On Argument as Practice:

What Makes Claims Evidence and Conclusions Trustworthy

James Paul Gee
Mary Lou Fulton Presidential Professor of Literacy Studies
Arizona State University
James.gee@asu.edu

Argument

Arguments are meant to persuade. Some arguments seek to persuade through rhetorical tricks. Such arguments are common in advertising and politics. Other arguments seek to persuade through evidence. Such arguments are the staple of science at its best, but they are important and necessary across many other areas of life and society.

We tend to call both of these types of arguments "persuasive arguments". I will call the first type a "manipulative argument" (in such arguments, even when evidence appears it is only in the service of manipulation). I will call the second type an "empirical argument" (where persuasion is a goal, but is subordinated to evidence). There are, of course, in between cases where we are not quite sure whether manipulation or evidence has the upper hand.

In education (schools, tests, and standards movements) even empirical arguments are not well understood as a practice. Textbooks on teaching or assessing arguments, regularly mistake reasons for evidence. They are different, as we will see.

Here is a rather typical statement about what an argument is:

> An argument is a connected series of statements or propositions, some of which are intended to provide support, justification or evidence for the truth of another statement or proposition. Arguments consist of one or more premises and a conclusion. The premises are those statements that are taken to provide the support or evidence; the conclusion is that which the premises allegedly support. For example, the following is an argument:

> The death penalty should be adopted only if it deters murder. However, it could only do this if murderers understood the consequences of their actions before acting, and since this is not so, we must reject adopting the death penalty.

> The conclusion of this argument is the final statement: "we must reject adopting the death penalty." The other statements are the premises; they are offered as reasons or justification for this claim. The premises of an argument are sometimes also called the "data," the "grounds" or the "backup" given for accepting the conclusion.

> http://www.iep.utm.edu/argument/

This source, like many others, confounds "having a reason" and "having evidence". If I attempt to argue that "Evolution is not true, because the eye is so well designed for its purpose that it could not have evolved step by step across a long period of time", I have given a reason (or the reason) for my belief: "The eye is so well designed for its purpose that it could never have evolved step by step". But this reason is not evidence. I have offered no support for it as a true claim. It is just a claim and, in fact, a false one to most biologists.

We can rewrite the argument above about the death penalty as:

> Premise 1: The death penalty should be adopted only if it deters murder

Premise 2: The death penalty could only deter murder if murderers understood the consequences of their actions before acting

Premise 3: Murders do not understand the consequence of their own actions before acting

Conclusion: Therefore we should not adopt the death penalty

While the above argument is, indeed, what philosophers call a "valid argument", the view of argument in the passage above is incoherent as a view of argument *as a practice*. The conclusion in the argument above is a claim (something someone purports to be true) and so is each of the premises. How can one claim be evidence for another claim? Why would I believe Premise 3, for example, any more or less than I would believe the Conclusion? They are both just claims from the same person.

The whole notion of what "support" or "evidence" means needs to be explicated if we are to understand argument as a practice. Claims are something we *assert* (a speech act). Evidence is something we *have* (a possession). And we *have* evidence because we have earned it thanks to *work* we have done.

The premises in the above argument are, in fact, evidence for the conclusion if and only if they can serve not just as claims (purported truths) but also as evidence, that is, *trustworthy grounds for belief or action*. So the crucial question is what makes a claim able to serve as evidence (trustworthy grounds for belief or action)?

Nothing in the above "argument" tells us whether the premises can bear the weight of being evidence. Nothing in the argument tells us whether these premises are likely to be true, how likely they are to be true, and whether and how they are relevant to the conclusion in the right way to offer it support as evidence. Until we know whether and how these premises can bear such weight, we do not know if the above "logical argument" is actually a real argument and not just a chain of claims.

A claim is evidence only if it is the product (outcome) of a type of work I will call "evidential inquiry". Evidentiary inquiry is, at its heart, always some form of observing (seeing) or listening or reading (studying). So, in the case of the argument above, it is a real argument if and only if each premise is associated with a statement of the work of evidential inquiry that went into establishing a firm foundation on which the claim can rest as evidence.

Evidential Work

Now we need to know what "evidential inquiry" is. Evidential inquiry is the work people have done that would make other people trust that a claim they make has a good probability of being true. Evidential inquiry makes a claim (a premise) trustworthy. Claims (premises) in arguments are evidence if and only if they are trustworthy. A real argument, as a form of practice, has the following form (with as many premises as we need):

Premise 1: You can trust me on X because of the evidential inquiry work I have done.

Premise 2: You can trust me on Y because of the evidential inquiry work I have done.

Conclusion: So, if you trust me on X and Y (as you should), then you must believe (trust me on) the conclusion Z, since if X and Y are trustworthy ("true"), then Z is trustworthy ("true") as well.

In the argument about the death penalty with which we started, the premises are evidence if and only if arguer can fill in for each premise a statement of the work of evidential inquiry that went into making the premises trustworthy. Each premise must be accompanied by what I will call a "certificate of trust" based on evidential inquiry work the person making the argument has done.

