1A Formal Methods
Shyane Siriwardena
ss2032@cam.ac.uk

Worksheet 2

A. Here is a limerick by George Boolos (1995):

> According to W. Quine,
> Whose views on quotation are fine,
> Boston names Boston,
> and Boston names Boston,
> But 9 doesn't designate 9.

Unfortunately, the quotation marks have gone missing from the third, fourth and fifth lines, rendering it both false and unamusing. Rectify this by adding quotation marks, to restore the proper distinction between use and mention.

As you may have noticed, this question concerns the use/mention distinction. We use quotation marks in order to clarify when we mean to talk about a *word* or the thing the word *represents* or *denotes*. So, for instance, if I want to talk about the Trinity philosopher who wrote *The Problems of Philosophy* (suppose I want to say of him that he was very clever), then I would do this:

(1) *Bertrand Russell was very clever.*

Now suppose I want to remark of the particular collection of letters that make up Russell's name that it contains four vowels. If I did this:

(2) *Bertrand Russell contains four vowels.*

Then I would be saying of the *man* that he (perhaps by some bizarre Scrabble-related accident) has four vowels inside of him. What I need to do here is use single quotation marks to make clear that I mean to be talking about the *words* themselves. So, here's what I should've done:

(3) 'Bertrand Russell' contains four vowels.

This sentence, i.e. (3), is true, where the previous one, (4), (probably) is not.

So far, I've been speaking loosely, though. When I say I want to "talk about" some thing or another in this case, I mean that I want to "pick it out" or *refer* to it. (This is also called *denoting* or *naming*.) So, to explain what just happened with the Russell sentences more precisely, in (1) and (2), the first two words denote or refer to the human being, whereas in (3) the bit in quotations denotes or refers to the words. Put another way, I'm ***using*** the name in (1) and (2) and only ***mentioning*** it in (3).

Here's why we should worry about the difference between use and mention: ambiguity. If we don't make explicit when we're using and when we're mentioning, we'll end up making invalid arguments look valid (there are other problems that arise too--this is just one of them). Observe:

1A Formal Methods
Shyane Siriwardena
ss2032@cam.ac.uk

Russell is the author of "On Denoting".
Russell is seven letters long.

∴ The author of "On Denoting" is seven letters long.

This inferences might appear valid; as is, it is seems to be of the form: A=B, A=C, ∴ B=C. But, it turns on a use/mention ambiguity. In Premise 2, I don't mean to talk about or refer to the person, but rather the *name* itself. I mean to be *mentioning* rather than *using* 'Russell'. Correcting the ambiguity, the argument should look like this:

Russell is the author of "On Denoting".
'*Russell*' *is seven letters long.*

∴ The author of "On Denoting" is seven letters long.

Now the argument is of the form: A=B, C=D, ∴ B=D, and is clearly *invalid*.

Okay! Now let's return to the original limerick:

According to W. Quine,
Whose views on quotation are fine,
Boston names Boston,
And Boston names Boston
But 9 doesn't designate 9.

The first two lines of this are fine--we don't need to add quotation marks there. But, as is, Line 3 is false, so we'll try to fix that first. Right now, since both instances of 'Boston' are without quotation marks, both denote or name the city in Massachusetts. And it is false that the city in Massachusetts *names* the city in Massachusetts since places aren't the kinds of things that can name. What we want to do is say that the *name* names (or denotes) the place! In other words, we want to say of a collection of letters that they name or denote Boston (the city). As we learned above, we can pick out or name a word by using single quotation marks. So,

Boston names Boston

should read:

‘*Boston*’ *names Boston*

Now the line says that ‘Boston’ (the name) names Boston (the place).

On to Line 4! Here, we have almost a duplicate of the previous line. We just need to correct it in a way that makes the claim true. That means, you *could* just repeat what we did to Line 3. But, the hope is that you fixed it in a different way. If not, don’t worry! But in my explanation here, I’ll be doing something slightly different with Line 4. Have a look at the following variation on Line 4:

And ‘*Boston*’ *names* ‘*Boston*’

Is that right? That is, is the statement there true? Answer: no it’s not. Here’s why: now both instances of ‘Boston’ denote the same collection of letters. In particular, they both denote the collection of letters that names the city (Boston). So, how would we say what names ‘Boston’ (the collection of letters)? By putting it in quotations, as before. That means, we’re putting the bit in quotations itself between a further set of quotation marks. When we do this, the convention is to use single quotations for one instance and double quotations for the other, just to make it clear that there are two different sets of quotation marks in use. So, Line 4 corrected would look like this:

And “ ‘*Boston*’ ” *names* ‘*Boston*’

That brings us to Line 5:

But 9 doesn’t designate 9.

