Logic for Lawyers

Prof. Daniel Pi

Contents

1 Argument Structure: Validity and Soundness

In order to understand the law, it is imperative that you firmly grasp the structure of argumentation. This brief article introduces the architecture of argumentation.[1] There is considerably more to argumentation than what is contained in this brief introduction, and students interested in exploring the topic further are strongly encouraged to explore the subject more extensively.[2] However, this précis should suffice for the purposes of our course.

An argument consists of *propositions*. A proposition is simply a claim about some state of affairs. For example, it is proposition *that grass is green* and *that pigs can*

[1] This subject matter is sometimes referred to as "informal logic" to distinguish it from mathematically rigorous "formal" logics.

[2] An excellent starting point to a more advanced study of the topic is the Stanford Encyclopedia of Philosophy entry on informal logic, which contains citations to further resources.

fly. Propositions can be either true or false. A false proposition (e.g., that pigs can fly) is no less a proposition.

Arguments are a way of structuring propositions. The purpose of the structure is to establish that a particular proposition, which we call the "conclusion," is true. In order to accomplish this, we begin with propositions that we *already* take to be true. These starting points are called "premises." We use the premises to "infer" the truth of the conclusion. These are the components of every argument: premises, inferences, and conclusion.

Premises are propositions which are purported to be true *ex ante*. Inferences are rules that tell us something about the relationship between propositions. An inference is "valid" if the truth of the premises requires that the conclusion be true also. A valid inference is what philosophers call a "truth-preserving" move. In other words, a valid inference is a logical rule which connects the premises to the conclusion so that the conclusion must *necessarily* be true *if* the premises are true.

If an argument's inferences are valid *and* its premises are true, then we call the argument "sound." This is our final goal. A sound argument necessarily entails a true conclusion. No matter how unlikely a conclusion may seem, no matter how much one might wish it were not so, if a conclusion is supported by a sound argument, then it *must* be true. No reasonable person can doubt the conclusion of a sound argument. Sound arguments are the *sine qua non* of good lawyering. They are the wares of our trade.

Before proceeding further, it is worth lingering on a terminological point, which is a frequent cause of sloppy reasoning. It is an improper use of language to characterize an argument as "true" or "false." To characterize an argument as "true" or "false" is a category error. Arguments are not the *kinds* of things that can be true or false. That would be like calling a sound "blue," or a pigment "loud." *Propositions* are the only things that can be true or false. Arguments are not propositions, but rather a way of structuring propositions. Premises can be true or false. Conclusions can be true or false. But an *argument* can only be "valid" (or "invalid") or "sound" (or "unsound").

Some examples may help to make the concepts intuitive.

ARGUMENT A

A1.	Smith killed Jones with malice aforethought.	premise
A2.	If Smith killed Jones with malice aforethought, then Smith is guilty of murder.	premise
A3.	Smith is guilty of murder.	*from* A1, A2

ARGUMENT A is a valid argument. It is valid, because *if* it were the case that

A1 and A2 were true, then it *must necessarily* be the case that A3 is true also. We do not know whether it is a *sound* argument, because we don't know whether the premises are true. For example, it could be the case that Smith did not kill Jones, or that Smith did not have malice aforethought, in which cases possibly Smith would not be guilty of murder.

A word of caution. It is important to recognize that if ARGUMENT A were unsound, then we would not be entitled to infer the contrary claim: that Smith is *not guilty* of murder. If ARGUMENT A were unsound, then Smith may or may not be guilty of murder. We simply don't know (absent further argumentation). A sound argument guarantees that its conclusion is true. An unsound argument cannot make that guarantee. But that does not mean the conclusion is necessarily false. The unsoundness of an argument does not foreclose the possibility that the conclusion could be true *for other reasons*.

Next, a valid argument can have a false conclusion. Consider:

ARGUMENT B

B1.	Socrates is a man.	premise
B2.	If all men are mortal, then Socrates is a man.	premise
B3.	All men are mortal.	*from* B1, B2

Clearly ARGUMENT B is invalid, because the conclusion B3 does not follow from the premises B1 and B2. To see that it is invalid, suppose it were the case that there existed an immortal man somewhere in the world. It could still be true that Socrates is a man (premise B1). And it could still be true that *if* all men were mortal, then Socrates would be a man (premise B2). Therefore, it is *possible* for B1 and B2 to be true, but for B3 to be false. Therefore, the inference is not truth-preserving.