Certificates of trust come in different grades. There are different types of evidential inquiry work and they lead to different degrees (stronger and weaker) of trust. For simplicity's sake we can distinguish three types of evidential inquiry work, three different sorts of certificates of trust to which such work can give rise. I will name these three types in terms of everyday language, but I mean my terms here as technical terms that I will define below. These three types are:

1) "I saw" (primary work and trust);
2) "I heard from someone who saw it" (secondary work and trust); and
3) "I heard from someone who heard" (tertiary work and trust).

"I saw" (primary work), "I heard from someone who saw it" (secondary work), and "I heard from someone who heard" (tertiary work) mean different things in different "semiotic domains". A semiotic domain recruits one or more modalities (e.g., oral or written language, images, equations, symbols, sounds, gestures, graphs, artifacts, and so forth) to communicate distinctive

types of messages. Here are some examples of semiotic domains: cellular biology, postmodern literary criticism, first-person-shooter video games, advertisements, Roman Catholic theology, modernist painting, midwifery, and so on and so forth through a nearly endless and motley list.

It is too bad we don't have a better term of art "semiotic domain" here. But the essential insight I am trying to capture is this: domains like vampire romance fan fiction for a young fan-fiction writer; Japanese anima manga (comic books) for an otaku (expert); "literal readings" of the Bible for a fundamentalist Christian; and cellular biology for a cellular biologist are each domains of specialized representations, modalities, knowledge, and practices. In their own ways, each is quite complicated and each is grounded in a group of people who have cognitive and social interests and help uphold a set of standards and norms. Each domain allows people to communicate distinctive sorts of messages (information, values, ideas) to each other. In each of these domains "I saw" and "I heard" mean different things.

Of all the domains in which we humans can make meaning and communicate, one is special. This is what I will call the "life world". For each of us, our life world is composed of all the times and areas in life in which we make claims as everyday people, not as specialists. These claims are based on "common sense" or the experiences we have had as everyday (non-specialist) people, experiences that do not require special training. When a person makes a claim as a doctor, *Yu-Gi-Oh* expert, or policy wonk, that person is speaking as an "expert" and often recruits a specialist style of language associated with his or her domain of expertise. But all of us can make claims based on our everyday knowledge and experiences as non-experts. Of

course, people's life worlds, and how they think and speak in them, differ in terms of their cultures.

In the life world, "I saw" means I saw something. If I make a claim like "A woodpecker made the hole in this tree" on the basis of having seen the woodpecker do it, this is primary work and in the life world can gain primary trust. If I make a claim like "A woodpecker made the hole in this tree" on the basis of having been told this by the person who saw it happen, this is secondary work and in the life world can gain secondary trust. If I make the claim that a woodpecker made the hole based on having heard it from someone else who heard it (third-hand, so to speak, or worse depending on how long the chain of "hear-say" is), this is tertiary work and in the life world can gain tertiary trust (or worse, depending on how long the chain of hearings is, how far back we have to go to get back to the person who saw/observed).

In experimental sciences, "I saw" means I carried out an experiment. That is primary work and can gain, in these domains, primary trust. "I heard from someone who saw" means I read or heard from the experimenter(s) about the results of the experiment. In these domains, this is secondary work and can gain secondary trust. "I heard from someone who heard" means I got my information from some source that was not the experimenter, but two or more steps removed from the experimenter's observations. Again, the length of this chain can quickly cause trust to dissipate.

In Biblical studies, "I saw" means I have read the text in the original ancient language in which it was written (primary work/trust). "I heard from someone who saw" means I read or heard from

someone who has read the text in the original language (secondary work/trust). And "I heard from someone who heard" means I read a text from or listened to someone who has not read a text or heard information from someone who has actually read the text in the original language (tertiary work/trust).

In other domains, "I saw". "I heard from someone who saw," and "I heard from someone who heard" mean different things and in the last of these (tertiary work) trust weakens as the chain of hearings lengthens. However, in no domain, including our life world domain, is trust based *only* on seeing and hearing. To generate trust, seeing and hearing have to have involved a certain amount of effort (which is why I call them "work").

Each level of work and trust requires two additional things beyond mere seeing and hearing. First, competence is required: Trust requires that, in a given domain, we judge the "see-er" as a competent see-er, that is, someone who is competent to say correctly what he or she has seen. We must also judge people in the chain of hearers as competent hearers, that is, people who are competent to report correctly what they have heard.

If Joe says he saw a coyote run down the street, but we know he cannot tell a coyote from a dog, then he cannot gain primary trust from his observation. If an experimenter has gone "rogue" and thinks that fairy dust could contaminate his samples, we do not trust his experiments and he cannot gain primary trust from them. If someone always reports things based on his or her own fantasies, then the person is not a competent hearer and reporter. Primary, secondary, and tertiary trust each requires judgments about competence. Such judgments, in specialist domains,

are made in terms of the (formal and informal) standards and norms of the people in the domain. In our life worlds, such judgments are made based on common sense and shared histories.

