Types of arguments
Kareem Khalifa

CONTENTS

General rules of thumb about short arguments

1. Distinguish premises from conclusions
    a. Whether you're composing an argument (writing, speaking) or trying to understand one (reading, listening), it's best to start by identifying the conclusion.
        i. In the act of composing this will help you from going off on tangents, and will thus help you to be focused, relevant, and where necessary, concise and succinct. Note that you do not actually have to assert your conclusion immediately (though this is often very effective), but you should at least know what your conclusion is before you start thinking of information that could be used as premises.
        ii. In the understanding an argument, it's best to start with the conclusion because while there can be an infinite number of premises in an argument, there can be only one conclusion. Each conclusion marks a new argument.
2. Present your ideas in a natural order.
    a. Generally, the conclusion should come first or it should come last.
    b. Short arguments should be no longer than two paragraphs (approximately a page).
3. Start from reliable premises.
    a. Recall that Nolt concluded our previous readings by identifying the truth of the premises as one of two conditions of a good argument.
    b. Reliable premises are just those that are likely to be true.

        i. What makes a premise reliable depends partly on what can be taken for granted by your audience. Who is a trusted authority? How large is the pool of common knowledge and stock examples?

4. Be concrete and concise.
    a. There is a common misconception that a deep understanding of something involves abstract, lengthy, complicated, and generally opaque language. I would argue just the opposite: the true mark that you have a deep understanding of something is that you can communicate a complicated idea in very simple terms.
5. Avoid loaded language.
    a. Once again this is a good rule both when you're trying to compose or understand an argument.
        i. If you're composing an argument, people may think that you're simply engaging in name-calling or polemics rather than offering a genuine argument.
        ii. If you start understanding the person's argument in loaded terms, you're liable to misrepresent his/her views.
            1. For example, check out this Sean Hannity and Ron Paul exchange (about 2:46 into the video clip): http://youtube.com/watch?v=xEZO7MPxJIs
6. Use consistent language.
    a. Recall that one thing we sought was a tight connection between premises and conclusion. Using consistent terminology helps to tighten up these connections. For example it is far harder to follow this argument:

        Playing in the mud makes for dirty clothes. Kareem gamboled in the loam. Therefore, the vestments of the son of the German teacher at Onondaga Central High School soon will suffer besmirching.

        then it is to follow this one:

        If you play in the mud, your clothes will get dirty. Kareem played in the mud. Therefore, Kareem's clothes will get dirty.
7. Stick to one meaning for each term.
    a. Using consistent language forges a tight connection between the premises and the conclusion. However, you can use consistent language to forge an illusory connection between the premises and the conclusion. This happens when you equivocate on meanings of a term throughout an argument.

Deductive arguments

The big gap in the previous reading was what it meant for premises to "support" their conclusion in an argument. What we want is for the premises to provide some "guarantee" of the truth of the conclusion. The strongest guarantee is that it is impossible that the premises are true and the conclusion is false. Any argument having this property is called a *deductively valid* argument.

DEDUCTIVE VALIDITY IS THE SINGLE MOST IMPORTANT CONCEPT IN THIS COURSE.

Memorize the following definition verbatim:

> **<u>A deductive argument is valid when, if its premises are true, its conclusion *must* be true.</u>**

Failure to memorize this will result in disastrous consequences for you (-5 extra points on any assignment where it comes up, plus general cluelessness throughout much of the semester).

Let's look at a valid argument to see how this works:

> All rabbits are mammals.
> Bugs Bunny is a rabbit.
> ∴ Bugs Bunny is a mammal.

When we say that it is *impossible* for the premises of this argument to be true and the conclusion false, we mean that a denying the conclusion will yield a <u>contradiction</u> in the premises. For example, suppose that Bugs is not a mammal. Then, by the first premise, he is not a rabbit. But this contradicts the second premise that he is a rabbit.

This kind of impossibility is called *logical impossibility*. There are other notions of impossibility, such as *physical impossibility*, that do not figure in assessments of validity. For example, the following argument makes assumptions that are physically impossible but not logically impossible:

> All rabbits are reptiles.
> Daffy Duck is a rabbit.
> ∴ Daffy Duck is a reptile.

