What the Lottery Paradox Tells Us About Default Reasoning

David Poole
Department of Computer Science
University of British Columbia
Vancouver, B.C., Canada V6T 1W5
poole@cs.ubc.ca

Abstract

In this paper I argue that we do not understand the process of default reasoning. A number of examples are given which serve to distinguish different default reasoning systems. It is shown that if we do not make our assumptions explicit we get into trouble with disjunctive knowledge, and if we make our assumptions explicit, we run foul of the lottery paradox. None of the current popular default reasoning systems work on all of the examples. It is argued that the lottery paradox does arise in default reasoning and can cause problems. It is also shown that some of the intuitively plausible requirements for default reasoning are incompatible. How different systems cope with this is discussed.

1 Introduction

Default reasoning is the ability to jump to a conclusion based on the lack of evidence to the contrary. Deduction in standard logic does not allow such reasoning; if some proposition follows from a set of axioms, it follows from a superset of the axioms. There have been many proposals for incorporating default reasoning in logic [Reiter, 1980, McCarthy, 1986, Moore, 1985, Delgrande, 1987, Poole, 1988]. I assume we use default reasoning to predict what is true.

In this paper we consider the problem of default reasoning, and discuss different choices that could be made in developing a default reasoning system. A set of examples is presented which indicates that current default reasoning systems do not work properly.

This is argued in two parts. In the first part (section 2), it is argued that if we do not commit to implicit assumptions made (for example, acknowledging that we assumed Tweety is not an emu when we concluded Tweety could fly) we get into trouble. In the second part I show how the lottery paradox arises when we do commit to assumptions.

When considering the lottery paradox, the main intuition I rely on is the "one step default" property: if "birds fly" (however we represent it) is a default and all we know about an individual is that it is a bird (in particular, we don't know it doesn't fly), we conclude it flies. This seems like a minimal property the default "birds fly" should have.

2 Commitment

Suppose we have d as a default, with exception e. The first question I want to consider is whether we should conclude $\neg e$ as a side effect of concluding d.

Consider the classic example of birds flying:

Example 1 Suppose we want to use the default "birds fly", with emus as exceptions. Suppose we know Polly is a bird, and know nothing else about Polly. As "Birds fly" is an assumption, it seems reasonable to conclude Polly flies. Should we conclude Polly is not an emu? There have been three different solutions to this suggested by different systems.

2.1 Non-committal

The first of the possible answers is that we should not conclude d at all. We should rather conclude only the disjunct $d \vee e$. The rationale is that we do not know whether the exception e is true, so we do not know whether d is true. If we cannot say whether e is true, we should not allow any side effect to the value of e.

This is exactly the situation with circumscription [McCarthy, 1986] with the exception being "fixed" during the minimisation.

I would argue this non-committalness loses the very reason for default reasoning: we can never conclude a default, but only the disjunct of the possibilities. We have lost the ability to jump to conclusions. Such a system is not doing default reasoning at all; we have just invented a new syntax for disjunctions.

We can never use the "birds fly" default to do what was originally intended, namely to conclude something flies from just knowing it is a bird. We would instead conclude either the bird is an emu or flies. Somehow we changed the meaning of "birds fly, but emu's are exceptions" to mean the logical statement "birds are either emus or fly". With many exceptions we could only conclude Polly flies if we could prove polly is not an emu, is not a roast duck, is not in the shell, etc.

2.2 Non-commitment

An alternate view is we should conclude default d, but make no commitment as to whether e is true or not. That is we conclude d, and not conclude $\neg e$. This occurs, in autoepistemic logic [Moore, 1985] when we use the formula[1]:

$$\neg L \neg d \wedge \neg L e \Rightarrow d$$

(where the operator L means "know"), to mean that if we don't know d is false and we don't know e is true, conclude d.

Similarly, we can use Reiter's semi-normal defaults (as advocated in [Reiter and Criscuolo, 1981]):

$$\frac{: M(d \wedge \neg e)}{d}$$

to mean if $d \wedge \neg e$ is consistent, conclude d.

