

Semantics for Classical Predicate Logic (Part I)*

Hans Halvorson

Formal logic begins with the assumption that the validity of an argument depends only on its logical form, and not on its content. In particular, if two arguments have the same logical form, then either they are both valid, or they are both invalid. But what is logical form? Our first approach (before the midsemester break) was to identify logical form with sentence connecting words, such as “and”, “or”, and “if ... then ...”. But we found that there are some arguments whose validity is not due to sentence connecting words; instead, these arguments are valid because of certain quantifier phrases, such as “all ...” or “some ...”.

We have now proposed four rules of inference for quantifier phrases: U-Elimination, U-Introduction, E-Introduction, and E-Elimination. It should seem reasonably clear that these rules are acceptable, i.e. will not lead you from true premises to a false conclusion. But since the rules are somewhat subtle, we should provide an independent method for checking their validity. Furthermore, it is not at all clear whether or not these rules form a “complete” system — i.e. if there is a sense in which they can be used to reproduce *any* valid argument.

Let’s consider first the question of the safety of these rules, and compare with the case of the sentence connective rules. How do we know that the rule &E is safe? Roughly speaking, we know that in every situation where $P \& Q$ is true, P is also true. So, the inference “ $P \& Q$ therefore P ” can never lead from truth to falsity. But what do we mean by “every situation”? We made the notion of a “situation” precise by means of a truth assignment to propositional variables P, Q, R, \dots . We stipulated that these propositional variables do not have any logical relations (e.g. P does not

*Version 0.1, revised April 14, 2006. Please report typos to: hhalvors@princeton.edu.

logically imply $\neg Q$), and so any combination of their truth values is a possibility we must consider.

Conversely, how do we know that the argument form " $P \vee Q$ therefore P " is *invalid*? We know that this argument form is invalid because there are *instances* of that form that have true premises and false conclusions, for example:

Tilghman is either Harvard's or Princeton's President. Therefore Tilghman is Harvard's President.

In fact, the point of this example is that there are sentences P and Q such that $P \vee Q$ is true, while P is false. So, to give a counterexample, we only need specify the corresponding row in the truth table for " \vee ".

It is obvious how to give "informal" counterexamples to predicate logic arguments: just replace predicate letters with real predicates, and replace name letters with real names. For example, to show that " $(\exists x)Fx$ therefore Fm " is invalid, simply note that somebody is a Princeton student ($Fx \equiv x$ is a Princeton student) but John Tesh is not a Princeton student ($m \equiv$ John Tesh). The main question of this Chapter is whether this heuristic method for finding counterexamples can be transformed into something systematic and efficient, something like truth tables.

So far the following is clear: to give a counterexample to a predicate logic argument, one must replace the predicate letters with real predicates, and the name letters with real names. It might then seem that once we replace the predicate letters with predicates and the name letters with names, then each sentence is either true or false. But actually, there is one small complication: unless we specify a "domain of quantification," then it is not clear what the quantifiers mean. For example, is it true or false that there was a Danish prince who loved Ophelia? Is it true or false that there are infinitely many prime numbers? Is it true that the color red exists?

Now you might say that " (x) " should always be translated as "For all x ", where the "all" is taken to apply to just those things that actually do exist. So then " $(x)Fx$ " would simply mean that all actually existing things are F . But this seems too strict. For example, it seems that we should also be able to use quantifiers to translate sentences such as "all of Jane Austen's characters are complex."

In order to deal with this and other related issues, we require that each time an argument is translated into quantifier form, one must provide

a “domain of quantification,” abbreviated “DoQ”. For example, we can show that the argument form

$$\frac{(x)(Fx \vee Gx)}{(x)Fx \vee (x)Gx} \quad (*)$$

is invalid by specifying the natural numbers as the DoQ, and by setting “ $Fx \equiv x$ is even” and “ $Gx \equiv x$ is odd.”

1 From Predicates to Sets

In order to give counterexamples in propositional logic, we employed a simplifying strategy: replace sentences with their truth values. In predicate logic, we will also employ a simplifying strategy: replace predicates with their extensions.

