belief and upbringing

belief and evolution

biased seemings

the challenge of disagreement

Does what we should believe depend on why we believe it?

Last time we were considering the following rule of belief:

Seems → Belief
If it seems to you that P is true, and you have no argument against P, you should believe P.

Right at the end we considered an objection to this idea. The objection is that what seems true to us might be due to factors in our past which make our "seemings" systematically unreliable.

Here are some real life examples of this phenomenon, and the line of thought to which they might lead.

I was raised a Catholic, and still am. But I know that other people were raised in different religions, and that people tend to believe the religion in which they were raised. On reflection, I think that it is probably true that, if I had been raised a Muslim, I would probably still be a Muslim. But it is just an accident that I was born into a Catholic family. So it is an accident that I think that Catholicism is true. So, I should give up my belief in Catholicism.

Here are some real life examples of this phenomenon, and the line of thought to which they might lead.

I was raised a Catholic, and still am. But I know that other people were raised in different religions, and that people tend to believe the religion in which they were raised. On reflection, I think that it is probably true that, if I had been raised a Muslim, I would probably still be a Muslim. But it is just an accident that I was born into a Catholic family. So it is an accident that I think that Catholicism is true. So, I should give up my belief in Catholicism.

Lets call this

the problem of dependency of belief on upbringing

the problem of dependency of belief on upbringing

Here's a related example:

> I have all kinds of moral beliefs: beliefs about what is good and bad, and beliefs about what I should and should not do. But all of our moral beliefs have been shaped by a long process of evolutionary change, which is itself a process of random genetic mutations and natural section operating in tandem. But there is no reason to think that this kind of evolutionary process would favor creatures which have true moral beliefs. So, since my beliefs are the result of this kind of process, I should think that my moral beliefs are probably false, and I should discard those beliefs.

the problem of dependency of belief on upbringing

I have all kinds of moral beliefs: beliefs about what is good and bad, and beliefs about what I should and should not do. But all of our moral beliefs have been shaped by a long process of evolutionary change, which is itself a process of random genetic mutations and natural section operating in tandem. But there is no reason to think that this kind of evolutionary process would favor creatures which have true moral beliefs. So, since my beliefs are the result of this kind of process, I should think that my moral beliefs are probably false, and I should discard those beliefs.

Let's call this:

the problem of dependency of belief on evolution

the problem of dependency of belief on upbringing

the problem of dependency of belief on evolution

A third case:

> How things seem to people are due to those people's belief systems. For example, almost all people have some implicit bias with respect to race or gender, and those biases affect how things seem to those people. But (as the example shows) the belief systems which affect the way things seem can be full of false beliefs. My belief system may be full of false beliefs. So, I should not trust the way things seem to me, and should not form beliefs on that basis.

the problem of dependency of seemings on belief

the problem of dependency of belief on upbringing

the problem of dependency of belief on evolution

the problem of dependency of seemings on belief

Our last case is somewhat different from the above cases, but (as we'll see) is related to them.

I believe that God exists. But I know lots of smart people who believe that God does not exist. On reflection, I think that those people are as smart as me, have thought about the issue as much as me, and have all of the evidence that I have. So I have no particular reason to think that I am right and they are wrong; either seems just as likely. So, I should give up my belief that God exists.

the problem of disagreement

the problem of dependency of belief on upbringing

the problem of dependency of belief on evolution

the problem of dependency of seemings on belief

the problem of disagreement

We've been asking the question: under what circumstances should I discard some belief that I have? The four types of cases we've just given examples of are four cases in which it is at least tempting to think that we should discard some belief of ours.

The cases are especially interesting because they are not far-out counterexamples: each case would seem to apply to a huge number of our most important beliefs. So, if these really are cases in which we should discard our beliefs, the practical consequences are enormous.

Let's consider these kinds of cases one by one.

the problem of dependency of belief on upbringing

We've all considered the thought that we would have different beliefs if we were raised in a different society, or by a different family. And it is a familiar idea that this can, and perhaps should, lead us to doubt those beliefs.

