Predictable Irrationality: The Fallacies and Other Irrationalities.

Human beings often reason fairly well, at least for some cases. But we are also subject to many errors in reasoning. We make errors in our own reasoning and we are persuaded by bad arguments given to us by others. Many of our errors fall into more or less clearly defined groups of *fallacies*. Philosophers use the term "fallacies" to refer to common errors in reasoning that, although erroneous, can convince. There are two reasons to study the fallacies. First, if you are interested in truth, you should avoid them, whether in your own reasoning or when someone tries to pull one on you. But second, you might want to employ them on others. For better or worse, the fallacies are used by any one in the business of persuading, whether lawyers, sales people (and the corporations they often work for) or politicians. Indeed, one of the best places to study the fallacies in action is by listening to what comes out of the mouths of such folk.

The fallacies are erroneous forms of reasoning that often persuade. Their persuasive power turns on a number of things. For one thing, a fallacious argument might resemble a good one. Indeed, as we will see, it is often difficult to draw a line between a fallacious argument and a pretty good one. Also, fallacious arguments often push psychological buttons. If you have those buttons, you are easily tricked. And again, it is often the case that people want to be tricked. They want to agree with you and only require the barest shadow of a reason to rise up on their hind paws and cheer. (See the section later of confirmation bias.) In any event, it is often hard to tell whether an argument is fallacious, especially in the hands of a skillful practitioner. There is often no mechanical way to distinguish good from fallacious arguments and reasonable people might disagree about them.

In this brief introduction to the fallacies, I explain some of the ones most commonly found in philosophical discussions, as well as in everyday discussions of controversial issues. You can consult any standard logic book, or any one of thousands of websites, for further information. A word of warning before plunging in. My guess is that some of you will read my descriptions of many of these fallacies and think to yourselves "What sort of idiot will fall for that one? Surely not me!" If you sincerely feel that way, than there is a fair chance that you are, indeed, just that sort of idiot! Even a full understanding of the fallacies is only partial protection against being taken in. When it comes to the fallacies, even the forewarned are not always forearmed. In a real sense, we were born to be fooled, and masters of the fallacies are hiding everywhere waiting for us, *homo gulliblis*, to step into their traps.

Hasty generalization: Generalization arguments are inductive arguments that involve forming generalizations from samples. There is nothing necessarily wrong or irrational about generalizing in this way, if it is carefully done. Indeed, generalizing from samples is an essential part of learning! This apple fell from the tree and hit Newton on the head, and that apple fell from the tree and smacked him yet again. "My goodness," he thinks, "all apples fall!" And thus gravity was discovered -- well, not really, but close enough. Such a clever man was Newton, though any reasonably bright squirrel was there ahead of him. But generalizing becomes fallacious when poorly done. I may generalize based on a sample that is too small, or one that is atypical. For example, I want to discover what percentage of students at Montana State University are Republican. I enter a room, barely paying attention to the sign on the door reading "Meeting of the Young Republicans 3PM today." I handout my survey and discover all of them are

Republicans! I generalize and conclude "Pretty much every student at MSU is a Republican." Well, admittedly this one is a bit far fetched. But it is nevertheless true that we often generalize from small and atypical samples. And in complex cases, it can be difficult to know whether one's sample is large enough or typical enough to provide grounds for generalizations. Indeed, even the expert pollsters, for whatever reason, often get it wrong. Consider two passages from a recent news report.[1]

> In this week's British elections, pollsters unanimously declared that no party would get a majority — until the Conservatives exceeded most predictions by a whopping 50 or more seats.
>
> Two months ago in Israel, pollsters were broadly convinced that Prime Minister Benjamin Netanyahu's coalition was in deep trouble — until it pulled off a sweeping victory.

So much for pollsters and their generalizations!

Weak analogy: We just considered inductive generalizations from examples, as occurs in polling. There is another kind of inductive reasoning that begins from examples. But rather than going from examples to a generalization, it involves a more modest kind of inference, one from examples to examples. This is called "arguing from analogy." The idea is that since something, A, has a certain characteristic, we can infer that something else, B, which is very much like A, probably has that characteristic as well. So I might reason that since this Honda Civic was quite reliable, this other Honda Civic is likely to be as well. There is nothing wrong with such arguments in themselves. They can be quite strong inductive arguments. But there is much that can go awry. The analogy between A and B might be weak. This occurs when A and B are not really all that much alike. The

