An Analogy

Imagine that you are at a baseball game. You are sitting high in the upper deck, right next to the scoreboard. Since you can't see the players, watching the scoreboard is the way you follow the game. Here is what the scoreboard says:

	Runs	Hits	Errors
Yankees	**2**	**10**	**0**
Twins	**15**	**14**	**0**

According to this scoreboard, the sentence 'The Twins are losing' is false. (Say: a team is losing if and only if it has less runs than its opponent.) And, according to this scoreboard, the sentence 'Neither team has any errors' is true.

(Note: to say that it is true according to *this* scoreboard that the Twins are winning is not to say that it is true relative to *every possible* scoreboard that the Twins are winning. Obviously, that latter claim is false. For example, let me describe a different scoreboard, according to which is it false that the Twins are winning: switch '2' and '15' in the scoreboard above, and leave everything else as it is.)

Anyway, imagine that I wanted to construct a scoreboard according to which the following sentence was true:

> *The Twins have more runs than hits.*

By constructing the scoreboard displayed above, I would satisfy my desire. (If you think that the scoreboard above couldn't represent a real baseball game, you are failing to consider quite a number of possibilities.)

The analogy here is supposed to be to interpretations. An interpretation is like a scoreboard relative to which some sentences are true and some false. Given an adequate interpretation, it is possible to calculate the truth-value of sentences *relative to that interpretation*. That was the first part of HW7. The second part involved constructing interpretations to make particular sentences true and particular sentences false. The third part of the assignment was more complicated.

Answers to Part A

When I'm trying to figure out whether a sentence is true relative to a particular interpretation, I find it helpful to first translate the sentence into quasi-English in an effort to get more clear about what the sentence 'is saying', and thus what it would take for an interpretation to make the sentence true (false).

As a result, here is what I'll generally do in what follows. First, I'll translate the sentence into quasi-English, then give the answer, then explain the answer in intuitive terms, and then give a complicated explanation in set-theoretic jargon. You might find the complicated explanation unhelpful, which is fine.

A. Let I denote the interpretation given by:

DoQ = {1, 2, 3, 4, 5}
Ext(Fx) = {1, 2, 3} Ext(Gx) = {4} Ext(Hx) = {5}
Ref(m) = 2 Ref(n) = 4

Determine if the following sentences are true or false relative to I.

1. (x)(~Fx ∨ ~Gx)

Everything is either not-F or not-G.

1 is true relative to I.

Consider each thing in the domain of I. Each of those things is either not-F or not-G. So, the claim that everything is either not-F or not-G is true.

Complicated explanation: True, because everything in the domain is a member of ~Ext(Fx)∪~Ext(Gx). Why? First, note that ~Ext(Fx) = {4, 5}, since ~Ext(Fx) consists of the members of the domain that are not members of Ext(Fx). Similarly, ~Ext(Gx) = {1, 2, 3, 5}. And from this it follows that ~Ext(Fx)∪~Ext(Gx) = {1, 2, 3, 4, 5}, since, in general, the *union* of a set X and a set Y, X∪Y, is the set that consists of everything that is a member of X or a member of Y or both—i.e., α is a member of X∪Y if and only if α is a member of X or α is a member of Y or both.

2. (x)((Gx & Fx) → Hx)

Everything is such that: if it is both G and F, then it is H.

2 is true relative to I.

Consider each thing in the domain. For each of those things, it is true to say of it that *if* it is both G and F, then it is H. So, it is true that everything is such that if it is both G and F, then it is H.

Complicated explanation: True, because everything in the domain that is a member of Ext(Gx)∪Ext(Fx) is also a member of Ext(Hx). In other words, true because Ext(Gx)∪Ext(Fx) is a subset of Ext(Hx). (Since Ext(Gx)∪Ext(Fx) = ∅, and ∅ is a subset of every set, it follows that Ext(Gx)∪Ext(Fx) is a subset of Ext(Hx).) Or, to put it another way, true because everything is a member of ~(Ext(Gx)∪Ext(Fx))∪Ext(Hx).

