

**The Effect of Knowledge on Belief:
Conditioning, Specificity and the Lottery
Paradox in Default Reasoning**

by

David Poole

Technical Report 90-19
June 1990

Department of Computer Science
The University of British Columbia
Vancouver, B.C. V6T 1W5
Canada

Abstract

How should what one knows about an individual affect default conclusions about that individual? This paper contrasts two views of "knowledge" in default reasoning systems. The first is the traditional view that one knows just what is in one's knowledge base. It is shown how, under this interpretation, having to know an exception is too strong for default reasoning. It is argued that we need to distinguish "background" and "contingent" knowledge in order to be able to handle specificity, and that this is a natural distinction. The second view of knowledge is what is contingently known about the world under consideration. Using this view of knowledge, a notion of conditioning that seems like a minimal property of a default is defined. Finally, a qualitative version of the lottery paradox is given; if we want to be able to say that individuals that are typical in every respect do not exist, we should not expect to conclude the conjunction of our default conclusions.

This paper expands on work in the proceedings of the First International Conference on Principles of Knowledge Representation and Reasoning [35].

1 Introduction

Default reasoning can be seen as jumping to conclusions about some individual based on knowledge about that individual.

Many papers have considered solutions to the so called “multiple extension problem” [38, 14, 24, 33], where conclusions of different defaults are in conflict. These solutions usually consider the case where the antecedents of the defaults both happen to be true. For example, if we have defaults that Quakers are pacifists and Republicans are not pacifists, we have to consider what to do when we have someone who is both a Quaker and a Republican [38]. Solutions to these problems usually consist of designing mechanisms for blocking defaults under appropriate conditions.

This paper concentrates on instances of the multiple extension problem, with the following property: there is some default such that whenever the antecedent of the default is true, there are extensions in conflict with the conclusion of the default. Unless the multiple extension problem is solved in a satisfactory way such defaults will never be used. Many of the standard ways to solve the multiple extension problem, for example deriving conclusions that are in all extensions, render such defaults useless; the defaults can effectively never be used.

There are two cases where this phenomena occurs. When the competing defaults have equivalent antecedents we have a qualitative version of the lottery paradox [17]. In the other case we have more specific knowledge competing with more general knowledge, and need to prefer more specific defaults if we want them to be usable.

In this paper we appeal to and expand on an intuition of “conditioning” that says “if p ’s are q ’s by default, and all we know about individual C is $p(C)$, we should conclude $q(C)$ ”. This is considered to be a minimal property of a default.

For example, suppose there is the default “birds fly”, and someone phones me up and says “Tweety is a bird”, and that is all I have ever heard about Tweety, then, using this default, I should conclude that Tweety flies. If I were not to use the default in this case, it seems as though this default would never be used. However, if I knew more about Tweety, then maybe I would not want to conclude that Tweety flies. What needs to be in a system to make sure that such a property holds is the basis of this paper.

The results of this paper can be summarised as:

1. Requiring one to know an exception (e.g., Reiter's Default Logic [37], Autoepistemic Logic [26, 21]) is too strong a condition to capture the naive intuition behind defaults. Similarly the idea of explicitly cancelling defaults cannot be used by itself to solve the multiple extension problem in a satisfactory way.
2. If we want a local interpretation of defaults, we need to have specificity; that is, we need to prefer more specific defaults over more general ones when they compete.
3. We should distinguish "background" and "contingent" knowledge in order to handle specificity. This is more than a syntactic distinction.
4. The lottery paradox arises naturally. If we want the ability to conclude that individuals that are typical in every respect do not exist, conditioning is incompatible with logical closure.

When we talk about a system getting the wrong answer, we have two possible meanings. For the "brave systems" that rely on membership in one extension [37, 26, 33], we mean that there is a wrong conclusion in one of the extensions (as opposed to not being able to derive the correct conclusion in one extension). For the more skeptical systems, such as those that require membership in all extensions [25, 24, 34] we mean that the desired result cannot be concluded.

2 Current Logic-based systems

Consider the following example:

Example 2.1 Suppose we are given:

All emus are birds.
Birds fly, by default.
Emus don't fly, by default.

Should we conclude that an arbitrary individual is not an emu?

This is more than a consideration of whether contrapositives should be allowed (e.g., ϵ -semantics [29] does not allow contrapositives in general, but

does answer “yes” to this question). The argument for concluding that an arbitrary individual is not an emu goes something like:

If the individual were a bird, we would conclude that the individual flies, and so we are implicitly assuming it is not an emu (as emus don’t fly). If it is not a bird, it is not an emu. Thus in either case it is not an emu.

We divide systems into classes as to whether they answer “yes” or “no” to this question. Virtually all of the systems considered do conclude that an arbitrary individual is not an emu [24, 33, 29, 19]. In some systems [37, 26, 21] the default can be represented so that the answer to the question is either “no” or “yes”. The representation that leads to the answer “no”, does so by explicitly blocking the side effect. In section 2.1 we argue that this explicit blocking is unintuitive in most cases. In section 2.2 it is shown how problems arise with the side effects (e.g., of concluding that arbitrary individuals are not emus in example 2.1) interacting.

2.1 Having to know an exception

There are a number of representation systems for which one has to “know” an exception before a default is blocked [37, 26, 21].

As an example, consider a representation of “birds fly, but emus are exceptional” using Reiter’s semi-normal defaults where we don’t want to conclude that typical birds are not emus ¹:

$$\frac{bird(x) : flies(x) \wedge \neg emu(x)}{flies(x)}$$

An equivalent formulation in autoepistemic logic [26, 21] can be given by the axiom schema:

$$L\ bird(x) \wedge \neg L\neg(flies(x) \wedge \neg emu(x)) \Rightarrow flies(x)$$

If we augment this with

¹Throughout this paper the convention of having variables, function symbols and predicate symbols in lower case and constants in upper case is used.

bird(Tweety)
bird(Polly)

we can conclude

flies(Tweety) ∧ flies(Polly)

using the default twice (once for $x = Tweety$, and once for $x = Polly$). There is no conclusion about the emuness of Tweety or Polly.

This may be considered as an appropriate answer. If, however, we add

emu(Tweety) ∨ emu(Polly),

we still have the same conclusion: both Tweety and Polly fly, even though the disjunct tells us that one of them is exceptional.