Definition. The *extension* of a predicate is just the collection of those things to which the predicate truly applies.

Example. The extension of the predicate “is even” (relative to the domain of natural numbers) is just the collection of even numbers. But if the DoQ is the set $\{1, 2, 3\}$ then the extension of the predicate “is even” is the singleton set $\{2\}$.

In order to give *formal* counterexamples to invalid arguments, we will make use of the mathematical theory called “set theory.” Roughly speaking, set theory is the study of collections of things.

We use curly braces “{” and “}” around a list of things in order to denote the *set* of those things. For example, the set consisting of the numbers 2, 6, and 17 is denoted by:

$$\{2, 6, 17\}$$

Axiom (Axiom of Extensionality:). Two sets A and B are equal, written $A = B$, just in case they have the same members.

For example $\{2, 4\}$ and $\{4, 2\}$ are two different ways to write the unique set whose elements are 2 and 4. Similarly, $\{2, 2, 4\} = \{2, 4\}$ since the former set contains exactly the same elements as the latter.

Definition (Member). The number 2 is a member of the set $\{2, 6, 17\}$. We can indicate this fact by writing $2 \in \{2, 6, 17\}$, where “ \in ” is shorthand for “is a member of.”

Definition (Subset). Every member of the set $\{2, 17\}$ is also a member of the set $\{2, 6, 17\}$. So, we say that $\{2, 17\}$ is a *subset* of $\{2, 6, 17\}$. And if we don’t feel like writing out the words, we might just write $\{2, 17\} \subseteq \{2, 6, 17\}$ as a shorthand.

Note that it does *not* make sense to write $2 \subseteq \{2, 6, 17\}$, because 2 is not itself a set. On the other hand, it does make sense to write $\{2\} \subseteq \{2, 6, 17\}$. The number “2” and the singleton set “ $\{2\}$ ” are two different things.

Definition (Empty Set). There is a set that does not contain anything at all. This set is called the *empty set*, and it is denoted by \emptyset . The empty set is the one and only set that is contained in all other sets.

Definition. If A and B are sets, we use $A \cap B$ to denote the set that consists of those things that are in both A and B . We use $A \cup B$ to denote the set that consists of those things that are in either A or B . Finally, we use $B - A$ to denote those things that are in the set B but not in the set A .

2 Formal Counterexamples

Now we revisit the concepts of counterexamples and interpretations, this time with an eye toward precision and rigor.

Definition. An *interpretation* of a sentence (or sentences) of the predicate calculus consists of:

- (a.) A set X , called the *domain of quantification*.
- (b.) An assignment Ref of names to elements of X .
- (c.) An assignment Ext of predicate letters to subsets of X .

Definition. A *formal counterexample* to an argument in predicate logic is an interpretation where the premises of the argument are true, and the conclusion of the argument is false.

I have not yet given a precise definition of when a sentence is true relative to an interpretation. I have not done so because the definition turns out to be quite complicated, and your intuitions will be sufficient for most cases.

Example. Consider the interpretation with domain $X = \{1, 2, 3, 4\}$ and

$$\begin{aligned} \text{Ref}(m) &= 1, & \text{Ref}(n) &= 1, & \text{Ref}(o) &= 2, \\ \underline{\text{Ext}}(Fx) &= \{1, 2, 3\} & \underline{\text{Ext}}(Gx) &= \{3, 4\} & \underline{\text{Ext}}(Hx) &= \emptyset. \end{aligned}$$

Then, it follows that:

- $(\exists x)(Fx \ \& \ Gx)$ is true relative to this interpretation, since the element 3 is in both $\underline{\text{Ext}}(Fx)$ and $\underline{\text{Ext}}(Gx)$.
- $(x)(Fx \rightarrow Gx)$ is false relative to this interpretation, since the element 1 is in $\underline{\text{Ext}}(Fx)$ but not in $\underline{\text{Ext}}(Gx)$.
- $(x)(Fx \vee Gx)$ is true relative to this interpretation, since every element in the domain is in either $\underline{\text{Ext}}(Fx)$ or $\underline{\text{Ext}}(Gx)$.
- $(x)(Hx \rightarrow Gx)$ is true relative to this interpretation, since $\underline{\text{Ext}}(Hx)$ is the empty set.