3. ~(x)Fx →(∃x)(Fx & Gx)

If it is not the case that everything is F, then something is both F and G.

3 is false relative to I.

Explanation: The sentence as a whole is a conditional, since the main operator is '→'. The antecedent is true, because '(x)Fx' is false. The consequent is false, because there isn't anything in the domain that is both F and G. (In other words, the consequent is false because Ext(Fx)∪Ext(Gx) = ∅.) Since the antecedent is true and the consequent is false, the conditional as a whole is false.

4. ~Gm & (∃x)(Fx & Gn)

m is not G and something is such that it is F and n is G.

4 is true relative to I. (This one was hard.)

Complicated Explanation: The sentence as a whole is a conjunction, since '&' is the main operator. The first conjunct is true, because 'Gm' is false. ('Gm' is false because the referent of 'm', 2, is not a member of Ext(Gx).) The second conjunct is true because, first, 'Gn' is true ('Gn' is true because the referent of 'n', 4, is a member of Ext(Gx)) and, second, because there is something in the domain that is a member of Ext(Fx). Or, to say it in another way, the second conjunct is true because there is something in the extension of the complex predicate 'Fx and Gn', where the extension of that predicate is identical to Ext(Fx) unless Ref(n) is not in Ext(Gx), in which case it is the empty set (so, in this case it is identical to Ext(Fx)).

5. **(∃x)(~Fx & ~Gx)**

Something is both not-F and not-G.

5 is true relative to I.

Explanation: True, because there is something in the domain that is a member of ~Ext(Fx)∩~Ext(Gx). (Namely, 5.) (In general, the *intersection* of a set X and a set Y, X∩Y, is the set that consists of everything that is a member of *both* X *and* Y—i.e., α is a member of X∩Y if and only if α is a member of X *and* α is a member of Y.)

6. **(∃x)(Fx → Gx)**

Something is such that if it is F, then it is G.

6 is true relative to I.

Look at everything in the domain. Can you find something such that it is true to say of it that *if* it is F, then it is G? Yes, you can; so, it is true that something is such that *if* it is F, then it is G.

As an example, let's consider 4. (5 would also work.) Now consider the claim: if 4 is F, then 4 is G. That's true, because the antecedent is false. The antecedent is false because 4 is not in the extension of F. Since it is true that *if 4 is F, then 4 is G*, it follows that *something is such that if it is F, then it is G*.

Complicated Explanation: True, because there is something in the domain that is a member of ~Ext(Fx)∪Ext(Gx). (Namely, 4 and 5.)

Answers to Part B

In this part, for each question, there were a number of legitimate strategies that you could have used to find the answer. I will try to illustrate each of the main strategies at some point in what follows.

B. Give counterexamples to the following invalid argument forms. You do not need to show your work, but you do need to present an interpretation, and explain how it solves the problem.

1. **(x)(Fx → Hx) │— (x)((Fx ∨ Gx) → Hx)**

In solving this problem, I will use a strategy I call the *Low-tech Method*. (This is the strategy that I recommend trying first.)

Our goal is to find an interpretation that shows that the argument form above is invalid. To do that, we need to find an interpretation—call it 'I'—such that the premise of the argument above is true relative to I but the conclusion is false relative to I. So, we want the truth-values of the sentences to come out as follows, relative to I:

(x)(Fx → Hx) │— (x)((Fx ∨ Gx) → Hx)
T F

Now we ask, "what, at a minimum, must our interpretation I be like if it makes the conclusion above false?". The intuitive answer is that *there must be something that is either F or G, but that is not-H*. How did we know that? Well, the conclusion above will be false relative to an interpretation just in case the following sentence is true relative to that interpretation:

(∃x)~((Fx ∨ Gx) → Hx)

And this sentence will be true relative to an interpretation I just in case there is something in the domain of I that is either in the extension of F or in the extension of G, but that is not in the extension of H—in other words, just in case there is something that is either F or G, but that is not-H.

