Skepticism and Floating Conclusions
(Preliminary Version)

John F. Horty

Philosophy Department and
Institute for Advanced Computer Studies
University of Maryland
College Park, MD 20742

horty@umiacs.umd.edu
www.umiacs.umd.edu/users/horty

Abstract: The purpose of this paper is to question some commonly accepted patterns of reasoning involving nonmonotonic logics that generate multiple extensions. In particular, I argue that the phenomenon of floating conclusions indicates a problem with the view that the skeptical consequences of such theories should be identified with the statements that are supported by each of their various extensions.

1 Introduction

One of the most striking ways in which nonmonotonic logics can differ from classical logic, and even from standard philosophical logics, is in allowing for multiple sanctioned conclusion sets, known as *extensions*. The term is due to Reiter [10], who thought of default rules as providing a means for extending the strictly logical conclusions of a knowledge base with plausible information. Multiple extensions arise when a knowledge base contains conflicting default rules, suggesting different, often inconsistent ways of supplementing its strictly logical conclusions.

The purpose of this paper is to question some commonly accepted patterns of reasoning involving theories that generate multiple extensions. In particular, I argue that the phenomenon of floating conclusions indicates a problem with the view that the skeptical consequences of such theories should be identified with the statements that are supported by each of their various extensions.

2 Multiple extensions

The canonical example of a knowledge base with multiple extensions is the Nixon Diamond, depicted in Figure 1. Here, the statements Qn, Rn, and Pn represent the propositions that Nixon is a Quaker, a Republican, and a pacifist; statements of the form $A \Rightarrow B$ and $A \rightarrow B$ represent ordinary logical implications and "default" implications respectively, with $A \not\Rightarrow B$ and $A \not\rightarrow B$ abbreviating $A \Rightarrow \neg B$ and $A \rightarrow \neg B$; and the special statement $\top$ represents truth. What the knowledge base tells us, of course, is this: Nixon is both a Quaker and a Republican, the fact that he is a Quaker provides a good reason for concluding that he is a pacifist, and the fact that he is a Republican provides a good reason for concluding that he is not a pacifist.

This example can be coded into default logic as the theory $\Delta = \langle \mathcal{W}, \mathcal{D} \rangle$, with $\mathcal{W} = \{Qn, Rn\}$ representing the basic facts of the situation and $\mathcal{D} = \{(Qn : Pn \;/\; Pn), (Rn : \neg Pn \;/\; \neg Pn)\}$ representing the two defaults. The theory yields two extensions: $\mathcal{E}_1 = Th(\mathcal{W} \cup \{Pn\})$ and $\mathcal{E}_2 = Th(\mathcal{W} \cup \{\neg Pn\})$. The first results when the basic facts of the situation are extended by an application of the default concerning Quakers; the second results when the facts are extended by an application of the default concerning Republicans.

In light of these two extensions, what are we to conclude from the initial information: is Nixon a pacifist or not? More generally, when a default theory leads to more than one extension, what should we actually infer from that theory—how should we define its set of consequences, or conclusions?

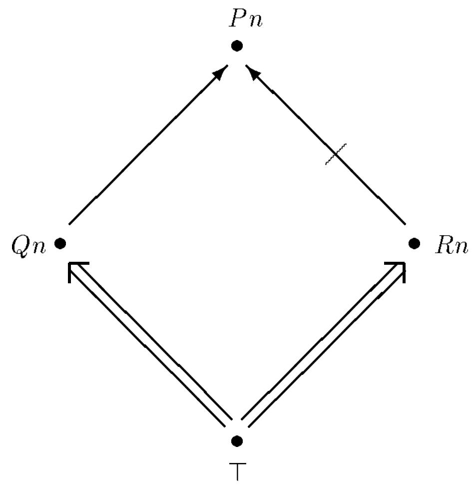

Figure 1: The Nixon Diamond

Several proposals have been discussed in the literature. One option is to suppose that we should arbitrarily select a particular one of the theory's several extensions and endorse the conclusions contained in it; a second option is to suppose that we should be willing to endorse a conclusion just in case it is contained in some extension of the theory. These first two options are sometimes said to reflect a *credulous* reasoning policy. A third option, now generally described as *skeptical*, is to suppose that we should endorse a conclusion just in case it is contained in every extension of the theory.[1]

Of these three options, the first—pick an arbitrary extension—really does seem to embody a sensible policy, or at least one that is frequently employed. Given conflicting defeasible information, we often simply adopt some internally coherent point of view in which the conflicts are resolved in some particular way, regardless of the fact that there are other coherent points of view in which the conflicts are resolved in different ways. Still, although this reasoning policy may be sensible, it is hard to see how it could be codified in a formal consequence relation. If the choice of extension really is arbitrary, different reasoners could easily select different extensions, or the same reasoner might select different extensions at different times. Which extension, then, would represent the real conclusion set of the original theory?

The second of our three options—endorse a conclusion whenever it is contained in some extension—could indeed be codified as a consequence relation, but it would be a peculiar one. According to this policy, the conclusion set associated with a default theory need not be closed under standard logical consequence, and might easily be inconsistent, even in cases in which the underlying default theory itself seems to be consistent. The conclusion set of the theory representing the Nixon Diamond, for example, would contain both Pn and $\neg Pn$, since each of these formulas belongs to some extension of the default theory, but it would not contain $Pn \wedge \neg Pn$, since this formula is not contained in any extension.