Second, quality is required. Even if a see-er or hearer is competent, how good was their work? Was it (the seeing or the hearing) thorough, consistent, careful, considered, and carried out in terms of high standards or not? Casual observations or quick and dirty experiments do not gain the same level of primary trust as do careful ones. Careless listening or reading does not gain the same level of secondary or tertiary trust as careful listening and reading does. Quality counts.

What makes a premise evidence?

We can now, at last, say what makes a premise evidence. To be evidence a premise must have, so to speak, an attached (assumed or actually stated) "certificate of trust". It must be marked as the result of primary, secondary, or tertiary work. If tertiary, it must be marked in terms of (perhaps roughly) how far back in the chain of hearers we have to go to get back to the see-er (observer).

For a premise to be evidence, we must also specify how competent, as see-ers or hearers, the see-ers or hearers involved in producing the premise are. For a premise to be evidence, we must, as well, specify the quality of the seeing or hearing that went into producing the premise. In a chain from many hearers back to a see-er we need to know or believe that each person in the chain is competent and has achieved a certain level of quality.

It is important to realize that judgments of trust, competence, and quality are always on a continuum. They are judgments about more or less. Judgments about the competence of see-ers and hearers and the quality of seeing and hearing range from high to low. They are not binary "yes" or "no" judgments.

To see what I mean by having a "certificate of trust" attached to each premise, consider the simple argument below:

Premise 1: I have seen lots of holes in trees just like the hole in this tree.

Certificate of trust for Premise 1:

Type of work: Primary, I saw the holes.

Competence: How competent am I to see holes in trees? I am fully competent in vision and in everything else I need to see holes in trees, as are most of us. I am somewhat less confident that I can be sure that this hole is "just like" others I have seen, but I hope you will give me the benefit of the doubt.

Quality: I am honest and, thus, you can trust me when I say I have seen lots of such holes, not just a few. I have sometimes looked at these sorts of holes carefully, more often rather casually.

Premise 2: Those holes were made by woodpeckers.

Certificate of trust for Premise 2:

Type of work: Combination of primary and secondary. A naturalist once told me that woodpeckers made such holes in trees and I have occasionally seen a woodpecker pecking on trees, though I have never seen one make a whole hole.

Competence: I am reliable hearer and reporter on matters like this, since I am interested in birds and nature, have devoted work to reading about them and observing them, and can recognize woodpeckers.

Quality: I care about birds and nature and pay attention to such things as well or better than the average person.

Conclusion: This hole was made by a woodpecker.

How trustworthy is this conclusion? It is based on primary and secondary work with adequate to good competence and quality. It is still possible the hole was made in some other way, of course, but it is likely that a woodpecker made it and so you can treat the conclusion as a very plausible assumption.

The premises in the above argument are only evidence when the information about trust (what we have called a "certificate of trust") has been filled in and when it is adequate, by whatever standards we apply or a domain applies.

A "certificate of trust" can be information the arguer explicitly gives. Sometimes all or part of the certificate of trust is not explicitly given, but assumed or taken for granted. In this case, hearers or readers have the choice to fill in the missing parts of the certificate of trust for themselves; simply to accept that the arguer could give the information if required; or to demand it. A certificate of trust often results, at least in part, from overt questioning and dialogue and not only from information given in the original statement of the argument. In this sense, certificates of trust are often co-constructed by arguers and hearers in the process of argument as dialogue.

Note however—and this is crucial—that the conclusions of empirical arguments (as opposed to purely logical arguments) are not really "true" or "false", but, in reality, more or less plausible and reliable. It is much as if we give premises grades in terms of trust. The conclusion inherits its trust grade from the premises (that is what makes it a conclusion of an argument).

As mentioned earlier, often in arguments certificates of trust are left tacit or unexpressed. Sometimes this is so because the arguer feels they are obvious and sometimes because the arguer wants to dupe us. In any case, it is part of the practice of argumentation for us to ask for certificates of trust if we feel we need them and for arguers to supply them or face rejection.

Trust is made from competent and high quality seeing and hearing, remembering that what seeing and hearing means is different in different domains. In some domains, seeing and hearing are far removed from literal seeing and hearing in the life world. Seeing and hearing can become quite technical practices. Therefore, it always crucial that we ask what domain an argument is connected to, what domain's language and standards it is committed to. However, it may not always be clear what domain is relevant, especially if the argument is taken out of context.

Types of arguments

In almost any domain, arguments can be about four different types of things:

1. A fact: "X is a fact" ["X is likely/probable"]
2. A value: "X is good"
3. A policy: "X is a successful policy or course of action"
4. A word (or concept): "X is the right way to use the word (apply the concept)"

Thus, we argue over what is true (fact), what is good (value), what is practical (policy), or what things mean (words/concepts). For each of these four types of argument, we develop and ask for certificates of trust work differently.

An argument about fact involves a conclusion about a factual statement being true (e.g., "This hole was made by a woodpecker"). An argument about value involves a conclusion about what is good, valuable, moral, or good taste, and so forth (e.g., "Lying is immoral", "This wine is outstanding"). I am using the term "value" here only for matters about what is good that cannot be settled solely by facts. An argument about policy involves a conclusion about what actions or

policies are most likely to succeed (e.g., “We should tax the rich more”; “Consult an expert before you buy insurance”). Note that “policy” here does not mean just official policies, but less official courses of action a well (like “Go bird watching early in the morning for best results”). An argument about words involves a conclusion about what a given word or concept means (or should mean) or how it should be used or applied (e.g., “Democracy means elections that are not determined by money”; “A real democracy does not allow money to determine the outcome of an election).