These get funny (and I would argue, incorrect) results, because they are being non-committal about the assumptions they are making. They do not allow us to conclude anything about the exception e. Consider the following example:

Example 2 Suppose by default people's left arms are usable, but a person with a broken left arm is an exception, and similarly people's right arms are, by default, usable, but broken right arms are an exception. In Reiter's notation (ignoring variables, which are irrelevant to this example) this is

$$\frac{: M(\textit{left-arm-usable} \wedge \neg \textit{left-arm-broken})}{\textit{left-arm-usable}}$$

$$\frac{: M(\textit{right-arm-usable} \wedge \neg \textit{right-arm-broken})}{\textit{right-arm-usable}}$$

If we know nothing about Matt's left arm, we conclude (correctly as to what we assumed a default was) his left arm is usable. If we know his left arm is broken, we (correctly again) do not conclude his left arm is usable.

Suppose we remember seeing him with a broken left arm or a broken right arm (we can't remember which). We add

$$\textit{left-arm-broken} \vee \textit{right-arm-broken}$$

In this case we cannot conclude he has a broken left arm and so conclude his left arm is usable. We also cannot conclude he has a broken right arm so we conclude his right arm is usable. We thus conclude both his left arm and his right arm are usable.

I would argue that this is definitely a bug, being able to conclude both arms are usable given we know one of his arms is broken. The problem is we have implicitly made an assumption, but have been prevented from considering what other assumptions we made as a side effect of this assumption. Somehow we needed to commit to the implicit assumption that his left arm was not broken when we used the first default.

This problem of disjunctive exceptions is endemic to the use of non-normal defaults[2].

2.3 Commitment to Assumptions

A third alternative is to conclude d, and as a side effect conclude $\neg e$. This reflects the idea that in concluding d, we are assuming e is not true, because if e were true we could not conclude d.

This is what happens in circumscription when we, for the "birds fly" example, minimise "ab", with "emu" varying and specify

$$\forall x(bird(x) \wedge \neg ab(x) \Rightarrow flies(x))$$
$$\forall x(emu(x) \Rightarrow ab(x))$$

or in Theorist [Poole, 1988] make "$birdsfly(X)$" a possible hypothesis and specify as facts

$$\forall x(bird(x) \wedge birdsfly(x) \Rightarrow flies(x))$$
$$\forall x(emu(x) \Rightarrow \neg birdsfly(x))$$

When we specify Tweety is a bird, we conclude Tweety is not an emu[3]. In the next section I consider the question as to whether such side effects can cause problems.

3 The Lottery Paradox

There is a famous problem which arises if we assume a proposition is false when its probability falls below some threshold. The problem arises because the conjunction of a number of likely propositions make become very unlikely or even impossible. This is known as the lottery paradox [Kyburg, 1961].

Suppose we have a threshold of ϵ. If there is a lottery with $> 1/\epsilon$ tickets, we assume each of these will not win. The conjunction of the assumptions is inconsistent. This is usually translated in probability theory as indicating that commitment is a bad idea.

In this section I show how the lottery paradox naturally arises in default reasoning systems[4] and can

[1] This analysis does not change if we use the more modular abnormality notation or use Gelfond's [1988] methodology for using autoepistemic logic.

[2] [Poole, 1988] shows how the use of preconditions in Reiter's defaults can lead to errors with disjunction. This example shows non normal defaults lead to errors with disjunctive knowledge. It is interesting to note that the simpler, and differently motivated Theorist system corresponds exactly to Reiter's normal defaults without preconditions [Poole, 1988, theorem 4.1].

[3] In both of these systems we conclude $\neg emu(c)$ for any constant c in our language that we do not know is an emu.

[4] [Kyburg, 1988, Perlis, 1987] also discuss how the lottery paradox can arise in a default reasoning system, but from a very different perspective.

potentially cause severe problems for current default reasoning systems. I then examine some possible responses to this problem.