This is a valid argument, even though its premises and conclusion are false, because there is no way that the conclusion can be false without contradicting one of the premises.

Examples like this may make you stop and wonder what's so great about validity. However, recall from Nolt that the two marks of a good argument are that its premises are true and that they are valid. In this case, the argument is lacking true premises. However it provides a tight connection between the premises and conclusion—indeed, we'll see that we can get a very precise notion of "tight connection" using the concept of deductive validity.

Note that the definition of validity does not apply to arguments that simply have true premises and a true conclusion. You might get lucky or string together true but unconnected statements. For example,

> 2+2 =4
> All rabbits are mammals.
> ∴ Tony Blair was the prime minister of the United Kingdom.

All of the propositions in the argument are true but the argument is not valid. This is because the conclusion *can* be false and the premises true, i.e., it is not a logical impossibility. It could be the case that Tony Blair never became prime minister of the UK, but this would have little effect on basic truths of arithmetic and zoology. Some arguments can even look like they have a tight connection, have all true premises, but aren't valid, e.g.,

> If it's raining, then the streets are wet.
> It's not raining.
> ∴ So the streets aren't wet.

This is <u>invalid.</u> Why? Suppose that someone opened up a fire hydrant. In this case, it would still be true that *if* it's raining then the streets are wet, it's not raining, but the streets are wet.

INVALID. Thus, it's possible that all of the premises are and the conclusion is false. We'll develop this idea more as the course progresses.

Types of deductive arguments

While Weston's presentation of the material might make you think that there he's providing *rules* for deductive arguments, his Rules 24-29 are really *types* of deductive arguments. Of course one rule you should follow is to use these types of arguments! It should be noted that the types of deductive arguments he cites are among the most pervasive in propositional logic.

Modus ponens

By far the most frequently used deductive argument, this runs as follows:

If p, then q.
p.
$\therefore$ q.

For example:

If it's raining, then the streets are wet.
It's raining.
$\therefore$ The streets are wet.

In case this isn't clear to you, keep the following two points in mind:

a. An argument is invalid if it can be shown that its premises can be true at the same time its conclusion is false.
b. As we discussed briefly in discussion sections, a hypothetical statement is false only if its antecedent is true and its consequent is false.

Now try your hardest to invalidate the argument:

1) By a, the conclusion would need to be false, so let's assume that the streets are NOT wet.
2) However, you still need both the premises to be true in order to invalidate the argument.
3) So it must NOT be raining; otherwise the hypothetical statement "If it's raining, then the streets are wet," will have a true antecedent and a false consequent, which by b, would make it false, and by a, would preclude the possibility of invalidating the argument.
4) However if it is NOT raining, then this contradicts the second premise.
5) Thus, it is impossible to have both premises of this argument be true while the conclusion is false.

Modus tollens

Alternatively, we might have a very similar inference:

If it's raining, then the streets are wet.
The streets are not wet.
$\therefore$ It's not raining.

More generically,

If p then q.
Not q.

∴ Not *p*.

A similar chain of reasoning as the previous section on modus ponens shows why modus tollens is a valid form of inference.

Hypothetical syllogism

This is a slightly different form of inference. Its form is:

If *p* then *q*.
If *q* then *r*.
∴ If *p* then *r*.

For example:

If Fido is a dog, then Fido is a mammal.
If Fido is a mammal, then Fido is warm-blooded.
∴ If Fido is a dog, then Fido is warm-blooded.

We can use principles a and b from our "Modus Ponens" section (yet again!) to show that this argument is valid. Try your best to invalidate it:

1) The conclusion must be false: So "If Fido is a dog, then Fido is warm-blooded" is false, which happens only if Fido is a dog and Fido is NOT warm-blooded.
2) However all of the premises must be true.
3) Since the falsity of the conclusion requires that Fido is a dog, then the first premise is only true if Fido is also a mammal.
   i. If Fido were a dog and not a mammal, then, for the hypothetical, "If Fido is a dog, then Fido is a mammal," the antecedent would be true and the consequent false, thus falsifying the whole first premise.
4) However, if Fido is a mammal, then for analogous reasons, Fido IS warm-blooded.
5) But this contradicts what we required in step 1.