So we always also need to ask or the arguer needs to be prepared to answer the question “Is this argument about a fact, a value, a policy, or a word?” This question is necessary because the “rules of the (argument) game” are different, in any domain, for each of these.

Burden of Proof

The conclusion to an argument gains its trust level from the trust levels the argument’s premises have earned. Different domains require different levels and kinds of trust. But any argument, by their very nature, must meet the basic burden of proof appropriate to the type of argument it is (an argument about fact, value, policy, or words). “Proof” here does not mean mathematical proof (“certainty”), but the likelihood of, or our trust or confidence in, the conclusion. The basic burden of proof is the basic bar an argument must pass for its premises to be relevant to the conclusion and, if true, for the argument to be valid.

Arguments about values are essentially about what is “good” in the sense of being ethically/morally good or highly valued in some other way that is not just a matter of brute facts

(e.g., the actual cost of gold is settled by facts; what the price of gold will be next month is a prediction, still a matter of fact, it is a guess about facts; whether gold is worth more than love is a judgment not settled only by facts). The burden of proof an argument about a value must meet is that it must answer the question: "Why is what the argument says is good actually good?" Answering this question will often involve giving criteria for what counts as good in a given domain or context. If an argument about a value has not answered this question (by the standards of the domain), then it fails. It has not met its burden of proof condition.

Arguments about policy (action) are essentially about what is "right to do" in the sense of being the right action to take to reach our stated goals successfully. The burden of proof an argument about policy must meet is that it must answer this question: "Why should we believe (trust) that this policy or action will succeed and not give rise to damaging unintended consequences?". If an argument about a policy has not answered this question (by the standards of the domain), then it fails. It has not met its burden of proof condition.

Arguments about words are essentially about the "right" way to use a word or apply a concept in specific contexts. The burden of proof an argument about policy must meet is that it must answer this question: "What good (either morally or practically) will it do to use the word this way this way?". If an argument about a word has not answered this question (by the standards of the domain), then it fails. It has not met its burden of proof condition. Note that arguments about words or concepts are always a species of argument about either a value (it would be morally good to use the word in the proposed way) or a policy (it would be practically good to use the word in the proposed way).

Arguments about facts are essentially about which facts are "right" in the sense that they are highly likely to true. The burden of proof an argument about a fact must meet is that it must answer this question: "How likely (probable) do the facts in the premises make the conclusion?". If an argument about a fact has not answered this question (by the standards of the domain), then it fails. It has not met its burden of proof condition

To meet their burden of proof condition arguments about a value must contain one or more premises about why a claimed good is actually good. When such an argument is directly about a good (e.g., Lying is bad/immoral; This wine is very good"), then it is obvious that the argument must contain some premises about why what is claimed to be good (or bad) is good (or bad). To do this the arguer must, at the least, give criteria for what constitutes value, good, or morality in such cases and, at the most, state a whole theory of value, morality, or taste in the requisite domain.

Sometimes arguments about a value are more indirect. They are about what we should or should not do as a matter of morality or taste, not policy ("should" can also have a policy meaning: what we should do practically). Such claims about "should" are implied claims about what is good or proper. Here, too, at least one premise in the argument must be a statement about why the thing claimed to be good (and, thus, what we should do or are obligated to do) is, indeed, good. To see this, consider the argument below:

Premise 1: You have more money than you need

Premise2: If you have more money than you need, you are morally obligated to help the poor

Conclusion: You are morally obligated to help the poor

Let's assume that Premise 1 has been shown to be true (assume we have a lot of trust in it)—say that the "you" referred to is a billionaire. This factual claim does not lead to the conclusion that you are obligated to help the poor; it is not evidence for that conclusion, unless the arguer can defend why helping the poor is the morally right thing to do (is a morally good thing). Why is it good? The argument does not say. Not everyone accepts that helping the poor is a good thing. They may think that helping the poor just deskills the poor or they may be social Darwinists who believe that allowing the poor to die off makes society stronger. An argument about a value must state (or be able to supply on demand) its criteria for what is good (moral, ethical, valuable, correct, etc.) and convince us we should accept those criteria. In an argument about a value, factual premises become relevant only when such criteria are offered or assumed.

To meet their burden of proof condition arguments about policy or a course of action require one or more premises about the effectiveness of the policy of course of action. Consider an argument that we should tax the rich more (Conclusion) because they have lots of money (Premise 1) and the government needs more revenue (Premise 2). The factual claims in the two premises do not lead to the conclusion that we should tax the rich more (they are not evidence for that conclusion) unless the arguer can say why and/or how the policy or action is likely to accomplish our goals without bad unintended consequences that will undermine those goals (and in some

cases and for some people it is important, as well, that the policy or action is moral or ethical). Such a statement is a distinctive form of factual claim, a claim about the likely effectiveness of a policy or action.

