An Introduction to Formal Logic

Richard L. Epstein

Advanced Reasoning Forum

ARF

Advanced Reasoning Forum
P.O. Box 635
Socorro, NM 87801
www.ARFbooks.org

ISBN 978-1-938421-27-3 paperback

ISBN 978-1-938421-28-0 e-book

An Introduction to Formal Logic

Preface to the Instructor

Logic is a tool we use to investigate how our language connects to the world in order to reason better and, we hope, to understand the world better. This book tells that story. It has a beginning, a middle, and an end.

We begin by setting out what formal logic is: the study of inferences for validity based on their form. Classical propositional logic is then presented as the simplest formal logic. In the development of that logic, the most important tools of formal logic are presented: a formal language, realizations, models, formal semantic consequence, and an axiom system. Examples of formalizing ordinary-language propositions and inferences show how to use classical propositional logic and also show some of its limitations.

The middle of the book relates form to the world. Predicate logic is motivated as a way to widen the scope of classical propositional logic to investigate more kinds of inferences whose validity depends on form. The large assumption that the world is made up of individual things is the basis for both its syntax and semantics. Besides form and the truth or falsity of atomic propositions, only the idea of assigning reference to terms as a kind of naming is needed. Many examples of formalizing show both the scope and limitations of classical predicate logic. Those depend on establishing criteria for what counts as a good formalization. The emphasis in those discussions is how the assumption that the world is made up of individual things is both useful and limiting.

In the last part of the book, second-order predicate logic is developed as an example of how metaphysics shapes formal logic. The final chapter reflects on the success of the formal methods we've developed as means to reason to truths. Our formal logics circumscribe what we mean by "individual thing", namely, what can be reasoned about in predicate logic. Our formal logics give us a way to be precise about how we understand possibilities, though only relative to the assumptions we make about form and meaning and what there is in the world.

This story gives the basics, the fundamentals of formal logic. Along the way I point out how the work here can be extended and modified to apply to a wider scope of what we can formalize from ordinary-language reasoning. The story is not finished. Not here, not elsewhere.

* * * * * * *

Some Points about the Organization and Content

- Appendices

  The appendices contain material that is either more technical than many students want, or is too philosophical for many students, or is supplementary to the main line of the story.

- The Formal Language

  In most logic texts, the definition of a formal language for predicate logic allows for superfluous quantifications. The rationale for including such formulas is that it simplifies the definition of the formal language, allowing a definition of bound and free variables to be made later. But the disadvantage is that a formula such as "$\forall x$ (Ralph is a dog)" that would correspond to the nonsensical "For everything, Ralph is a dog" is deemed acceptable. The semantics for superfluous quantifiers treat that formula as equivalent to "Ralph is a dog", which can be true. That is not consonant with our normally treating nonsense as false in our reasoning, as you can see in "Truth and Reasoning" in my *Reasoning and Formal Logic*. The advantages of not allowing superfluous quantification, beyond ridding our semi-formal languages of nonsense, are significant: we need no axiom schemes for superfluous quantification, and many proofs about the language are simplified by no longer having to treat cases of superfluous quantification separately.

- Proof theory
  Hilbert-style axiomatizations of classical propositional logic, classical predicate logic, and classical predicate logic with equality are presented in the text, and their completeness proofs appear in an appendix. Natural deduction and other methods of proof can be left until the skills are needed.

- Functions
  Only in some parts of mathematics are functions used, and even in much of mathematics partial functions are essential. To extend classical predicate logic to allow for formalizing reasoning that involves partial functions , we would need to analyze how to reason with non-referring names, descriptive names, and descriptive functions. That's too much. I do that in *The Internal Structure of Predicates and Names*.

- The History of Logic
  It's a big subject. It's complicated. And it's not illuminating at this level. I present a lot in my books *Propositional Logics*, *Predicate Logic*, *Computability*, and *Classical Mathematical Logic*.

- English and Formal Logic
  Some might object that the development of formal logic here is too closely tied to motivations and examples from English. But, perhaps in slightly modified form, the examples and motivation here will apply to reasoning in many other languages. If they do not serve, then that would be evidence that the notion of thing is not as deeply embedded in the other language.