Thus, it's impossible to invalidate this argument and any other Hypothetical Syllogism.

Disjunctive Syllogism

The form of this inference is:

Either *p* or *q*.
Not *p*.
∴ *q*.

For example,

Either you eat fries with your sandwich or you eat a salad with your sandwich.
You do not eat fries with your sandwich.
∴ So you eat a salad with your sandwich.

This requires a slightly different proof for validity. As before, we need the following:

a. An argument is invalid if it can be shown that its premises can be true at the same time its conclusion is false.

However, we now need to recall the following:

b. A disjunctive "or" statement is false if and only if both of its disjuncts are false.

Now try your best to invalidate a disjunctive syllogism:

1. By a, the conclusion, "You eat salad with your sandwich" must be false.
2. However, if you do not eat salad with your sandwich, then by b, you eat fries with your sandwich.
3. However, this contradicts the second premise that you do not eat fries with your sandwich.

Thus, it is logically impossible for the premises of this argument to be true while the conclusion is false. The argument is VALID.

Another point: There are two senses of "or" in English: exclusive or inclusive. *Exclusive* disjunctions are like what you find with sides on a menu (as in the example above): either *p* or *q* but not both. *Inclusive* disjunctions are ones in which both disjuncts can be true at the same time. It's shorthand for *p* or *q* or both *p* and *q*. For example, my Advanced Logic course requires that you have taken Introduction to Logic or that you have a strong background in mathematics. Clearly, a person who has taken Intro to Logic and has a strong math background is still allowed to take the course.

Dilemma

A dilemma, or "constructive dilemma" as it's sometimes called, is of the following form:

> Either *p* or *q*.
> If *p* then *r*.
> If *q* then *s*.
> ∴ Either *r* or *s*.

For example,

> You can take either the blue pill or the red pill.
> If you take the blue pill, you will become taller.
> If you take the red pill, you will become shorter.
> ∴ You will become taller or shorter.

Reductio ad absurdum

Technically, the form is:

> If not *p*, then *q* and ~*q*.
> ∴ *p*.

Implicitly the form is:

> If *p*, then *q* and ~*q*.
> It can never be the case that both *q* and ~*q*. (Because this is a contradiction).
> ∴ *p*.

Weston then highlights the procedure/reasoning lurking beneath this proof:

> To prove: *p*
> Assume the opposite: Not-*p*.
> Argue from that assumption we'd have to conclude: *q*
> Show that *q* is false (silly, contradictory, absurd)
> Conclude: *p* must be true after all.

For example,

> If $2+2 \neq 4$, then $4 - 2 \neq 2$ and yet, $4 - 2 = 2$.
> ∴ $2 + 2 = 4$.

Inductive arguments

Due to both practical and conceptual constraints, the premises of an argument may support the premises in a manner slightly weaker than that required by deductive validity. For example, the following is a generally regarded as a good inference though it is not deductively valid:

My parents have told me that my name is Kareem Khalifa.
∴ My name is Kareem Khalifa.

However, the conclusion can be false and the premises true. For example, suppose my parents are deceitful people and really named me Bobo Laughingstock, forging documents at every possible turn to hide my true name from me. Thus, we might distinguish two types of arguments. The first, deductively valid arguments are of the form:

Premise 1
Premise 2
…
Premise n.
∴ (Necessarily) Conclusion

The second, what we call *inductively strong* arguments are of the form:

Premise 1
Premise 2
…
Premise n.
∴ (Probably) Conclusion

In other words, whereas deductively validity is a property of argument such that the truth of the premises makes the conclusion necessarily true, inductive strength is a property of an argument such that the truth of the premises makes the conclusion highly probable. For example, my inferring that my name is Kareem Khalifa on the basis of my parents telling me that my name is Kareem Khalifa is more amenable to an inductively strong argument. Note that unlike necessity, probability comes in *degrees*. Specifically, probability can range from 0 to 100%. A deductively valid argument always yields a 100% probable conclusion on the basis of its premises[1].