An invalid argument can also have a true conclusion. Consider:

ARGUMENT C

C1.	Fido is a dog.	premise
C2.	Fido likes to play fetch.	premise
C3.	Fido has hair.	*from* C1, C2

Let us assume that C1, C2, and C3 are all true propositions. Notice that despite ARGUMENT C having true premises and a true conclusion, the inference is not valid. Even though the premises are true, and the conclusion is true, the premises C1 and C2 do not *logically force* the conclusion C3. The inference is not valid, therefore the argument is not valid, therefore the argument is unsound.

It is instructive to consider what *would* make an argument like Argument C valid. We would need a different premise to force the conclusion. Something like this would be a valid revision of Argument C:

Argument D

D1.	Fido is a dog.	premise
D2.	If Fido is a dog, then Fido has hair.	premise
D3.	Fido has hair.	*from* D1, D2

Now, assuming Fido were in fact a dog, Argument D would be a sound argument. The premises are all true (by assumption), and the inference is valid (it could not be the case that D1 and D2 were true, but that D3 be false). An argument with true premises and a valid inference is a *sound argument*. Thus, the logic forces us to accept the conclusion that Fido has hair.

Before introducing the substantive inference rules we will be using, it may be useful to reiterate the main points about argument structure exposited thus far:

Summary

- Propositions are claims about states of affairs.
- Propositions can be true or false (but not both).
- An argument is a way of structuring propositions. An argument consists of:
  - Premises
  - Inferences
  - Conclusion
- Premises are propositions.
- A conclusion is a proposition.
- Inferences are rules that relates the truth of the conclusion to the truth of the premises.
- A valid inference guarantees that a conclusion is true if the premises are true.
- An argument is valid if every inference contained in it is valid.
- An argument is sound if it is a valid argument with true premises.
- If an argument is sound, then its conclusion must be true.

2 Modus Ponens

I have thus far spoken about inferences as rules relating premises to conclusions. I will now describe the two main inference rules that we will need for the purposes of this course. The first rule is what logicians call "*modus ponens*." All the argument examples we have considered up to now have already used the rule of *modus ponens*. The structure goes like this:[3]

Modus Ponens Generalized

m1.	x	premise
m2.	If x then y.	premise
m3.	y	*from* m1, m2

In the study and practice of law, all arguments will have this form. Sometimes it will not be superficially obvious how to fit an argument into this form, but with a bit of word-juggling, you should be capable of rephrasing every legal argument into this form. If an "argument" really *cannot* be restated in this form, then it is not really an argument.

The premises required for *modus ponens* can be further categorized into two types: (i) fact premises and (ii) legal premises. Fact premises are propositions that make claims about the world. For example, it is a fact premise *that Smith killed Jones*. It is a fact premise *that Smith possessed a kilogram of cocaine*. It is a fact premise *that Smith's blood alcohol level was 0.2*. These are all examples of fact premises.

Note that a fact premise is not necessarily a fact. A fact premise is a proposition that claims some possible fact to be the case. It could be a true proposition or a false proposition. It is up to the trial courts to determine whether a fact premise is true (i.e., whether it *is* a fact).

Legal premises have the form: *if* [fact] *then* [outcome]. Legal premises state "laws." For example, it is a legal premise *that if Smith kills Jones with malice aforethought, then Smith is guilty of murder*. It is a legal premise *that if Smith possessed a kilogram of cocaine, then he intended to distribute cocaine*.

Note that a legal premise is not necessarily the law. A legal premise is a proposition that claims what the law is. It is up to appeals courts to determine whether a legal premise is true (i.e., whether it *is* the law).

It is critically important to recognize the conceptual distinction between fact premises and legal premises. I cannot overstate the importance of this point. The division between fact premises and legal premises is crucial to understanding the

[3]Let x and y be variables denoting propositions.

entire structure of our legal system. Indeed, we do not even use the same forums to litigate these different categories of premises: trial courts adjudicate disputes about fact premises; appeals courts adjudicate disputes about legal premises.