So, here we want something that doesn’t designate 9 (the number). As you may have spotted, after our discussion on Line 3, 9 (the number) doesn’t designate (or name, or denote, or refer--these are all synonymous) 9! So, in fact, Line 5 is true and can be left as is. That said, there are other correct answers (in fact, this is the case for Lines 3 and 4 as well). For instance, you might have done this:

But “ ‘9’ ” doesn’t designate 9.

This would also be correct since, as we saw in Line 4, ‘ “*9*” ’ (the name of the name) in fact designates ‘9’ (the name), and not 9 (the number).

1A Formal Methods
Shyane Siriwardena
ss2032@cam.ac.uk

B. ⋆ Determine whether the following arguments are valid or invalid. You may use any shortcuts you like, and can use either complete or partial truth tables, as appropriate.

1. $A \to B, B \therefore A$
2. $L \to (L \wedge S), \neg S \therefore L \to R$
3. $(B \leftrightarrow \neg R) \wedge \neg\neg Q \therefore \neg((B \wedge Q) \to R)$
4. $(A \leftrightarrow \neg B) \leftrightarrow C \therefore \neg(A \leftrightarrow (B \leftrightarrow C))$
5. $A \vee (B \wedge C) \therefore (A \vee B) \wedge (A \vee C)$
6. $(A \wedge B) \vee (C \wedge D) \therefore (A \wedge C) \vee (B \wedge D)$

Here, I will give an explanation of partial truth tables and shortcuts. The keys to using these tools are as follows: (1) an understanding of the definition of in/validity, (2) a good grasp of the truth tables for each connective.

First, a reminder: An argument is **valid** if and only if it is not possible for the all the premises to be true and the conclusion false.

Now that we have the language of truth values at our disposal, we can put this more precisely, in the following terms:

An argument is valid if and only if there is no valuation (i.e. line of the truth table) on which all the premises are true and the conclusion false.

From which it follows that: *An argument is invalid if and only if there is some (at least one) valuation on which all the premises are true and the conclusion false.*

As we were saying last class, it would be a good idea to memorise these definitions. The more you stick to the letter of these definitions, the better you'll do on questions that involve in/validity of arguments in some way. In this way, your application of the concept of validity should be somewhat mechanical.

The next thing to get under your belt as soon as possible is the truth table for each connective. Knowing these by memory is just about essential. This is necessary for doing truth-tables without referring to the text. But, it will also be very helpful for understanding some of the material later in the course. All of that to say, it's worth spending time to make sure that you've got a solid knowledge of the truth-tables for each connective. You can find these in the textbook on pages 32 and 33.

1A Formal Methods
Shyane Siriwardena
ss2032@cam.ac.uk

When we're using truth tables to figure out whether an argument is valid or invalid, we need to keep the exact definitions in mind. Let's start with validity. We said that an argument is valid if and only if there is no valuation on which all the premises are true and the conclusion false. In other words, when we consult each line of the truth table, on all of the lines of that table where all of the premises are true, the conclusion had better be true as well. Or, put differently, on all of the lines where the conclusion is false, it had better not be that all of the premises are true. Let's start with some simple examples:

$$(1) \quad A \wedge B \therefore B$$

First, let's construct our truth table. I'll need a column for each of the atomic propositions, one for each of the premises, and one for the conclusion.

Atom	*Atom*	*Premise 1*	*Conclusion*
A	**B**	**A ∧ B**	**B**
T	T	T	T
T	F	F	F
F	T	F	T
F	F	F	F

This truth table is simple since it just involves the truth-table for conjunction as well as all the possible truth values of A and B. I've filled out the whole thing, but I will show you in a moment how to use a kind of shortcut to cut out some of the work.

We wanted to know if the argument given in (1) is valid. To determine this using the truth-table just constructed, we need to make sure that *whenever the conclusion is false, there is also at least one false premise*. So we are interested in the following lines (in green):

Atom	*Atom*	*Premise 1*	*Conclusion*
A	**B**	**A $\wedge$ B**	**B**
T	T	T	T
T	F	F	**F**
F	T	F	T
F	F	F	**F**

These are the only lines on which the conclusion is false. Now we need to check all of the premises on those lines. In this case, there is only ONE premise, so there's only one other column to check, namely column 3.

Looking at column 3, the premise is false on both of the lines where the conclusion is false. Therefore, there is no line on which the premises are true and the conclusion false. So, the argument is VALID.