There are other interesting differences between deductive and inductive arguments. For example, deductive arguments are *monotonic* while inductive arguments are *non-monotonic* or *defeasible*. The acceptability of a conclusion of a monotonic argument cannot be affected by the introduction of additional information; for a non-monotonic argument it can. For example consider the following inductive argument:

Jim is bright and studies hard.
∴ Jim will probably get an A in Khalifa's Logic course.

This may be an inductively strong inference, but suppose after finding out that Jim is bright and studies hard, we subsequently discover the following:

Jim is in the third grade.

---

[1] For you stats nerds, a deductively valid argument will have the following property: P(Conclusion | Premises$_{1\ldots n}$) = 1.

This will lower the probability of Jim getting an A in Khalifa's Logic course. However, suppose we then find out the following:

Jim is a boy genius who could solve complex differential equations at age 5.

This will then raise the probability of Jim getting an A in Khalifa's Logic course, etc. In contrast, let's go back to our deductively valid Bugs Bunny argument.

All rabbits are mammals.
Bugs Bunny is a rabbit.
∴ Bugs Bunny is a mammal.

Nothing you can add to this argument will alter the validity of it. Try your best! (We'll see that there are some interesting consequences of this in the coming weeks).

Types of inductive argument

There is a wide variety of inductive arguments. Some of the most frequently used inductive arguments include the following:

Arguments by example

These arguments are typically of the following form.

Some (representative) *F's* are *G's*.
∴ (Probably) Most *F's* are *G's*.

There are several versions of this. One of the most common arguments by example is called *enumerative induction*. This has the following form:

All observed *F's* are *G's*.
∴ (Probably) All *F's* are *G's*.

Here is another:

This (random) *F* is a *G*.
∴ (Probably) All *F's* are *G's*.

In all of these cases, the biggest danger arises from choosing unrepresentative examples. As we saw, inductive arguments are defeasible or non-monotonic: the unrepresentative example is the big "defeater" for an argument by example. For instance, the first argument by example only works if the *F's* in the premise aren't anomalies. For example,

Some people in Vermont earn over $100,000
∴ Most people in Vermont earn over $100,000.

This might work if you looked, e.g., only at people in the Administration at Middlebury, but note that the median income (per capita) in Vermont is less than $21,000! You can easily extrapolate this to the other kinds of arguments by example.

Thus, if you argue by example, you should try to safeguard against the unrepresentative example. If you are encountering an argument by example, you should make sure the person offering that argument has taken the appropriate safeguards. These include:

8. Giving more than one example.

A single example is liable to look like something the arguer "cherry-picked" to support his/her conclusion. It could be that a majority of the relevant cases out there provide negative support for the conclusion, but the one cited is an anomaly, i.e., is unrepresentative.

The smaller the set of relevant cases, the more imperative it is to consider all of the cases, as each case makes a huge difference in a small set. For example, compare the following arguments.

> Khalifa has a PhD.
> ∴ Most philosophy professors at Middlebury have PhD's.

> Khalifa has a PhD.
> ∴ Most professors at Middlebury have PhD's.

Now imagine that there are three professors in the philosophy department that do not have PhD's[2]. Clearly, the first argument is undermined to a far greater extent by this evidence than the second argument, since it is clearly possible that the philosophy department does not have many PhD's while the rest of the College's departments do.

9. Use representative examples

As we said, this is the greatest danger of arguments by example. So what makes an example representative?

Here's a very useful rule of thumb: when arguing by example, use at least two examples: one that is as typical an example as you can find, and then choose at least one example that is *least likely* to prove your conclusion and show that, despite the odds, it does. Often this second kind of example can be prefaced with the word "Even." For example,

> The median income in Vermont is less than the United States. As might be expected, much of this has to do with the large rural areas in Vermont, e.g., Orleans county which is 82% rural, has a median household income of $35,000. <u>Even</u> in Burlington, which one might expect a deviation in state income numbers, the median household income is $41,000, less than the $48,000 national median.