All arguments using *modus ponens* must involve both fact premises and legal premises. However, an argument can have multiple fact premises and multiple legal premises. For example, to prove that Smith murdered Jones, we typically require *two* fact premises and *one* legal premise:

ARGUMENT E

E1.	Smith killed Jones.	fact premise
E2.	Smith had malice aforethought.	fact premise
E3.	If Smith killed Jones, and Smith had malice aforethought, then Smith is guilty of murder.	legal premise
E4.	Smith is guilty of murder.	*from* E1, E2, E3

E1–E3 entail E4. To make this point explicit, we can extend the diagram of *modus ponens* to represent indefinitely many conjunctions thusly:

Modus Ponens EXTENDED

M1.	x	premise
M2.	x'	premise
M3.	x''	premise
...	...	...
Mn.	x^n	premise
M$n+1$.	If x and x' and x'' and ... and x^n then y.	premise
M$n+2$.	y	*from* M1, M2, M3, ..., Mn, M$n+1$

Here, premises M1–Mn are fact premises, and premise Mn+1 is the legal premise. In a criminal statute, the fact premises are usually referred to as "elements." When we talk about the *elements of an offense*, we are referring to the fact premises required for the argument to be valid.

There are additional complications that can arise when we extend *modus ponens* to involve more logically complicated premises (for example, using logical connectives like "or" or "neither"). We will observe how to deal with this logical complexity as the course progresses. For now, it suffices to understand the structure of *modus ponens* in its most generalized form.

3 Universal Instantiation

The second inference rule we will need for the purposes of this course is what logicians call "universal instantiation." Laws are not usually stated in specific terms. For example, you will not find it stated in any criminal code *that if Smith kills Jones with malice aforethought, then Smith is guilty of murder.* Legal codes do not typically reference specific individuals like "Smith" or "Jones." Rather, laws are typically stated in a general form. For example, *for any x and y, if x is a person, and y is a person, and x kills y with malice aforethought, then x is guilty of murder.*

The structure of universal instantiation goes like this:[4]

Universal Instantiation

U1.	For any given x, $P(x)$	premise
U2.	$P(a)$	*from* U1

The rule of universal instantiation means that a general statement applies to particulars. For example:

Argument F

F1.	For any given x, if x is a person and x possesses more than a kilogram of cocaine, then x intends to distribute cocaine.	premise
F2.	If Smith is a person and Smith possesses more than a kilogram of cocaine, then Smith intends to distribute cocaine.	*from* F1

F1 entails F2. In most instances, the use of universal instantiation is obvious and unstated. However, to be clear and rigorous lawyers, we must be explicitly aware of the logical moves we are making. It is important to make this inference rule explicit, because (as we will discover later in this course) some cases do hinge upon the validity of this move.

4 Arguments, Counterarguments, and Nesting

If we understand the structure of arguments using *modus ponens* and universal instantiation, then we have all the logical machinery we will ever need to make legal arguments. It is therefore of utmost concern that you acquire a deep familiarity with

[4]Let $P(x)$ denote a proposition which contains the variable term x. Let a denote a particular.

these logical structures. The preceding sections conveyed a great deal of content in a short space. It is worth re-reading them until you feel very comfortable with the concepts that were introduced.

In this final section, I will briefly touch upon some things we can do with arguments. You have learned that the truth of an argument's conclusion depends upon the truth of its premises (and the validity of its inferences). You may have wondered though how we come to know that a premise is true in the first place. In the argument examples above, we simply *assumed* some of the premises were true. But we cannot always rely on assumptions. Some premises will be disputed.

When the truth of a premise is disputed, then it must be supported by *more argument*. We do this by nesting an argument within another argument—like Russian nesting dolls. For example, suppose we have the following argument:

ARGUMENT G

G1.	Fido is a mammal.	premise
G2.	For any x, if x is a mammal, then x has hair.	premise
G3.	If Fido is a mammal, then Fido has hair.	*from* G2
G4.	Fido has hair.	*from* G1, G3

This simple argument uses both of the inference rules we learned in the previous sections. We infer G3 from G2 using the rule of universal instantiation. And we infer the conclusion G4 from G2 and G3 using *modus ponens*. We know that *modus ponens* and universal instantiation are valid inferences. Therefore, ARGUMENT G is a valid argument.

But now suppose that someone disputed our premise that Fido is a mammal. What are we to do? Our clients would be quite justifiably outraged if we simply shrugged our shoulders and declared our readiness to "agree to disagree." Rather, we should construct *another* argument to convince the skeptic of the truth of G1.