7 Formalizing in Classical Predicate Logic

A. Rewriting English Sentences

We've been very stingy with the grammatical tools we allow ourselves in analyzing propositions. For instance, where is the predicate and where the name in "All dogs bark"? Where is the name in "A dog is a dog"?

In this chapter we'll see that following the assumptions we used in devising classical predicate logic we can establish conventions for how to rewrite these examples and a wide variety of other propositions in such a way that the new English sentence is equivalent to the original for all our logical purposes and can be easily formalized.

But why should we be concerned with English sentences anymore? After all, it's the vagueness and imprecision of English that drove us to create a formal language. Why not jettison reasoning in English completely and use only semi-formal languages? If we do, then we're committed to reasoning only under the assumptions of classical predicate logic, and we will not be able to see the limitations of those or compare them to other ways of seeing the world.

I can't give rules for rewriting every proposition we'll want to reason with, and often there's more than one way to rewrite a sentence leading to equally reasonable choices for a formalization. A tradition has developed of many explicit and implicit agreements and assumptions. What I'll do in this and the next two chapters is set out explicitly many of the major conventions about formalizing sentences in the current tradition of predicate logic. In this chapter I'll be concerned mostly with the basic division of a proposition into names and predicate.

To begin, we have the criteria of formalization that we adopted with classical propositional logic (p. 29). The first, as modified for predicate logic, is:

- The formalization respects the assumptions that govern our choice of syncategorematic vocabulary and definition of truth in a model. The constraints we work under when we adopt classical predicate logic must be observed.

The other two are:

- If a proposition is informally true due to its form or false due to its form, then its formalization should be a tautology or anti-tautology, respectively.
- If one proposition follows informally from another or a collection of other propositions, then its formalization should be a formal semantic consequence of the formalizations of the other(s).

When I say in what follows that a proposition is true or is false, I mean with respect to an implicit model in which predications are assigned their "obvious" truth-values. When I use x, y, and z, they're meant to stand for x_1, x_2, and x_3.

B. Common Nouns as Subject or Object

Relative quantification with "all"

How can we formalize the following valid inference?

(1) (a) All dogs bark.
(b) Ralph is a dog.
Therefore, (c) Ralph barks.

Propositions (1b) and (1c) are clearly atomic. How shall we formalize the first premise? The grammatical subject of (1a) is "dogs". That's a common noun, not a name of an individual or a pronoun that we can construe as a place-holder. There's no one thing in the world called "dogs" that barks. We do not mean by (1a) that the collection of all dogs barks but that each individual dog does. So let's rewrite (1a) as:

(2) Every dog barks.

Perhaps we should take a realization with universe all dogs and formalize this as:

$\forall x$ (x barks)

But then we'd have no way to formalize (1b) to show the validity of the inference.

In a realization in which there are objects other than dogs in the universe, the quantification in (2) is restricted. We're not asserting that all things bark but only each dog. So let's rewrite (2) as:

Everything which is a dog barks.

If we try to read this taking "thing" in "Everything" and "which" as pronouns, we get "$\forall x$ (x is a dog barks)", which is nonsense. We need to supply a clearer subject for "barks". It's not the collection of dogs (in this new guise) that barks, but a dog, any dog. We rewrite this last version as:

(3) Everything which is a dog is a thing which barks.

We have two predicates "— is a dog" and "— barks" with a quantification that governs both. A formalization of (3), then, should have the form:

$\forall x$ (x is a dog $*$ x barks)

Here $*$ should link these two open sentences in a way faithful to "is" in (3). Our only resources for $*$ are the propositional connectives $\rightarrow$, $\wedge$, $\vee$, or some combination of those and $\neg$.

We don't want $\vee$ because "$\forall x$ (x is a dog $\vee$ x barks)" would be false if the universe contained a mute cat, whereas (1a) could be true. For the same reason, "$\forall x$ (x is a dog $\wedge$ x barks)" can't be right. So consider:

(4) $\forall x$ (x is a dog $\rightarrow$ x barks)

This can be understood as formalizing:

For each thing, if it is a dog then it barks.