So raising taxes on the rich can only be the right (practical) thing to do to raise revenue if it actually will raise revenue and will not have unintended consequences that harm our interests and goals. Some will argue that the rich will just evade the taxes or park their money in tax havens and revenue will not go up. It may even go down, they will say. Other will argue that taxing the rich too heavily will discourage people from getting rich, thereby lowering potential revenue in the long haul, however much it improves it in the short term. What this shows is that policy arguments need premises (assumed or state) about the likely effectiveness of actions.

To meet their burden of proof condition, arguments about words or concepts require one or more premises about what good value or effective policy or action will result from using the word or applying the concept in the way the conclusion suggests. An argument like: Money determines votes in the United States (Premise 1); In a democracy, all the people and not just the rich should determine the outcome of an election (Premise 2); So the United States is not a democracy (Conclusion) works only if the two factual premises are supplemented by some premise or premises about what good it will do us (morally or practically) to use the term "democracy" in this way.

For instance, one could add any of the following premises (and their associated certificates of trust) to the argument: "Calling something a democracy when votes can be bought, will lead to a

lack of trust in government"; "Using the word 'democracy' only for countries where elections are free of big money will improve civil society and clarify political debate"; "The founders did not intend elections to be influenced by the rich". Remember, though, that each of these premises themselves has to be defended and is open to question and further debate. Note that arguments about a word are always a species of either an argument about a value (it is morally good to use the word this way) or policy (it is practically good to use the word this way).

To meet their burden of proof condition arguments about facts require one or more premises about how much trust the other premises in the argument should give us about the likelihood of the conclusion being true. Such premises about how likely the other premises make the conclusion are often not explicitly stated, but, rather, assumed based on the standards and practices of the domain in which the argument is made. For example, a premise like "I saw John holding a gun" makes the conclusion "John had a gun" quite likely by the standards of the life world, based on the usually unstated life world premises that "seeing is believing" and "seeing is usually accurate". The premise that "I saw John holding a gun" makes the conclusion that John had a gun less likely (trustworthy) in the domain of law in courtrooms. Here the conclusion requires some further shoring up by a premise like "Eyewitness testimony tends to be accurate only under certain circumstance and those circumstances were met here".

When we analyze an argument, we need always to carefully ask whether it met its burden of proof condition. If the premises together to not answer the required burden of proof question, the arguments is not trustworthy (in that sense, not "valid"). It is surprising how many things printed

as examples of "good" arguments in texts on argumentation do not meet the required burden of proof condition and, thus, are not valid (trustworthy, successful) arguments.

Questions

So for any argument here are the questions we should always ask and the arguer should always be prepared to answer, if the premises in the argument do not already answer them:

1. DOMAIN QUESTION: What domain is the argument in?
2. TYPE QUESTION: Is this argument about a fact, a value, a policy, or a word?
3. WORK QUESTION: For each premise, we need to ask whether it is the product of primary, secondary, or tertiary work.
4. COMPETENCE QUESTION: For each premise, we need to ask how competent the arguer is as a see-err or hearer in the requisite domain in regard to the claim made by that premise.
5. QUALITY QUESTION: For each premise, we need to ask how high (or low) the quality of seeing or hearing was in regard to that premise, given the standards in the requisite domain.
6. BURDEN OF PROOF QUESTION: The premises of the argument must answer one of the following questions depending on what type of argument it is (fact, value, policy, or words):
    A. Why is the value expressed or implied in the conclusion actually good?
    B Why should we believe (trust) that this policy will succeed and not give rise to damaging unintended consequences?.

C. What good (either morally or practically) will it do to use the word this way?

D. How likely (probable) do the facts in the premises make the conclusion?

7. TRUST QUESTION: For the conclusion, we need to ask ourselves how much trust it is has inherited from the premises and their associated certificates of trust, given the standards of the requisite domain (how much trust a given domain requires for arguments to work or "win"). We need to make a judgment about how much trust we will place in the conclusion.

Complexities

The work of making and assessing arguments is very hard work. This is something that is not at all apparent from the simple cases I have used so far to make my point. Each of the seven questions above can be very hard to answer in many cases. Let's take each question and, at least, display what sorts of complexity can and regularly do arise. First, we will consider the domain question. It is not always clear what domain an argument is meant to be in or a part of, what the domain it is that determines the standards of language, reasoning, and evidence the argument needs to meet. Even an argument about whether a woodpecker made the hole in the tree could be in the life world domain or a scientific domain dealing with birds or forests. People often argue over what domain a claim should be argued within. For example, is a claim about whether some expression is "correct English" a claim that should be argued in (the language, standards, and values of) the life world, linguistics, sociology, education, or some other area? Each domain

means something different by "correct English", uses different methods, has different ways to argue with different standards, and looks at evidence in different ways.

Let's turn now to the type question. It is not always clear—sometimes it is a matter of contention—whether a claim is a matter of fact, value, policy, or words. Some cases are clear, others are not. For example, some people think the claim "Life begins at conception" is a factual claim, others that it is a values claim, and others (myself included) think it is an argument about words and how they should be used in or applied to the world. The claim could even be seen as a policy matter mixed with an argument about words—thus, an argument about how we should use the word "life" in formulating laws or governmental policies.