The problem is in the semi-normal nature of the default; having to “know” or “prove” an exception is much too strong. The disjunction is not strong enough to cancel either default (or even cancel the use of both defaults together).

Consider how other information in the knowledge base could prevent the conclusion of the conjunction. To block the conjunction we have to use the disjunction in some way, as without the disjunction we want to conclude the conjunction. To block one of the default instances, we need to conclude that one of the birds is an emu (or does not fly). Consider how to block the default for Tweety. To use the disjunction to conclude *emu(Tweety)* we need to conclude $\neg emu(Polly)$ (for example, by having $\forall x emu(x) \Rightarrow \neg flies(x)$ as a fact). This is precisely the side effect that the semi-normal defaults do not allow.

Note that this problem is endemic to the use of non-normal defaults. If we have the semi-normal defaults (or instances of semi-normal defaults):

$$\frac{: \alpha \wedge \beta}{\beta} \quad \frac{: \gamma \wedge \delta}{\delta},$$

the fact $\neg \alpha \vee \neg \gamma$ does not block the conclusion $\beta \wedge \delta$.

This problem does not require the explicit statement of disjuncts:

Example 2.2 Consider the semi-normal defaults:

$$\frac{bird(x) : flies(x) \wedge \neg dead(x)}{flies(x)}$$

$$\frac{of_ancient_species(x) : fossilised(x) \wedge dead(x)}{fossilised(x)}$$

If we are given the facts

$$bird(Fred) \wedge of_ancient_species(Fred)$$

we can conclude

$$flies(Fred) \wedge fossilised(Fred).$$

The semi-normal nature of the defaults does not recognise the implicit assumptions that Fred is both dead and not dead.

Example 2.3 This problem also manifests itself in a different way if we follow Brewka's [4] suggestion of using semi-normal defaults of the form²:

$$\frac{: M flies(x)}{bird(x) \Rightarrow flies(x)}$$

to allow case analysis on the antecedents [33], and to also block contrapositives. This representation allows us to conclude $flies(Tweety) \vee flies(Polly)$ from $bird(Tweety) \vee bird(Polly)$, and blocks the conclusion of $\neg bird(Sylvester)$ from $\neg flies(Sylvester)$. However, given $\neg flies(Fred) \vee \neg flies(Mary)$, the contrapositive is not blocked and $\neg bird(Fred) \vee \neg bird(Mary)$ is concluded.

Example 2.4 Although this may seem like peculiar behaviour for these default examples, there are examples where this behaviour does seem appropriate [7]. Consider the default that someone who has a motive, and may be guilty, should be a suspect:

$$\frac{has-motive(x) : suspect(x) \wedge guilty(x)}{suspect(x)}$$

This default can be blocked for an individual *John* if we knew $\neg guilty(John)$.

The disjunctive exceptions also seem reasonable for this example; if we know both Pete and Mary have motives and we know one is not guilty, it is reasonable to conclude that they both are suspects.

This example perhaps shows the distinction between default reasoning and autoepistemic reasoning that was pointed out by Moore [26]. There are cases where not knowing the particular counterexample is important, but these seem to be the exception rather than the rule.

²This is also the form suggested by Levesque [21, p. 291] to represent "birds fly".

2.2 Concluding Exceptions are False

The alternative answer to the question posed at the start of section 2 is to conclude that exceptions are false (i.e., that the arbitrary individual in example 2.1 is not an emu).

Example 2.5 Consider the following elaboration of example 2.1:

All emus are birds.
Birds fly, by default.
Emus don't fly, by default.
If something looks like an emu, it is an emu, by default.

Suppose we are also told two facts

Tweety is an emu.
Polly looks like an emu.

Intuitively we would like to conclude that Tweety does not fly (as we have the direct default that emus don't fly, and thus are exceptional birds), and, similarly, we would like to conclude that Polly is an emu. This example is considered as two separate cases in the following two sections.

2.2.1 Specificity

There is a very strong intuition that, based on the information in example 2.5, we should be able to conclude that Tweety does not fly. Although there are two potentially applicable defaults, the one applicable to emus is more specific and thus should be preferred over the more general default about birds (it is a more specific default as it is about a more specific class). This notion of preference for more specific knowledge has been advocated by many authors [42, 32, 24, 22, 11, 41, 29].

If we don't want to conclude Tweety does not fly in example 2.5, it seems as though the default "emus don't fly" can never be used. Whenever it is able to be used, the "birds fly" default is also applicable, and competes with this default. Thus, unless we want a default to be useless, we should prefer to use the more specific default.

There are three basic approaches that have been considered to ensure that we conclude that Tweety does not fly:

1. Force the user to add “cancellation axioms” to stop the use of the more general default [24, 33]. For example, we could name the first default “ $\neg ab(Birds\ fly, x)$ ” by writing

$$\forall x\ bird(x) \wedge \neg ab(Birds\ fly, x) \Rightarrow flies(x)$$

and adding a cancellation axiom

$$\forall x\ emu(x) \Rightarrow ab(Birds\ fly, x)$$

In example 2.5, we would conclude that Tweety does not fly, as we can use the cancellation axiom to prove the “birds fly” default is not applicable to Tweety.

2. Build a general priority system, and make the user add priorities³ [24, 3]. The user would make the “emu’s don’t fly” default have higher priority than the “birds fly” default; when they compete, as in this example, the higher priority default would prevail.
3. Incorporate specificity into the default reasoning system automatically, [42, 32, 22, 11, 41, 27, 1]. This is discussed further in section 3.

2.2.2 Inheritance of cancellation

Based on the information in example 2.5, we also want to conclude that Polly, who looks like an emu, is an emu. This is similar to the previous specificity case, in that the direct default “if it looks like an emu, it is an emu” competes with the conclusion that Polly is not an emu using the first two defaults. If we ever want the third default to be used we have to counter the conclusion of Polly not being an emu.