Generally look for a common factor that all of the examples have that might give the arguer false confidence in his/her conclusion. For example:

> In a recent survey, 70% people in New York City said they would vote for Hillary Clinton for president.
> ∴ Hillary Clinton will probably win the presidency.

Of course, New York City's constituency is a Democratic stronghold, and thus not a strong barometer for how the election will go. Better to look at a "swing" state that is not consistently red or blue, such as Florida, Pennsylvania, or Ohio. Of course, if a traditionally red state votes Democrat, this is very salient evidence—subject to an "Even" argument.

10. Background information is crucial

As we saw, inductive arguments are defeasible or non-monotonic. Thus, additional background information is crucial.

---

[2] Incidentally, this is false. ☺

It's generally a good idea to consider what kind of evidence might undermine an otherwise inductively strong argument. The most important question to ask when thinking about these sorts of things is, "What is the larger population being referenced?" (This is called the *reference class.*)

For example, the assertion that 30% of people have blue eyes might be true if the reference class is Europe, but is certainly false if the reference class is the world (only about 5-8% of the world population has blue eyes, as birth rates in Europe are much lower than Asia, Africa, and South America—where there is an overwhelming majority of brown-eyed people). Thus, you should make sure that the reference class doesn't bias the inference.

11. Consider counterexamples

Arguments example typically involve an extrapolation from some examples to a more general statement. Any statement of the form "All *F's* are *G's*" is making a very strong claim about *F's*: a single *F* that is NOT a *G* proves it false. This is called a *counterexample* to the generalization that all *F's* are *G's*.

Whenever you argue from examples, you should be mindful of counterexamples. This is part of a more general practice in critical thinking that I call *reverse role-playing* or *adversarial thinking*. You'll hear me say this more than once: Imagine that your worst enemy wrote your paper, and endeavor to humiliate him/her.

Considering counterexamples might cause you to *reject* a general claim. Alternatively, it might cause you to *qualify* a general claim. For example:

> Generalization: No mammals lay eggs.
> Counterexample: Platypuses and echidnae, members of the monotreme family, lay eggs.
> Revised generalization: No non-monotreme mammals lay eggs.

Other times you might hold fast to the generalization and *reinterpret the counterexample* as consistent with the generalization. For example:

> Generalization: All professional athletes are overpaid.
> Alleged counterexample: Minor league athletes earn very little money.
> Reinterpreted counterexample: Even a minor league athlete earns more than the average American.

Note that we use the "Even" argument here too!

Arguments by Analogy

An argument by analogy runs as follows:

> *a* is an *X*.
> Both *a* and *b* are similar with respect to *F*, *G, H, …*
> $\therefore$ So *b* is (probably) an *X*.

In this argument, *a* and *b* are often called *analogues*. *F, G, H* are *points of similarity*, and *X* the *target*. For example:

> Fido likes Alpo.
> Both Fido and Spot are dogs.

∴ So Spot (probably) likes Alpo.

The chief danger with arguments by analogy is that nearly anything can be analogous with anything else *in some respect*. For example,

Fido likes Alpo.

Both Fido and my desk have four legs.

∴ My desk (probably) likes Alpo.

To avoid bad arguments by analogy, a *relevant analogy* is required in the second premise.

12. Analogy requires relevantly similar analogues

As we saw with arguments by example, different kinds of inductive arguments have different kinds of "defeaters." A typical defeater of an argument by analogy highlights that the analogues being compared are different in respects that are more relevant than their similarities.

So how can you discern if two analogues are relevantly similar? Here is a useful heuristic to follow:

a. In general, try to be as precise as possible about the points of similarity (what was labeled as *F, G, H*…in the schema above) and the target (X above).
b. Once you've figured out exactly what *F, G, H, etc.* are, ask yourself whether things having *F, G, H, etc.* necessarily or typically have the target *X* as well.
c. If so, then the analogy works; if not, it needs to be revised or rejected.

Let's return to our examples of good and bad arguments by analogy to see how this three-step process works. In the Fido-Spot example, we see that (a) the point of similarity is that both Fido and Spot are dogs and the target is that they like Alpo; (b) since dogs typically like Alpo; (c) this is relevant similarity.