One way of avoiding these peculiar features of the second option is to think of the conclusions generated by a default theory as being shielded by a kind of modal operator. Where A is a statement, let $\mathcal{B}(A)$ mean there is good reason to believe that A; and suppose a theory provides us with good reason to believe a statement whenever that statement is included in some extension of the theory, some internally coherent point of view. Then we could define the initial conclusions of a default theory $\Delta = \langle \mathcal{W}, \mathcal{D} \rangle$ as the set that extends $\mathcal{W}$ with a formula $\mathcal{B}(A)$ whenever A belongs to some extension of Δ, and we could go on to define the theory's conclusion set as the logical closure of its initial conclusions.

This variant of the second option has some interest. It results in a conclusion set that is both closed under logical consequence and consistent as long as $\mathcal{W}$ itself is consistent. And Reiter's original paper on default

[1] The use of the *credulous/skeptical* terminology in this context was first introduced by Touretzky et al. [13], but the distinction itself is older than this; it was already implicit in Reiter's paper on default logic, and was described explicitly by McDermott [8] as the distinction between *brave* and *cautious* reasoning. Makinson [6] refers to the first of the two credulous options described here as the *choice* option.

logic [10, Section 4] provides a proof procedure, sound and complete under certain conditions, that could be used in determining whether $\mathcal{B}(A)$ belongs to the conclusion set as defined here. Unfortunately, however, this variant of the second option also manages to sidestep our original question. We wanted to know what conclusions we should actually draw from the information provided by a default theory—whether or not, given the information from the Nixon Diamond, we should conclude that Nixon is a pacifist, for example. But according to this variant, we are told only what there is good reason to believe—that both $\mathcal{B}(Pn)$ and $\mathcal{B}(\neg Pn)$ are consequences of the theory, so that there is good reason to believe that Nixon is a pacifist, but also good reason to believe that he is not. This may be useful information, but it is still some distance from telling us whether or not to conclude that Nixon is a pacifist.[2]

Of our three options for defining a notion of consequence in the presence of multiple extensions, only the third, skeptical proposal—endorse a conclusion whenever it is contained in every extension—seems to hold any real promise. This option leads to a single conclusion set, which is both closed under logical consequence and consistent as long as the initial information is. And it provides an answer to our original question concerning proper conclusions that is at least initially attractive. In the case of the Nixon Diamond, for example, since neither Pn nor $\neg Pn$ belongs to every extension, this third option tell us that we should not conclude that Nixon is a pacifist, but that we should not conclude that Nixon is not a pacifist either. Since there is a good reason for each of these conflicting conclusion, we should remain skeptical.

3 Floating conclusions

Default logic defines a direct, unmediated relation between a particular default theory and the statement sets that form its extensions. Another class of formalisms—known as *argument systems*—takes a more roundabout approach, analyzing nonmonotonic reasoning through the study of interactions among competing defeasible arguments.[3]

Although the arguments themselves that are studied in these argument systems are often complex, we can restrict our attention here entirely to linear arguments, analogous to the reasoning paths studied in theories of defeasible inheritance.[4] These linear arguments are formed by starting with a true statement and then simply stringing together strict and defeasible implications; each such argument can be said to *support* the final statement it contains as its conclusion. As an abstract example, the structure $\top \Rightarrow A \rightarrow B \not\Rightarrow C$ can be taken to represent an argument of the form "A is true, which defeasibly implies B, which strictly implies $\neg C$," supporting the conclusion $\neg C$. As a less abstract example, we can see that the Nixon Diamond provides the materials for constructing the two arguments $\top \Rightarrow Qn \rightarrow Pn$ and $\top \Rightarrow Rn \not\rightarrow Pn$, supporting the conflicting conclusions Pn and $\neg Pn$.

Where α is an argument, we will let $^*\alpha$ represent the particular conclusion supported by α. Where Φ is a set of arguments, we will let $^*\Phi$ represent the set of conclusions supported by the arguments in Φ—that is, the set containing the statement $^*\alpha$ for each argument α belonging to Φ.

The primary technical challenge involved in the development of an argument system is the specification of the coherent sets of arguments that an ideal reasoner might be willing to accept on the basis of a given body of initial information. We will refer to these coherent sets of arguments as *argument extensions*, to distinguish them from the *statement extensions* defined by theories such as default logic. Again, the actual definition of argument extensions is often complicated in ways that need not concern us here. Without going into detail, however, we can simply note that, just as theories like default logic allow multiple statement extensions, argument systems often associate multiple argument extensions with a single body of initial information. In the case of the Nixon Diamond, for example, an argument system patterned after multiple-extension theories of defeasible inheritance would generate the two extensions $\Phi_1 = \{\top \Rightarrow Qn, \top \Rightarrow Rn, \top \Rightarrow Qn \rightarrow Pn\}$ and $\Phi_2 = \{\top \Rightarrow Qn, \top \Rightarrow Rn, \top \Rightarrow Rn \rightarrow \neg Pn\}$. The first results from supplementing the initial information with the argument that Nixon is a pacifist because he is a Quaker, the second from supplementing this information with the argument that Nixon is not a pacifist because he is a Republican.

[2]Note that this objection is directed only against the use of modal operators to capture the epistemic interpretation of default logic. Other interpretations, involving other modal operators, are possible; it is shown in [2], for example, that a deontic interpretation, with default conclusions wrapped inside of deontic operators, generates a logic for normative reasoning corresponding to that originally suggested by van Fraassen [14].