Much of the domain of law is about stipulating the meaning of a word (like "theft", "assault", "property", "murder", "speech") that has different meanings in different domains and contexts of use. Thus, often in the law, arguments are not really a matter of arguing about what words should mean, but, rather (after acknowledging that they have differing, contextually changing, and often vague meanings) stipulating what they will mean for the purposes of laws, courts, and justice, largely matters of public policy. Stipulations are conventional and open to change as times change or policy issues change.

Even authentic arguments about how a word should be used or how a concept ought to apply often have a values or a policy orientation. If I argue that any country that allows money to determine the outcome of elections should not be considered a democracy, I am arguing over the

meaning of a word and I am expressing a value (democracies should not advantage the rich, it is not fair) or advocating a policy (constrain campaign funding).

Consider now the work question. Sometimes it is clear what constitutes primary work ("I saw"). For example, it is clear in experimental science and in the case of many life world claims about the world. But what constitutes primary work and generates primary trust can be a hard question in certain domains and for certain types of arguments. In policy matters—arguments about what to do—it is clear that having engaged in the action, observed its results, or built models or experiments to view its results indirectly are all forms of seeing, forms of primary work. Hearing about such things is secondary at best.

While it is clear that in the case of a value judgment like "This is a great pinot noir wine" having tasted the wine is a primary form of work, a type of "seeing" (namely tasting) in the domain of wine connoisseurship, the issue is often much less clear for claims about moral values. What is "seeing" in the case of a moral judgment like "Lying is morally wrong"? What sort of evidence can gain primary trust here? People do not actually see moral values and they do not usually carry out experiments on them.

In the case of moral values, in some religious domains, the primary work is the word of the priest or minister who has some sort of direct or privileged access to God or God's word. The priest or minister has heard from someone who has seen (namely God), but this form of hearing is as close to "seeing" (direct observation) that we humans can get when God is involved. A form of secondary work (hearing from someone who has seen, namely God) becomes here a form of

primary work (seeing, observing, primary access, as close as we can get). Unfortunately, this leads to an "argument from authority" (the priest or minister) and not evidence per se.

Claims about moral values can, of course, have other foundations than the authority of God and his/her surrogates. They can, for example, be founded on claims that every human, barring mental illness, accepts (so we are all, in the case, see-ers). For example, we could argue that all humans, barring mental illness, accept the premise that "if something would hurt someone that is a reason not to do it, unless there are compelling over-riding reasons in a given situation". In fact, we could argue that if some human did not accept this premise, we at least would not bother to reason with him or her about moral values. Then we can use this premise and further premises (e.g., that lying hurts people and other premises about what sorts of over-riding reasons count, if they should, or why they shouldn't if they shouldn't) to argue for a further moral claim like "lying is morally wrong".

Arguments about words (concepts) can be equally vexed. Who is to say what a word or expression like "democracy" or "the good life" is to mean in different contexts, especially in the life world where there are no specialist stipulations? A primary observation here may be imagining (a form of seeing) what a life, society, or the world would look like if a word was used in a certain way or a concept was applied in a certain way. And, of course, imaginings can be backed up by more or less reflection, study, and modeling.

Finally, the competence, quality, and trust questions can each become very complex because criteria for competence, quality, and necessary level of trust can vary across different domains.

And, too, they can be matters of contention, negotiation, and compromise even within a given domain.

Answering the burden of proof question can often get so vexed that it deserves its own section below.

Confirmation Basis

Empirical arguments in all domains are prone to a very common human mental "bug" called the "confirmation basis". Humans tend to consider only information and evidence that confirms their beliefs. They do not, unless actively pushed, look for possible negative evidence and counter-examples. In many scientific domains corrections for confirmation basis are built in, often through colleagues who actively want to defeat other colleague's arguments and who thus look for negative evidence and counter-examples. This is part of what makes science inherently social and collaborative (competition is, oddly, here a form of collaboration). In many other cases, the role of defeating confirmation basis falls to people who are listening to or reading an argument. As critical listeners or readers they need to seek out negative evidence and counter-examples and confront the arguer with them as part and parcel of the dialogic practice of argumentation.

A good empirical arguer anticipates critical listening and reading and seeks to deal with and forestall negative evidence and counter-examples beforehand. This is to say that good arguers are critical listeners and readers of their own arguments. They actively combat the dangers of confirmation bias.

Burden of Proof: Going Meta: Defending Theories

Ensuring that the premises of an argument answer the required burden of proof question can be vexed enough sometimes that it leads to what I will call "going meta". We can be asked to defend our criteria for good (arguments about a value); for the probability of success of a proposed action (arguments about policy); for the moral or practical need for using a word a certain way (arguments about a word); or for the probability our conclusion is true given our premises (arguments about a fact). Indeed, in the case of arguments about facts, in certain technical domains, we may have to defend which theory of probability we are using (there are several).