Consider the three proposed solutions to the specificity problem above:

1. If we are using cancellation, we have to cancel defaults that argue against Polly being an emu. The most direct counterargument uses the cancellation axiom introduced for specificity in section 2.2.1. By assuming the first default (i.e., $\neg ab(Birds\ fly, Polly)$), we can use the cancellation axiom to conclude that Polly is not an emu (even given

³If we want to automatically add these priorities, this is considered to be the third case.

no explicit facts about Polly). We must block this conclusion to allow only the conclusion that Polly is an emu. To block the conclusion we have to specify something like

$$\forall x \text{ looks_like_emu}(x) \Rightarrow ab(\text{Birds fly}, x)$$

In other words, objects that look like emus must inherit the cancellation of the emu class. Similarly some properties that, by default, allow us to conclude that an individual looks like an emu, must also inherit the cancellation of the “birds fly” default.

This is still not right. Suppose we have an individual Fred that is a bird and looks like an emu, but is not an emu. With the above cancellation axioms in effect, we cannot conclude that Fred flies, even though we know that Fred is not a member of the only exceptional class of birds given.

Thus it seems as though “cancellation axioms” do not provide the tools we need to treat even this simple example in a satisfactory way⁴.

2. The second case is to use some form of prioritisation. For Polly there are three defaults that together are in conflict (no two of which are in conflict). We need the default about looking like an emu to have higher priority than at least one of the other defaults. As the “birds fly” default has lower priority than the “emus do not fly” default, the default about looking like an emu must have higher priority than the “birds fly” default.

Such arguments about the relative priority of defaults can lead to “counter examples” to the universal applicability of static prioritisation. The following example is constructed in a manner similar to the example for cancellation. It may look complicated, but the idea is simple. We want to create an example where, in order for the defaults to

⁴Note however, that the use of cancellation advocated here (and in [33]) can solve the problems that motivated the development of prioritised circumscription in [24]. Rather than cancelling the cancellation axioms, as in [24], a simpler idea is to add new defaults for the subclasses. The intuition is that emus, because they are birds, are exceptional beings with respect to flying, and because they are emus, are exceptional birds. They have their own reason for not flying. See [33] for more details. When we try to use these defaults in practice however, they break because of the reasons in the text.

be used, we need some minimal ordering of defaults. We then set into competition two defaults that are not closely related.

Example 2.6 Suppose we want to represent the following defaults:

- (i) Canadians speak English as a native language.
- (ii) Quebecois are Canadians who do not speak english as a native language.
- (iii) If someone is in Quebec they are Quebecois.
- (iv) If someone says that they are "en Quebec" they are in Quebec.
- (v) If Fred says someone said they are "en Quebec", they said they were "en Quebec".

In order for the second default to be applicable, it has to have higher priority than the first default. In order for the third default to be applicable, it has to have higher priority than one of the first and the second, and so must have higher priority than the first. Similarly the fourth and fifth defaults must also have priority over the first in order for them to be applicable.

Suppose we have the facts:

- (a) Mary is a Canadian.
- (b) Mary is not Quebecois.
- (c) Mary is in Quebec.
- (d) Fred says that Mary said she was "en Quebec".
- (e) If Mary is a native English Speaker, she would not have said she is "en Quebec".

We have created a competition between the first and the fifth defaults. Because of the previous considerations, the fifth default must have priority over the first, and thus we conclude that Mary said she was "en Quebec" and so is not a native English speaker. This is very peculiar: the logic tells us that the default about Fred's literal reliability should have priority over the typicality of non-Quebecois Canadians.

There are many questions that arise as to where priorities come from, how do we add priorities, and how does a user know where to add priorities (this is particularly important when we have recursive rules, such as in frame axioms [14]). If we want to be able to automatically infer priorities (e.g., [13]), we need to consider the next case.

3. The last case is where the system can automatically handle specificity. This is discussed further in section 3. It is important to note that, as the inheritance of cancellation of example 2.5 shows, specificity can be more complicated than the (conceptually simple, but still tricky) case discussed in section 2.2.1. When we have an object that looks like an emu, we need some form of specificity to override the natural tendency to conclude that the object is not an emu.

2.2.3 A Qualitative Lottery Paradox

The previous example and discussion considered cases where there is a default, such that whenever the antecedent of the default is true there are extensions that run counter to the conclusion of the default. In the previous examples, at least some of the competing defaults were more general (there are cases where the more general default is applicable, and the more specific is not), and it was these defaults that needed to be blocked to force the more specific default to apply. There is one case where the problem is not to do with specificity. This is where the competing defaults all have the same (or equivalent) antecedents. This turns out to form a qualitative lottery paradox. As we can't use all of the defaults, which can or should we use?

The answer to this problem is closely related to the question of whether individuals that are typical in every respect really exist. Are there birds that are typical in all respects? Are there houses that are typical in all respects? While this may be an arguable point, most default reasoning systems take a very strong stand on such questions: not only do they exist, but they are the typical individual. Many of these systems break if we try to say that there are no individuals of a certain type that are typical in every respect.

Saying that typical individuals of a particular class do not exist, far from being an exceptional situation, would seem to be the norm for large knowledge bases. Consider the following elaboration of our familiar ornithological example:

Example 2.7 Suppose we want to build a knowledge base about birds. Suppose also that all we are told about Tweety is that Tweety is a bird.

We first state knowledge about the different birds we are considering:

$$\forall x \text{ bird}(x) \equiv \text{emu}(x) \vee \text{penguin}(x) \vee \text{hummingbird}(x) \vee \\ \text{sandpiper}(x) \vee \text{albatross}(x) \vee \dots \vee \text{canary}(x)$$

We now state defaults about birds (e.g., they fly, are within certain size ranges, nest in trees, etc.). For each sort of bird that is exceptional in some way we will be able to conclude that Tweety is not that sort of bird:

- We conclude that Tweety is not an emu or a penguin because they are exceptional in not flying.
- We conclude that Tweety is not a hummingbird as hummingbirds are exceptional in their size (consider making a bird cage for Tweety; we have to make assumptions about the size of birds).
- We conclude that Tweety is not a sandpiper as sandpipers are exceptional in nesting on the ground. This assumption would be made if we are walking our robot in the outdoors and someone says “look at that bird nest”; the robot would have to make assumptions of where to look first.
- We conclude that 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 that Tweety is a canary. If the canary is also exceptional, then in all of the systems considered, we can no longer conclude that Tweety flies (or we conclude it and its negation in different extensions). We have thus lost effectively the use of the default “birds fly”.