[3]A recent survey of a variety of argument systems can be found in Prakken and Vreeswijk [9].

[4]The development of the path-based approach to inheritance reasoning was initiated by Touretzky [12]; a survey can be found in [3].

When a knowledge base leads to multiple argument extensions, there are, as before, several options for characterizing the appropriate set of conclusions to draw on the basis of the initial information. Again, we might adopt a credulous reasoning policy, either endorsing the set of conclusions supported by an arbitrary one of the several argument extensions, or perhaps endorsing a conclusion as believable whenever it is supported by some extension or another. In the case of the Nixon Diamond, this policy would lead us to endorse either ${}^*\Phi_1 = \{Qn, Rn, Pn\}$ or ${}^*\Phi_2 = \{Qn, Rn, \neg Pn\}$ as the conclusion set of the original knowledge base, or perhaps simply to endorse the statements belonging to the union of these two sets as believable.

As before, however, we might also adopt a kind of skeptical policy in the presence of these multiple argument extensions, defining the appropriate conclusion set through their intersection. In this case, though, since these new extensions contain arguments rather than statements, there are now two alternatives for implementing such a policy. First, we might decide to endorse an argument just in case it is contained in each argument extension associated with an initial knowledge base, and then to endorse a conclusion just in case that conclusion is supported by an endorsed argument. Formally, this alternative leads to the suggestion that the appropriate conclusions of an initial knowledge base Γ should be the statements belonging to the set

$$^*(\bigcap\{\Phi : \Phi \text{ is an extension of } \Gamma\}).$$

Or second, we might decide to endorse a conclusion just in case that conclusion is itself supported by each argument extension of the initial knowledge base Γ, leading to the formal suggestion that the appropriate conclusions of the knowledge base should the statements belonging to the set

$$\bigcap\{{}^*\Phi : \Phi \text{ is an extension of } \Gamma\},$$

where the order of * and $\bigcap$ is reversed.

Of course, these two alternatives for implementing the skeptical policy come to the same thing in the case of the Nixon Diamond: both lead to $\{Qn, Rn\}$ as the appropriate conclusion set. But there are other situations in which the two alternatives yield different results. A well-known example, due to Ginsberg, appears in Figure 2, where Qn and Rn are interpreted as before, Dn and Hn are interpreted to mean that Nixon is a dove or a hawk respectively, and En as meaning that Nixon is politically extreme (regarding the appropriate use of military force). What this diagram tells us, then, is that Nixon is both a Quaker and a Republican, that there is good reason to suppose that Nixon is a dove if he is a Quaker, a hawk if he is a Republican, and that he is politically extreme if he is either a dove or a hawk.

Again, a system patterned after multiple-extension inheritance theories would associate two argument extensions with this knowledge base, as follows:

$$\begin{aligned}
\Phi_1 \quad = \quad \{ & \top \Rightarrow Qn, \;\; \top \Rightarrow Rn, \\
& \top \Rightarrow Qn \rightarrow Dn, \\
& \top \Rightarrow Qn \rightarrow Dn \not\Rightarrow Hn, \\
& \top \Rightarrow Qn \rightarrow Dn \Rightarrow En \}, \\
\\
\Phi_2 \quad = \quad \{ & \top \Rightarrow Qn, \;\; \top \Rightarrow Rn, \\
& \top \Rightarrow Rn \rightarrow Hn, \\
& \top \Rightarrow Rn \rightarrow Hn \not\Rightarrow Dn, \\
& \top \Rightarrow Rn \rightarrow Hn \Rightarrow En \}.
\end{aligned}$$

Since no arguments except for the trivial $\top \Rightarrow Qn$ and $\top \Rightarrow Rn$ are contained in both of these extensions, the first of our two alternatives for implementing the skeptical policy, which involves intersecting the argument extensions themselves, would lead to $\{Qn, Rn\}$ as the appropriate conclusion set, telling us nothing more than the initial information that Nixon is a Quaker and a Republican. On the other hand, each of these two argument extensions supports the statement En—one through the argument $\top \Rightarrow Qn \rightarrow Dn \Rightarrow En$, the other through the argument $\top \Rightarrow Rn \rightarrow Hn \Rightarrow En$. The second of our two alternatives for implementing the skeptical policy, which involves intersecting supported statements rather than the arguments that support them, would therefore lead to the conclusion set $\{Qn, Rn, En\}$, telling us also that Nixon is politically extreme.

Statements like En, which are supported in each extension associated with a knowledge base, but only by different arguments, are known as *floating conclusions*. This phrase, coined by Makinson and Schlechta

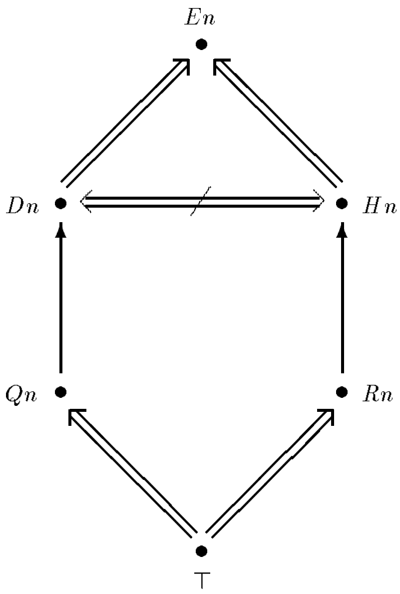

Figure 2: Is Nixon politically extreme?