When an argument goes meta, we stop arguing about a claim (the conclusion) and begin arguing about the theory that produced our claim and the criteria we are using to meet the required burden of proof. Claims and criteria are always the results of theories, whether these are claims and criteria about facts, values, policies, or words. Arguments over whether the burden of proof has been met often lead to arguments over theories.

A theory is a set of principles and procedures for generating claims and evaluating their likelihood of being true or trustworthy. For example, in the life world most humans hold a theory something like "Seeing is believing as long as the conditions for seeing are good enough". So we tend to trust claims made on the basis of seeing. If we deem someone competent and honest, we rarely question the truth of what he or she claims to have seen under good conditions. However, a whole host of work in psychology and law has shown that eye witness testimony is

highly unreliable. This work has shown that people are not, in fact, very good at accurately reporting what they have seen. Their minds, desires, and memories distort what they have seen and make them misreport what they have seen even when they fully believe that they have seen what they claim to have seen.

Normally when someone claims to have seen something in the life world, we constrain ourselves to arguing (or making judgments) about the claimer's competence (Is he honest? Does he have 20/20 vision? Does he need glasses?) and the quality of his seeing (Were the conditions good? Did he pay attention?). Of course, usually these matters just go "without saying". But imagine I bring up the literature on the unreliability of reports of seeing (eye witness testimony) and question the very theory that "seeing is believing". This is what I will call "going meta". I am no longer just arguing over a claim (like "John saw Jane go into the house next door late last night"), but questioning the whole theory that underlies claims to have seen or witnessed and their trustworthiness. This could well be an assault on the life world by a specialist in psychology or law. In mundane cases, it would be quite rude. In court, it would be quite common. Courts have much higher standards for seeing than does the life world.

Going meta and questioning theories is common in cross domain disputes. A creationist does not only argue specific factual claims about evolution, but frequently goes meta and questions the whole theory of evolution, the theory that tells biologists how to generate and evaluate claims about the past development of species. However, sometimes people go meta within one domain. Two biologists that hold different but related theories of evolution—say one thinks natural selection happens at the level of bodies and the other thinks it happens at the level of genes—

may well argue over which theory is correct when they find they cannot really agree on how to interpret a factual claim until they get straight whether they should, in general, be talking about bodies or genes. Of course, fights about facts will play a role in arguments about theories, but the real goal is to get straight how we should generate, state, and evaluate claims in the first place.

Because arguments about values are often so hard and complex, we regularly go meta in such cases. If you claim "Sex outside or marriage is immoral" I am quite likely to ask you to defend your theory of morality, the theory that generated this specific moral claim. If you claim that something is true because your minister said so, I am quite likely to ask you to defend what makes your minister an authority in such matters, why I should trust an argument from authority in his case—that is, what is your theory of how authority of this type gives rise to trustworthy claims (or why, indeed, we should trust authority at all).

When an argument has gone meta and we are now arguing over a theory, how do we defend a theory? What sorts of claims can gain primary trust when we are defending a theory? What constitutes primary work here? We defend theories by explaining them and showing their track record of success in generating trustworthy claims in an appropriate domain. Primary work is a matter of developing or actively using a given theory. Secondary work is having consulted such primary work competently and carefully. Tertiary work is having consulted secondary work competently and carefully--and, again, in disputes about theories, work (hearing or reading) that gets too far away from the primary sources is less and less to be trusted.

An Argument

Let's look at the argument about the death penalty with which we started. I reprint my version of it below:

Premise 1:	The death penalty should be adopted only if it deters murder
Premise 2:	The death penalty could only deter murder if murderers understood the consequences of their actions before acting
Premise 3:	Murders do not understand the consequence of their own actions before acting
Conclusion:	Therefore we should not adopt the death penalty

DOMAIN QUESTION: What domain is the argument in?

What domain is this argument in? Which domain's standards, values, practices, theories, and language is it meant to be using? It does not appear to be the life world domain, since Premises 1 and 3 in the argument are surely not widely shared common everyday understandings. It does not appear to be a legal argument. Perhaps, it is founded in philosophy or psychology, or both, since its premises do seem tied to claims made in these domains.

TYPE QUESTION: Is this argument about a fact, a value, a policy, or a word?

The conclusion is a values statement ("should not"), so this appears to be an argument about a value. I assume the conclusion is a claim about morality or ethics. I suppose it could be taken as a claim about policy—what actions are practical and liable to be successful for specific

purposes—but the argument does not appeal to the language and theories typical of policy discussions about the death penalty (which tend to appeal to statistics about things like deterrence, not claims about the minds of murderers).

WORK QUESTION: For each premise, we need to ask whether it is the product of primary, secondary, or tertiary work.

We have no idea whatsoever whether the work behind any of the premises is primary (direct study), secondary (hearing or reading about direct studies), or tertiary (hearing or reading sources that themselves are not based on reports on direct studies). We have no idea how much trust to place in each premise. Premise 2 could be taken to be "obvious" to anyone who has a mind and understands what "understanding the consequences of ones actions" means (thus, the result of primary work)—but see more below on this.