ARGUMENT H

H1.	For any x, if x is a dog, then x is a mammal.	premise
H2.	If Fido is a dog, then Fido is a mammal.	*from* H1
H3.	Fido is a dog.	premise
G1.	Fido is a mammal.	*from* H2, H3
G2.	For any x, if x is a mammal, then x has hair.	premise
G3.	If Fido is a mammal, then Fido has hair.	*from* G2
G4.	Fido has hair.	*from* G1, G3

First, we infer H2 from H1 using the rule of universal instantiation. Then we use H1 and H3 to prove the disputed premise G1 using *modus ponens*. Thus, our naked premise *that Fido is a mammal* is now supported by a further argument. And what if our opponent disputed whether premise H2 were true? Simple. We provide yet another argument to prove H2. And we continue nesting arguments to prove the premises of premises until we arrive at premises that are undisputed.

Clearly, we can continue nesting *ad infinitum*. This infinite nesting is a source of frustration for philosophers seeking to ground knowledge in some unobtainable "absolute truth." But for lawyers, the infinitely nested structure of knowledge is a great boon. It means that we can never run out of arguments. We can always dispute our opponent's premises, and the premises which support their premises, and the premises which support the premises which support their premises, and so forth. The only constraint limiting how deeply we argue is how far our clients are willing to pay us to dig.

Finally, a precise understanding of logic is useful not only in constructing sound arguments, but also in understanding how to dismantle our opponent's unsound arguments. When our opponent presents an adverse argument, using what we have learned, we are now well equipped to rebut it.

At the level of utmost generality, we understand that our task is to reveal how our opponent's argument is unsound. The first step is to analyze their words into the formal structures we learned in this introduction. It is worth repeating that this is not always a trivial task, and it may involve some work to get our opponent's words to "fit" into the structures we have studied. It is challenging work. But it gets easier with practice. When your opponent purports to be giving an argument, you must always be looking to identify the fact premises, the legal premises, and the inference.

If the inference is invalid, then our work is easy. We simply point out that the conclusion is not a necessary consequence of the premises, and our opponent's argument is instantly refuted.

However, if the inference is valid, then we must next look to the truth of the premises. There is a common lawyer's adage: "If the facts are against you, then argue the law. If the law is against you, then argue the facts." If our opponent's argument is valid, then we are left to dispute their fact premises or their legal premises (or both). You can dispute fact premises using a variety of evidentiary tools like eyewitnesses, expert testimony, documents, and recordings. You can dispute legal premises using a variety of economic, historical, philosophical, and institutional theories.

Yet these empirical and theoretical tools are never the end. Lest you be blindsided by your opponent's clever counterarguments (or overlook clever counterarguments available to your side), always remember that these resources are themselves merely

the components of further argumentation. You can *always* contest the premises of any argument.[5] For example, suppose your opponent provides an eyewitness purporting to having seen your client Smith killing Jones. Embedded in their presentation of their witness is the implicit argument:

ARGUMENT I

I1.	The witness says he observed Smith killing Jones.	premise
I2.	If the witness says he observed Smith killing Jones, then Smith killed Jones.	premise
I3.	Smith killed Jones.	*from* I1, I2

Premises I1 and I2 are both challengeable. Is the witness *really* claiming that he observed Smith killing Jones, or did he merely perceive something that *seemed* like Smith killing Jones? Is the witness credible? It does not follow that an event occurred simply because someone claims to have observed it. Make your opponent provide more arguments to support I1 and I2. Construct your own arguments to prove that I1 or I2 are false.

Another example: suppose your opponent interprets the term "malice aforethought" as being equivalent to *intention*. Embedded in this interpretation is the argument:

ARGUMENT J

J1.	If "malice aforethought" is an element of murder, then *intent* is an element of murder.	premise
J2.	"Malice aforethought" is an element of murder.	premise
J3.	*Intent* is an element of murder.	*from* J1, J2

Again, premises J1 and J2 are both challengeable. Does the statute say "with malice aforethought"? When was the statute enacted? How was the statute enacted? Was it merely a codification of common law or was it given full consideration by the legislature? Did the legislature intend "malice aforethought" to be equivalent to *intent*? What effect would equating "malice aforethought" with *intent* have on incentives? What effect would equating "malice aforethought" with *intent* have on the fairness of outcomes? Construct more arguments!

[5]It must be conceded that in some instances, the nested arguments can get to the point where things start to sound a bit ridiculous. However, for practical purposes it is almost always the case that you should argue more rather than stopping too early.