[7], nicely captures the picture of these conclusions as floating above the sea of different and conflicting arguments that might be taken to support them.

The phenomenon of floating conclusions was first investigated in the context of defeasible inheritance reasoning, particularly in connection with the theory developed by Thomason, Touretzky, and myself in [4]. That theory was based on a two-step architecture. First, in contrast to the multiple-extension accounts considered so far, a single argument extension was defined that was thought of as containing the "skeptically acceptable" arguments based on a given inheritance network. Second, the appropriate conclusions were defined simply as the statements supported by those skeptically acceptable arguments.

Ginsberg's political extremist example was meant to show that no such two-step architecture could correctly represent skeptical reasoning. The single argument extension could not consistently contain both the arguments $\top \Rightarrow Qn \rightarrow Dn \Rightarrow En$ and $\top \Rightarrow Rn \rightarrow Hn \Rightarrow En$, since the strict information in the knowledge base shows that each of these arguments conflicts with an initial segment of the other. The argument extension could not contain either of these arguments without the other, since that would involve the kind of arbitrary decision appropriate only for credulous reasoning. And if the argument extension contained neither of these two arguments, it would not support the conclusion En, which Ginsberg considers to be an intuitive consequence of the initial information: "given that both hawks and doves are politically [extreme], Nixon certainly should be as well" [1, p. 221].[5]

Both Makinson and Schlechta [7] and Stein [11] also consider floating conclusions in the context of defeasible inheritance reasoning, also focusing on the theory developed in [4]. Makinson and Schlechta share Ginsberg's view that the appropriate conclusions to derive from a knowledge base are those that are supported by each of its argument extensions:

> It is an oversimplification to take a proposition A as acceptable ... iff it is supported by some [argument] path α in the intersection of all extensions. Instead A must be taken as acceptable iff it is in the intersection of all *outputs* of extensions, where the output of an extension is the set of all propositions supported by some path within it [7, pp. 203–204].

From this they likewise argue, not only that the particular theory of [4] is "just wrong," but more generally, that any theory based on the same two-step architecture—attempting to define the acceptable statements

[5]Although first published in the textbook cited here, this example had previously been part of the oral tradition for many years—I first heard the example during the question session after the AAAI-87 presentation of [4], when Ginsberg raised it as an objection to that theory.

through a single set of acceptable arguments—will be mistaken. And Stein reaches a similar judgment, for similar reasons:

> The difficulty lies in the fact that some conclusions may be true in every credulous extension, but supported by different [argument] paths in each. Any path-based theory must either accept one of these paths—and be unsound, since such a path is not in every extension—or reject all such paths—and with them the ideally skeptical conclusion—and be incomplete [11, p. 284].

What lies behind these various criticisms, of course, is the widely-held assumption that the first, rather than the second, of our two skeptical alternatives is correct—that floating conclusions should be accepted, and that a system that fails to classify them among the consequences of a defeasible knowledge base is therefore in error. The purpose of this paper is to question that assumption.

4 An example

Why not accept floating conclusions? Their precarious status can be illustrated through any number of examples, but we might as well choose a dramatic one.

Suppose, then, that my parents have a net worth of one million dollars, but that they have divided their assets in order to avoid the United States inheritance tax, so that each parent currently possess half a million dollars apiece. And suppose that, because of their simultaneous exposure to a fatal disease, it is now settled that both of my parents will die within a month. This is a fact: medical science is certain.

Imagine also, however, that there is some expensive item—a yacht, say—whose purchase I believe would help to soften the blow of my impending loss. Although the yacht I want is currently available, the price is good enough that it is sure to be sold by the end of the month. I can now reserve the yacht for myself by putting down a large deposit, with the balance due in six weeks. But there is no way I can afford to pay the balance unless I happen to inherit at least half a million dollars from my parents within that period, and if I fail the pay the balance on time, I will lose my large deposit. Setting aside any doubts concerning the real depth of my grief, let us suppose that my utilities determine the following conditional preferences: if I believe I will inherit half a million dollars from my parents within six weeks, it is very much in my benefit to place a deposit on the yacht; if I do not believe this, it is very much in my benefit not to place a deposit.

Now suppose I have a brother and a sister, both of whom are extraordinarily reliable as sources of information. Neither has ever been known to be mistaken, to deceive, or even to misspeak—although of course, like nearly any source of information, they must be regarded as defeasible. My brother and sister have both talked with our parents about their wills, and feel that they understand the situation. I have written to each of them describing my delicate predicament regarding the yacht, and receive letters back. My brother writes: "Father is going to leave his money to me, but Mother will leave her money to you, so you're in good shape." My sister writes: "Mother is going to leave her money to me, but Father will leave his money to you, so you're in good shape." No further information is now available: the wills are sealed, my brother and sister are trekking together through the Andes, and our parents, sadly, have slipped into a coma.

Based on my current information, what should I conclude? Should I form the belief that I will inherit half a million dollars—and therefore place a large deposit on the yacht—or not?