Is that an appropriate reading of (3)? In the classical interpretation of $\rightarrow$, things in the universe that are not dogs are "irrelevant" to the truth-value of (4): if x does not refer to a dog, then "x is a dog $\rightarrow$ x barks" is vacuously true. This is just what we want, and indeed it was in part to be able to model propositions such as (3) that we chose this interpretation of $\rightarrow$ in classical logic. So we formalize (1) as:

$\forall x$ (x is a dog $\rightarrow$ x barks)
Ralph is a dog.
Therefore, Ralph barks.

This is valid in classical predicate logic, an example of universal instantiation.

Similarly, it seems we should formalize

(5) All cows are white.

as

(6) $\forall x$ (x is a cow $\rightarrow$ x is white)

But suppose the universe of our realization contains no cows. Then (6) is true. Would (5) be true, too? Some say no. They'd formalize (6) as:

(7) $\forall x$ (x is a cow $\rightarrow$ x is white) $\wedge$ $\exists x$ (x is a cow)

Most logicians do not want to ascribe an existential assumption to a proposition unless compelled to do so by the role of the proposition in inferences, and they do not consider “All cows are white, therefore there is a cow” compelling. The general convention now is to use (6) as a formalization of (5), adopting (7) only when the context demands it. This is only a convention. We’re not asserting anything about the deep meaning of “all” in English. Nor would a statistical survey of speakers of English be relevant in deciding which convention to adopt, for if there is a disagreement, we need only test which of the formalizations best respects the assumptions of the context in which the proposition is used.

Relative quantification with “some”

Consider:

(8) Some cats are feral.

Rewriting this in the manner we used for (1a), we get:

Some thing which is a cat is a thing which is feral.

But we should not formalize this as:

$\exists x$ (x is a cat $\rightarrow$ x is feral)

That would be true in a universe that contained only domestic cats and a dog, for any dog (or wombat, or stick) makes the conditional “x is a cat $\rightarrow$ x is feral” true if taken as reference for x, while (8) would be false. Rather, we want:

(9) $\exists x$ (x is a cat $\wedge$ x is feral)

Some say that because of the plural “cats” in (8), the proposition is true only if there is more than one cat that is feral. We’ll see how to formalize that reading in Chapter 9. The tradition now is to formalize (8) as (9) unless context suggests otherwise. Certainly (9) is a good formalization of both “There is a cat that is feral” and “There exists a feral cat”. That we use the same formalization for each of these and (8) leads some to say that these sentences all “express the same proposition”, but what that proposition is I cannot say.

Consider now:

(10) A cat is mewing.

Do we mean by this that some one cat is mewing or every cat? Surely just one will do, taking (10) as elliptical for “There is a cat that is mewing”. So on the pattern established above, we can formalize the proposition as:

$\exists x$ (x is a cat $\wedge$ x is mewing)

The article “a” should be understood as “there is a”.

But we also have:

(11) A dog is a dog.

This is surely true due to its form only, for we don’t need to know anything about dogs to evaluate it. So it should be formalized as a tautology. If we convert the subject noun "dog" into a predicate and supply an indefinite subject, we get in analogy with our earlier rewritings:

A thing which is a dog is a thing which is a dog.

But one thing? No, any thing. We formalize (11) as:

$\forall x$ (x is a dog $\rightarrow$ x is a dog)

This is a tautology. Despite the similarity of "a cat is" and "a dog is", we formalize those phrases differently by reflecting on when the propositions in which they appear would be true. Though we certainly want to formalize propositions of "the same form" in the same way, that constraint must be secondary to ensuring that we have a formalization that will have the same truth-value as the original with respect to any universe.

C. Common Nouns and Multiple Quantifiers

Consider:

Some dogs bark and every cat meows.

We formalize this as:

$\exists x$ (x is a dog $\wedge$ x barks) $\wedge$ $\forall y$ (y is a cat $\rightarrow$ y meows)

We respect the grammar of the original as made up of two propositions, and each of those we know how to formalize. This illustrates a general convention.

We formalize parts of propositions in accord with our established agreements whenever possible, then formalize the way those parts are put together in the original.