Consider the following example where the circumscription convention of using named abnormality is used, as above. Assume all we are told about Tweety is that Tweety is a bird.

We start off by writing the sort of birds we may encounter in our domain and have a formula like[5]:

$$\forall x\ bird(x) \equiv emu(x) \vee penguin(x) \vee hummingbird(x) \vee sandpiper(x) \vee albatross(x) \vee ... \vee canary(x)$$

Now add defaults about birds. For each sort of bird that is exceptional in some way we will be able to conclude Tweety is not a bird of that sort.

- We conclude that Tweety is not an emu or a penguin because they are exceptional in not flying.
- We conclude Tweety is not a hummingbird as hummingbirds are exceptional in their size (consider for example the case of making a bird cage for Tweety; we have to make an assumption about the size of birds),
- we conclude Tweety is not a sandpiper as sandpipers are exceptional in nesting on the ground (for example, when bush walking and someone says "look at that bird nest", we have to look somewhere first; we look up by default if all we know is the nest belongs to a bird);
- we conclude Tweety is not an albatross as albatrosses are exceptional in some other way.

If every sort of bird is exceptional in some way, except for, say, the canary, we conclude Tweety is a canary (as we have ruled out all the other alternatives). This may or may not be a bad side effect. When we add the fact canaries are abnormal in being a bright colour, suddenly nothing works. We can no longer conclude Tweety flies! $flies(Tweety)$ is no longer in all minimal models. There is one model in which Tweety does not fly and in which all of the other abnormalities are false.

The problem is that local, seemingly irrelevant information (namely information about how different sorts of birds are abnormal in different ways) can interact to make nothing work. When we follow the advertised way to use these default reasoning systems, we find we get very strange behaviour. For seemingly unrelated statements to interact to produce such disastrous side effects is a bad technical problem.

Unlike McDermott [Hanks and McDermott, 1986, McDermott, 1987], I do not suggest this is evidence to give up on the programme of formalising commonsense reasoning using logic, but rather use this problem to shed more light on the phenomenon we are trying to formalise.

[5]This sort of statement naturally arises in systems where we assume complete knowledge.

4 Possible Responses

There a number of possible responses to this problem:

4.1 Denial

The first response is denial that this problem will ever arise in practice. Unfortunately this is an empirical question and not a theoretical question. We can argue about this forever, but until we actually go and build real systems and find out what does happen, the argument will be as irrelevant as trying to determine how many angels can fit on the head of a pin.

The problem outlined here was discovered by using our Theorist system [Poole et. al., 1987, Poole, 1988], and noticing funny side effects and obscure reasons why we should not predict (membership in all extensions) certain expected outcomes. We are currently building larger systems to determine whether such problems do arise. Unfortunately we will never be able to say this problem does not arise in practice, but only be able to determine it does.

I do not believe the scenario above, considering each type of bird as being exceptional in some way, is so far fetched. I would not be surprised, in a large database, if each subclass of bird is indeed exceptional in some way. All we need for the above problem to arise is some way to determine there is no completely typical individual. Once we can determine this, none of the formalisms (that commit to an assumption) work correctly. In large knowledge bases, not only would I expect such situations to arise, but they would be normal. For example, the "normal" person (who is 175cm tall, has an IQ of 100, has a grade 12 education and has 2.2 children), does not exist, although we may want to make these assumptions so we can point out to others how someone is different to that "normal" person.