Problem. Give a formal counterexample to the argument:

$$(x)(Fx \vee Gx) \vdash ((x)Fx \vee (x)Gx).$$

Solution: Let $X = \{1, 2\}$, let $\underline{\text{Ext}}(Fx) = \{1\}$ and $\underline{\text{Ext}}(Gx) = \{2\}$. Since all elements of X are either in $\underline{\text{Ext}}(Fx)$ or $\underline{\text{Ext}}(Gx)$, $(x)(Fx \vee Gx)$ is true relative to this interpretation. But since $\underline{\text{Ext}}(Fx)$ is not equal to X , $(x)Fx$ is not true relative to this interpretation. Similarly, $(x)Gx$ is not true relative to this interpretation, and so $(x)Fx \vee (x)Gx$ is not true relative to this interpretation.

Problem. Give a formal counterexample to the argument with no premises and conclusion $(x)(y)(Fx \rightarrow Fy)$.

Solution: Let $X = \{1, 2\}$, and let $\underline{\text{Ext}}(Fx) = \{1\}$. Then it is not true that for all x, y in the DoQ, if Fx then Fy . In particular, $1 \in \underline{\text{Ext}}(Fx)$ but $2 \notin \underline{\text{Ext}}(Fx)$.

2.1 Specifying interpretations with tables

It can be boring to write out the interpretations of a bunch of predicates. So, in order to spice up life, we give here an alternative way to specify the interpretation of predicate letters. (This method only works when the domain X is a finite set.)

First we list the predicate letters in the top row of the table, and we list the elements of the domain in the first column of the table. We then write “+” in cell (i, j) if the predicate in column j applies to the object on row i . Otherwise, we write “-” in cell (i, j) .

Problem. Give an interpretation for a problem involving the predicate letters F, G, H , and the names m, n, o .

Solution:

$$X = \{1, 2, 3\}, \quad \text{Ref}(m) = 1, \quad \text{Ref}(n) = 1, \quad \text{Ref}(o) = 2.$$

	F	G	H
1	+	-	-
2	+	+	-
3	+	-	-

Instead of drawing the table, we could have just written:

$$\underline{\text{Ext}}(Fx) = \{1, 2, 3\}, \quad \underline{\text{Ext}}(Gx) = \{2\}, \quad \underline{\text{Ext}}(Hx) = \emptyset.$$

But tables are fun.

When giving an interpretation, you must be clear and explicit about which subsets are assigned to which predicate letters. But given that you satisfy the requirements of clarity and explicitness, you can use whatever method you want to specify your interpretation.

Note. An interpretation must assign the same subset of the domain to both Fx and Fy , even though they have different variables. In other words, the interpretation gives the extension of the predicate letter F , and doesn't care about what variable we put after F . We will deal later with the tricky case of sentences that involve the same predicate letter with different variables — e.g., $(x)(y)(Fx \rightarrow Fy)$ and $(\exists x)(Fx \rightarrow (y)Fy)$.

3 Semantic Properties and Relations

Recall that within propositional logic, we defined a bunch of special kinds of sentences (e.g., tautologies, inconsistencies), and a bunch of special logical relationships between sentences (e.g., implies, is subcontrary to). Actually, these concepts are defined in terms of interpretations, and so they extend naturally to predicate logic. For example, we would define a predicate logic sentence as *tautologous* just in case it is true relative to every predicate logic interpretation.

Definition.