The situation is depicted in Figure 3, where the statement letters are interpreted as follows: F represents the proposition that I will inherit half a million dollars from my father, M represents the proposition that I will inherit half a million dollars from my mother, $BA(\neg F \wedge M)$ represents the proposition that my brother asserts that I will inherit my mother's money but not my father's, and $SA(F \wedge \neg M)$ represents the proposition that my sister asserts that I will inherit my father's money but not my mother's. The defeasible links $BA(\neg F \wedge M) \rightarrow \neg F \wedge M$ and $SA(F \wedge \neg M) \rightarrow F \wedge \neg M$ reflect the fact that any assertion by my brother or sister provides good reason for concluding that the content of that assertion is true. The strict links in the diagram record various implications and inconsistencies. Notice that, although the contents of my brother's and sister's assertions—the statements $\neg F \wedge M$ and $F \wedge \neg M$—are jointly inconsistent, the truth of either entails the disjunctive claim $F \vee M$, which is, of course, all I really care about. As long as I can conclude that I will inherit half a million dollars from either my father or my mother, I should go ahead and place a deposit on the yacht.

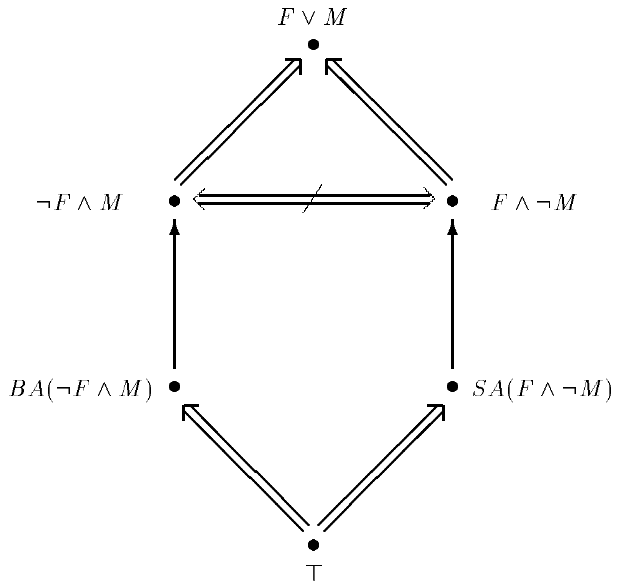

Figure 3: Should I buy the yacht?

A multiple-extension approach would associate the following two argument extensions with this knowledge base:

$$
\begin{aligned}
\Phi_1 \;=\; \{ \; & \top \Rightarrow BA(\neg F \wedge M), \\
& \top \Rightarrow SA(F \wedge \neg M), \\
& \top \Rightarrow BA(\neg F \wedge M) \rightarrow \neg F \wedge M, \\
& \top \Rightarrow BA(\neg F \wedge M) \rightarrow \neg F \wedge M \quad \not\Rightarrow F \wedge \neg M, \\
& \top \Rightarrow BA(\neg F \wedge M) \rightarrow \neg F \wedge M \Rightarrow F \vee M \; \}, \\
\Phi_2 \;=\; \{ \; & \top \Rightarrow BA(\neg F \wedge M), \\
& \top \Rightarrow SA(F \wedge \neg M), \\
& \top \Rightarrow SA(F \wedge \neg M) \Rightarrow F \wedge \neg M, \\
& \top \Rightarrow SA(F \wedge \neg M) \Rightarrow F \wedge \neg M \quad \not\Rightarrow \neg F \wedge M, \\
& \top \Rightarrow SA(F \wedge \neg M) \Rightarrow F \wedge \neg M \Rightarrow F \vee M \; \}.
\end{aligned}
$$

Again, the first of our two alternatives for implementing the skeptical reasoning policy, which involves intersecting arguments, would lead to $\{BA(\neg F \wedge M), SA(F \wedge \neg M)\}$ as the appropriate conclusion set, telling me only that my brother and sister asserted what they did. But since each of the two extensions contains some argument supporting the statement $F \vee M$, the second alternative, which involves intersecting supported statements, leads to the conclusion set $\{BA(\neg F \wedge M), SA(F \wedge \neg M), F \vee M\}$, telling me also—as a floating conclusion—that I will inherit half a million dollars from either my father or my mother.

In this situation, then, there is a vivid practical difference between the two skeptical alternatives. If I were to reason according to the first, I would not be justified in concluding that I am about to inherit half a million dollars, and so it would be foolish for me to place a deposit on the yacht. If I were to reason according to the second, I would be justified in drawing this conclusion, and so it would be foolish for me not to place a deposit.

Which alternative is correct? I have not done a formal survey, but most of the people to whom I have presented this example are suspicious of the floating conclusion, and so favor the first alternative. Most do not feel that the initial information from Figure 3 would provide sufficient justification for me to conclude, as the basis for an important decision, that I will inherit half a million dollars. Certainly, this is my own opinion—that the example shows, contrary to the widely-held assumption, that it is at least coherent for a skeptical reasoner to withhold judgment from floating conclusions. Although both my brother and sister are reliable, and each supports the conclusion that I will inherit half a million dollars, the support provided by each of these reliable sources is undermined by the other; there is no unopposed reason supporting the conclusion. Since either my brother or sister must be wrong, it is therefore easy to imagine that they might both be wrong, and wrong in this way: perhaps my father will leave his money to my brother and my mother will leave her money to my sister, so that I will inherit nothing.

5 Comments on the example

First, in case this example does not yet seem convincing, it might help to modify things a bit. Suppose, then, that I had written only to my brother, and received his response—that my father had named him as sole beneficiary, but that my mother would leave her money to me. That is, suppose my starting point is the information depicted in the left-hand side of Figure 3. In this new situation, should I conclude that I will inherit half a million dollars, and therefore place a deposit on the yacht?