With that in mind, let's sketch one such interpretation, the simplest one we can think of.

DoQ = {1,
Ext(Fx) = {1
Ext(Gx) = {
Ext(Hx) = {

If we were to 'close all the brackets' above, the resulting interpretation would make the conclusion false. But what about the premise? We have to make sure that our interpretation makes the premise above true in addition to making the conclusion above false if it is going to show that the argument above is invalid.

This exposes a problem with the interpretation we've started to sketch. As it stands, it doesn't make the premise above true. And if we were to add '1' to the extension of G so that it did make the premise true, this would have the effect of making the conclusion false relative to the resulting revised interpretation.

There are two simple solutions to this problem (among others) that will get us what we want. I'll outline one and then the other.

First, we might see that if we were to go ahead and add 1 to the extension of G, we could also add a new thing to the domain to make the conclusion come out false, while not affecting the truth of the premise, as long as we added the new thing to the extension of G and not of F. Here is how this would go:

DoQ = {1, 2}
Ext(Fx) = {1}
Ext(Gx) = {1, 2}
Ext(Hx) = { } = ∅

This interpretation provides a counterexample to the argument above.

A more crafty interpretation results from the observation that we could have initially made the conclusion false by putting 1 in the extension of G rather than in the extension of F. If we had done that, then we would have avoided the problem we ran into above. In other words, the following interpretation also provides a counterexample to the argument:

DoQ = {1}
Ext(Fx) = ∅
Ext(Gx) = {1}
Ext(Hx) = ∅

**Important:* when you come up with an interpretation that you think provides a counterexample to an argument, make sure you go back and *double-check your work* to make sure that you are right. The easiest way to do this is to separately calculate the truth-value relative to your interpretation for each of the premises and for the conclusion, and then make sure that all the premises come out true but the conclusion false relative to your interpretation.

2. **(∃x)(Fx → Gx), (∃x)Fx ∨ (∃x)Gx |— (∃x)Fx → (∃x)Gx**

I recommend the Low-Tech Method for this problem too. To show that this argument is invalid, we need an interpretation relative to which all the premises are true but the conclusion is false. If the conclusion—which is a conditional—is to be false relative to our interpretation, then its antecedent must be true and its conclusion false relative to our interpretation. So, relative to our interpretation, it must turn out that '(∃x)Fx' is true but '(∃x)Gx' is false. And this means that the extension of G must be empty and the extension of F non-empty. This suggests the following starting point:

DoQ = {1
Ext(Fx) = {1
Ext(Gx) = ∅

Now, we need to ask whether this interpretation, as it stands, makes all the premises true but the conclusion false. Unfortunately, it does not—because, although it does also make the second premise true and the conclusion false, it fails to make the first premise true. So, we have a problem. Fortunately, however, we can solve this problem by adding a new thing to the domain and not placing it in the extension of F—thereby making the initial premise true without affecting the truth-value, relative to the resulting interpretation, of either the other premise or the conclusion. Here is the resulting interpretation:

DoQ = {1, 2}
Ext(Fx) = {1}
Ext(Gx) = ∅

This interpretation works as a counterexample to the argument.

Let's now consider how we could have used Algorithm C to solve this problem. We begin by conjoining all the premises of the argument and the negation of the conclusion (see handout.)

(Why do we do this? Because if we want to use Algorithm C to figure out whether the argument is valid, our first step is to form a conjunction of the premises and the negation of the conclusion, and then use Algorithm C to figure out whether that conjunction is consistent. If it is consistent, then that means that there is an interpretation that makes it true—which, if you think about it, also means that there is an interpretation that makes all the premises of the argument true, but the conclusion false, which means that the argument is invalid. If the conjunction is inconsistent, then the argument is valid.

(*Important*: Make sure you understand this point—in other words, make sure you understand why figuring out whether the sort of conjunction we've described is consistent is sufficient for figuring out whether the argument it represents is valid.)