But the situation is more complicated with relations. Consider:

(12) Every dog hates every cat.

We have a relation "x hates y" and ask which x and which y. As with (1a), it's not all objects x, but only those that are dogs that are relevant. And it's not all objects y, but only those that are cats. We need both these restrictions as antecedents, and since no order is intended or needed for the pair of universal quantifiers, we can put them both at the start:

(13) $\forall x\, \forall y$ ((x is a dog $\wedge$ y is a cat) $\rightarrow$ x hates y)

The formalization is good: both (12) and (13) are true or both false for any universe, and the categorematic words appearing in (13) are exactly those of (12).

Alternatively, we could follow the earlier examples by taking care of one quantifier at a time:

(14) $\forall x$ (x is a dog $\rightarrow$ $\forall y$ (y is a cat $\rightarrow$ x hates y)

This is semantically equivalent to (13) and puts the formalization of the "any" for cats later in the proposition as in the original. There is no principled choice to be made here. But if we are to have some regularity in formalizing, some rules that we can use as a guide rather than improvising with each new example, we need to choose one for our standard method. Before we do that, let's look at other pairs of quantifiers.

Consider:

(15) Some dog hates some cat.

We have two ways to formalize this in accord with what we did with a single existential quantifier and in analogy with the choices for a pair of universal quantifiers. We can put all the quantifiers at the start, or add them as they apply:

(16) $\exists x\, \exists y\, ((x \text{ is a dog} \land y \text{ is a cat}) \land x \text{ hates } y)$

(17) $\exists x\, (x \text{ is a dog} \land \exists y\, (y \text{ is a cat} \land x \text{ hates } y))$

These are equivalent, and both would be true in a model iff for that universe (15) is true.

Now consider:

(18) Every dog hates some cat.

In accord with what we did with a single existential quantifier and a single universal quantifier, we now have three choices:

(19) $\forall x\, \exists y\, ((x \text{ is a dog} \rightarrow (y \text{ is a cat} \land x \text{ hates } y)$

(20) $\forall x\, \exists y\, ((x \text{ is a dog} \land y \text{ is a cat}) \rightarrow x \text{ hates } y)$

(21) $\forall x\, (x \text{ is a dog} \rightarrow \exists y\, (y \text{ is a cat} \land x \text{ hates } y))$

These are semantically equivalent, as you can show, and each is true in a model iff for that universe (18) is true.

Consider also:

(22) Some dog hates all cats.

Again we have three semantically equivalent choices for formalizing this in accord with what we've seen so far:

(23) $\exists x\, \forall y\, ((x \text{ is a dog} \land y \text{ is a cat}) \rightarrow x \text{ hates } y)$

(24) $\exists x\, \forall y\, (x \text{ is a dog} \land (y \text{ is a cat} \rightarrow x \text{ hates } y)$

(25) $\exists x\, (x \text{ is a dog} \land \forall y\, (y \text{ is a cat} \rightarrow x \text{ hates } y))$

As a general method, putting all the quantifiers at the start, giving the quantifier for each conversion maximal scope, still leaves us choices: (19) or (20), and (23) or (24). There is even less to motivate a choice between those. However, making the conversions in order, one at a time, with the quantifiers for the conversions having minimal scope, does determine the formalizations: (21) and (25). It gives us a clear inductive method. So let's opt for the latter method. Thus, we take (14) as "the" formalization of (12), and (17) as the formalization of (15), and (21) as the formalization of (18), and (25) as the formalization of (22).

Summarizing, we have the following convention:

> *To formalize a proposition containing a common noun (phrase) used as a subject or object, we convert the noun (phrase) into a predicate and supply some form of quantification.*
>
> *Universal quantification over just those objects covered by a common noun is modeled by taking the converted predicate as (possibly one conjunct of) the antecedent of a conditional that is universally quantified. Existential quantification is formalized by taking the predicate as a conjunct.*
>
> *When there is more than one common noun in the proposition, we apply these rules sequentially, using minimal scope for the quantifications.*

For the cases we've seen, this method can be expressed rather succinctly.