Some might say no—that even in this simplified situation I should not make such a important decision on the basis of my brother's word alone. But this objection misses the point. Most of what we know, we know through sources of information that are, in fact, defeasible. By hypothesis, we can suppose that my brother is arbitrarily reliable, as reliable as any defeasible source of information could possibly be—as reliable as perception, for instance, or the bank officer's word that the money has actually been deposited in my account. If we were to reject information like this, it is hard to see how we could get by in the world at all. When a source of defeasible information that is, by hypothesis, arbitrarily reliable tells me that I will inherit half a million dollars, and there is no conflicting evidence in sight, it is reasonable for me to accept this statement, and to act on it. Note that both of the two skeptical alternatives yield this outcome in our simplified situation, since the initial information generates only a single argument extension, in which the conclusion that I will inherit half a million dollars is supported by a single argument.

Now suppose that, at this point, I hear from my equally reliable sister with her conflicting information—that she is my mother's beneficiary, but that my father will leave his money to me. As a result, I am again in the situation depicted in Figure 3, with two argument extensions, and in which the statement that I will inherit half a million dollars is supported only as a floating conclusion. Ask yourself: should my confidence in the statement that I will inherit half a million dollars be diminished in this new situation, now that I have heard from my sister as well as my brother? If it seems that my confidence can legitimately be diminished—that this new information casts any additional doubt on the outcome—then it follows that floating conclusions are somewhat less secure than conclusions that are uniformly supported by a common argument. And that is all we need. The point is not that floating conclusions might be wrong; any conclusion drawn through defeasible reasoning might be wrong. The point is that a statement supported only as a floating conclusion seems to be less secure than the same statement when it is uniformly supported by a common argument. As long as there is this difference in principle, it is coherent to imagine a skeptical reasoner whose standards are calibrated so as to accept statements that receive uniform support, but to reject floating conclusions.

As a second comment, notice that, if floating conclusions pose a problem, it is not just a problem for argument systems, but also for traditional nonmonotonic formalisms, such as default or model-preference logics. Indeed, the problem is even more serious for these traditional formalisms. With argument systems, where the extensions generated are argument extensions, it is at least possible to avoid floating conclusions by adopting the first of our two skeptical alternatives—endorsing only those arguments belonging to each extension, and then endorsing only the conclusions of the endorsed arguments. Since arguments are represented explicitly in these systems, they can be used to filter out floating conclusions. In most traditional nonmonotonic logics, arguments are suppressed, and so the materials for carrying out this kind of filtering policy are not even available.

To illustrate, a natural representation of the information from our yacht example in default logic is the theory $\Delta = \langle \mathcal{W}, \mathcal{D} \rangle$, where $\mathcal{W} = \{BA(\neg F \wedge M), SA(F \wedge \neg M)\}$ describes what my brother and sister said and $\mathcal{D} = \{(BA(\neg F \wedge M) : \neg F \wedge M \ / \ \neg F \wedge M), (SA(F \wedge \neg M) : F \wedge \neg M \ / \ F \wedge \neg M)\}$ reflects the defaults that whatever my brother and sister say should be taken as true. This theory has two extensions: $\mathcal{E}_1 = Th(W \cup \{\neg F \wedge M\})$ and $\mathcal{E}_2 = Th(W \cup \{F \wedge \neg M\})$. The extensions of default logic are statement extensions, and so the only possible policy for skeptical reasoning appears to be: intersect the extensions. Since the statement $F \vee M$ belongs to both extensions, skeptical reasoning in default logic tells me, immediately and without ambiguity, that I will inherit half a million dollars.

Of course, default logic is essentially a proof-theoretic formalism, and it is easy to see how it could be modified so that the extensions defined would contain proofs rather than statements; such a modification would then allow for floating conclusions to be filtered out by a treatment along the lines of our first alternative. It is harder to see how floating conclusions could be avoided in model-preference logics. In a circumscriptive theory, for instance, the yacht example could naturally be expressed by supplementing the facts $BA(\neg F \wedge M)$ and $SA(F \wedge \neg M)$ with the statements $(BA(\neg F \wedge M) \wedge \neg Ab_b) \supset \neg F \wedge M$ and $(SA(F \wedge \neg M) \wedge \neg Ab_s) \supset F \wedge \neg M$, and then preferring those models in which as few as possible of the propositions Ab_b and Ab_s—the abnormalities associated with the rare situations in which my brother or sister is mistaken—are true. Of course, there can be no models in which neither Ab_b not Ab_s is true. The most preferred models will therefore be those in which only one of these abnormalities holds. The statement $F \vee M$ is true in all of these models, and would therefore follow as a circumscriptive consequence.[6]

6 Other examples

Once the structure of the yacht example is understood, it is easy to construct other examples along similar lines: just imagine a situation in which two sources of information, or reasons, support a common conclusion, but also undermine each other, and therefore undermine the support that each provides for the common conclusion.