Can we formulate a possible rule of belief which would be a candidate to explain this?

Here's a natural suggestion:

Social Dependency → No Belief
If you believe P, but would not have believed P if you had been raised in a different society, you should not believe P.

Can we formulate a possible rule of belief which would be a candidate to explain this? Here's a natural suggestion:

Social Dependency → No Belief
If you believe P, but would not have believed P if you had been raised in a different society, you should not believe P.

This has some plausibility to it. For if the only reason why you have some belief is that you were raised to believe it, shouldn't you then think that the belief is just an accident of your upbringing, and should be discarded?

But on closer examination this leads to some pretty implausible results. For the following all look reasonably plausible:

If you had been raised in the family of Genghis Khan, you would have thought that torture is permissible.

If you had been raised in ancient Greece, you would have thought that slavery is permissible.

If you had been raised in the middle ages, you would have thought that some animals came to life by spontaneous generation.

Social Dependency → No Belief
If you believe P, but would not have believed P if you had been raised in a different society, you should not believe P.

If you had been raised in the family of Genghis Khan, you would have thought that torture is permissible.

If you had been raised in ancient Greece, you would have thought that slavery is permissible.

If you had been raised in the middle ages, you would have thought that some animals came to life by spontaneous generation.

But should these facts cause you to abandon your belief in the wrongness of torture or slavery, or your belief in the truth of the theory of biogenesis? Surely not.

And that just seems to show that Social Dependency → No Belief is false.

But should these facts cause you to abandon your belief in the wrongness of torture or slavery, or your belief in the truth of the theory of biogenesis? Surely not.

And that just seems to show that Social Dependency → No Belief is false.

But that might seem to be a somewhat unsatisfactory stopping place. Can't reflection on the fact that your beliefs are just a product of certain kinds of socialization give you good reason to doubt those beliefs?

If you think that it can, then there are two options. The first would be to try to find a way to modify the above rule of belief into one which does not have consequences like the ones just listed.

But there is also another option. Here's an example:

> On reflection, I realize that the only reason why I believe that God exists is that my parents told me this. So I was trusting in my parents' reliability. But I now see no reason to trust in their reliability on this matter — I don't think that they had good reasons for their beliefs. So, I now think that the reason why I believed that God exists was not a good reason. So, I should abandon this belief.

But there is also another option. Here's an example:

> On reflection, I realize that the only reason why I believe that God exists is that my parents told me this. So I was trusting in my parents' reliability. But I now see no reason to trust in their reliability on this matter — I don't think that they had good reasons for their beliefs. So, I now think that the reason why I believed that God exists was not a good reason. So, I should abandon this belief.

This kind of reasoning seems perfectly legitimate. But it seems that it has nothing special to do with the dependence of your beliefs on your upbringing.

Someone carrying out the above line of reasoning seems to be relying on something like the following rule of belief:

> **Bad Reasons → No Belief**
> If the only reason why you believe P is that you believe Q, and you come to believe that Q is a bad reason to believe P, you should not believe P.

Someone carrying out the above line of reasoning seems to be relying on something like the following rule of belief:

Bad Reasons → No Belief
If the only reason why you believe P is that you believe Q, and you come to believe that Q is a bad reason to believe P, you should not believe P.

There are at least two ways in which you might come to find that your believing Q is a bad reason for your believing P.

First, you might discover that Q is false. For example, you might believe that South Bend has great weather only because you believe that South Bend is in California, and discover that the latter belief is false. That would be a good reason to give up your belief that South Bend has great weather.

Second, you might discover that Q does not make P likely to be true. For example, you might believe that South Bend has great weather only because you believe that South Bend is in Indiana, and discover that being in Indiana is not likely to make a city have great weather. That too would be a good reason to give up your belief that South Bend has great weather.

Social Dependency → No Belief
If you believe P, but would not have believed P if you had been raised in a different society, you should not believe P.