Example 3 As a natural example, take the well known default in the legal system namely "people are presumed innocent unless proven guilty", and the knowledge that someone is guilty (as there was a crime committed). This could be represented as

$$\forall x\ \neg ab(innocence, x) \Rightarrow \neg guilty(x)$$
$$\exists x\ guilty(x)$$

For any particular individual we do not conclude they are not guilty. I would not like to be the one to explain to judge Jones that we do not conclude

$$\neg guilty(judge\_jones)$$

but do conclude

$$\neg guilty(judge\_jones) \vee \neg guilty(jack\_the\_ripper)$$

4.2 Technical Patches

I described this as a "technical problem", and as such it seems as though it should have a technical solution. I believe the problem is endemic to current ideas about how defaults work (see section 5 below). There is good evidence to suggest that any solution to the lottery paradox above will not work for the broken arm problem (the structure of both of them is remarkably similar, but we expect different answers). I see this as a challenge to those who like to find technical solutions, but I feel as though the problem is we do not understand the phenomena we are trying to formalise.

Suggestions such as prioritised circumscription [McCarthy, 1986] will not work. There is a symmetry about this example. The side effect will affect whatever is the lowest priority default.

One interesting thing can be seen in this example. If we predict what is in one extension rather than what is in all extensions (following the definition of extension in [Reiter, 1980] or [Poole, 1988]), we find "Tweety flies" is in one extension. If we add the exceptions (emu's are abnormal with respect to flying) as facts, we can also predict Tweety doesn't fly. Poole [1988] suggests this problem (of having the side effect explicit) could be solved by using "constraints" to prune the scenarios without being part of the scenarios. This works if we equate prediction with being in one extension. We can explain Tweety flying, but cannot explain the negation. We can explain each of the other typical properties of birds; we cannot predict the conjunction of the properties. The use of constraints also does not let us conclude both of Matt's arms are usable.

This seems to be a good "technical patch" of the type we were looking for. However, equating prediction with membership in an extension leads to the peculiar property of predicting a proposition and also predicting its negation. Careful structuring of the knowledge base may help this (see section 4.3), but this is not a general solution.

If, instead we equate prediction with membership in all extensions, the use of constraints again does not work. The conjunction of all of the normal assumptions about birds is inconsistent; removing the assumption Tweety is normal with respect to flying is a way to make an extension from which we cannot conclude Tweety flies.

4.3 Breaking Conventions

Let us now consider one sort of knowledge it is claimed defaults capture. This is the idea that default reasoning models a notion of conventional reasoning. The reason "birds fly" is a default is that if I tell you Tweety is a bird, and I do not tell you Tweety cannot fly, I am telling you Tweety can fly. This is the motivation for autoepistemic reasoning [Moore, 1985, section 2].

If we take this meaning of default reasoning seriously, not only does the lottery paradox above not arise, but the multiple extension problem in general does not arise.

According to the defaults as conventions view, the default "birds fly" means if I add knowledge about a particular bird, I must assert it doesn't fly if it doesn't fly. With this convention, if I assert Tweety is a bird, and I do not assert Tweety does not fly, I am saying Tweety flies. If Tweety does not fly I have broken the convention. The knowledge based should be fixed up just as if I had asserted something that is false.

If there are multiple extensions they are mutually inconsistent, so at most one can be true of the world under consideration (the intended interpretation). Thus one of the extensions must be false in the intended interpretation. So either something I added explicitly is false in the intended interpretation, or else there is a default which is not applicable in the world under consideration. In the latter case, to follow our convention, I must tell the system about that exception. Multiple extensions indicate I did not follow the convention.

In the lottery paradox example above, Tweety is exceptional in at least one of the properties, so to follow the convention, I should tell you that property. Thus the lottery paradox example cannot arise. Moreover, none of the multiple extension problems can arise. Multiple extensions are thus not a problem to be solved, but indicate a convention has been broken; we need to patch up our buggy knowledge base rather than solve the multiple extension problem.

Automatically enforcing such constraints is not as difficult as it may, at first, seem. In the Theorist system [Poole et. al., 1987, Poole, 1988], we can maintain a knowledge base with only one extension [Poole, 1989a] by ensuring that:

1. When a new default is added to the knowledge base, if we can explain an instance of the negation of the default and cannot prove that instance, this default introduces multiple extensions. If not, we still only have one extension.
2. When a new fact is added, if we can explain the negation of the fact with a explanation containing more than one default and cannot explain it with a subset of that explanation containing only one default, the new fact introduced multiple extensions.