The problem is that lots of seemingly irrelevant information (namely about how different sorts of birds are exceptional in different aspects) can interact to make none of the defaults applicable. For seemingly unrelated statements to interact to produce such side effects seems like a very bad problem.

The reason that we divide the class of birds into subclasses is because each subclass is exceptional in some way. Rather than being a pathological

example, this would seem to be the general rule, typical of hierarchies with exceptions.

This problem is analogous to the lottery paradox of Kyburg [17]. In the lottery paradox we have the default that each ticket will not win (as we want to make plans assuming our ticket will not win, and only seriously plan on what to do with the money if we actually win), however we also have the knowledge that one ticket will win. If we conjoin all of the default conclusions, we end up with a contradiction. Most of the default logic systems “solve” this problem by ignoring all of the defaults⁵, rather than the arguably more intuitive idea [17, 18] of not conjoining the conclusions to get a contradiction. Given that the user added the defaults, the system is being very presumptive to ignore the explicit defaults, presumably deciding the user was not rational in adding them (as has been advocated by Shoham [40, p. 392]).

One possible patch⁶ to fix the problem in this example is to disjoin the class “typical birds” to the other sorts of birds. We thus conclude that Tweety is a typical bird and has all of the typical properties. The disjunction, however, is a strange statement, as the “typical bird” is not another sort of bird like emus and sparrows, but rather is an artifact of the representation. This solution ignores the fact that all birds, even typical birds, are some sort of bird (the bird that is not of some type would indeed be exceptional!). It does not allow us to reason by cases as to properties of birds. Also, the resulting knowledge base would not allow us to reason to the identity of a bird by ruling out other cases.

It is also not clear how to expand this “solution” to cases where the exceptions are not as homogeneous as in the previous example, and where the lack of a typical bird is derived not from case analysis, but from, say, physical constraints as in the following example:

Example 2.8 Consider making assumptions about houses (this is done by real-estate sales-people so they can advertise the “features” of a particular house). We need to make an assumption about the size of a house to interpret

⁵Membership in one extension systems also effectively ignore the defaults. For each default there is an extension containing the negation of the conclusion of the default (by assuming the other defaults). Depending on how the single extension is chosen, each default can have its conclusion or the negation of its conclusion in the extension.

⁶This was suggested to me by Matt Ginsberg, May 1989.

statements such as “a large house”, or “a normal sized house”. We would also make assumptions about the number of bedrooms, the number of other rooms and the size of each typical room. The analogous situation to the lottery paradox occurs if we can conclude that the “typical house” does not exist because the sizes do not add up. We would not expect the sizes to add up, as the typically sized house has some room larger than normal (as a selling point).

In this example we derive the non-existence of a typical house, not by case analysis, but rather by physical constraints. One could imagine adding some buffer in the description, but this would entail that the typical house has some space not in any room, which is physically impossible, or as “ghost” rooms that no house in fact has.

In section 4, this qualitative lottery paradox is formalised.

3 Different sorts of “knowing”

In much of AI there is an assumption that there is one sort of “knowing”; one knows something if and only if it is in the knowledge base [20]. In this section I argue that such a definition is inadequate for the intuition behind defaults and for any formalisation of default reasoning that incorporates specificity and standard logical connectives (material implication, in particular).

Example 3.1 highlights some differences between different intuitions behind the statement “all I know”.

Example 3.1 Suppose we have a large knowledge base that includes the defaults

Cats purr.

Mammals live in the wild.

and also includes the fact that cats are mammals, but does not include any directly contradictory knowledge (e.g., that cats don’t purr or that mammals don’t live in the wild).

Suppose we have never heard about “Fred”, and all someone tells us is “Fred is a cat”; there is some notion that “Fred is a cat” is all that we know about that individual.

It seems as though we should be able to conclude

purrs(Fred)

but not necessarily

lives_in_wild(Fred)

as cats may be exceptional with respect to living in the wild. It is presumable that, depending on the other knowledge in the system, the “mammals live in the wild” default could be blocked for cats, but the “cats purr” default could not be blocked for cats without rendering the default useless.

In another sense [20, 21] we “know” other things about Fred:

- We “know” all tautologies mentioning the constant *Fred* are true, for example,

$green(Fred) \wedge (green(Fred) \Rightarrow sick(Fred)) \Rightarrow sick(Fred)$

- We “know” all general knowledge about Fred (i.e., all the knowledge that we know is true for all individuals is true of Fred), for example:

$square(Fred) \Rightarrow rectangle(Fred)$

$cat(Fred) \Rightarrow mammal(Fred)$

- We also “know” inferred knowledge about Fred, for example, we can derive

$mammal(Fred)$

from $\forall x cat(x) \Rightarrow mammal(x)$, as we assumed this is in the knowledge base.

There seems to be two very different forms of “know” here. The first is everything that is in the knowledge base (including *cat(Fred)* and *mammal(Fred)*). The second is that *cat(Fred)* is different to the other sorts of knowledge; as far as the defaults are concerned, all we know about Fred is *cat(Fred)* (so the “cats purr” default should be applicable to Fred), but we know more than *mammal(Fred)* (so the “mammals live in the wild” default is not necessarily applicable to Fred).

This is a very different form of “all I know” than that formalised by Levesque [21].

Example 3.2 It seems as though there is enough information in the defaults:

Birds fly, by default.
Emus don't fly, by default.

and the facts

Emus are birds.
Edna is an emu.

to conclude that Edna does not fly, using the intuition of specificity. The default that Edna should fly because she is a bird should not be applicable as we have more specific information about Edna.

The facts involved are

$$\begin{aligned} \forall x \text{ emu}(x) \Rightarrow \text{bird}(x) \\ \text{emu}(\text{Edna}) \end{aligned}$$

Example 3.3 Suppose we change example 3.2 by swapping the role of *emu* and *bird* in the facts. We end up with the facts:

$$\begin{aligned} \forall x \text{ bird}(x) \Rightarrow \text{emu}(x) \\ \text{bird}(\text{Edna}) \end{aligned}$$

With the defaults as in example 3.2 and these facts we would, by symmetry, want to conclude *flies*(Edna), which is the opposite of the conclusion in example 3.2.