Bad Reasons → No Belief
If the only reason why you believe P is that you believe Q, and you come to believe that Q is a bad reason to believe P, you should not believe P.

If you find this plausible, that provides a kind of indirect reason for doubting Social Dependency → No Belief. On this view, discovering that your beliefs are due to your upbringing might well be an occasion for you examining the reasons for those beliefs.

And, when you do that, you might find that your reasons for holding those beliefs are bad. Then, plausibly, you should ditch them. But here the dependence of your beliefs on society is just the occasion for re-examination — it is not the **reason** why you should ditch them.

After all, there is no reason why you could not discover that your beliefs are due to your upbringing but, on examination, find that you have good reasons for keeping those beliefs.

the problem of dependency of belief on evolution

Let's turn to our second example of a way in which the causes of your beliefs might lead you to doubt those beliefs.

Here is a striking thought experiment from the contemporary philosopher Sharon Street:

"Evolutionary forces have played a tremendous role in shaping the content of human evaluative attitudes. ... [But] allowing our evaluative judgments to be shaped by evolutionary influences is analogous to setting out for Bermuda and letting the course of your boat be determined by the wind and tides: just as the push of the wind and tides on your boat has nothing to do with where you want to go, so the historical push of natural selection on the content of our evaluative judgments has nothing to do with evaluative truth."

"Evolutionary forces have played a tremendous role in shaping the content of human evaluative attitudes. ... [But] allowing our evaluative judgments to be shaped by evolutionary influences is analogous to setting out for Bermuda and letting the course of your boat be determined by the wind and tides: just as the push of the wind and tides on your boat has nothing to do with where you want to go, so the historical push of natural selection on the content of our evaluative judgments has nothing to do with evaluative truth."

Suppose that you tried this: you set out for Bermuda, and just let your boat go where the winds take it. Would you be very likely to end up in Bermuda? Of course not. It is not **impossible** that you would end up there — but the odds that you would are vanishingly small.

But, Street suggests, this is somewhat like the attitude of someone who takes their beliefs about right and wrong to reflect the truth about morality. This is because she thinks two things:

Our beliefs about right and wrong are largely shaped by our evolutionary history.

The odds of a process of random mutation and natural selection yielding creatures with true moral beliefs are vanishingly small.

Our beliefs about right and wrong are largely shaped by our evolutionary history.

The odds of a process of random mutation and natural selection yielding creatures with true moral beliefs are vanishingly small.

Let's look at these two assumptions in turn.

First, consider the claim that our beliefs about morality are shaped by our evolutionary history. A prominent part of our moral belief system is a set of views about our obligations to other human beings — e.g., situations in which we should help them.

Because human beings are evolved from simpler organisms, it is natural to try to explain this part of human life in terms of simpler mechanisms in these simpler organisms. And in fact we find plenty of examples of simpler organisms helping other organisms — for example, bees sacrificing themselves for the sake of the hive.

The ways in which these traits can evolve in simpler organism is itself a large topic of study in evolutionary biology. A bee's sacrifice of its life can seem puzzling from an evolutionary point of view — isn't the point of evolution that the fittest survive, and isn't a bee who gives up its own life the very paradigm of an organism whose traits will not be passed on?

Because human beings are evolved from simpler organisms, it is natural to try to explain this part of human life in terms of simpler mechanisms in these simpler organisms. And in fact we find plenty of examples of simpler organisms helping other organisms — for example, bees sacrificing themselves for the sake of the hive.

The ways in which these traits can evolve in simpler organism is itself a large topic of study in evolutionary biology. A bee's sacrifice of its life can seem puzzling from an evolutionary point of view — isn't the point of evolution that the fittest survive, and isn't a bee who gives up its own life the very paradigm of an organism whose traits will not be passed on?

But this is too simple, for a few reasons. One is that one organism helping another might increase the chances of the latter helping the former — which might increase the former's chances of survival and reproduction. This is **reciprocity**.