> *Conventions on converting nouns into predicates* Let γ (gamma) and δ (delta) stand for common or collective nouns, and $\gamma(x)$, $\delta(x)$ stand for the conversions of those into predicates. We formalize:
>
> All γ are δ.
> $\forall x\,(\gamma(x) \rightarrow \delta(x))$
>
> Some γ are δ.
> $\exists x\,(\gamma(x) \wedge \delta(x))$
>
> All γ are P to all δ.
> $\forall x\,(\gamma(x) \rightarrow \forall y\,(\delta(y) \rightarrow \mathrm{P}(x, y)))$
>
> Some γ is P to some δ.
> $\exists x\,(\gamma(x) \wedge \exists y\,(\delta(y) \wedge \mathrm{P}(x, y)))$
>
> All γ is P to some δ.
> $\forall x\,(\gamma(x) \rightarrow \exists y\,(\delta(y) \wedge \mathrm{P}(x, y)))$
>
> Some γ is P to all δ.
> $\exists x\,(\gamma(x) \wedge \forall y\,(\delta(y) \rightarrow \mathrm{P}(x, y)))$

With these conventions we can formalize propositions with any number of common nouns, converting them into predicates in succession

Exercises

1. Why formalize ordinary language propositions instead of reasoning entirely in semi-formal languages?

2. Why require that the proposition we're formalizing and the formalization have exactly the same categorematic words?

3. What does our choice of one of the following depend on as a formalization of "All donkeys bray"?

   $\forall x$ (x is a donkey $\rightarrow$ x barks)

   $\forall x$ (x is a donkey $\rightarrow$ x brays) $\wedge$ $\exists x$ (x is a donkey)

4. Show that the following are equivalent. (Hint: See Chapter 5.K.)

   $\forall x \exists y$ (x is a dog $\rightarrow$ (y is a cat $\wedge$ x hates y))

   $\forall x$ (x is a dog $\rightarrow$ $\exists y$ (y is a cat $\wedge$ x hates y))

5. Show that the following are equivalent:

   $\exists x \forall y$ (x is a boy $\wedge$ (y is a dog $\rightarrow$ x loves y))

   $\exists x$ (x is a boy $\wedge$ $\forall y$ (y is a dog $\rightarrow$ x loves y))

6. Show that the following are equivalent:

   $\forall x \exists y$ ((x is a dog $\rightarrow$ (y is a cat $\wedge$ x hates y)

   $\forall x \exists y$ ((x is a dog $\wedge$ y is a cat) $\rightarrow$ x hates y)

   $\forall x$ (x is a dog $\rightarrow$ $\exists y$ (y is a cat $\wedge$ x hates y))

7. Show that the following are equivalent:

   $\exists x \forall y$ ((x is a dog $\wedge$ y is a cat) $\rightarrow$ x hates y)

   $\exists x \forall y$ (x is a dog $\wedge$ (y is a cat $\rightarrow$ x hates y)

   $\exists x$ (x is a dog $\wedge$ $\forall y$ (y is a cat $\rightarrow$ x hates y))

8. Explain why we should or should not formalize all of the following the same:

   Dogs hate cats.
   All dogs hate all cats.
   Every dog hates any cat.
   Any dog hates any cat.
   All dogs hate every cat.

9. Explain why we should or should not formalize all of the following the same:

   Some cat is smarter than some dog.
   Some cats are smarter than some dogs.
   There's a cat that's smarter than some dog.
   There are a cat and a dog such that the cat is smarter than the dog.
   There is at least one cat that's smarter than at least one dog.
   There exist a cat and a dog with the cat smarter than the dog.
   There's a cat that's smarter than a dog.

10. Explain why we should or should not formalize all of the following the same:

    Every dog hates some cat.
    For any dog there's a cat that it hates.
    Any dog hates at least one cat.
    All dogs hate some cats.

D. Adjectives

The following inference is not valid:

> Bon Bon is a small donkey.
> Therefore, Bon Bon is small.