When we detect we have multiple extensions we can ask the user to debug the knowledge base by cancelling one of the defaults [Poole, 1988]. These detection procedures are, in general, undecidable. However it seems appropriate to assign these to low priority background processes, which report when they find an inconsistency or a multiple extension. Just as people do not immediately (if at all) realise they have been mislead (or lied to), these background processes may or may not return to report a breaking of a convention.

The importance of this section is that "defaults as conventions" is a consistent view of defaults; whether it corresponds to the use of the term default is a different question.

4.4 We don't understand the Phenomena.

The fourth response is we do not understand the phenomena we are trying to formalise. If we mean some sort of "typically", the response in section 4.3 does not seem to be appropriate. If this is the case we must recognise that the lottery paradox can arise in the formal systems defined so far. If we claim the lottery paradox does not arise in the "commonsense" view of a default, then the formal systems do not capture our normal sense of "default". Thus we do not understand the phenomena we are trying to characterise.

In the next section I show that one intuitive reading of "birds fly" is incompatible with many of the formal models of non-monotonic reasoning.

5 One Step Default Property

The property underlying the intuition in the lottery paradox example is what I call the "one step default property"[6].

I will use the notation "$p(x) \rightarrow q(x)$" is a default to mean "p's are q's by default". No meaning should be placed in this notation. Different systems use different notations and have different semantics. I intend this discussion to include all of these notations.

Definition 1 A default reasoning system has the **one step default property** if whenever "$p(x) \rightarrow q(x)$" is a default and all that is given about constant c is "$p(c)$" (in particular we do not know the truth of $q(c)$), it concludes "$q(c)$".

For example, under this property if I tell you "birds fly", and all I tell you about Tweety is Tweety is a bird, if a system has the one step default property it concludes Tweety flies. This seems like a minimal property "birds fly" should have.

The following theorem puts a constraint on the type of systems with this property.

Theorem 1 A default reasoning system cannot have all of the following properties:

(i) The one step default property.

(ii) If it concludes two answers, it concludes their conjunction. That is, if it concludes "a" and concludes "b", it concludes "$a \wedge b$".

(iii) The ability to represent disjunctive knowledge, and to allow arbitrary (not directly conflicting) defaults.

(iv) It does not conclude anything known to be false[7].

Proof: To prove this it suffices to give one set of inputs which follow the constraints given in (iii). By showing that properties (i) and (ii), lead to a contradiction with (iv), we demonstrate that a system with all four properties cannot exist.

Suppose

$$p(x) \rightarrow q_i(x)$$

is a default for $i = 1..n$, and

$$\forall x\ \neg q_1(x) \vee \neg q_2(x) \vee ... \vee \neg q_n(x)$$

is a fact, and we are given

$$p(c)$$

By (i) we conclude each "$q_i(c)$", and by (ii) we conclude their conjunction, which is inconsistent, and so must be false, contravening (iv). □

Given these four intuitive properties are inconsistent, it is interesting to consider which property different systems have given up.

(i) is given up in circumscription [McCarthy, 1986], in any minimal model solution [Shoham, 1987] and systems which require membership in all extensions [McDermott and Doyle, 1980]. This is because they want the expressiveness that property (iii) gives, they need property (ii) by their very nature, and always reject having inconsistent extensions or reducing to no models.

(ii) is given up in many probability-based systems [Neufeld and Poole, 1988, Bacchus, 1989], and in systems which, for prediction, only require membership in one extension [Reiter, 1980, Moore, 1985, Poole, 1988]. These latter systems seem to get the one step default property for the wrong reason, namely by being able to predict a proposition and also predict its negation.

(iii) is given up in inheritance systems [Thomason and Horty, 1988]. These allow (i), (ii) and (iv), however they lack the expressiveness of the richer logic-based formalisms.