Observation 3.4 If we just consider the instances of the facts that are relevant to Edna, we find something interesting:

$$\text{emu}(\text{Edna}) \wedge (\text{emu}(\text{Edna}) \Rightarrow \text{bird}(\text{Edna}))$$

is logically equivalent to

$$\text{bird}(\text{Edna}) \wedge (\text{bird}(\text{Edna}) \Rightarrow \text{emu}(\text{Edna}))$$

As far as the semantics are concerned, the above formulae are indistinguishable (they have the same truth value in every model; they are true if and only if both $emu(Edna)$ and $bird(Edna)$ are true), but they elicit different answers with respect to specificity.

This observation can be summarised in the claim:

Result 3.5 ⁷ If in a default reasoning system that uses classical logic for the facts, we

1. treat defaults modularly (i.e., their representation does not depend on the facts),
2. consider only the instances of the facts for the individuals under consideration, and
3. treat logically equivalent facts as the same,

then we cannot have specificity.

Proof: If the system incorporates specificity, it gets opposite answers in examples 3.2 and 3.3. However, under the conditions of the result, examples 3.2 and 3.3 are identical, and so cannot elicit different answers. \square

The problem is that there is not enough information in the semantic content of the instances of the facts to handle specificity. Note that the use of cancellation axioms or user-defined priorities violates the modularity of defaults.

If we want to use classical logic for the background facts, we have to be able to exploit some difference between the facts of examples 3.2 and 3.3 to account for the opposite answers. There seems to be two possible answers:

1. The difference between $emu(Edna)$ and $\forall x emu(x) \Rightarrow bird(x)$ is syntactic. We know Edna is both an emu and a bird, but we have to take into account the universally quantified formula, and somehow the fact that all the other emus are also birds is crucial. This has been advocated by Bacchus [2]. He argues that we need to randomise over the name *Edna*, in order to consider just the typical emu.

⁷This is not called a theorem, because it deliberately uses undefined terms. In particular we do not define what it means for a default reasoning system to "have specificity". The one necessary condition is that it concludes that Edna does not fly in example 3.2.

2. There is a difference in kind between the fact $emu(Edna)$ and the fact $\forall x emu(x) \Rightarrow bird(x)$. The latter is always true in the domains under consideration ("background knowledge") and the former only happens to be true ("contingent knowledge").

The following example shows that two syntactically identical formulae can produce opposite answers, thus showing that the distinction is more than just syntactic.

Example 3.6 Suppose we are building an expert system with defaults;

Professors are sedentary.
Runners are not sedentary.

This expert system has nothing to say about professors who are runners. Suppose also that we are providing a facility to ask the user for particular knowledge about the case under consideration [5].

Suppose that we ask the user "who are runners?", and the user replies "all of the professors". The user is thus saying

$$\forall x professor(x) \Rightarrow runner(x)$$

is true about their particular world. Suppose they also tell us that Maria is a professor. Should we conclude that Maria is sedentary? Given that we had nothing to say about professors who are runners, we should not conclude that Maria is sedentary just because all of the professors in the world under consideration happen to be runners.

If, however, we had designed the knowledge base taking into account the fact that all professors are runners, then we should conclude by specificity that Maria is sedentary. In some sense, the default "Professors are sedentary" would have already taken into account the fact that professors are runners.

What is important about this example is that it shows that there is no syntactic distinction between background and contingent knowledge. It is rather a distinction that must be explicit in building the knowledge base.

This distinction between necessary and contingent facts does not reflect differences in the world being represented, but rather differences in the knowledge bases. There is nothing in the domain that prescribes whether "all the professors are runners" is background or contingent. This choice reflects

whether the knowledge base has been constructed taking this fact into account or not.

Background knowledge is about all possible worlds. To say that “all professors are runners” in background knowledge says that there could not be a professor who is not a runner. Contingent knowledge is about the individuals in one particular world. The contingent statement “all professors are runners” is about the individuals in one particular world (the particular world that the user is in, presumably). This distinction is similar to the distinction between propositional and statistical probabilities of Bacchus [2].

In summary, this section has argued for the following claim:

Claim 3.7 We need to distinguish explicitly between background knowledge and contingent knowledge in a default reasoning system.

This distinction is very common; it can be seen in the following:

1. the distinction between the network and markers in marker passing systems such as NETL [9];
2. the difference between the probabilistic knowledge (such as $p(A|B) = 0.345$) and the conditioning knowledge (B in the preceding equation) in probability theory [29] (see example 3.8 below);
3. the difference between background knowledge and observations in abduction [36, 34].
4. the distinction between the general knowledge provided by a knowledge engineer and the particular knowledge provided by a user in a typical expert system architecture [5].

This distinction was first used with respect to nonmonotonic reasoning in Poole’s 1985 paper work on comparing explanations for specificity [32], and more recently in the work of Delgrande [6] and Geffner [11].

This distinction arises very clearly in Bayesian probability theory [28]. There are two ways to say that A is true. The first is to say $p(A) = 1$. The second is to condition on the knowledge A .

Example 3.8 To show how this distinction arises in probability theory consider the following conditional probabilities:

$$\begin{aligned} p(\textit{flies}|\textit{bird}) &= 0.87 \\ p(\textit{flies}|\textit{emu}) &= 0.056 \\ p(\textit{bird}|\textit{emu}) &= 1 \end{aligned}$$

Suppose we want to determine the probability that an individual who is an emu flies. If we want to say that *emu* is true, we cannot write $p(\textit{emu}) = 1$. This statement is logically inconsistent with the conditional probabilities above [28].

We need to condition on *emu*, and ask $p(\textit{flies}|\textit{emu})$.

It is important to realise how important and subtle this distinction may be. For example, Pearl in [29] does not distinguish between background and contingent knowledge, and cannot represent the “birds fly” example where it is a fact that emus are birds. He only distinguishes between facts and defaults. Because he does not have the category of background knowledge, his formalism does not work when the statement “emus are birds” is a fact; he needs to make it a default. In a closely related paper, Geffner and Pearl [12], can solve the “birds fly” example, in part because they distinguish between the “background context”, and the given knowledge on the left hand side of their “ \models ” relation.