In contrast, our Fido-desk example (a) the point of similarity is that both Fido and my desk have four legs and the target is that they both like Alpo; (b) however, many four-legged things don't like Alpo (furniture, other animals, etc.); (c) so the analogy is not relevant.

Arguments from Authority

Recall the example from above:

My parents have told me that my name is Kareem Khalifa.

∴ My name is Kareem Khalifa.

The example of inferring my name from what my parents' tell me is called an *argument from authority* or an *argument from testimony*. Expert testimony, what you read in textbooks and newspapers, even readings from instruments are all inferences of this sort. Their general form is:

Some source asserts or indicates that *p*.

∴ (Probably) *p*.

The defeater to these kinds of arguments is that the source is *unreliable* about making judgments that *p*. For example:

Amy Winehouse says that rehab clinics do not effectively treat addiction.

∴ (Probably) Rehab clinics do not effectively treat addiction.

13. Sources should be cited.

Knowing *who* or *what* a source is essential to evaluating the reliability of he/she/its testimony. With the Winehouse example, one could cite the following evidence:

> Amy Winehouse is a pop singer with no medical expertise on drug addiction.

You'll often see statistics cited without a source or hear an argument from authority start with something like "8 out of 10 experts say…" Ask yourself who is providing these stats or who these "experts" are,

14. Seek informed sources

Typically, a more informed source is a more reliable source. All else being equal, it would be better to get advice about physics from Susan Watson or Jeff Dunham than from a freshman in Intro to Physics.

Importantly, being an informed source about one thing doesn't make you an expert about something else. For example, a now notorious commercial begins with "I'm not a doctor, but I play one on TV." Now, if the person were then to proceed to talk about the acting techniques involved in playing a doctor in a television show, he would be a reliable source. However, the commercial proceeds to tell you about the value of Vick's Cough Syrup as a remedy for colds. Is a television actor really qualified to tell you about these things?

In writing, your audience determines how much of a person's credentials you get to take for granted. Academic discourse has developed a sophisticated system of citation that makes this a bit easier than in everyday discussion.

15. Seek impartial sources.

You may know a source's name, and he/she/it may seem quite informed, but still not be reliable because of certain biases. For example,

> A Harvard-Smithsonian astrophysicist asserts that the average global temperature did not increase during the twentieth century.
>
> ∴ (Probably) the average global temperature did not increase during the twentieth century.

However, here is an important piece of background information:

> The Harvard-Smithsonian astrophysicist is Willie Soon, who is funded by the American Petroleum Institute, a lobbying group for the oil industry.

Clearly this changes the reliability of the testimony.

16. Cross-check sources.

In this day and age, many professional experts are funded by groups with particular interests. This does not automatically make them biased or unreliable. One way to check these things is to see how other experts, who are not funded by someone with these interests (or who are funded by different interests), evaluate the claims of the expert cited. If there is general consensus in an expert community that the source is reliable or that the claim is probable, this provides additional support for the claim.

Certain claims are not amenable to arguments by authority, and cross-checking is a good litmus test: if there is no consensus on a claim, then probably citing one person's testimony on the matter will be unconvincing.

17. Personal attacks do not disqualify a source.

The most reliable ways of correctly assessing a source's reliability are to examine how informed and impartial he/she/it is and to cross-check this with the broader expert community of which he/she/it is a part. Anything else might easily slide into a personal attack or *ad hominem argument*. For example:

> Linnaeus claimed that every organism can be classified according to its genus and species.
>
> ∴ (Probably) every organism can be classified according to its genus and species.

However, suppose the following information is revealed:

> Linnaeus was a creationist.

Should we infer from this that it is less probable that organisms can be classified by genus and species? Note that being a creationist has profound effects on a person's attitude toward biology! However, look at how remarkably well Linnaeus's classification system does when you cross-check it: all contemporary biologists, who are generally not creationists, accept his classificatory scheme, suggesting that it is impartial and informed. Furthermore, it isn't quite clear how this classification is made more/less probable by being a creationist. For example, Linnaeus might have held that God, in His infinite wisdom, made species amenable to such classification to further our understanding.