If we do that, we get the following:

$$(\exists x)(Fx \rightarrow Gx) \ \& \ ((\exists x)Fx \vee (\exists x)Gx) \ \& \ \sim((\exists x)Fx \rightarrow (\exists x)Gx)$$

Since this is a *pure monadic sentence* (see handout), the next step is to put it into *disjunctive normal form* (see Lemmon pp. 190-5).

The task of putting the sentence above into disjunctive normal form may be intimidating. However, it may help to note that the sentence above has the following form:

$$P \ \& \ (Q \vee R) \ \& \ \sim(Q \rightarrow R)$$

The task of putting this sentence in disjunctive normal form is less intimidating. Here is what I propose to do. I will put the simplified sentence immediately above into disjunctive normal form, and then substitute '(∃x)(Fx→Gx)', '(∃x)Fx', and '(∃x)Gx' back in for 'P', 'Q', and 'R' respectively. This should make the task easier for me and easier for you to follow. I will cite the rules on page 191 that allow me to make the various transformations as I go.

First, the simplified formula is equivalent to:

P & (Q ∨ R) & ~(~Q ∨ R) (Rule 8)

And that is equivalent to:

P & (Q ∨ R) & (~~Q & ~R) (Rule 11)

And that is equivalent to:

P & (Q ∨ R) & (Q & ~R) (Rule 7)

And that is equivalent to:

((P & Q) ∨ (P & R)) & (Q & ~R) (Rule 13)

And that is equivalent to:

((Q & ~R) & (P & Q)) ∨ ((Q & ~R) & (P & R)) (Rule 13)

And that is equivalent to:

(~R & P & Q) ∨ (Q & ~R & P & R) (Rule 1 and the convention of omitting redundant parentheses)

And this, finally, is in disjunctive normal form.

Next, let's substitute '(∃x)(Fx→Gx)', '(∃x)Fx', and '(∃x)Gx' back in for 'P', 'Q', and 'R' respectively. This gives us the following:

(~(∃x)Gx & (∃x)(Fx → Gx) & (∃x)Fx) ∨ ((∃x)Fx & ~(∃x)Gx & (∃x)(Fx → Gx) & (∃x)Gx)

The next step is to change each negated simple monadic sentence into an equivalent simple monadic sentence using the quantifier negation rules (see handout). If we do that, we get:

((x)~Gx & (∃x)(Fx → Gx) & (∃x)Fx) ∨ ((∃x)Fx & (x)~Gx & (∃x)(Fx → Gx) & (∃x)Gx)

Now, we take each disjunct, one at a time, and put it into Algorithm B (see handout). Let's take the left disjunct first.

What we do now is reorder the sentences (in the left disjunct) so that those starting with an existential quantifier occur first (see handout). Here is the set of sentences we get when we do that:

{ $(\exists x)(Fx \rightarrow Gx)$, $(\exists x)Fx$, $(x){\sim}Gx$ }

The next thing we do is to choose an arbitrary name for each sentence beginning with an existential quantifier (see handout). In this case, we need two, since there are two sentences beginning with existential quantifiers. Let's pick 'a' and 'b' for our arbitrary names. Substituting the relevant sentences in gives us the following set:

{ $Fa \rightarrow Ga$, Fb, $(x){\sim}Gx$ }

The next thing we do is to substitute in, for each universally quantified sentence, a substitution instance of that sentence for each name we previously introduced (see handout). If we do that, we get the following set:

{ $Fa \rightarrow Ga$, Fb, ${\sim}Ga$, ${\sim}Gb$ }

Now we put these sentences into Algorithm A, and see if there is an interpretation that makes them all true—in other words, we now see whether those sentences form a consistent set of sentences. It turns out that they are consistent, and so there is an interpretation that makes them all true—for example, consider an interpretation relative to which v(Fa) = T, v(Fb) = F, v(Ga) = F, and v(Gb) = F.