[1] http://www.politico.com/story/2015/05/british-election-polling-flop-conservatives-surveys-117772.html

less alike they are, the weaker the analogy and the weaker the inductive inference. If it is too weak, the argument commits the fallacy of weak analogy. For example, suppose a particular Honda Civic made in a particular year in a particular factory is quite reliable. That gives reason to think another Civic made in the same year in the same factory is probably reliable. (If I crashed my Civic, I might feel safe picking up the other as a used car.) But the analogy weakens if we are considering a Civic made in a different year, in a different factory, and so on. As the differences add up, the analogy, and its corresponding inference, weakens. Weak analogies are quite common, though it is often hard to know how weak an analogy is. Consider the following argument I picked up from a logic text. "Guns are like hammers—they're both tools with metal parts that could be used to kill someone. And yet it would be ridiculous to restrict the purchase of hammers—so restrictions on purchasing guns are equally ridiculous." Obviously this is not a great analogy, but is it a *terrible* one? And some years back there was a rather curious argument going around for requiring the federal government to balance its budget. "A government is like a home. Just as a home must balance its budget, so must the federal government." Putting aside the obvious fact that it often makes perfect sense for home owners to go into debt -- I myself have a rather nice mortgage, we need to ask how strong an analogy we have here. Is a nation state like ours really all that much like a home? Do the rules that apply to home economies apply to national economies? Well, maybe, and maybe not. Analogies are also common, say, in the abortion debate. A standard pro-choice argument turns on an analogy between the fetus and a tissue. Just as the tissue has no right to life, neither does the fetus. But is the fetus really so similar to a mere tissue, say, a bit of skin?

False Cause: We often engage in causal reasoning, that is, reasoning about causes and effects. There is nothing inherently wrong with such reasoning, indeed, it is crucial if we are to lead our lives well and even to survive. But there are certain common mistakes that we make with causal reasoning. These errors are based on the assumption that because things, or kinds of things, are associated in some way, there is a causal connection between them. The classic false cause fallacy is the "post hoc, ergo propter hoc" fallacy, which translates as "after this, therefore because of this." The idea is that because one event came after another, the later event must have been caused by the earlier. Of course, sometimes later events are caused by earlier one. But the mere fact of temporal order proves nothing. I may walk under a ladder and some time later come to grief. But it would be an error, if that is all I have, to conclude that walking under the ladder caused the grief. In the same way, I might be ill and take some one's medical advice and recover. But if that is all I have, I cannot conclude that the treatment caused me to get better. If you doubt this, just consider: a huge number of things occurred between my walking under the ladder and my coming to grief, and a huge number of things happened between my following someones medical advice and my getting better. If all we have to go on is "after this, therefore because of this" we have little to go on.

A version of the fallacy of false cause involves trying to infer causation from statistical correlation. Suppose that certain sorts of events are correlated with other sorts of events. One might think that this is enough to establish a causal link between the two. But it is perfectly possible that the correlation is a coincidence. As statisticians sometimes say, "correlation does not prove causation." Indeed, even when we have a very strong statistical correlation, it can be very difficult to prove causation. For example, there has

been a general decrease in crime of many sorts in the United States going back many years. But what is the cause? Lots of things have happened in this time period. Some argue that the drop in crime in New York City, for example, is the result of more active policing, including such policies as "stop and frisk." But crime has fallen even in places where such policies have not been introduced. Perhaps it is the end of the crack epidemic, or the aging of the population (younger folk are more likely to cause crimes than older folk). Indeed, if all we focus on is correlation, the drop in crime correlates with the rise of the World Wide Web, not to mention with my growing older! ("Look here! Levy was 40 and now he is over sixty and during all that time, crime has been decreasing! Obviously, things will only improve has the old guy turns truly ancient.") Clearly, correlation leaves a lot of questions about causation open.

Slippery slope: Slippery slope arguments are in a sense *conservative* arguments. They warn us against the dangers of changes. They have the following form. "Here is our current situation, say, a situation with respect to the law. It might not be perfect, but it is not so bad. Here is a proposal to change the situation. The proposal might make things better in one way or another, but it risks starting a chain of events, a sort of snowballing, which can ultimately lead to disaster. So we should not take the first step, no matter how attractive it might seem." Slippery slope arguments are inductive. A prediction is made: take this step, and a series of future events are likely which are likely to lead to disaster. There is nothing inherently wrong with such arguments. But much can go wrong, and when enough goes wrong, we have the slippery slope fallacy.

Slippery slope arguments involve a series of predictions. If we do A, then B is likely to occur, and if B occurs, C is likely to occur, and so on. There are at least two possible

problems here. First, people who give this sort of argument often do nothing to show that any, let alone all, of the individual steps are likely to occur. They rely on some sort of gut reaction that each step is at least *possible* and that that we will confuse the (often only remotely) possible with the likely. Yah, we are that dumb. We make that mistake all the time. The second problem is that even if a reasonably high probability can be assigned to each step in a chain of events, the probability of *the entire chain* of events occurring can be low. Keeping things simple, reflect on the following. Suppose that event A has a nice high .9 probability, and that event B also has a .9 probability. Then the probability of both A and B occurring is .9 X .9 or .81. As we add more and more high probability events to the chain, the lower the probability of the entire sequence. If there are enough steps in the sequence from our initial action to the disaster, the probability of the entire chain might be very low.

Here are two sample slippery slope arguments. I leave it to you to ask whether either is fallacious.