Of course, this won't apply in the case of the bee. But the bees in a hive are all genetically related; so if a bee sacrifices itself to help other members of the hive survive and reproduce, that will help the bee's genetic material to live on. This is **kin selection**.

This is a very superficial introduction to a complex topic, but will be enough to get us going. Reciprocity and kin selection exhibit ways in which organisms might evolve the capacity for **altruism** — helping others even when doing so is not to one's own immediate benefit.

When we look to more complex organisms, like primates, we find more complex forms of altruism. There are various theories about how simple examples of kin selection and altruism might have evolved into the kind of altruistic behavior we find in primate groups, which sometimes has no immediate explanation in terms of reciprocity or kin selection.

In larger social structures composed of psychologically complex organisms, these simpler altruistic tendencies could have evolved into more complex attitudes towards things like loyalty towards members of your group and cooperation with others, which jointly would have improved the likelihood of the group's survival.

The idea is then that our own ideas about cooperation and what we owe to each other are evolutionary descendants of these. If this scientific research program is broadly on the right track, that would seem to support Street's claims that our own attitudes about right and wrong are largely due to evolutionary processes.

The idea is then that our own ideas about cooperation and what we owe to each other are evolutionary descendants of these. If this scientific research program is broadly on the right track, that would seem to support Street's claims that our own attitudes about right and wrong are largely due to evolutionary processes.

Our beliefs about right and wrong are largely shaped by our evolutionary history.

The odds of a process of random mutation and natural selection yielding creatures with true moral beliefs are vanishingly small.

Let's look at the second claim.

Consider how natural selection works. It is just a series of random processes taking place over millions of years. Now suppose that there are a bunch of facts about what is objectively right and wrong, good and bad. What are the odds that this evolutionary process would culminate in beings who happened to evolve in such a way that they believed these claims about right and wrong?

Surely about as high as the chances of getting to Bermuda just by following the winds and tides.

Our beliefs about right and wrong are largely shaped by our evolutionary history.

The odds of a process of random mutation and natural selection yielding creatures with true moral beliefs are vanishingly small.

This argument is importantly different from our first argument about the dependence of certain beliefs on upbringing. For here we have additional reason to believe that the beliefs in question are false. This comes from the antecedent improbability of an evolutionary process giving us the right moral beliefs.

Once one sees how the argument works, it can be expanded to other domains. For example, similar evolutionary arguments have been given for the conclusion that we should think that it is very unlikely that our religious beliefs are true.

We can, by modifying our two assumptions, turn this into a valid argument for the conclusion that our beliefs about morality are very likely to be false.

We can, by modifying our two assumptions, turn this into a valid argument for the conclusion that our beliefs about morality are very likely to be false.

The evolutionary argument against moral beliefs

```
1. My moral beliefs are due to evolutionary
   processes.
2. If a moral belief is due to evolutionary
   processes, it is very unlikely to be true.
------------------
C. My moral beliefs are very unlikely to be true.
```

This is a highly simplified argument — but it is enough to bring out a few different ways in which the defender of moral beliefs might try to reply.

The evolutionary argument against moral beliefs

```
1. My moral beliefs are due to evolutionary
   processes.
2. If a moral belief is due to evolutionary
   processes, it is very unlikely to be true.
------------------
C. My moral beliefs are very unlikely to be true.
```

A first line of reply is that the argument is **self-undermining**.

One might argue as follows: this line of argument does not just apply to moral beliefs; it applies to **any** beliefs. No belief produced by random evolutionary forces is very likely to be true.

But then consider my belief that the theory of evolution is true. It follows that that belief is itself very unlikely to be true. So, the first premise of the argument is very unlikely to be true.

This objection is best understood as a dilemma. Either we say that the second premise is false, in which case the argument fails; or we say that the second premise is true, in which case it applies to all beliefs, and the first premise is false.