(iv) is not given up by any system I know, although it is argued [Israel, 1980, Perlis, 1987, Kyburg, 1988] that commonsense reasoning does indeed require reasoning under inconsistency.

[6] This discussion is in terms of parametrized (open) defaults as is it most natural for this case. However the argument is purely propositional, and covers propositional systems as well as systems allowing defaults with free variables.

[7] We do not want it to be inconsistent if the facts are consistent. This property does not constrain the system at all if the facts given are inconsistent.

The ϵ-semantics of [Pearl, 1988] fits into this analysis in a very interesting way. For this theorem it fails in property (iii). There is no consistent probability assignment for the defaults and facts given in the proof. This could be translated as meaning it solves the problem nicely, but I would claim it means we must treat seriously the semantics saying there are only infinitesimally few exceptions. It shows we cannot use the system if the proportion of exceptions does not have measure zero. In particular this system does not seem appropriate to represent "birds fly", as it is not true there are infinitesimally few birds that don't fly. His semantics means accepting the "convention" view of defaults (section 4.3).

Shoham [1987] rejects the one step default property in his discussion on the lottery paradox. However his discussion indicates that we would not want to write such defaults, but explicitly rejects the view of defaults as autoepistemic statements (section 4.3). Rather than indicating to the user that the knowledge base is inconsistent, he would rather [Shoham, 1987, p. 392] the system decide that the user was not rational in adding the default that each lottery ticket would not win, and so not allow the one step default conclusion.

6 Where to look for a solution

I think there are two areas to look for a solution to this problem: these are in the areas of probability theory, and in comparing logical arguments as to why we should believe some proposition or not.

6.1 Probability

Pearl [1988] and Cheeseman [1985] argue very logically and convincingly that probability theory is the correct way to consider reasoning under uncertainty.

The one step default property is ingrained at the very foundation of probability theory. $p(A|B) = v$ only tells us information about A when all we know is B[8] [Pearl, 1988]. Not unsurprisingly, default reasoning systems based on probability theory (eg. [Geffner, 1988]) end up with different properties than those based on minimal models or other logical formalisms which do not have the one step default property.

According to probability theory the lottery paradox is a problem with commitment to assumptions. The problem is concluding a proposition is true without being certain of the proposition. Instead of concluding Tweety flies we could conclude the probability of Tweety flying is high. The conjunction of the conclusions would have probability zero, but we know we cannot conclude the conjunction of propositions is likely just because the proposition is likely.

One of the promising ideas in this area is to use qualitative probabilities [Aleliunas, 1988], where instead of using numbers we can use more linguistic probability values in a probability algebra. The relationship between this and notions of default reasoning is not clear.

Another promising idea is that of Neufeld [1988], where "birds fly" means the probabilistic statement Tweety being a bird increases our belief in Tweety flying:

$$p(flies|bird) > p(flies)$$

The lottery paradox is overcome by not allowing us to conclude the belief in the conjunction is increased just because belief in each proposition is increased.

6.2 Arguments

The second promising area is to consider the role of logical arguments.

Logic can be seen as the study of arguments; it is the study of when we should believe an argument based on the truth of the premises. A valid logical argument is one in which the conclusion must be true if the premises are true. It has been shown [Poole, 1988] that defaults can be treated as possible hypotheses that can be used in the premise of a logical argument. The defaults are the premises of a logical deduction; we do not defeat the argument, but defeat the premises (by showing they are inconsistent). All of the arguments are standard logical proofs. Multiple extensions indicate there is an argument for a proposition and an argument against a proposition.

A natural way to consider default reasoning is to compare the arguments for and against some proposition. Poole [1989a] shows how membership in all extensions can be seen as a process of dialectics. A goal is in all extensions if and only if there is a set of explanations for the proposition such that there is no scenario inconsistent with all of the explanations. This can be modelled at two agents having an argument; one agent finds arguments for the goal and the other agent tries to find a scenario in which all of the first agent's arguments fall down [Poole, 1989a].