- If ϕ is false relative to some interpretation, then ϕ is said to be *falsifiable*.
- If ϕ is false relative to every interpretation, then ϕ is said to be *inconsistent*.
- If ϕ is true relative to some interpretation, then ϕ is said to be *consistent*.
- Let Γ be a set of sentences. If there is an interpretation relative to which every sentence in Γ is true, then Γ is said to be *consistent*. Otherwise, Γ is said to be *inconsistent*.
- A set Γ of sentences *logically implies* a sentence ψ if there is no interpretation that makes all of Γ true while making ψ false. That is, there is no *formal counterexample*. In that case, the argument with premises Γ and conclusion ψ is *valid*.
- Two sentences ϕ, ψ are *logically equivalent* if they have the same truth value relative to every interpretation.

Example. $(x)(Fx \vee \neg Fx)$ is tautologous, because in any interpretation, each object in the domain is either in $\text{Ext}(Fx)$, or is in the complement of $\text{Ext}(Fx)$.

Example. $(\exists x)(Fx \ \& \ \neg Fx) \rightarrow Gm$ is tautologous, because $(\exists x)(Fx \ \& \ \neg Fx)$ is inconsistent.

4 Semantic Problems and Answers

An answer to a problem requiring the presentation of an interpretation is best seen as having two parts, as follows: (1) state the domain of your proposed interpretation, and present (using one of the above methods) the interpretation of the names and predicate symbols; (2) state the truth values of the various sentences, defend your claim that they have those truth values, and (most importantly) be explicit as to how this information solves the problem you began with. Let's consider a couple of examples of solved problems.

Problem. Show that $(\exists x)Fx \rightarrow (x)Fx$ is not tautologous.

Solution:

1. Let $X = \{1, 2\}$, and let $\text{Ext}(Fx) = \{1\}$.
2. The sentence is false relative to this interpretation: Since 1 is in $\text{Ext}(Fx)$, $(\exists x)Fx$ is true relative to this interpretation. However, since 2 is not in $\text{Ext}(Fx)$, $(x)Fx$ is false relative to this interpretation. Therefore (by truth tables) $(\exists x)Fx \rightarrow (x)Fx$ is false relative to this interpretation. Since the sentence is false relative to *some* interpretation, it is not tautologous. \square

Problem. Show that the sentence from the previous problem is consistent (true relative to some interpretation).

Solution:

1. Let $X = \{1\}$, and let $\text{Ext}(Fx) = \{1\}$.
2. Since $\text{Ext}(Fx) = X$, it follows that $(x)Fx$ is true relative to this interpretation, and therefore (by truth tables) $(\exists x)Fx \rightarrow (x)Fx$ is true relative to this interpretation. Since the sentence is true relative to some interpretation, it is consistent.

5 Systematic Searches for Interpretations

Solving a semantic problem (e.g. “is the argument with premises Γ and conclusion ϕ valid?”) requires one to check all possible interpretations. Now in propositional logic, there are only finitely many interpretations of a finite set of sentences (namely, 2^n interpretations, where n is the number of propositional variables). But in predicate logic, there are infinitely many interpretations of any sentence. So, it is simply impossible to search through all interpretations of a predicate logic sentence.

In fact, it has been proven that the task of deciding if a predicate logic argument is valid cannot be “automated” — i.e. there is no algorithm for solving the validity problem for predicate logic arguments. (This interesting result would be discussed in detail in a more advanced logic course, such as PHI 312 or PHI 321.) However, in the special case of sentences containing only predicates with a single variable (i.e. no relation symbols), there are algorithms for solving the validity problem. These algorithms are the topic of the current section.

We will consider three successively stronger algorithms for testing the consistency of a finite set of sentences. If you feed a finite set Δ of sentences into the algorithm then either it will return the verdict “Inconsistent”, or it will return an interpretation that makes all the sentences true. Of course, if we can solve the consistency problem then we can solve the validity problem. So, these algorithms solve the validity problem for arguments with certain types of sentences.

5.1 Algorithm A

The first algorithm works for sentences that contain no quantifiers.