This tells us that the pure monadic sentence that we started with (namely, '$(\exists x)(Fx \rightarrow Gx)$ & $((\exists x)Fx \vee (\exists x)Gx)$ & ${\sim}((\exists x)Fx \rightarrow (\exists x)Gx)$') is true relative to the following interpretation:

DoQ = {1, 2}
Ext(Fx) = {1}
Ext(Gx) = ∅

In other words, this tells us that there is an interpretation relative to which all the premises of the argument we started with are true but the conclusion false. In other other words, this means that there is a counterexample to the argument we started with—namely, the interpretation we've just considered. (Note that this is the same interpretation we arrived at via the Low-Tech Method.)

3. **$(x)Fx \leftrightarrow (x)Gx \;\vdash\; (x)(Fx \leftrightarrow Gx)$**

Let's use the Small Domain Method to solve this problem.

In general, if we want to use the small domain method to check whether an argument is valid, we begin by conjoining all the premises of the argument and the negation of the conclusion. In other words, if we want to use the small domain method our first step is to form one (possibly long) sentence that is a conjunction of all the premises and the negation of the conclusion. We will then use the small domain method to figure out whether the resulting conjunction is consistent. If it is consistent, then that means that there is an interpretation that makes it true—which, if you think about it, also means that there is an interpretation I such that all the premises of the relevant argument are true relative to I, but the conclusion is false relative to I, which means that the argument is invalid. If the conjunction is inconsistent, then the argument is valid.

Here is what we get when we conjoin the premise of the argument above with the negation of the conclusion:

$$((x)Fx \leftrightarrow (x)Gx) \ \& \ {\sim}(x)(Fx \leftrightarrow Gx)$$

In order to make it easier to find a quantifier-free sentence that is equivalent to this sentence relative to domains of various sizes, let's use the quantifier negation rules to derive the following equivalent sentence:

$$((x)Fx \leftrightarrow (x)Gx) \ \& \ (\exists x){\sim}(Fx \leftrightarrow Gx)$$

Next, we follow the instructions on the handout and consider a quantifier-free sentence that is equivalent relative to a domain of one object to the pure monadic sentence above. (See handout.) If we follow those instructions, here is what we get:

$$(Fa \leftrightarrow Ga) \ \& \ {\sim}(Fa \leftrightarrow Ga)$$

This is an inconsistent sentence, as a truth-table test will show. (It is a contradiction.) Given that, the next thing we are supposed to do is to consider a quantifier free sentence that is equivalent relative to a domain of two objects to the pure monadic sentence above. (See handout.) If we follow those instructions, we get the following:

$$(Fa \ \& \ Fb) \leftrightarrow (Ga \ \& \ Gb)] \ \& \ [{\sim}(Fa \leftrightarrow Ga) \vee {\sim}(Fb \leftrightarrow Gb)]$$

This sentence is consistent, as the following truth-assignment shows: Let v(Fa) = T, v(Fb) = F, v(Ga) = F, and v(Gb) = F. This tells us that the pure monadic sentence above is true relative to the following interpretation:

DoQ = {1,2}
Ext(Fx) = {1}
Ext(Gx) = ∅

In other words, this tells us that there is an interpretation relative to which all the premises of the argument we started with are true but the conclusion false. In other other words, this means that there is a counterexample to the argument we started with—namely, the interpretation we've just considered.

Answers to Part C

C. Determine if the following arguments are valid. If you say that an argument is valid, explain how you know that fact. If you say that an argument is invalid, provide a counterexample.

1. **(x)(Fx → Gx), (∃x)(Gx&Hx) |— (∃x)(Fx&Hx)**

Let's use Algorithm B to solve this problem. (Of course, you could also use either the Low-Tech Method or the Small Domain Method.)