Suppose you are a military commander pursuing an enemy that currently holds a strong defensive position. It is suicide to attack while the enemy occupies this position in force, but you have orders to press ahead as quickly as possible, and so you send out your reliable spies. After a week, one spy reports back that there can now be only a skeleton force remaining in the defensive position; he has seen the main enemy column retreating through the mountains, although he also noticed that they sent out a diversionary group to make it appear as if they were retreating along the river. The other spy agrees that only a skeleton force remains in the defensive position; he has seen the main enemy column retreating along the river, although he notes that they also sent out a diversionary group to make it appear is if they were retreating through the mountains. Based on this information, should you assume at least that the main enemy force has retreated from the defensive position—a floating conclusion that is supported by both spies—and therefore commit your troops to an attack? Not necessarily. Although they support a common conclusion, each spy undermines the support provided by the other. Perhaps the enemy sent out two diversionary groups, one through the mountains and one along the river, and managed to fool both your spies into believing that a retreat was in progress. Perhaps the main force still occupies the strong defensive position, awaiting your attack.

Or suppose you attend a macroeconomics conference during a period of economic health, with low inflation and strong growth, and find that the community of macroeconomic forecasters is now split right down the middle. One group, working with a model that has been reliable in the past, predicts that the current strong growth rate will lead to higher inflation, triggering an economic downturn. By tweaking a few parameters in the same model, the other group arrives at a prediction according to which the current low inflation rate will actually continue to decline, leading to a dangerous period of deflation and triggering an economic downturn. Both groups predict an economic downturn, but for different and conflicting reasons—higher inflation versus deflation—and so the prediction is supported only as a floating conclusion. Based on this information, should you accept the prediction, adjusting your investment portfolio accordingly? Not necessarily. Perhaps the extreme predictions are best seen as undermining each other and the truth lies somewhere in between: the inflationary and deflationary forces will cancel each other out, the inflation rate will remain pretty much as it is, and the period of economic health will continue.

There is no need to labor the point by fabricating further examples in which floating conclusions are suspect. But what about the similar cases, exemplifying the same pattern, that have actually been advanced as supporting floating conclusions, such as Ginsberg's political extremist example from Figure 2?

I have always been surprised that this particular example has seemed so persuasive to so many people. The example relies on our understanding that individuals adopt a wide spectrum of attitudes regarding the appropriate use of military force, but that Quakers and Republicans tend to be doves and hawks, respectively—where doves and hawks take the extreme positions that the use of military force is either never

[6] This form of circumscription, which involved minimizing the truth of statements rather than the extensions of predicates, is a special case of the more usual form; see Lifschitz [5, pp. 302-303] for a discussion.

appropriate, or that it is appropriate in response to any provocation, even the most insignificant. Of course, Nixon's own position on the matter is well known. But if I were told of some other individual that he is both a Quaker and a Republican, I would not be sure what to conclude. It is possible that this individual would adopt an extreme position, as either a dove or a hawk. But it seems equally reasonable to imagine that such an individual, rather than being pulled to one extreme of the other, would combine elements of both views into a more balanced, measured position falling toward the center of the political spectrum—perhaps believing that the use of military force is sometimes appropriate, but only as a response to serious provocation. Given this real possibility, it might be appropriate to take a skeptical attitude, not only toward the questions of whether this individual would be a dove or a hawk, but also toward the question whether he would adopt a politically extreme position at all.

Another example appears in Reiter's original paper on default logic, where he suggests [10, pp. 86–87] defaults representing the facts that people tend to live in the same cities as their spouses, but also in the cities in which they work, and then asks us to consider the case of Mary, whose spouse lives in Toronto but who works in Vancouver. Coded into default logic, this information leads to a theory with two extensions, in one of which Mary lives in Toronto and in one of which she lives in Vancouver. Reiter seems to favor the credulous policy of embracing a particular one of these extensions, either concluding that Mary lives in Toronto or concluding that Mary lives in Vancouver. But then, in a footnote, he also mentions what amounts to the skeptical possibility of forming only the belief that Mary lives in either Toronto or Vancouver—where this proposition is supported, of course, as a floating conclusion.

Given the information from this example, I would, in fact, be likely to conclude that Mary lives either in Toronto or Vancouver. But I am not sure this conclusion should follow as a matter of logic, even default logic. In this case, the inference seems to rely on a good deal of knowledge about the particular domain involved, including the vast distance between Toronto and Vancouver, which effectively rules out any sort of intermediate solution to Mary's two-body problem.

By contrast, consider the happier situation of Carol, who works in College Park, Maryland, but whose spouse works in Alexandria, Virginia; and assume the same two defaults—that people tend to live in the same cities as their spouses, but also tend to live in the cities in which they work. Represented in default logic, this information would again lead to a theory with multiple extensions, in each of which, however, Carol would live either in College Park or in Alexandria. Nevertheless, I would be reluctant to accept the floating conclusion that Carol lives either in College Park or in Alexandria. Just thinking about the situation, I would consider it equally likely that Carol and her spouse live together in Washington, DC, within easy commuting distance of both their jobs.

7 Skepticism

Why is it so widely thought that floating conclusions should be accepted by a skeptical reasoner, so that a system that fails to generate these conclusions is therefore incorrect? It is hard to be sure, since this point of view is generally taken as an assumption, rather than argued for, but we can speculate.

Suppose an agent believes that either the statement B or the statement C holds, that B implies A, and that C also implies A. Classical logic then allows the agent to draw A as a conclusion; this is a valid principle of inference, sometimes known as the principle of constructive dilemma. The inference to a floating conclusion is in some ways similar. Suppose a default theory has two extensions, $\mathcal{E}_1$ and $\mathcal{E}_2$, that the extension $\mathcal{E}_1$ contains the statement A, and that the extension $\mathcal{E}_2$ also contains the statement A. The standard view is that a skeptical reasoner should then draw A as a conclusion, even if it is not supported by a common argument in the two extensions.