> Capital punishment reduces our respect for life. This, in turn, bit by it, one execution at a time, will create a society more tolerant of violence. One day our society will degenerate to the point of a war of all against all. The only way to stop this slide is to end capital punishment.

And on the flip side, we have the following argument.

> Capital punishment is currently practiced in a number of states. Were we to end capital punishment, we would signal a willingness to tolerate even the most horrible crimes. We would signal that murder, torture, rape and so on, are just not so bad. This would inevitably lead to a lower respect for life, and the law generally. We can expect this to

lead to disaster as people come to lose all respect for law, life, and everything else we hold dear. A war of all against all is the likely result.

So we have two slippery slope arguments pointing in opposite directions. Somehow, I doubt they can both be right. And perhaps neither one is!

Appeal to false authority: Often we argue by appeal to respected authorities. This can be perfectly legitimate. What could be more reasonable than appealing to a well-qualified historian about an historical matter, or a well-qualified biologist about a biological matter, or a well-qualified mechanic about a mechanical matter? But such appeals are fallacious under several conditions. First, many of the authorities people appeal to are not really authorities at all. Sometimes we know this, as when a famous athlete appears in an advertisement for something he knows nothing about. You might think you cannot be fooled by such nonsense, but large corporations would not pay celebrities huge sums of money to represent them if there were not good evidence it worked, that is, if there were not good evidence that we really are that dumb. Second, someone might actually be an expert but be paid to express an opinion that they may not actually endorse, or have any reason to endorse. How many athletes are paid to endorse a particular shoe they know and care nothing about? More importantly, there is evidence that researchers are paid by corporations (thinking tobacco and pharmaceuticals) to endorse, and even put their names on, questionable research. Third, often experts in a field disagree. This is often the case in, say, economics. When you can find there are significant numbers of economists arguing that a particular economic policy is a good one, and significant numbers arguing it is a bad one, it is not much of an argument for (or against) that policy to say "Hey everyone, I got a nice economist who agrees with me!"

Woopty doo.

Ad populum (Appeal to the People): There are several versions of this fallacy. They all appeal to the desire most have to be in accord with others. The most iconic version is *the bandwagon fallacy* in which the arguer tries to convince the audience to do or believe something because many others do. For example, not all that long ago it was common to argue that gay marriages are immoral since polls showed that a significant majority of Americans thought it was. Recently, the flip argument has been employed, that gay marriage is not immoral since a steadily growing majority do not find it immoral. Another version of the ad populum fallacy is known as *playing to the gallery.* Here, a speaker seeks to persuade by arousing relevant prejudices and emotions in his audience, perhaps by using colorful and emotion laden language. Of course, the speaker must work with the prejudices of his actual audience. If you have the wrong audience for your prompts, it could well backfire.

Ad hominem (against the person). Suppose I am in a debate with someone about some issue. For example, I might believe in changing the tax system to a flat tax while my opponent believes in progressive taxation. In a proper debate, we will both argue in favor of our preferred policies and against that of our opponent. But anyone with much experience with debate about taxes will know that this is extremely ineffective. Stick to the issue too closely, and it can backfire. You may even get the dreaded title of "wonk" if you actually show that you know something about the issue, a clear debate looser! .

So what is a poor debater to do? Go after your opponent! Be as in your face and personal as you like, or as your audience will enjoy. Enter the Ad Hominem. This is the fallacy of arguing "against the person" rather than against the claims the person

makes. There are three basic forms of ad hominem arguments. The most basic form is the *abusive ad hominem* which consists of name calling. In my youth, a great deal of political "argument" from the left consisted of shouting "fascist" (often screamed by people who did not know what a fascist really was). And a longstanding, very effective form of "argument" from the right is the shriek of "socialist" (often shrieked by people who do not actually know what a socialist is). Again, for example, in the 2015 debates among prospective Republican candidates for president, for lack of much of substance to say, the fearsome Donald Trump spends time declaring his opponents are stupid, of low energy, and have unattractive faces. Of course, he gets his share of abuse, particularly about his hair. These are, of course, rather shallow versions of the ad hominem, but seem to be pretty much what people want to hear. An interesting twist on the abusive ad hominem is that the name calling might very well be accurate: there are fascists and their are socialists and there are all sorts of other folk a particular audience might find despicable. But even someone who truly is despicable might be right on some point and might provide convincing arguments for that point. The second kind of ad hominem is the *circumstantial ad hominem*. Here it is argued that we should not take someone's views or arguments seriously because of the *circumstances* of the arguer, often because it is thought the person's arguments are in one way or another self-serving (or perhaps in-group serving). So, it is common to argue that *of course* the urban poor would support expanding social welfare programs since such programs are in the interests of the urban poor. Hence we can assume that nothing such lazy folk (nothing like a little abusive ad hominem to spice the circumstantial ad hominem) say in favor of such programs can be trusted. It is also common to argue that *of course* the wealthy want to reduce their own

taxes so *of course* any arguments they give that doing so will benefit society as a whole should be dismissed as self-serving. No doubt we seek to serve the interests of ourselves and our group. But it might still be true that we are right about something. The third version of the ad hominem fallacy is the *tu quoque* or "you too" argument. It involves rejecting, say, a recommendation because the person making the recommendation does not herself live up to it. "Why should I do as you say, you hypocrite!" Well, even hypocrites can be right!