Let's do the same thing again, with a new example:

(2) $A \rightarrow B \therefore A$

Once again, I'll start by sketching out a truth table. Again, I'll need a column for each of the atomic propositions, one for each of the premises, and one for the conclusion

Atom	*Atom*	*Premise 1*	*Conclusion*
A	**B**	**A $\rightarrow$ B**	**A**
T	T		
T	F		
F	T		
F	F		

1A Formal Methods
Shyane Siriwardena
ss2032@cam.ac.uk

This time, instead of filling out the entire truth table right away, I'm going to start by filling out the column for the conclusion. Observe:

Atom	*Atom*	*Premise 1*	*Conclusion*
A	**B**	**A $\rightarrow$ B**	**A**
T	T		T
T	F		T
F	T		F
F	F		F

As before, I need only be concerned with the lines of the truth table where the conclusion is false. I.e. the following lines in green:

Atom	*Atom*	*Premise 1*	*Conclusion*
A	**B**	**A $\rightarrow$ B**	**A**
T	T		T
T	F		T
F	T		**F**
F	F		**F**

The reason for this is that, it is the lines on which the conclusion is false that are relevant to the argument being valid or not. If the argument is valid, one of the premises will be false on each of these lines. If the argument is invalid, all of the premises will be true on at least one of these lines. So, with this in mind, I will only fill out the green rows of the truth table for the premise column (in this case there's just one premise column since there's only one premise).

Atom	*Atom*	*Premise 1*	*Conclusion*
A	**B**	**A → B**	**A**
T	T		T
T	F		T
F	T	**T**	F
F	F	**T**	F

As you can see, the premise is true on *both* of the lines where the conclusion is false. Thus, there is at least one valuation on which all the premises are true and the conclusion is false. Therefore, this argument is *invalid*.

Notice, I can know this without even filling out the remaining spaces in the truth table. Those are lines where the conclusion is true, so they are not going to tell me anything about the validity or invalidity of the argument. Why is this the case? It is because of the definitions of validity and invalidity; as I said before (though it's worth belabouring the point) the definitions of validity and invalidity are concerned with the situations in which the conclusion is false. The argument could still be valid or invalid, regardless of what's happening on the lines of the truth table where the conclusion is true.

So, this is all you need to know to successfully complete a truth table for an argument of *any length*. Here are a couple questions you might have:

Q: How do I know how many lines to include in my truth table.

A: We need to make sure that the truth table for any given argument includes *<u>all possible combinations</u>* of truth values for each of the *<u>atomic propositions</u>*. That is to say, the truth table should exhaust all of the possible combinations of 'true' and 'false' for each distinct 'letter' (so to speak) in the premises and conclusion. So, if you have only 1 atomic proposition A, your truth table will need to have 2 lines -- one for when A is true, and one for when A is false. If you have 2 atomic propositions, you'll need 4 lines. And, in general, if you have *n* atomic propositions, you'll need 2^n lines.

1A Formal Methods
Shyane Siriwardena
ss2032@cam.ac.uk

Q: What is a partial truth table? How do I know when to use one?

A: A partial truth table is one that only includes a single line. This is different from the shortcut I was demonstrating above, where you are just leaving some blanks in the complete table. <u>You are allowed to use a partial truth table **only** when you are showing an argument to be **invalid**.</u> The reason for this, again, lies in the definitions of validity and invalidity. In order for an argument to be *invalid* it is enough for there to be a <u>single</u> line (i.e. valuation) on which the premises are all true and the conclusion false. So, if you know an argument to be invalid, you only need to give an instance of one valuation on which the premises are true and the conclusion false.

However, in order to show that an argument is *valid*, one line will never be enough. You need to show that on *all* of the lines on which the conclusion is false, some premise is also false. This means you need to write out *all* of the lines for the conclusion column of the truth table. You cannot just write out the rows on which the conclusion is false because that leaves open the possibility that there is some other line that isn't written out on which the conclusion is also false. So, to show that the false lines included exhaust *all* of the false lines, you need to fill in the true lines too. Once *that's* done, as we saw in the question above, you can ignore the lines on which the conclusion is true, and fill out the other rows accordingly.

Okay. So that's all well and good, but how do you *know* when to use a partial truth table? Unfortunately, the only way to know that all you need is a single line of a truth table is if <u>you already know that the argument is invalid</u>! I, personally, find this to be very hard to spot unless I use a truth table. The safest bet is always to start by filling out the entire column for the conclusion, and then to work out the rest based on the lines on which the conclusion is false. If, however, you can tell an argument is invalid before you've done the truth table, then you are definitely allowed to just write out one of the relevant lines of the truth table (i.e. one of the lines on which, given the truth values of the atomic propositions, the premises are true and the conclusion is false). But don't feel like you have to do this.