First, we consider the set of sentences that consists of all the premises together with the negation of the conclusion:

{ (x)(Fx → Gx), (∃x)(Gx&Hx), ~(∃x)(Fx&Hx) }

Now, in order to use Algorithm B we need to have a set of simple monadic sentences—and the set above is *not* a set of simple monadic sentence (see handout). However, we can transform the set above into a set of simple monadic sentences by using the quantifier negation rule on the last sentence above. If we do that, we get the following:

{ (x)(Fx → Gx), (∃x)(Gx&Hx), (x)~(Fx&Hx) }

And this resulting set *is* a set of simple monadic sentences. The next thing we do is to reorder the sentences so that the existentially quantified sentences occur first. Here is what we get when we do that:

{ (∃x)(Gx&Hx), (x)(Fx → Gx), (x)~(Fx&Hx) }

The next thing we do is to choose an arbitrary name for each sentence beginning with an existential quantifier (see handout). In this case, we only need one, since there is only one sentence that begins with an existential quantifier. Let's pick 'a' for our arbitrary name. Substituting the relevant sentences in gives us the following set:

{ Ga&Ha, (x)(Fx → Gx), (x)~(Fx&Hx) }

The next thing we do is to substitute in, for each universally quantified sentence, a substitution instance of that sentence for each name we previously introduced (see handout). If we do that, we get the following set:

{ Ga&Ha, Fa → Ga, ~(Fa&Ha) }

Now we put these sentences into Algorithm A, and see if there is an interpretation that makes them all true—in other words, we now see whether those sentences form a consistent set of sentences. It turns out that they are consistent, and so there is an interpretation that makes them all true—for consider the interpretation relative to which v(Fa) = F, v(Ga) = T, v(Ha) =T.

This tells us that there is an interpretation relative to which all the premises of the argument we started with are true but the conclusion false—namely, the following interpretation:

DoQ = {1}
Ext(Fx) = ∅
Ext(Gx) = {1}
Ext(Hx) = {1}

And this means that there is a counterexample to the argument we started with—namely, the interpretation we've just considered.

For the following questions, I recommend using Low-Tech Method to 'figure out' what the answer is, and then answering the questions with the following informal proofs.

2. (x)((Fx&Gx) → Hx), (∃x)Fx |— ~(x)Gx v (∃x)Hx

Assume, for reductio, that there is an interpretation I such that the premises are both true but the conclusion false relative to I.

Since the conclusion is false relative to I, it follows that both '(x)Gx' and '~(∃x)Hx' are true relative to I.

Thus, some member of the domain of I—call it 'α'—is both F and G, since '(x)Gx' is true relative to I and, by assumption, '(∃x)Fx' is true relative to I.

But then α is H, since '(x)((Fx&Gx) → Hx)' is true relative to I and, as we've noted, α is both F and G.

Since α is H and α is a member of the domain of I, it follows that '(∃x)Hx' is true relative to I.

But this means that both '(∃x)Hx' and '~(∃x)Hx' are true relative to I, which is absurd.

So, the assumption must be false—which is to say that there is no interpretation I such that the premises are both true but the conclusion false relative to I—and that means that the argument is valid.

3. (x)Fx ↔ (∃x)Gx |— (x)Fx ∨ (x)~Gx

Assume, for reductio, that there is an interpretation I such that the premise is true but the conclusion false relative to I.

Since the conclusion is false relative to I, it follows that both '~(x)Fx' and '~(x)~Gx' are true relative to I.

Thus, '(∃x)Gx' is also true relative to I, since '~(x)~Gx' is true relative to I.

From the fact that '(∃x)Gx' is true relative to I together with the fact that, by assumption, the premise '(x)Fx ↔ (∃x)Gx' is true relative to I, it follows that '(x)Fx' is also true relative to I, since a biconditional is true if and only if its primary constituents have the same truth-value.

But this means that both '(x)Fx' and '~(x)Fx' are true relative to I, which is absurd.

So, the assumption must be false—which is to say that there is no interpretation I such that the premise is true but the conclusion false relative to I—and that means that the argument is valid.