This distinction can be compared with the distinction between necessary and contingent knowledge in modal logics [15]. There is one important difference: in typical modal logics of necessity the necessary propositions have a more important status than the contingent ones. If something follows from a contingent proposition it follows from the necessary proposition (as $Lp \rightarrow p$ is an axiom in all modal logics I know of where L is interpreted as “necessary”, as opposed to say, “belief”). In the distinction presented here, the contingent facts have a more important status than the necessary (background) facts; making a fact necessary tends to reduce its impact (as example 3.6 shows). Background facts are passive and can be ignored unless they are needed in the reasoning process; contingent facts demand to be taken into consideration and accounted for.

4 Closure and the Lottery Paradox

In this section, some of the properties outlined in the previous examples are formalised.

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 notation and have different semantics; I intend this discussion to include every notation.

The property that seems to be a minimal property for a default is what I call “conditioning”⁸.

Property 4.1 (Conditioning) A default reasoning system has the **conditioning property** if whenever “ $p(x) \rightarrow q(x)$ ” is a default and the contingent knowledge is “ $p(C)$ ” (where constant C does not appear in the background knowledge base), it concludes “ $q(C)$ ”.

Thus if “ p ’s are q ’s” by default, and all we know contingently about some object is that p is true of it, we should conclude q is true of the object.

For example, suppose a system has the default “birds fly” and all we tell it about some object is that it is a bird. If a system has the conditioning property it concludes that the bird flies. This seems like a minimal property “birds fly” should have.

Note that this is an extremely weak property. If we know anything else about C this property, by itself, does not sanction us to use the default.

This property is the simplest property of many of the recent conditional accounts of default reasoning [29], namely the, seemingly uncontroversial

$$p \vdash_{\Delta} q \text{ if } p \rightarrow q \in \Delta$$

Property 4.2 (Finite Conjunctive Closure) A system has the (finite conjunctive) closure property if it concludes finite conjunctions of its conclusions.

This property says that if a system concludes α and concludes β then it concludes $\alpha \wedge \beta$. The “finite” condition means that we do not demand that the system can prove $\forall x p(x)$ if it can prove $p(C)$ for all C .

The third property is a restriction of minimal representational power:

⁸This discussion is in terms of parametrized (open) defaults as it is most natural for this case. However the argument is purely propositional, and covers propositional systems as well as systems allowing defaults with free variables. A similar notion (without the background–contingent distinction) was called the “one step default property” in [35].

Property 4.3 (Horn representability) The system can at least represent Horn clauses. That is, it can represent implications of the form $a_1 \wedge \dots \wedge a_n \Rightarrow b$ and restrictions of the form $\neg(a_1 \wedge \dots \wedge a_n)$ ⁹.

Property 4.4 (Consistency) Defaults do not introduce inconsistencies. If the facts are consistent, the system doesn't conclude anything at odds with the facts.

Property 4.5 (Arbitrary defaults) Beyond perhaps making a restriction on non-directly contradicted defaults, whether a default is acceptable does not depend on other facts and defaults¹⁰.

It is presumed that the defaults make sense to the person adding them. A system with the arbitrary defaults property allows users to add any defaults they think are appropriate. This property makes no claims as to whether $p(x) \rightarrow q(x)$ should be acceptable as a default if either $p(x) \rightarrow \neg q(x)$ is a default or $\forall x p(x) \Rightarrow \neg q(x)$ follows from the facts (it does not seem reasonable to want to conclude q as well as $\neg q$ if a p is encountered).

The following is a constraint on systems with these properties.

Result 4.6 A default reasoning system cannot have all of the following properties:

- (i) Conditioning.
- (ii) Finite conjunctive closure.
- (iii) Horn representability.
- (iv) Consistency.
- (v) Arbitrary defaults.

⁹Note that we are using Horn clauses in a way different to how Prolog uses them: the negated conjunctions are used as facts rather than as queries.

¹⁰This is set up as a "straw man" in order to consider what constraints on arbitrary defaults are implied by the other conditions.

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

Suppose

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

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

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

is a background fact, and we are given the contingent fact

$$p(C)$$

By (i) we conclude each " $q_i(C)$ ", and by (ii) we conclude their conjunction, which is inconsistent, contravening (iv). \square

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

- (i) Conditioning is given up in circumscription [24], in any minimal model solution [40] and in systems which require membership in all extensions [25]. 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. This means that they cannot guarantee that "birds fly" can be used when all they are told is that something is a bird.

Pollock, in his defeasible reasoning system [31], explicitly gives up this property using his "principle of collective defeat", as he wants the property of finite conjunctive closure. Similarly Gabbay [10] gives up conditioning by his "compatibility of the $>$ rules" property, which is essentially finite conjunctive closure.

- (ii) Finite conjunctive closure is given up in many probability-based systems [27, 1], and in systems which, for prediction, only require membership in one extension [37, 26, 33]. These latter systems get the conditioning property for the wrong reason, namely by being able to conclude a proposition and also conclude its negation (albeit in different extensions).

- (iii) Horn representability is given up in inheritance systems [42, 41]. These allow (i), (ii), (iv) and (v), however they lack the expressiveness of the richer logic-based formalisms.
- (iv) Consistency is given up by thresholding probability as a basis for acceptance [17, 23], but is not given up by any of the default reasoning systems I know of. It is, however, argued [16, 30, 18] that commonsense reasoning does indeed require reasoning under inconsistency.
- (v) Arbitrary defaults is given up in ϵ -semantics [29, 12]. There is no consistent probability assignment for the defaults and facts given in the proof of result 4.6. 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.

Consider now the implication imposed by the other conditions on “arbitrary defaults”. We get into problems when the conjunction of the conclusions of the defaults directly following from some contingent knowledge are inconsistent. Requiring the other four conditions is like imposing the condition that not only does the individual that is normal in every respect exist, but the individual that is normal in every respect is the normal individual.

This section relied on the use of conditioning in formalising our version of the lottery paradox. It is not that the lottery paradox only arises when we use conditioning (example 2.7 was not stated in terms of conditioning), but rather that the use of conditioning helps us understand what is violated in the lottery paradox. Without casting it in terms of background and contingent knowledge it is difficult to give a condition under which it is unreasonable that a default not be applicable (c.f., example 3.1).