Arguing against a person is not always fallacious. There are at least two kinds of situation in which we legitimately argue against a person. First, the issue might actually be something about the person. Is this person a jerk, a war criminal, a fascist or a socialist? I do not want to work with a jerk, marry a war criminal or vote for people of certain political persuasions so I need to consider evidence specifically about people and their characters. I may be intolerant for not wanting to work with a jerk, but I am not committing any fallacy by considering evidence that this person is a jerk. Second, suppose that we are supposed to believe something someone says just because they are supposed to have some kind of knowledge or wisdom. Then the issue is whether that person really is a source of knowledge and wisdom. For example, in a court of law, it is perfectly legitimate, at least sometimes, to try to undermine a supposed witness's testimony by impugning his character. "This supposed eyewitness," says the defense attorney, "is a criminal spending twenty years behind bars who has negotiated with the authorities to testify against my client for a significant reduction in his sentence." I would begin to wonder, quite reasonably, about that testimony. In the same way, it is often perfectly legitimate to question a supposed scientific study on some controversial topic by

impugning the character of the person who generated it. I may know nothing about the science behind, say, pharmaceuticals. But suppose I discover that a scientist is on the payroll of a large pharmaceutical manufacturer, and before that was on the payroll of large tobacco companies. I would quite legitimately be skeptical of his research, research I cannot actually understand. It is all good and well to say I must address the research and not the character of the person, but only an idiot would take off from her life for five or ten years to become a specialist in a field she cares little about just to evaluate a scientific study that purports to show that tobacco really is good for you. And it makes a good bit of sense to reason that someone paid by the tobacco industry to tout the health benefits of tobacco is not to be trusted, whatever his credentials. Of course, his research *might* be legitimate. But where should you place your bet?

Appeal to Emotion: There are a variety of fallacies that involve trying to sway an audience by appeal to emotion rather than reason. Sometimes the author or speaker seeks to *arouse fear,* a common form of argument in American politics. (Indeed, in American politics, appeals to fear are often combined with a the spicy ad populum of talking about how we are the "land of the brave". Telling em they are so so brave and then scaring the bleep out of em. Very effective.) Fear is often quite irrational and once aroused has the power to swamp reason. By arousing fear, it is often easy to get people to agree to many otherwise unsupported things. There is also the *appeal to pity.* This is when an arguer tries to get people to accept a conclusion by making them feel sorry for someone. So, for example, someone might argue that she *deserves* a better grade, a pay increase, or whatever. But rather than giving reasons relevant to *desert*, she might instead try to arouse pity. "I have been sick, my life is hard, I had all these problems, I am a recovering

Dr. Pepper addict and will probably never recover from this thing I have for Oprah." All this might be true, but not relevant to whether the person actually *deserves* the grade, pay increase, or whatever.

Of course, not all appeals to emotion are irrational. Indeed, emotion can itself be thought of as a kind of detection device, a device for detecting genuine reasons for action and belief. Fear can detect things worthy of fear. Pity can detect things worthy of pity. And often we should reach conclusions specifically because something really is fearsome or worthy of pity. Nuclear war is a fearsome thing, and this fact should affect our policies. Appealing to fear might make our actions more rational rather than less in such cases. And arguably, that something is pitiful can be a reason to act, a reason best brought home to us by arousing our pity. When do such appeals turn into fallacies? It is not always easy to know.

Appeal to ignorance: One commits the fallacy of appeal to ignorance when one, in effect, argues that since there isn't very good evidence on an issue, you should accept my view on the issue. There are two general forms. "No one has proven X, so you should accept not-X," and "No one has proven not-X, so you should accept X. The problem with this sort of argument can been seen from the possibility of "dueling appeals to ignorance." For example, some have argued that since no one has proven that God exists, we should conclude he does not. Others have argued that since no one has proven prove God does not exist, we should accept that he does. Hm . . . Not much progress there! Indeed, it would seem that the most rational response to ignorance on a topic is to *withhold* judgment rather than to commit to either side. There may be exceptions to this prescription, but anyone who thinks there is an exception has to make the case for it.