The premise is true and the conclusion is false: my donkey Bon Bon is huge compared to a mouse or even a cocker spaniel. The adjective "small" is a *relative adjective*, that is, it can be evaluated only relative to a kind of thing: a small elephant, a small mouse, a small skyscraper. There is no absolute idea of what's small. So we can't detach the adjective from the noun to formalize "Bon Bon is a small donkey" as "Bon Bon is small $\wedge$ Bon Bon is a donkey". We have to take "Bon Bon is a small donkey" as atomic. But then we can't respect that the following inference is valid:

(26) Bon Bon is a small donkey.
Therefore, Bon Bon is a donkey.

Perhaps we could add an *ad hoc* assumption:

> $\forall x$ (x is a small donkey $\rightarrow$ x is a donkey)

We'd formalize (26) only relative to that, as if it were an axiom. But these inferences are valid, too:

> Bidú is a big dog.
> Therefore, Bidú is big.

> Birta is an old dog.
> Therefore, Birta is a dog.

They have the same form as (26), but we could formalize them only by adding axioms governing the use of "small", "big", and "old" with every common noun. That's just to say that these inferences are outside the scope of classical predicate logic, which is supposed to codify forms of valid inferences. We can extend classical predicate logic to formalize these by taking account of the internal structure of atomic predicates viewing, for example, "— is a small donkey" as a simple predicate plus a modifier of it that: "(— is a donkey) / small"; but that's another story you can see in my *The Internal Structure of Predicates and Names*.

In contrast, it seems that we can formalize:

> Bon Bon is a beautiful donkey.
> Therefore, Bon Bon is a donkey.

as

> Bon Bon is beautiful $\wedge$ Bon Bon is a donkey.
> Therefore, Bon Bon is a donkey.

That's because "beautiful" is an *absolute* adjective. It doesn't matter what kind of thing we call beautiful: the same standard is used for all things. If you don't agree that "beautiful" is absolute, try "sweet tasting" or "black". If these or any other adjective is absolute, then we can formalize uses of it in predicate logic.

Uses of adjectival clauses are more amenable to formalizing. Consider:

Every dog which yelps also barks.

Suppose we rewrite this along the lines of our earlier examples:

(27) Every thing which is a dog which yelps is also one that barks.

By deleting "also" we get what I take to be an equivalent proposition, so we can ignore that word. We don't want to take "x is a dog which yelps" as atomic because then we can't respect the validity of "Ralph is a dog which yelps, therefore Ralph is a dog". We can formalize (27) as:

$\forall x$ ((x is a dog $\wedge$ x yelps) $\rightarrow$ x barks)

The adjectival clause is added as a conjunct to the predicate that formalizes the common noun (phrase). That conjunct need not be atomic:

Every dog which yelps or whimpers also barks.

I'll let you provide the justification for formalizing this as:

$\forall x$ ([x is a dog $\wedge$ (x yelps $\vee$ x whimpers)] $\rightarrow$ x barks))

Nor are we restricted to only propositional connectives in formalizing an adjectival clause. Consider:

Every person who owns a dog is happy.

We can formalize this as:

$\forall x$ ([x is a person $\wedge$ $\exists y$ (x owns y $\wedge$ y is a dog)] $\rightarrow$ x is happy)

I'll let you formulate a general convention for formalizing adjectival clauses.

E. Adverbs

Can we formalize the following valid inference?

(28) Juney is barking loudly.
Therefore, Juney is barking.

We can't detach "loudly" from "barking" because the adverb "loudly" describes the barking, not Juney: "Juney is barking $\wedge$ Juney is loudly" is nonsense. Adverbs modify verbs, and whatever verbs correspond to in the world, it's not individual things. In classical predicate logic, the premise of (28) has to be formalized as atomic. To respect the validity of (28) we'd have to add an axiom "Juney is barking loudly $\rightarrow$ Juney is barking". We'd also have to add axioms to ensure the validity of the following:

Dick is eating quickly.
Therefore, Dick is eating.

Tom ran fast.
Therefore, Tom ran.

Yet it seems clear that all of these are valid due to their form.