5 Related Work

5.1 Lottery Paradox and Default Reasoning

Perlis [30] has also discussed how the lottery paradox can arise in default reasoning. He shows that “omnithinkers” who are Socratic (admit that some of their beliefs are wrong), and recollective (can recall all of their default conclusions) cannot be consistent.

Perlis’s argument is very different to the one presented here. We do not require either of these assumptions. We are not talking about how a set of derived conclusions are in conflict, but rather about how one conclusion cannot be drawn because the defaults are in conflict.

Whereas Perlis’s argument depends on reasoners reasoning in time, and then admitting that they made a mistake, and so being inconsistent, we rely on the argument that a reasoner will not derive even a direct conclusion. Rather than finding that their conclusions lead to a contradiction, the systems under consideration go to extreme lengths to avoid inconsistencies, even to the point of not being able to use their defaults.

Notice also that the argument presented here consists of a “narrow scope” [8] application of defaults. We are only trying to use a default for one individual and not trying to derive conclusions about an entire population (as do typical instances of the lottery paradox [17, 30, 18]). The problem is not that there is a default that is not applicable for multiple individuals, but rather that there is an individual for which a set of multiple defaults is not applicable.

Thus, intuitively, restricting the scope of individuals [8] is not a viable solution to the instance of the paradox presented here. However we can solve instances of the qualitative lottery example by reifying the defaults [8]. This is done when we use the abnormality notation; we can use $\neg ab(Birdsfly, x)$ to make the constant *Birdsfly* denote the default. The lottery paradox of example 2.7 can be avoided for the conclusion about whether Tweety flies by restricting the scope to, say, the constants *Birdsfly* and *Tweety*. The question then arises as to how we knew that we wanted “*Birdsfly*” in the scope. Either it had something to do with a query that we were asking the knowledge base, or it didn’t. If it has something to do with a query, then, presumably we are giving up finite conjunctive closure. Different queries will have different scopes, and so it should not be expected that their

conclusions conjoin. If the scope does not depend on the query, then either all of the contradictory defaults are in the scope, in which case we still have the qualitative lottery paradox, or some are missing in which case there are some conclusions that have direct defaults, but cannot be concluded. This is very much like just ignoring some of the defaults.

5.2 Probabilistic Systems

One of the main features of probabilistic interpretations of defaults [29], whether they are based on probabilities arbitrarily close to one [12], an increase in probability [27] or simple majority [1], is the use of conditioning.

For example, the first, seemingly uncontroversial, axiom of ϵ -semantics [12] is

$$\text{If } (p \rightarrow q) \in \Delta \text{ then } p \models_{\langle L, \Delta \rangle} q$$

This corresponds to the conditioning property presented in section 4, but is, however, slightly stronger.

Example 5.1 Suppose, as background knowledge, we include

$$bird(Big_bird) \wedge \neg flies(Big_bird)$$

and have the default $bird(x) \rightarrow flies(x)$ (in [12] this means the set of all its ground instances is in Δ). The conditioning property does not sanction us to conclude anything given $bird(Big_bird)$. However the above axiom of ϵ -semantics lets us conclude $flies(Big_bird)$ from $bird(Big_bird)$.

The above fact and default is, in fact, ϵ -inconsistent [29]. Geffner and Pearl [12] do not allow explicit exceptions as part of the background knowledge in this way. Semantically, the reason is that the default implies the conditional probability $p(flies(Big_bird) | bird(Big_bird))$ is close to 1 and the background facts about *Big_bird* imply that it is zero. Syntactically it is because the conditioning rule allows the conclusion of $flies(Big_bird)$ from $bird(Big_bird)$, and a deduction axiom (we conclude what logically follows from the facts) lets us conclude $\neg flies(Big_bird)$.

To represent the fact that Big Bird is an exceptional bird Geffner and Pearl would make a predicate $big_bird(x)$ that is true if x is Big Bird, and write

$$big_bird(x) \Rightarrow bird(x) \wedge \neg flies(x)$$

as a background fact. Rather than using a constant to denote the individual, we need to use a predicate to say that a particular constant denotes Big Bird. Thus although the conditioning in this paper and the conditioning in ϵ -semantics are different, in practice this difference would not be encountered.

6 Conclusion

The unifying theme between the specificity and the lottery paradox problems, is to consider what happens when there are always multiple extensions when the antecedent of a default is true. That is, to consider the case when $p(x) \rightarrow q(x)$ is a default such that whenever $p(C)$ is true for any C , there are competing defaults that do not allow the conclusion of $q(C)$ (or also allow the conclusion $\neg q(C)$). If we do not handle the multiple extensions appropriately, the default becomes useless. We do not like defaults that can never be used; if a user didn't want a default to be used they would not have added the default in the first place.

It was argued that solutions to the multiple extension problem that rely on "knowing" exceptions do not work. It was shown why we need to distinguish between "background" and "contingent" knowledge, and why we should not expect to have conjunctive closure of our default conclusions.

The qualitative lottery paradox was discovered using our Theorist system [33] on (pseudo-real) problems. It was surprising to me to find out how naturally it arises in practice, and how difficult it is to get specificity working satisfactorily beyond trivial examples. There is much work that remains to be done on what problems arise in practice. We don't want to be like the intellectuals in Galileo's time [39, p.520], and mistakenly think we know what the phenomena is that we are trying to formalise. We need to look at real representational problems and build more experiments to determine what these things we call defaults really are.

Acknowledgements

Thanks to all of the people at the 2nd International Workshop on Nonmonotonic Reasoning, and the International Conference on Representation and Reasoning who were involved in sharpening my ideas on this problem. Spe-

cial thanks to Hector Geffner, Ray Reiter and the anonymous reviewers. This research was supported under NSERC grant OGPOO44121.