Why might such a ridiculous form of argument persuade people, besides the fact that people are not so bright? Perhaps because this form of argument actually can make sense, but only in a special situation. Suppose that we have tried to prove something, X, and have failed. And suppose further that X is the sort of thing *that we should be able to prove if it were true*. That could give us good reason to conclude that X is probably not true. Here is a simple case. No one has yet found a real Tyrannosaurs Rex alive and well in Montana. I take that to be pretty good evidence there are no Tyrannosaurs alive and well in Montana. After all, a live Tyrannosaurs is not easy to hide, especially if we consider that any such creature would have to be part of a breeding population. That none have been found is pretty good evidence there there are none to be found. And if that example seems rather fanciful, consider the following. Physicists sometimes predict that certain subatomic should exist given their best theories. Further, they might have good reason to think that if those particles do exist, a particular sort of experiment (say in the Large Hadron Collider at CERN) will find them. If the experiments do not find them, that is pretty good (though probabilistic) evidence that the particles do not exist. (I have read some commentators who say that after a point, failure to find the Higgs Boson would have been pretty good evidence that it does not exist, which is something of a bugger for certain theories in physics. Fortunately, as I understand it, the Large Hadron Collider has turned it up.)

Straw man: Anyone who has been in a fist fight knows that it can be hard to knock an opponent down. Slamming your fist into someone's jaw can hurt your hand more than their jaw, and there is always the risk that they might hit back. Ouch! It is much safer to knock around something that is soft and cannot defend itself, like a person made of straw.

Well, in the world of fisticuffs that is not likely to get you anywhere, but it might in the world of debate. Debate is not about truth, it is about winning. And this is true whether the debate takes place in a school debating competition, or among those running for president. Suppose you are in a debate and your opponent believes something. She argues for her position fairly well. . You want to refute it, but you are having a hard time of it. Winning is all and truth be damned! So you try to confuse your audience about what your opponent believes. Suppose you can make your audience think that your opponent believes something stupid, perhaps something far more extreme than she really accepts. That is the argumentative straw man. You can now knock that straw man to argumentative pieces and declare yourself the winner. If you are skillful enough, or your audience is stupid enough, you might convince them that you have actually refuted your opponent when you have only knocked down a straw man.

Straw men abound in political discussions, as well as elsewhere. Suppose your opponent has a very nuanced policy proposal on the table concerning a resolution to the conflict between Palestinians and Israelis. You might construct a straw man by grossly over simplifying it and then kicking the simplified version to pieces. “She wants to totally abandon Israel; surely you can all see how dangerous that is!” Or your opponent might have a fairly modest position on how to reform the penal system -- we have more folk in jails and prisons than any other nation on earth so, arguably, it is in need of reform. So you make your opponent seem very extreme. “He wants to do away with all punishment and to reward criminals. You must see how dangerous his proposal is!” Or suppose someone has a modestly restrictive abortion proposal. You argue “He just wants to keep all women barefoot and pregnant. Can’t you see how dangerous that is!” Cheap tricks

perhaps, but very effective in many situations.

Red herring: A red herring is a fish. I have never made the acquaintance of one, but I assume it is a rather smelly fish, especially if it is dead. Metaphorically, a red herring is something that is supposed to distract us from what is really going on. The story, if it is to believed, is that people training hunting dogs to follow a trail used to drag a read herring across the trail to see whether or not they could distract the dog. I take it that with enough abuse, dogs can learn to ignore the distraction. How does the metaphorical distraction work? At least in mystery novels, a criminal might plant a false trail, trying to get the detectives onto the wrong track. In logic, the *fallacy* of the red herring also involves a false trail. I am trying to establish a thesis, perhaps a philosophical one or perhaps one about how best to set up a school lunch program or how best to deal with Russian intervention in the Ukraine. But I find myself unable to persuade many people. So what am I to do? *Change the topic!* I try, ever so subtly, to move the discussion on to a slightly different topic. The subtly different topic is the argumentative red herring. If I do my job well, the audience might not even notice that I have shifted topics replacing the real one with a dead and rather smelly fish. I now say something persuasive about the fish. "As you can all agree, I am sure, this herring, like all herring, is really gross and stinky." The members of the audience start bobbing their heads in agreement with such an obvious truth. And , if I am careful, I can slide back to the main topic while their heads keep bobbing agreement but without their even noticing the argumentative slight of hand. Of course, it must be done subtly. Otherwise you will find people saying "That jerk did not answer the question, he shifted topics!"

Examples of red herrings are common in political debates. For example, when

arguing about particular economic policies such as a possible tax cut on capital gains, or increasing the minimum wage, politicians often shift from the value of the particular policy under consideration and start talking abstractly about the value of a strong economy generally, or family values, or this, that or the other thing. Who could disagree about a strong economy or family values or this or that? Our heads bob in sympathy and keep bobbing when the shift back is made. When someone is caught in a despicable action, say, cheating on a spouse, they may try to shift the discussion from their own craven ways to some more abstract and often imponderable issue such as "Well, what is morality anyway? Why should we be bound by mere convention?" Gee willikers, I don't know, so maybe cheating was not so bad after all? When debating the fairness of a system of grading in class, a topic close to many of your hearts, a student might well shift the topic to the importance of getting into medical school. Or here is one of my personal favorites, a student may shift from the topic of fair grading to , "Why should I have to take this course anyway?!." Well, maybe you shouldn't have to take this course. Maybe you should drop it. Talk to your adviser. But that is not relevant to the question of fair grading!

