shera, Jesse. 1970. "Library and Knowledge." Pp. 82-110 in *Sociological Foundations of Librarianship*. New York: Asia Publishing House.
- Sosa, Ernest. 2007. *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume 1*. Oxford: Oxford.
- Steup, Matthias and Ernest Sosa, eds. 2005. *Contemporary Debates in Epistemology*. Malden, Massachusetts: Blackwell.
- Stich, Stephen P. 1990. *The Fragmentation of Reason*. Cambridge: MIT.
- Suppe, Frederick R. 1985. "Current Epistemological Approaches to Observation and Interpretation." Pp. 187-92 in *Towards Foundations of Information Science*, ed. Laurence B. Heilprin. White Plains, New York: Knowledge Industry Publications.
- Svenonius, Elaine. 2004. "The Epistemological Foundations of Knowledge Representations." *Library Trends* 52:571-87.
- Symons, Ann K. and Charles Harmon. 1995. *Protecting the Right to Read*. New York: Neal-Schuman.

- Taylor, Shelly E. and Jonathon D. Brown. 1988. "Illusion and Well-Being: A Social Psychological Perspective on Mental Health." *Psychological Bulletin* 103:193-210.
- Thellefsen, Torkild. 2004. "Knowledge Profiling: The Basis for Knowledge Organization." *Library Trends* 52:507-14.
- Vaidhanathan, Siva. 2001. *Copyrights and Copywrongs*. New York: NYU Press.
- von Foerster, Heinz. 1980. "Epistemology of Communication." Pp. 18-27 in *The Myths of Information*, ed. Kathleen Woodward. Madison, Wisconsin: Coda Press.
- Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford: Oxford.
- Wilson, Patrick. 1983. *Second-Hand Knowledge*. Westport, Connecticut: Greenwood Press.
- Wolkoff, Kathleen N. 1996. "The Problem of Holocaust Denial Literature in Libraries." *Library Trends* 45:87-96.
- Zagzebski, Linda. 1996. *Virtues of the Mind*. Cambridge: Cambridge.
- Zagzebski, Linda. 2003. "The Search for the Source of Epistemic Good." *Metaphilosophy* 34:12-28.
- Zins, Chaim. 2006. "Redefining Information Science: From "Information Science" to "Knowledge Science"." *Journal of Documentation* 62:447-61.