We can formalize inferences that involve adverbs in the same way we can formalize inferences that involve relative adjectives by reading the atomic predicate "— is barking loudly" as a simple predicate and a modifier, as in "(— is barking) / loudly". That story, too, you can read in my *The Internal Structure of Predicates and Names*

F. Mass Terms

The following inference is valid:

(29) (a) Snow is white.
(b) All that's white is not black.
Therefore, (c) Snow is not black.

Can we formalize it in predicate logic?

In (29a) and (29c) we seem to be using "Snow" as a name. But what thing does it name? What individual is there that is snow? Snow is physical, but it's not in one place. Yes, snow is distinct from all else in the world, but we can't point to it, even in theory, and say that is snow. We can only use it to describe lots of masses or things that are made of snow. I can point and say, "That ball is made of snow". I can point to a tree when walking and say, "The stuff on that tree is snow". But I can't point to snow complete, all snow.

Don't we have the same problem with "Dogs"? Consider:

(30) Dogs bark.
All that barks is not a cat.
Therefore, Dogs are not cats.

We recognize this as valid, and "Dogs" is used as a subject much like a name. We can't point to dogs complete at one time. I can point to this dog or that dog, but not all dogs even in theory. Rather, "Dogs" describes lots of individual things we call "dogs". So we reformulate "Dogs" as a predicate and formalize (30) as:

(31) $\forall x$ (x is a dog $\rightarrow$ x barks)
$\forall x$ (x barks $\rightarrow$ $\neg$ (x is a cat))
Therefore, $\forall x$ ((x is a dog $\rightarrow$ $\neg$ (x is a cat))

This is valid. And in any model with universe that contains all dogs and cats and in which "dog", "cat", and "barks" are given their usual interpretation, (31) reflects what we intended in (30).

Why can't we do the same with snow? We might try formalizing (29) as:

$\forall x$ (x is snow $\rightarrow$ x is white)
$\forall x$ (x is white $\rightarrow$ $\neg$ (x is black))
Therefore, $\forall x$ (x is snow $\rightarrow$ $\neg$ (x is black))

But what in a model could be used as a reference to make "x is snow" true? What could we put into the universe of a model that is all snow? We would have to include all quantities of snow, whether a snowball, or a snowman, or what's lying on the branches of a tree. We would also have to include all parts of those, for every part of snow is snow. I can pick up a handful of snow and there are innumerable "quantities" of snow in it, as there are innumerable "quantities" of water in a glass of water: the water up to one-third of the way from the bottom of the glass, the water up to .01726825623 of the way from the bottom of the glass, ... Only with a description picking out a particular quantity of snow can we specify a thing that is snow because snow, unlike dogs, does not come in identifiable quantities, it does not come in distinct parts each of which is snow. Snow is a mass, not a collection of things. We have no conception of how we could pick out in a general way just any item that "snow" describes. So we don't say "All snow is white" but "Snow is white".

What holds for "snow" holds equally for other *mass terms* that we use to describe masses: "gold", "mud", "water", ... Every part of a mass is a mass of the same kind. In contrast, no part of a thing is a thing of the same kind. No part of my dog Birta is a dog.

At best, we can formalize propositions that use mass terms by using predicates like "— is made of snow", "— is a piece of ice", "— is a glass of beer". These would be true of what we do accept as things, such as snowballs, ice cubes, glasses of beer, and they would be false of other things, such as my dogs or the pen on my desk. We don't put snow or ice or beer in the universe of a model, nor do we put bits of snow, bits of ice, or bits of beer in the universe of a model, but only identifiable parts of snow, ice, and beer that we can pick out and distinguish. Masses and unidentifiable bits of masses are not things.

Aside: *Masses as collections of things*
Some say that water, which is a mass, can be treated as a collection of things. A quantity of water is a collection of the smallest parts of water, H_2O molecules. So to quantify over quantities of water is to quantify over collections of H_2O molecules. (We'll see how to quantify over collections in Chapter 10.) But if quantities of water are just collections of H_2O molecules, then no one ever drank a glass of water because, physicists tell us, there is no collection of only H_2O molecules. To identify water with collections of H_2O molecules mistakes an abstraction for our experience, as I explain in "Models and Theories" in *Reasoning in Science and Mathematics*. For a mass such as mud it's even more obvious that there are no smallest bits.