The fallacy of red herring is often committed intentionally, as an argumentative strategy. Watch politicians being interviewed as they evade questions by shifting topics. But we also pull red herrings on ourselves. A topic is hard, and your thesis is difficult to defend. You start out bravely, chin in the air and with a certain confidence in your prose But then, when the going gets tough, the tough get going with a rapid shift to a hopefully easier point. And when that one gets a bit gnarly, as it always does, another shift of topic, and then another. Finally, after ten pages, nothing has been accomplished since

there have been so many shifts in topic. And yet the paper ends with a confident assertion of victory. As if a pile of odorous fish is a victory.

False dichotomy: Arguments often proceed by a process of elimination. Suppose I know that there are three possibilities, A, B and C. If I can eliminate A and B, I am left with C. There is no fallacy here. If I later find that there was yet another alternative, say D, I can no longer be sure that C is the right answer without more work. But again, no fallacy has been committed. The fallacy of false dichotomy occurs when an arguer *intentionally* leaves out some of the alternatives. She might argue that "Since the killer is either John or Sue, and it is not John, it must be Sue," knowing full well that there are other possible killers.. Why might someone do this? In the case just given, I might actually want Sue to be convicted and executed, or I might want the real killer to escape, especially if I myself am the killer!

In public policy debates, politicians and commentators often simplify the alternatives and thereby commit the fallacy of false dichotomy. I remember an argument for our fighting in Vietnam. "Either we fully commit to that war or before we know it, the red army will be raising its flag over Washington D.C." Obviously there were other alternatives that were ignored, such as keeping out of the war and fighting off the nonexistent Vietnamese invasion force. (I suspect that a boy scout troop stationed in California could have held off the threat to D.C., and gotten merit badges for their efforts!) Throw in a few other fallacies such as slippery slopes (see below) a bunch of ad populums in the form of colorful patriotic slogans and off we go. How can we fall for such poor reasoning? Perhaps when the world gets too complex, and we are subject to threats and fear, there is something human about simplifying the world into black and

white, and maybe a touch of chartreuse. Either we build an impenetrable fence on the boarder with Mexico, or we will be swamped by a flood of illegal Mexican rapists and drug dealers. Either we end all use of fossil fuels or the world will be destroyed by climate change.

Begging the question (petitio principii): In current ordinary English, the expression "begging the question" does not seem to refer to a fallacy. Rather, it seems to mean something like "brings up the question." One might use the expression to point out that a further issue needs to be discussed. But in logic, the expression refers to fallacies. In its classic form, it involves *assuming what you are trying to prove*. In the simplest case it is called "circular reasoning." Suppose someone wants to establish something, P, and we ask them why we should accept it. And they say "Well, you see, P is true and since P is true, it follows that P." Well duh! Obviously, no progress is being made here. But isn't this downright stupid? Who would ever give such an argument, and who would be fooled by it. Well, look in the mirror. Guilty as charged! In its simplest form, as just described, it often takes the form of "argument by repetition." I have received many papers in my philosophy classes in which students adopt interesting and controversial theses but end up arguing for them by just repeating them over and over and over. They often do not even realize they are doing it. But it is also a fact that repeating something tends to convince people. If you hear something often enough, you tend to believe it. So listen to your favorite politician repeat over and over and over, often without any evidence save the repetition. And the question begger actually has ways to cover his tracks. Here is a slight modification that can make all the difference. The premise,which does nothing but repeat the conclusion, might nevertheless *be put into different words.* With a silver

tongue, and good audience connection, an orator might confuse you until you do not even know she is doing nothing but repeating the same thing in different words.

There is another form of begging the question. First, some warm up. Suppose you and I disagree on some topic and we are trying to resolve our disagreement by arguments. If any such argument is going to resolve our disagreement, the premises of the argument must embody common ground that is less controversial than the conclusion we are trying to settle on. We disagree about something. You say X and I say Y. You scratch your head and then say, in your best imitation of Socrates, "Look Levy, you accept A don't you, and B and C as well?" "Of course," I respond, "who doesn't?" Well, then you should accept X as well!" "Oh," says I, "you got me." So far there is nothing problematical here. In fact, this is what scientists often do, and do quite correctly. We have a scientific disagreement. We seek non-controversial starting points to resolve the disagreement. Often these are provided by experiments. Experimental results will sometimes provide common, non-controversial ground to resolve theoretical disputes. Fine. But a problem arises when the premises of your argument, A, B, and C, are things I do not accept. Even worse, A, B and C might be things almost no one would accept *who did not already accept the conclusion you are trying to prove.* Those premises are not common ground and are every bit as controversial as the conclusion. Appealing to them does nothing to establish the conclusion in a way that can settle the debate.