In the example of section 3, given Tweety is a bird, there is a very short argument that Tweety flies (namely because Tweety is a bird and "birds fly"). There is a long convoluted argument saying Tweety does not fly (namely by assuming other normalities of birds which eliminates all other possible types of birds Tweety can be, except for the non-flying ones).

It seems reasonable to view reasoning as a process of evaluating logical arguments (or more precisely the premises of logical arguments), and preferring more direct (in some sense) arguments. There is one consequence of this way to view the lottery paradox. We end up with a direct argument that Tweety flies. We end up with all of the other direct arguments about Tweety. The problem is the conjunction of these assumptions is inconsistent. Although we would predict a number of consequences of our knowledge, we

[8] In particular, if $v \neq 0$, it tells us nothing about the value if $p(A|B \wedge C)$.

may not want to predict the conjunction of these consequences. This is exactly the lottery paradox. We predict any particular lottery ticket is not going to win. When we conjoin many such predictions problems arise.

One of the reasons the lottery paradox example is so persuasive is because of our intuitions about wanting to prefer more specific knowledge [Touretzky, 1986, Thomason and Horty, 1988, Poole, 1985, Geffner, 1988]. One intuition behind specificity is exactly the one step default property (we prefer the one step default that emu's don't fly over the longer argument that emus are birds and birds fly). If this is so, any method that compares extensions or models (without regard to the question being asked) is not going to be the basis for a model of default reasoning incorporating specificity. Either it says "yes" to the conjunction of the predictions (which is inconsistent, and so would predict something known to be false), or it says "no" to flies(Tweety), and so must find very circuitous arguments to defeat the direct implication (which seems antithetical to the notion of preferring more specific knowledge).

This idea of solving specificity problems by comparing logical arguments is pursued further in [Poole, 1989b].

7 Conclusion

Rather than suggesting we give up on logic when we find the default reasoning formalisms do not give the answers I would like, I have argued that we need to reconsider the phenomena we are trying to formalise. The way the lottery paradox can easily arise shows the fragility of current default reasoning systems. I believe the right solution is to consider the role of dialectics; we must compare arguments for and against propositions. However, for logicists to defeat the arguments that probability theory is the appropriate framework in which to view this will not be easy. We both need to understand the problems we are trying to solve.

The other moral of this paper is that we must build systems to see how our reasoning systems work in practice. The instance of the lottery paradox was found while using our Theorist implementation [Poole et. al., 1987]. No one understands what other problems will arise when we start to solve non-trivial problems.

Acknowledgements

Thanks to all of the people at the 2nd International Workshop on Nonmonotonic Reasoning, at Grassau in June 1988 who were involved in sharpening my ideas on this problem, especially John McCarthy, Matt Ginsberg, Michael Gelfond, Brian Haugh, Scott Goodwin, Vladimir Lifschitz and Gerd Brewka. Thanks also to Randy Goebel, Eric Neufeld, Paul Van Arragon, Romas Aleliunas, and Greg Provan for valuable discussions on the points raised in this paper.

This research was supported under NSERC grant OGPOO44121.

References

[Aleliunas, 1988] R. Alelianas, "A new normative theory of probabilistic logic", *Proc. CSCSI-88.*

[Bacchus, 1989] F. Bacchus, "A Modest, but Semantically Well Founded Inheritance Reasoner", Tech. Report, University of Waterloo.

[Cheeseman, 1985] Peter Cheeseman. In defense of probability. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, pages 1002–1009, Los Angeles, California, August 1985. International Joint Committee on Artificial Intelligence.

[Delgrande, 1987] J. P. Delgrande, "A first-order conditional logic for prototypical properties", *Artificial Intelligence*, Vol. 33, No. 1, 105-130.

[Gelfond, 1988] M. Gelfond, "Autoepistemic Logic and Formalization of Commonsense Reasoning: Preliminary Report", in *Proceedings of the 2nd International Workshop on Non-Monotonic Reasoning*, Springer-Verlag, Lecture Notes in Artificial Intelligence, No. 346, pp. 176-186.