1. Perform a truth table test on the sentences in Δ . If there is no truth assignment that makes them all true, then return “Inconsistent.” If there is a truth assignment j that makes them all true, then proceed to Step 2.
2. Put the names that occur in the sentences in Δ in a list. Let the DoQ be $\{1, \dots, n\}$ where n is the number of names. For each predicate letter F , and for each $i = 1, \dots, n$, let $i \in \underline{\text{Ext}}(Fx)$ just in case $j(Fa_i) = T$, where a_i is the i -th item on the list of names.

5.2 Algorithm B

The second algorithm works for simple quantified sentences.

Definition. A *simple monadic sentence* is one that contains no relation symbols, and a single quantifier whose scope is the entire sentence.

For example, $(x)((Fx \ \& \ Gx) \rightarrow \neg Hx)$ is a simple monadic sentence.

Definition. If ϕ is a simple monadic sentence, we let ϕ^a denote the instance of ϕ that is obtained by taking off the initial quantifier and replacing all instances of the variable with the name a .

For example, if ϕ is the sentence $(\exists x)(Fx \ \& \ Gx)$ then ϕ^a is the sentence $Fa \ \& \ Ga$.

We now show how to determine if a collection ϕ_1, \dots, ϕ_n of simple monadic sentences is consistent.

1. Reorder the sentences if necessary so that the first m sentences begin with an existential quantifier, and the remaining sentences begin with a universal quantifier.
2. If none of the sentences begins with an existential quantifier, then choose one arbitrary name a . Now put the sentences $\phi_1^a, \dots, \phi_n^a$ into Algorithm A.
3. If there is at least one sentence beginning with an existential quantifier, then choose m arbitrary names a_1, \dots, a_m (one per existential sentence). Now put the following sentences into Algorithm A:

$\phi_1^{a_1}, \dots, \phi_m^{a_m}$ (instances of existential sentences)

$\phi_{m+1}^{a_1}, \dots, \phi_{m+1}^{a_m}$

\vdots

$\phi_n^{a_1}, \dots, \phi_n^{a_m}$

(instances of universal sentences)

That last step could use some explanation: for each distinct existential sentence ϕ_i , we choose a distinct name a_i and construct the corresponding

instance $\phi_i^{a_i}$. Then for each universal sentence ϕ_j , we construct m instances $\phi_j^{a_1}, \dots, \phi_j^{a_m}$, one for each name that was introduced for an existential sentence.

5.3 Algorithm C

Before you read this section, you need to know how to transform a sentence into an equivalent sentence in disjunctive normal form. For this, see Appendix A of Lemmon's book.

The third algorithm is a simple extension of Algorithm B; it works for truth functional combinations of simple monadic sentences.

Definition. A *pure monadic sentence* is a sentence that is a truth-functional combination of simple monadic sentences.

For example,

$$(\exists x)Fx \ \& \ (y)(Gy \rightarrow Fy),$$

is a pure monadic sentence. On the other hand,

$$(x)(Fx \rightarrow (\exists y)Fy)$$

is not simple monadic, since it has nested quantifiers.

1. Transform $\phi_1 \ \& \ \dots \ \& \ \phi_n$ into disjunctive normal form. The result will be of the form $\psi_1 \vee \dots \vee \psi_m$, where each ψ_i is a conjunction of simple monadic sentences and negated simple monadic sentences.
2. Within each disjunct ψ_1, \dots, ψ_m , change each negated simple monadic sentence into an equivalent simple monadic sentence using the quantifier-negation equivalences.
3. For $i = 1, \dots, m$ put the modified conjuncts in ψ_i into Algorithm B. If Algorithm B yields an interpretation, then output this interpretation. If Algorithm B says Inconsistent, then repeat this step until all disjuncts have been checked, and if all disjuncts fail the test, output Inconsistent.

5.4 The small domain method

Although Algorithm C always gives the correct answer, it has the drawback that it does not match the way that we normally reason when we try to decide if some sentences are consistent. The method in this section is perhaps less “clean” than Algorithm C (e.g. it might be more difficult to implement it in a computer program), but it might be somewhat more intuitive.