Something like this can occur in common moral disputes. I think that two of the classic arguments about abortion, at least without a lot of extra work, simply beg the question. Some one reasons that abortion is wrong because, after all, (here is the key premise) the fetus is morally the same as an adult with all the adult's rights. But this

premise is not likely to be accepted by anyone on the other side, by any one who did not already accept that abortion is wrong. The premise that the fetus is just like an adult and has all the rights of an adult is every bit as controversial as the anti-abortion conclusion. The argument begs the question. On the other hand, suppose someone starts from the premise "The fetus is nothing but a tissue, all the way till birth," and concludes that abortion is legitimate. This person is assuming a premise that is highly unlikely to be accepted by someone who is against abortion. It is every bit as controversial as the pro-choice conclusion it is used to defend. It is not in any way common ground so the argument begs the question. Rather than being independent grounds to accept conclusions, both these arguments just "preach to the choir." A great many moral arguments are like this. Here is a pair of textbook examples. "Active euthanasia is morally acceptable since it is an ethical thing to help another human being escape suffering through death." And again, "Murder is morally wrong so active euthanasia is morally wrong." Do the premises of these arguments really count as relatively non-controversial common ground? I doubt it.

Equivocation: One word can have multiple meanings. For example "right" can be a term of *evaluation*, as opposed to "wrong," it can refer to a sort of moral *claim*, as when we speak of a right to life, it can be a direction, as opposed to left, and it can refer to a general political stance, generally conservative. This phenomenon of multiple meanings does not usually cause a problem, and can lead to much amusement. (Do a search for the famous Abbot and Constello routine "Whose on First.") But it can cause problems when a an argument involves a deceptive slide between different meanings of a word. When such a slide occurs, and affects the validity or strength of an argument, the fallacy

of equivocation has occurred. Here are two textbook examples. One might reason that giving money to a charity is the right thing to do, so charities have a right to our money. The slide is between "right" meaning the right thing to do and "right" meaning a claim on us. These are related, but importantly different sense of the word "right." Again, one might reason (1) The end of life is death and (2) happiness is the end of life. So (3) death is happiness. "The end of life" in the first premise means ceasing to live. In the second, it refers to life's purpose. That the words are used with two different meanings undermines the argument. Might someone actually be fooled by these equivocations? Perhaps. A more serious, and common equivocation, is on two meanings of the word "theory." "Theory" can mean something like *a speculation* which has yet to be proven, as in the expression "that is only theory and not fact." But in science, mathematics and philosophy it sometimes has a different meaning. It can refer to a general explanation of something, an explanation that is often very well confirmed, widely accepted and not at all thought of as a mere speculation. So, when a scientist talks about general relativity theory, or quantum theory, or a mathematician talks about set theory, usually they are not using "theory" to mean some sort of speculation. This dual use of the word "theory" often leads to an unintentional equivocation, as when someone hears the expression "the theory of evolution" from a scientist. The scientist likely understands the expression in the second sense. It is a general account of how species come into and go out of existence. You would be hard pressed to find more than an handful of biologists who viewed the theory of evolution as a mere speculation, even though they know that there are speculative bits in it (When did dogs diverge from wolves anyway? That is a matter of speculation.) But many folk are not aware of this dual use of "theory." When they

here “theory of evolution,” they often respond “Even scientists admit it is just a theory, a mere speculation.” Not really.

The Modern Science of Predictable Irrationality.

Since the beginning of philosophy, philosophers have been interested in classifying some of the most common errors in reasoning. But in recent years, the list of predictable errors has been expanded, and our understanding of them has been improved, by psychologists. Interestingly, a lot of the research is done by psychologists in business schools. For who could be more interested in possible irrationality than those who wish to sell you stuff? A few more words on this at the end.

Much of the research I am interested in here is part of “heuristics and biases” research. Roughly, a heuristic is a mental shortcut used to solve problems, whether practical or cognitive. Some researchers speak of heuristics as “fast and frugal.” They are fast in that you do not have to go through a long thought process to come to a conclusion, and they are frugal in the sense that you do not require much processing power. Here is a commonly mentioned example. It might be thought that when you catch a fly ball, you must do a complex calculation dealing with everything from gravity to air resistance and which involves a great deal of math that you never learned. Not likely! Instead, you employ the fast and frugal “gaze heuristic.” The heuristic is to fix your gaze on the ball, start running, and modify your speed so that the angle of gaze is constant. Following this heuristic makes it likely that the ball will land in your hands, or on your head. There are also heuristics involved in thinking. For example, there is the “availability” heuristic. When thinking about complex issues, one does not, like a computer, process all the examples relevant to a particular issue that you have ever come

across or heard of. Rather, your mind zips to an easily available example, perhaps one from recent memory, and you draw conclusions based on that easily available example.

Heuristics often help us come to correct answers quickly and frugally. Of course, the use of heuristics is never anything but probabilistic. They are often compared to "rules of thumb." But like rules of thumb, they can lead us astray. The problem is that the heuristic might be designed to work well in some situations, but not in others. Think of a heuristic as being activated by a push of a psychological button. What happens when the button is pushed in a situation the heuristic was not designed for? No telling, but it is often bad! Think of your learned heuristics involving social interaction. If your main social interactions are with, say your high school or college friends, you might follow heuristics appropriate for that group in wildly inappropriate situations. Many a job interview has been blown by people following social heuristics that are perfectly good for interacting with college friends but not when dealing with possible employers. And I suppose many friendships have been weakened by following heuristics appropriate for business interactions when one is, instead, trying to make friends.