Notice the difference between these two cases, though. In the first case, the classical reasoning agent believes both that B and C individually imply A, and also that either B or C holds. In the second case, we might as well suppose that the skeptical reasoner knows that A belongs to both the extensions $\mathcal{E}_1$ and $\mathcal{E}_2$, so that both $\mathcal{E}_1$ and $\mathcal{E}_2$ individually imply A. The reasoner is therefore justified in drawing A as a conclusion by something like the principle of constructive dilemma—as long as it is reasonable to suppose, in addition, that either $\mathcal{E}_1$ or $\mathcal{E}_2$ is correct. This is the crucial assumption, which underlies the standard view of skeptical reasoning and the acceptance of floating conclusions. But is this assumption required? Is it necessary for a skeptical reasoner to assume, when a theory leads to multiple extensions, that one of those extensions must be correct?

Suppose that each of the theory's multiple extensions is endorsed by some credulous reasoner. Then the assumption that one of the theory's extensions must be correct is equivalent to the assumption that one of these credulous reasoners is right. But why should a skeptical reasoner assume that some credulous reasoner, following an entirely different reasoning policy, must be right? Of course, there may be situations in which it is appropriate for a skeptical reasoner to adopt this standard view—that one of the various credulous reasoners must be right, but that it is simply unclear which one. That might be the extent of the skepticism involved. But there also seem to be situations in which a deeper form of skepticism is appropriate—where each of the multiple extensions is undermined by another to such an extent that it seems like a real possibility that all of the credulous reasoners could be wrong. The yacht, spy, and economist examples illustrate situations that might call for this deeper form of skepticism.

As a policy for reasoning with conflicting defaults, the notion of skepticism was originally introduced into the field of nonmonotonic logic to characterize the particular system presented in [4], which did not involve the assumption that one of a theory's multiple extensions must be correct, and did not support floating conclusions. By now, however, the term is used almost uniformly to describe approaches that do rely on this assumption, so that the "skeptical conclusions" of a theory are generally identified as the statements supported by each of its multiple extensions, including the floating conclusions. Of course, there is nothing wrong with this usage of the term, as a technical description of the statements supported by each extension—except that it might tend to cut off avenues for research, suggesting that we now know exactly how to characterize the skeptical conclusions of a theory, so that the only issues remaining are matters concerning the efficient derivation of these conclusions. On the contrary, if we think of skepticism as the general policy of withholding judgment in the face of conflicting defaults, rather than arbitrarily favoring one default or another, there is a complex space of reasoning policies that could legitimately be described as skeptical, many of which involve focusing on the arguments that support particular conclusions, not just the conclusions themselves.

Acknowledgments

I owe many of the ideas in the paper to a series of conversations with Tamara Horowitz. I am grateful for valuable comments from Aldo Antonelli, David Makinson, and Richmond Thomason.

References

[1] Matt Ginsberg. *Essentials of Artificial Intelligence.* Morgan Kaufmann Publishers, 1993.

[2] John Horty. Moral dilemmas and nonmonotonic logic. *Journal of Philosophical Logic*, 23:35–65, 1994.

[3] John Horty. Some direct theories of nonmonotonic inheritance. In D. Gabbay, C. Hogger, and J. Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming, Volume 3: Nonmonotonic Reasoning and Uncertain Reasoning*, pages 111–187. Oxford University Press, 1994.

[4] John Horty, Richmond Thomason, and David Touretzky. A skeptical theory of inheritance in nonmonotonic semantic networks. *Artificial Intelligence*, 42:311–348, 1990.

[5] Vladimir Lifschitz. Circumscription. In D. Gabbay, C. Hogger, and J. Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming, Volume 3: Nonmonotonic Reasoning and Uncertain Reasoning*, pages 297–352. Oxford University Press, 1994.

[6] David Makinson. General patterns in nonmonotonic reasoning. In D. Gabbay, C.Hogger, and J. Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming, Volume 3: Nonmonotonic Reasoning and Uncertain Reasoning*, pages 35–110. Oxford University Press, 1994.

[7] David Makinson and Karl Schlechta. Floating conclusions and zombie paths: two deep difficulties in the "directly skeptical" approach to defeasible inheritance nets. *Artificial Intelligence*, 48:199–209, 1991.

[8] Drew McDermott. Non-monotonic logic II. *Journal of the Association for Computing Machinery*, 29:33–57, 1982.

[9] Henry Prakken and Gerard Vreeswijk. Logics for defeasible argumentation. In Dov Gabbay, editor, *Handbook of Philosophical Logic (Second Edition)*. Kluwer Academic Publishers, forthcoming.

[10] Raymond Reiter. A logic for default reasoning. *Artificial Intelligence*, 13:81–132, 1980.

[11] Lynn Stein. Resolving ambiguity in nonmonotonic inheritance hierarchies. *Artificial Intelligence*, 55:259–310, 1992.

[12] David Touretzky. *The Mathematics of Inheritance Systems*. Morgan Kaufmann, 1986.

[13] David Touretzky, John Horty, and Richmond Thomason. A clash of intuitions: the current state of nonmonotonic multiple inheritance systems. In *Proceedings of the Tenth International Joint Conference on Artificial Intelligence (IJCAI-87)*, pages 476–482. Morgan Kaufmann, 1987.

[14] Bas van Fraassen. Values and the heart's command. *The Journal of Philosophy*, 70:5–19, 1973.