Causal arguments

Causal arguments take many forms (see Appendix). Perhaps the most crucial one is:

> (Under appropriate, controlled conditions,) there is a strong correlation between X and Y.
>
> X precedes Y in time.
>
> ∴ X (probably) causes Y.

For example,

> Smoking correlates with lung cancer.
>
> Smoking precedes lung cancer, i.e., a person who smokes is more likely to get lung cancer later.
>
> ∴ Smoking causes lung cancer.

The big defeater for causal arguments is called the *spurious correlation*. Certain correlations don't entail causation. For example:

> Yellow-stained teeth correlate with lung cancer.
>
> Yellow-stained teeth precede lung cancer.
>
> ∴ Yellow-stained teeth cause lung cancer.

Clearly something is amiss here. Most spurious correlations can be traced back to whether the correlation obtains "under the appropriate, controlled conditions" in the first premise of the causal argument schema at the beginning of this section. To get a sense of the appropriate conditions under which a correlation entitles you to infer a causal claim, imagine how you might construct a machine that produces the effect. You'll find that the "appropriate, controlled conditions" are what allow the machine to work.

18. Explain how cause leads to effect

One reason to think that yellow teeth don't cause lung cancer is that there's no plausible answer to the question, "How do yellow teeth cause lung cancer?" Generally, these kinds of questions require you to cite a *mechanism* that shows how certain aspects of the putative cause interact together to yield the effect. For example, when doctors claim that smoking causes cancer, they cite the fact that smoking (the cause) puts tar in the lungs, tar chronically irritates the lungs, stimulating the growth of cancerous cells (the effect).

To use our "machine" exercise, if you needed to build a "cancer-making machine" (say as a biological weapon for the Department of Defense), you'd need to have a good blueprint or schematic for how to build that machine, so you'd want to specify all of the parts and their interactions.

19. Propose the most likely cause

Of course, simply concocting a mechanism is just a function of how clever and imaginative you are. For example, an advocate of the hypothesis that yellow stained teeth cause cancer might fabricate the following mechanism. The enamel of our teeth is color sensitive, and creates carcinogens when colored yellow.

To return to our machine exercise, we don't like this mechanism because it's not likely to produce the desired effect (cancer). In contrast, the "tar" hypothesis is well supported by empirical observation, our best-accepted medical hypotheses, etc.

Sometimes no cause is all that likely. In this case, it's probably a good idea to refrain from making causal inferences, and make more modest claims about correlation. For example, Steve Levitt, a famous economist has noticed a correlation between the legalization of abortions and the lowering of crime rates in the subsequent decades (Levitt and Dubner 2005). Many people, including Levitt, proposed the following mechanism: potential mothers aborted unwanted children, unwanted children are more prone to crime, so potential mothers aborted most of the people who would currently be criminals. The problem of course, is that it is very difficult to establish many parts of this mechanism, e.g.,

a. A child who is initially unwanted by his/her mother during pregnancy may not be unwanted once he/she is actually born.
b. It is not clear that being unwanted correlates (and thus causes) people to be criminals. Many rich suburban kids are unwanted and do not become criminals.
c. The women with the highest abortion rates are not representative of mothers most likely to give birth to/raise criminals.

These are all parts of Levitt's "crime-reduction machine" that don't seem to work too smoothly. All of this is to say that there may be an interesting correlation between legalizing abortion and the drop in crime rates, but perhaps it would be hasty to infer that one causes the other.

20. Correlated events are not necessarily causally related

One possible account of the correlation between abortion and crime is that it is merely *coincidental:* for example, in America, there are all sorts of economic, political, and social changes that happened in between the legalization of *Roe v. Wade* and the drop in crime during the 1990's that might "screen off" the abortion hypothesis as the cause of the crime drop.

Here's another example that nicely illustrates a coincidental correlation (heehee):

> John Jones's taking birth control pills correlates with his not getting pregnant.
> ∴ John Jones's taking birth control pills (probably) correlates with his not getting pregnant.