As with social heuristics, a cognitive heuristic which is useful in one sort of situation can backfire in others. We refer to some of these as *biases*. A psychological button is pushed, and a cognitive processes is set in motion in an inappropriate situation. Consider the availability heuristic mentioned above. We often reason on the basis of the most easily available examples. This is such a common heuristic that I assume it must be valuable somewhere, but it is also the source of many errors. For example, even today we find people reasoning, "I just don't believe smoking can be bad for you. My grandma smoked every day of her life and lived to 95." Here, an example which is

easily brought to mind overpowers whatever understanding of statistics one might have. Appeal to an easy to recall example is, well, easy. Thinking about difficult statistics and probabilities is, well, difficult.

Confirmation Bias. Though there are many cognitive biases that we bring to the table, there is one in particular I want to draw your attention to since it is relevant to everything you might do in a philosophy class. It is the confirmation bias. Suppose you have a particular thesis that you are more or less (often more) committed to. Perhaps against your will, say, because you have to write a philosophy paper on the subject, you try to evaluate this thesis. On a standard approach to rational thought, you would go about it like this. You would seek evidence, what philosophers call "arguments," for and against your thesis. You would state those arguments as strongly as possible so as to give each side its best shot. You would list and evaluate all the responses to these argument as fairly as possible, as well as the counter-responses, and so on. You would then do your best to evaluate all these arguments, responses and counter-responses so as to come to a reasonable conclusion. You are now in a position to judge whether the thesis is plausible. Except it is very unlikely you go through this process, or anything even remotely like it, very often. Odds are you did not even try, even if you thought you did.

The confirmation bias is a tendency for people to interpret, favor and recall information that confirms their preconceptions or hypotheses. For example.

1. We often only seek confirmation of our views and never seek, or we ignore, possible disconfirmation.
2. We tend to inflate the significance of anything that tends to confirm our views

> and we tend to deflate the significance of anything that disconfirms our views. This leads us to accept very weak evidence in favor of what we believe and to ignore serious problems with our views.

Sometimes this cashes out as a search for anything that confirms what we already believe, no matter how weak. And when this weak evidence is found, our effort to evaluate ends. If evidence against our view is forced upon you, it is often dismissed on weak grounds.

For example, suppose you know little about a subject and are therefore reliant on experts, as is often the case. But suppose that, even though you know next to nothing about the topic, you already "know" the answer to your question on that topic. (Yah, I am talking about YOU! And me too.) If pressed to defend your totally uninformed, but strongly held, position, you seek out an "expert," even if that person is not really an expert, who agrees with you. You then declare the matter settled. If other experts (perhaps real ones this time) are found who appose your expert, you find some way to undermine them, no matter how weak that undermining is. You then declare, chin aggressively stuck out, "If that is the best you guys can bring against me [Yah, this is personal] it only proves my position [which has nothing much to support it] must be right! And anyone who disagrees is just stooopid!" Wonderful! What would all those aliens zipping around in space ships and kidnapping humans for study think of our behavior? That might be why they think we are just fodder for their scientific experiments! (I still remember my conversation with an alien when I was abducted. I suggested that maybe it was wrong to kidnap and experiment on intelligent beings. It, the alien, stared at me for a moment and then responded, "Yes, we never experiment on intelligent beings" and then and then proceeded to remove my brain, which it discarded as not worth close inspection.

And in truth, to this day I seem to get along pretty well without it.)

At the beginning of this brief introduction to biases and heuristics, I mentioned that some of those most interested in this topic are businesses. The reason should be obvious. If you have buttons which, when pushed, get you to do irrational stuff (spend money you do not have on products you do not want, for example), they want to know about it. Ever wonder why, a number of years ago, many restaurants stopped putting dollar signs before the prices on their menus? Ever wonder why restaurant menus stopped simply telling you what was available and started explain your "options" with little blurbs bordering on poetry? "Handcrafted, Triple-Basted, Slow-Cooked, Golden-Brown, Juicy, Hand-Selected, Fresh from the Oven, Fork Tender, Delicate Mac and Cheese." Because research shows that you have buttons that can be pushed. Restaurant folk know that colorful words will make you spend far more for your Mac and Cheese, and leaving out the dollar signs makes you go for the higher priced items. Predictable, and for someone else, highly profitable irrationality.

For those interested in heuristics and biases. There are many popular works written by major figures working in this field. See Dan Ariely, *Predictably Irrational,* and Kahneman, *Thinking, Fast and Slow* for two accessible options. The Ariely is particularly fun and he has an even more accessible talk on the Ted.com website.