References

- [1] F. Bacchus, "A Modest, but Semantically Well Founded Inheritance Reasoner", *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence* 1989.
- [2] F. Bacchus, *Representing and Reasoning with Probabilistic Knowledge*, M.I.T. Press, 1990.
- [3] G. Brewka, "Preferred subtheories: an extended logical framework for default reasoning", *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, Detroit, pp. 1043-1048, 1989.
- [4] G. Brewka, *Nonmonotonic Reasoning: From Theoretical Foundation Towards Efficient Computation*, Ph.D. thesis, University of Hamburg, 1989.
- [5] B. G. Buchanan and E. H. Shortliffe, *Rule-based expert systems*, Addison-Wesley, Reading, MA 1984.
- [6] J. Delgrande, "An approach to Default Reasoning Based on a First-Order Conditional Logic: Revised Report", *Artificial Intelligence*, Vol. 36, No. 1, pp. 63-90, 1988.
- [7] D. Etherington, *Reasoning from incomplete information*, Research notes in AI, Pitman, London, 1987.
- [8] D. Etherington, S. Kraus and D. Perlis, "Nonmonotonicity and the Scope of Reasoning", to appear *Proc. AAAI-90*, Boston 1990.
- [9] S. Fahlman, *NETL: A system for representing and using real-world knowledge*, MIT Press, Cambridge, MA, 1979.
- [10] D. M. Gabbay, "Theoretical Foundations for non-monotonic reasoning systems", in K. R. Apt, *Logics and Models of Concurrent Systems*, Springer-Verlag Berlin Heidelberg, pp. 439-457, 1985.

- [11] H. Geffner, "On the logic of Defaults", *Proc. AAAI-88*, pp. 449-454, Minneapolis, 1988.
- [12] H. Geffner and J. Pearl, "A framework for reasoning with defaults", to appear in *Defeasible Reasoning and Knowledge Representation*, Kluwer Publishers, 1989.
- [13] H. Geffner, "Conditional entailment: closing the gap between defaults and conditionals", *Proc. 3rd Workshop on Nonmonotonic Reasoning*, Lake Tahoe, May 1990.
- [14] S. Hanks and D. McDermott, "Default Reasoning, Nonmonotonic Logics and the Frame Problem", *Proc. AAAI-86*, pp. 328-333, 1986.
- [15] G. E. Hughes and M. J. Cresswell, *An Introduction to Modal Logic*, Methuen, London, 1968.
- [16] D. J. Israel, "Whats wrong with non-monotonic logic", *Proc. AAAI-80*, pp. 99-101, 1980.
- [17] Henry E. Kyburg, Jr., *Probability and the Logic of Rational Belief*, Wesleyan University Press, Middletown, 1961.
- [18] Henry E. Kyburg, Jr., "Probabilistic Inference and Non-monotonic Inference", *Proc. Fourth Workshop on Uncertainty in Artificial Intelligence*, pp. 221-228, 1988.
- [19] D. Lehmann, "What does a conditional knowledge base entail?", *Proceedings of the First International Conference on Principles of Knowledge Representation and Reasoning (KR-89)*, pages 212-222, Toronto 1989.
- [20] H. J. Levesque, "Foundations of a functional approach to knowledge representation", *Artificial Intelligence*, 24, 155-212, 1984.
- [21] H. J. Levesque, "All I know: a study in autoepistemic logic", *Artificial Intelligence*, 43 (2-3), 263-309, 1990.
- [22] R. P. Loui, "Defeat amongst arguments: a system of defeasible inference" *Computational Intelligence* 3 (2), 100-106, May 1987.

- [23] R. P. Loui, "Real Rules of Inference: Acceptance and Non-monotonicity in AI", Technical Report TR-191, University of Rochester, 1986.
- [24] J. McCarthy, "Applications of Circumscription to formalising common-sense knowledge", *Artificial Intelligence*, Vol. 28, pp. 89-116, 1986.
- [25] D. V. McDermott and J. Doyle, "Non-monotonic logic 1", *Artificial Intelligence*, Vol. 13, pp. 41-72, 1980.
- [26] R. C. Moore, "Semantical Considerations on nonmonotonic logic", *Artificial Intelligence* 25 (1) 75-94, 1985.
- [27] E. Neufeld and D. Poole, "Probabilistic Semantics and Defaults", *Proc. Fourth Workshop on Uncertainty in Artificial Intelligence*, pp. 275-282, 1988.
- [28] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufmann, 1988.
- [29] J. Pearl, "Probabilistic Semantics for Nonmonotonic Reasoning: A Survey" *Proceedings of the First International Conference on Principles of Knowledge Representation and Reasoning*, Toronto, May 1989, pp. 505-516.
- [30] D. Perlis, "On the Consistency of commonsense reasoning", *Computational Intelligence*, Vol. 2, No. 4, pp. 180-190, 1987.
- [31] J. L. Pollock, "Defeasible Reasoning", *Cognitive Science*, Vol. 11, pp. 481-518, 1987.
- [32] D. Poole, "On the Comparison of Theories: Preferring the Most Specific Explanation", *Proc. IJCAI-85*, pp. 144-147, 1985.
- [33] D. Poole, "A logical framework for default reasoning", *Artificial Intelligence*, Vol 36, pp. 27-47, 1988.
- [34] D. Poole, "Explanation and Prediction: An Architecture for Default and Abductive Reasoning", *Computational Intelligence* Vol. 5, No. 2, pp. 97-110, May 1989.

- [35] D. Poole, "What the Lottery Paradox tells us about default reasoning", *Proceedings of the First International Conference on Principles of Knowledge Representation and Reasoning*, Toronto, May 1989, pp. 333-340.
- [36] H. E. Popl, Jr., "On the mechanization of abductive logic", *Proc. IJCAI-73*, 147-152, 1973.
- [37] R. Reiter, "A Logic for Default Reasoning", *Artificial Intelligence*, Vol. 33, pp. 81-132, 1980.
- [38] R. Reiter and G. Criscuolo, "On Interacting Defaults", *Proc. Seventh International Joint Conference on Artificial Intelligence*, pp. 270-276, 1981.
- [39] B. Russell, *A History of Western Philosophy*, Counterpoint, London 1984 (first published 1945).
- [40] Y. Shoham, "Nonmonotonic Logics: Meaning and Utility", *Proc. IJCAI-87*, pp. 388-393, 1987.
- [41] 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, 1988.
- [42] D. S. Touretzky, *The Mathematics of Inheritance Theory*, Pitman / Morgan Kaufmann, 1986.