COMPETENCE QUESTION: For each premise, we need to ask how competent the arguer is as a see-err or hearer in the requisite domain in regard to the claim made by that premise.

Having no information on domain and on the sort of work that went into establishing each premise as actual evidence, we cannot really answer the competence question. Knowing the identity of the arguer would not be enough. He or she may have a good reputation, we but we will still require to know that in this case that competence was applied.

QUALITY QUESTION: For each premise, we need to ask how high (or low) the quality of seeing or hearing was in regard to that premise, given the standards in the requisite domain.

Again, having no information on domain and the sort of work that went into establishing each premise as actual evidence, we cannot answer the quality question.

BURDEN OF PROOF QUESTION: If the argument is taken to be on about policy, it must answer the requisite burden of proof question: Why should I believe (trust) that this policy or action will succeed and not give rise to damaging unintended consequences? Nothing in the premises answers this question. Surely getting rid of the death penalty will accomplish the policy goal of getting rid of the death penalty, but, even if we agree with this policy, we still need an argument that getting rid of the death penalty will not have damaging unintended consequence like leading to murders. If this argument is taken as an argument about values, then its requisite burden of proof question is: Why is the value expressed in the conclusion good? This question is not answered either. We need to know why not having the death penalty is morally a good thing. There are, of course, people who would argue that the death penalty is moral (and not having it immoral), even if it does not deter murderers, because they believe retribution is a moral value. If the arguer ends up having to argue over whether or not retribution is a moral value, the original argument is going to go-meta and lead to an argument over competing theories of morality,

TRUST QUESTION: For the conclusion, we need to ask ourselves how much trust it is has inherited from the premises and their certificates of trust, given the standards of the requisite domain (how much trust a given domain requires for arguments to work or "win"). We need to make a judgment about how much trust we will place in the conclusion.

The trust question cannot be answered. Conclusions to arguments inherit their trust (or "trust grade") from their premises. Since we cannot grade the trust of the premises, we cannot assign any trust value to the conclusion.

So it is clear that this "argument" about the death penalty is not a real argument. It is just a collection of statements or claims that take the form of a valid logical argument (the conclusion is true if the premises are, but, sadly, we have no idea whether the premises are true).

Now let's consider what work would have to be done to make this "argument" a real argument. Let's start with the first premise: "The death penalty should be adopted only if it deters murder". If this claim is a moral one, then we need to know why we should accept or believe this moral value. We may well have to "go meta" here and ask the arguer to defend a moral theory from which such a claim follows. If this claim is about policy—what we should do for effective action—then we need to know why this would be a good, effective, or successful policy. Let's assume the premise is a moral claim and that then this argument is in the domain of ethics or morality. There are, of course, people who argue for the death penalty as a form of justice or retribution even if it does not deter murderers. The truth of the premise is not obvious, but needs to be established.

The second premise—"The death penalty could only deter murder if murderers understood the consequences of their actions before acting"—is probably meant to be a claim that is "obvious" to anyone and everyone. Thus, it would be a life world domain claim. Such claims can, of course, be carried over to more specialist domains. This would be primary work and engender

primary trust. The claim is plausible, but we could argue that the death penalty could deter murder by putting to death all those people who kill without understanding the consequences of their actions, leaving only people who do understand such consequences behind in the world. Thus, the claim is, in fact, not purely obviously true.

The third premise—"Murders do not understand the consequence of their own actions before acting"—is clearly false in the life world, where almost all of us assume some people, even murderers, understand the consequences of their actions and some do not. This claim, then, must be based in a more specialist domain with more esoteric knowledge. Perhaps, it is meant to be based in psychological studies, interviews, or some philosophical theory of the mind. The premise would be much more plausible had the subject of the claim been "Many murders" and "Murders", but then the whole argument would not work, since for those murderers who do understand the consequences of the actions, the death penalty might deter murder. There is also the problem that historically it seems some people have committed murder either because they wanted to be put to death or did not care if they were.

We can see that it is going to be an uphill battle to garner trust for these premises, with the possible exception of Premise 2. The "argument" as printed is not a real argument and it looks unlikely that it could be made into a trustworthy real argument (i.e., an argument where the trust placed in the premises is strong enough—by some domain's standards—to allows us to place a good deal of trust in the conclusion).

It is unfortunate, however, that textbooks on teaching argument often print such things as the death penalty argument above as an example of argument. Standards movements about argumentation—including the new Common Core Standards—often do likewise. Arguments like the one above are regularly printed in books about teaching argument or assessing arguments or argument writing. Claims (premises) that are reasons are confused with evidence and trust and the conclusions of such arguments are said to be "supported" when they are not, at least not evidentially supported. Arguments have their true homes in argumentation, which is a domain-specific social practice. Reasons are freely given, but evidence is the result of work and trust has to be earned.

Students learning about argument should always be learning about domains, inquiry, trust, dialogue (argument as a form of interaction between people or between a person and a text), and, yes, logic (consistency, inference, and validity). Since argument is a social and interactive practice it is always, too, a form of discourse (language in use) in the head and in the world. We recruit different forms of language in different domains to represent arguments and to carry them out in interaction with people and texts.