Finally, a note on doing truth tables for more complex propositions (i.e. propositions that include multiple connectives). Make sure that you're aware of the <u>main operator</u>! In general, you should be aware of the 'order' in which the proposition is put together. Here's what I mean by that. Consider the following proposition:

1A Formal Methods
Shyane Siriwardena
ss2032@cam.ac.uk

$(B \longleftrightarrow \neg R) \wedge \neg\neg Q$

We need to break this down into its constituent propositions in order to know how to evaluate the truth value of the whole. The main operator is always the *last connective* to go into the proposition (and the last one to come out when we're breaking a proposition down). So, we can imagine breaking down any given proposition into a kind of construction tree like below. At each step, we identify the main operator (in blue), and the proposition(s) next to that operator.

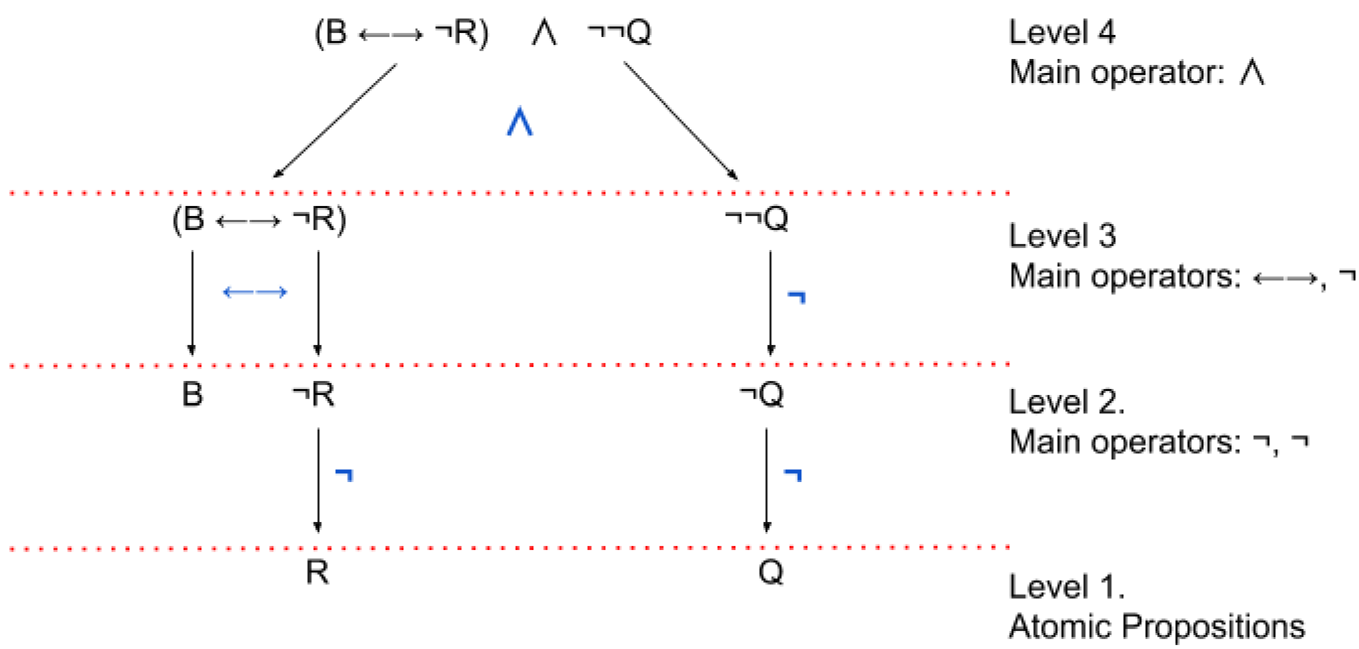

If we worked from the Level 1 of the tree, upwards, we would have the order in which we should consider the truth values for our truth table. First we need to consider the atomic propositions, then the next level of connectives (in this case the row of the tree where negations are added (Level 2)), then the next level, and so on until we get to the proposition at the top of the tree. The final step in the tree--i.e. the final connective to go in--is the main operator of the entire proposition (in this case, it's a conjunction). So, if you were doing a truth table for the proposition at the top of the page, in order to figure out the truth-value of the entire conjunction, you would first need to know the truth-value of the biconditional on the left-hand side, and of the negated proposition on the right-hand side (i.e. the propositions at Level 3). And to know those, you'd need the truth-values of the propositions at Level 2, and so on.