The “small domain method” is based on the following fact that has been proven by paid logicians:

Fact: If a pure monadic sentence is consistent, then it is true relative to some “small” domain. (In fact, the size of the domain needed is a function of the number of predicates and variables in the sentences.)

Thus, to test a pure monadic sentence for consistency, do the following: Find a quantifier-free sentence that is equivalent to the original sentence relative to a domain with one object; test this resulting sentence for consistency using ordinary truth tables. If the resulting sentence is consistent, you are done — the original sentence is true in that domain. If the resulting sentence is inconsistent, then start over again with a domain with two individuals. If the resulting sentence is consistent, you are done — the original sentence is true in that domain. If the resulting sentence is inconsistent, then repeat the procedure in a domain with three individuals, etc., until either you find an interpretation relative to which the sentence is true, or you conclude that there is no such interpretation. For when you are entitled to draw the latter conclusion, see the last section.

5.4.1 Finding equivalent quantifier-free sentences

Suppose that the domain X has only three objects 1, 2, 3. In this case, the universal statement “ $(x)Fx$ ” is equivalent to a conjunction: “1 is an F , 2 is an F , and 3 is an F .” Similarly, the existential statement “Something is a F ” is equivalent to the disjunction “Either 1 is an F , or 2 is an F , or 3 is an F .” In sum, when there are only finitely many things, we can (by naming each object) translate every simple monadic sentence into a sentence without quantifiers.

The same sort of equivalences also hold for truth-functional combinations of simple monadic sentences (i.e., pure monadic sentences). For example, if we take the standard association of names with numbers:

$$\text{Ref}(a) = 1 \quad \text{Ref}(b) = 2 \quad \text{Ref}(c) = 3,$$

then $(x)Fx \rightarrow (\exists x)Fx$ is equivalent to:

$$(Fa \& Fb \& Fc) \rightarrow (Fa \vee Fb \vee Fc).$$

Generally, in order to obtain a quantifier-free sentence that is equivalent (relative to some finite domain) to a sentence ϕ , you should do the following:

1. Disassemble ϕ into truth-functionally simple components;
2. If any of these components are quantified statements, expand them into equivalent conjunctions or disjunctions;
3. Put the expanded statements back together again using the original truth-functional connectives.

Problem. Find a quantifier-free sentence that is equivalent in a domain with three objects to $(x)(Fx \rightarrow Gx) \vee (\exists x)(Hx \& Mx)$.

Solution: Since the main operator is a disjunction “ \vee ”, we separate the original sentence into $(x)(Fx \rightarrow Gx)$ and $(\exists x)(Hx \& Mx)$. Since these two statements are quantified statements, we expand them as follows:

$$\begin{aligned} (x)(Fx \rightarrow Gx) &\equiv (Fa \rightarrow Ga) \& (Fb \rightarrow Gb) \& (Fc \rightarrow Gc) \\ (\exists x)(Hx \& Mx) &\equiv (Ha \& Ma) \vee (Hb \& Mb) \vee (Hc \& Mc) \end{aligned}$$

Finally, we put these back together again using \vee to obtain:

$$\begin{aligned} &[(Fa \rightarrow Ga) \& (Fb \rightarrow Gb) \& (Fc \rightarrow Gc)] \\ &\vee [(Ha \& Ma) \vee (Hb \& Mb) \vee (Hc \& Mc)]. \end{aligned}$$

□

Problem. For each of the sentences $(x)Fx$, $(\exists x)Fx \rightarrow (\exists x)Gx$, $(x)(Gx \vee Fx)$, find a quantifier-free sentence that is equivalent relative to a domain containing two individuals.

Solution:

$$\begin{aligned} (x)Fx &\equiv Fa \& Fb \\ (\exists x)Fx \rightarrow (\exists x)Gx &\equiv (Fa \vee Fb) \rightarrow (Ga \vee Gb) \\ (x)(Gx \vee Fx) &\equiv (Ga \vee Fa) \& (Gb \vee Fb) \end{aligned}$$

5.4.2 Testing for consistency

Relative to domains with a finite number of individuals, we can test the consistency of a sentence of monadic PC using ordinary truth tables: Just translate the sentence into an equivalent quantifier-free sentence.