Of course, assuming John is a man, we all know why he's not getting pregnant. ☺

In terms of our machine exercise, the key thing to note here is that certain correlations are like extraneous parts. For example, just as a car will function just the same regardless of its exterior paint job, John will remain non-pregnant regardless of his consumption of birth control pills.

21. Correlated events may have a common cause

Sometimes, two events are correlated because they are both effects of a common cause. For example, smoking is a common cause of yellow-stained teeth and lung cancer, and thus correlates with both, but if you control for smoking, e.g., by looking at NON-smokers who happen to have yellow-stained teeth (e.g., from a diet heavy in turmeric or coffee), then no correlation obtains between yellow teeth and lung cancer.

In terms of our machine exercise, you can see how the common cause is doing the "real work" of the machine. You might think about this in terms of switches. Suppose that you have a "Smoking" switch and a "Yellow Teeth" switch. If the Smoking switch is turned off and the Yellow Teeth switch is turned on, there will be no cancer. If, in contrast, the Smoking switch is turned on, then regardless of whether the Yellow Teeth switch is on or off, you'll get cancer.

22. Either of two correlated events may cause the other

Sometimes, we don't know which event came first. For example, there is a correlation between depression and alcoholism. You might think that depressed people are more inclined to "turn to the bottle," but alternatively, it may be that alcoholism, because of chemical dependency and imbalances, causes depression. Indeed, in this case, even if you could figure out which came first, it's not likely to tell you very much, since both alcoholism and depression are conditions that typically last for a very long time and interact with many other causal factors throughout a person's life. Insofar as you can make a causal inference, you should:

a. Go with causal hypothesis that relies on the most reliable mechanism.
b. Try to screen off one factor from another, e.g., look at depressed people who don't drink and see if they eventually turn to the bottle and look at alcoholics who aren't depressed and see if they eventually become depressed.

23. Causes may be complex

Causal relations can quickly become complex, often working in "loops" (people with high-self esteem tend to succeed, success then brings on greater self-esteem). Be mindful of these loops, as they might furnish the most plausible mechanism.

Appendix (optional): Other causal arguments

There are MANY other forms of causal argument: notably Mill's Methods and Inference to the Best Explanation.

The following five causal argument schemas are called *Mill's Methods*[3]. The first of Mill's Methods is the *Method of Difference:*

> Case *a* and case *b* are very similar, though *a* ends with the event *Y* and *b* does not.
> Case *a* has an event *X* in its history; case *b* does not.
> ∴ *X* (probably) caused *Y*.

The next, the *Method of Agreement*:

> Case *a* and *b* are very different, though both end with the event *Y*.
> Both *a* and *b* have event *X* in their histories.
> ∴ *X* (probably) caused *Y*.

The next, the *Joint Method of Agreement and Difference*:

> Cases *a*, *b,* and *c* are similar in some respects, different in others, but *a* and *c* end with the event *Y* and *b* does not.
> The only event common to *a* and *c's* histories and absent from *b's* is *X*.
> ∴ *X* (probably) caused *Y*

Fourth, Mill has the *Method of Concomitant Variation*.

> Cases *a* and *b* both end in event *Y*.
> The intensity of *Y* is much greater in *a* than in *b*.
> The intensity of *X* is much greater in *a* than in *b*.
> ∴ *X* and *Y* are causally related.

Finally, there is the *Method of Residues:*

> Case *a* ends with events *W* and *Y*.
> *T*, which is in *a's* history, causes *W*.
> *X* is the only other (relevant) event in *a's* history.
> ∴ *X* (probably) causes *Y*.

Another causal argument is called Inference to the Best Explanation (IBE):

> *Y*.
> *X* is the best explanation of *Y*.
> ∴ (Probably) *X*.

There are many more forms of inductively strong arguments, most of which we won't discuss in this course.

Levitt, Steven D., and Stephen J. Dubner. 2005. *Freakonomics : a rogue economist explores the hidden side of everything*. 1st ed. New York: William Morrow.

[3] Those of you who have taken ethics or social and political philosophy will recognize John Stuart Mill, the originator and namesake of these methods, as the author of "On Liberty" and "Utilitarianism."