[Geffner, 1988] H. Geffner, "On the logic of Defaults", Proc. AAAI-88, Minneapolis.

[Hanks and McDermott, 1986] S. Hanks and D. McDermott, "Default Reasoning, Nonmonotonic Logics and the Frame Problem", *Proc. AAAI-86*, pp. 328-333.

[Israel, 1980] D. J. Israel, "Whats wrong with nonmonotonic logic", *Proc. AAAI-80*, 99-101.

[Kyburg, 1961] Henry E. Kyburg, Jr., *Probability and the Logic of Rational Belief*, Wesleyan University Press, Middletown, 1961.

[Kyburg, 1988] Henry E. Kyburg, Jr., Probabilistic Inference and Non-monotonic Inference, *Proc. Fourth Workshop on Uncertainty in Artificial Intelligence*, pp. 221-228.

[McCarthy, 1986] J. McCarthy, Applications of Circumscription to formalising common-sense knowledge, *Artificial Intelligence*, Vol. 28, pp. 89-116.

[McDermott and Doyle, 1980] D. V. McDermott and J. Doyle, "Non-monotonic logic 1", *Artificial Intelligence*, Vol. 13, pp. 41-72.

[McDermott, 1987] D. McDermott, A Critique of Pure Reason, *Computational Intelligence*, Vol. 3, No. 3, August 1987, pp. 151-160.

[Moore, 1985] R. C. Moore, Semantical Considerations on nonmonotonic logic, *Artificial Intelligence* 25 (1) 75-94.

[Neufeld and Poole, 1988] E. Neufeld and D. Poole, Probabilistic Semantics and Defaults, *Proc. Fourth Workshop on Uncertainty in Artificial Intelligence*, pp. 275-282.

[Pearl, 1988] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufmann, San Mateo, 1988.

[Perlis, 1987] D. Perlis, "On the Consistency of commonsense reasoning", *Computational Intelligence*, Vol. 2, No. 4, 180-190.

[Poole, 1985] D. Poole, "On the Comparison of Theories: Preferring the Most Specific Explanation", Proc. IJCAI85, pp.144-147.

[Poole, 1988] D. Poole, A logical framework for default reasoning, *Artificial Intelligence*, Vol 36, pp. 27-47.

[Poole, 1989a] D. Poole, "Explanation and Prediction: An Architecture for Default and Abductive Reasoning", to appear, *Computational Intelligence.*

[Poole, 1989b] D. Poole, Dialectics and Specificity, in preparation.

[Poole et. al., 1987] D. L. Poole, R. G. Goebel, and R. Aleliunas, "Theorist: a logical reasoning system for defaults and diagnosis", in N. Cercone and G.McCalla (Eds.) *The Knowledge Frontier: Essays in the Representation of Knowledge*, Springer Varlag, New York, 1987, pp. 331-352.

[Reiter, 1980] R.Reiter, "A Logic for Default Reasoning", *Artificial Intelligence*, Vol. 33, pp. 81-132.

[Reiter and Criscuolo, 1981] R. Reiter and G. Criscuolo, "On Interacting Defaults", *Proc. Seventh International Joint Conference on Artificial Intelligence*, pp. 270-276.

[Shoham, 1987] Y. Shoham, "Nonmonotonic Logics: Meaning and Utility", *Proc. IJCAI-87*, pp. 388-393.

[Thomason and Horty, 1988] R. H. Thomason and J. F. Horty, "Logics for Inheritance Theory", in *Proceedings of the 2nd International Workshop on Non-Monotonic Reasoning*, Springer-Verlag, Lecture Notes in Artificial Intelligence, No. 346, pp. 220-237.

[Touretzky, 1986] D. S. Touretzky, *The Mathematics of Inheritance Theory*, Pitman / Morgan Kaufmann.