Problem. Could $(\exists x)Fx \& (\exists x) \neg Fx$ be true relative to a domain with only one object?

Solution: If there were only one thing in the universe, then $(\exists x)Fx \& (\exists x) \neg Fx$ would be equivalent to $Fa \& \neg Fa$, which is a contradiction. So, no; this sentence cannot be true if there is only one object in the domain. \square

Problem. Could $(\exists x)Fx \& (\exists x) \neg Fx$ be true if there were exactly two objects in the domain?

Solution: In this case $(\exists x)Fx \& (\exists x) \neg Fx$ would be equivalent to:

$$(Fa \vee Fb) \& (\neg Fa \vee \neg Fb).$$

A truth table test shows that this sentence is consistent, e.g. the truth assignment j such that $j(Fa) = T$ and $j(Fb) = F$. So, yes; $(\exists x)Fx \& (\exists x) \neg Fx$ could be true if there were two things in the domain. \square

5.4.3 Putting it all together

Problem. Use the small domain method to show that $(x)(Fx \rightarrow \neg Fx)$ is consistent.

Solution: In a domain with one individual, the original statement becomes $Fa \rightarrow \neg Fa$. This is true when Fa is false. So, $(x)(Fx \rightarrow \neg Fx)$ can be true

in a domain with one individual. In particular, here is an interpretation which shows explicitly that $(x)(Fx \rightarrow \neg Fx)$ is consistent:

$$X = \{1\}, \quad \text{Ext}(Fx) = \emptyset.$$

□

Problem. Use the small domain method to determine if the following argument is valid:

1. $(\exists x)Hx \rightarrow (x)(Fx \rightarrow Gx)$
2. $(\exists x)Fx \quad // \quad (\exists x)Hx \rightarrow (x)Gx$

Solution: Relative to a domain with one member, we get the argument:

1. $Ha \rightarrow (Fa \rightarrow Ga)$
2. $Fa \quad // \quad Ha \rightarrow Ga$

A truth table test shows that if the two premises are true, then the conclusion must also be true. So, we now try a domain with two members. Relative to a domain with two members, we get the argument:

1. $(Ha \vee Hb) \rightarrow ((Fa \rightarrow Ga) \& (Fb \rightarrow Gb))$
2. $Fa \vee Fb \quad // \quad (Ha \vee Hb) \rightarrow (Ga \& Gb)$

A truth table test shows that the premises can be true while the conclusion is false. For example, we could choose the truth-assignment:

$$\begin{aligned} j(Ha) = T & \quad j(Hb) = F \\ j(Ga) = T & \quad j(Gb) = F \\ j(Fa) = T & \quad j(Fb) = F \end{aligned}$$

Therefore, the original argument is invalid.

□

5.4.4 Decision procedures

Suppose that you need to determine whether a sentence ϕ is consistent. You check a domain with one individual, and ϕ is false. You check a domain with two individuals and ϕ is false. You check a domain with three individuals and ϕ is false. Surely you cannot check domains of all sizes! When, if ever, are you entitled to conclude that there is *no* interpretation relative to which ϕ is true? Amazingly, it has been shown that:

Suppose that ϕ is a pure monadic sentence with n predicate letters. If ϕ is consistent, then there is an interpretation \mathcal{I} whose domain has less than or equal to 2^n elements, and ϕ is true relative to \mathcal{I} .¹

So, to take a specific case, if ϕ has two predicate letters, then for ϕ to be consistent, it must be true in some domain with at most 4 objects! It follows that the small domain method is a “decision procedure” for the consistency of pure monadic sentences: it will answer any question you have about consistency in a finite amount of time.

¹Compare with Boolos and Jeffrey, *Computability and Logic*, p. 250.