The evolutionary argument against moral beliefs

```
1. My moral beliefs are due to evolutionary
   processes.
2. If a moral belief is due to evolutionary
   processes, it is very unlikely to be true.
------------------
C. My moral beliefs are very unlikely to be true.
```

This objection is best understood as a dilemma. Either we say that the second premise is false, in which case the argument fails; or we say that the second premise is true, in which case it applies to all beliefs, and the first premise is false.

Against this, the proponent of the evolutionary argument is likely to distinguish between beliefs about morality and empirical beliefs about one's environment. Evolution plausibly **does** select for the truth of the latter. But those are the beliefs on which our scientific beliefs are based.

By contrast (so the reply goes) there is no evolutionary pressure for groups to arrive at the truth about morality. The only evolutionary pressure is for them to find the rules which maximize their chances of survival and reproduction.

The evolutionary argument against moral beliefs

```
1. My moral beliefs are due to evolutionary
   processes.
2. If a moral belief is due to evolutionary
   processes, it is very unlikely to be true.
------------------
C. My moral beliefs are very unlikely to be true.
```

A second line of objection is a different way of objecting to the second premise.

This is to say that moral beliefs **are** relevant to survival. Maybe "morality" is just a term for the system of rules which is most conducive to the survival of a group.

One potential worry here is that this would seem to lead to a kind of moral relativism (which we'll discuss more next week). Suppose that slavery were conducive to the survival of some group. Would that make slavery the right practice for that society to engage in?

The evolutionary argument against moral beliefs

```
1. My moral beliefs are due to evolutionary
   processes.
2. If a moral belief is due to evolutionary
   processes, it is very unlikely to be true.
------------------
C. My moral beliefs are very unlikely to be true.
```

So accepting a kind of relativism would be one way of saving the truth of moral beliefs. But suppose that you are not a relativist. How might you reply to the argument?

A third line of reply is to deny the first premise. One argument against it is that many of our moral views don't fit especially well into at least the most obvious evolutionary explanations. Consider, for example, a firefighter sacrificing her life, or, for a less dramatic example, someone donating money to a charity which serves people in another part of the world.

Is either action likely to further the propagation of one's genes? Surely not.

A third line of reply is to deny the first premise. One argument against it is that many of our moral views don't fit especially well into at least the most obvious evolutionary explanations. Consider, for example, a firefighter sacrificing her life, or, for a less dramatic example, someone donating money to a charity which serves people in another part of the world.

Is either action likely to further the propagation of one's genes? Surely not.

Against this, the proponent of the evolutionary argument might say that our view that these actions are praiseworthy comes from a kind of overgeneralization of our evolved morality. During the millennia when morality evolved, humans lived in small groups. In those groups, giving to anyone, or saving anyone, might well indirectly benefit oneself. One might then see our current endorsement of these activities as akin to our present consumption of fatty foods: an activity which was once evolutionarily useful but no longer is.

But even if this leads to a stalemate, one might deny that our moral beliefs are **fully** determined by evolution. Perhaps our moral beliefs are akin in some ways to our mathematical or scientific beliefs. About the latter, we might say that we evolved a general capacity to reason, which was evolutionarily useful, and then were able to use that ability for things (like mathematics) which weren't a part of the reason why the capacity to reason was selected for.

Against this, the proponent of the evolutionary argument might say that our view that these actions are praiseworthy comes from a kind of overgeneralization of our evolved morality. During the millennia when morality evolved, humans lived in small groups. In those groups, giving to anyone, or saving anyone, might well indirectly benefit oneself. One might then see our current endorsement of these activities as akin to our present consumption of fatty foods: an activity which was once evolutionarily useful but no longer is.

But even if this leads to a stalemate, one might deny that our moral beliefs are **fully** determined by evolution. Perhaps our moral beliefs are akin in some ways to our mathematical or scientific beliefs. About the latter, we might say that we evolved a general capacity to reason, which was evolutionarily useful, and then were able to use that ability for things (like mathematics) which weren't a part of the reason why the capacity to reason was selected for.

Just so, could we say that we evolved certain altruistic traits, and then used our reason to form a more complex moral code, which was not itself selected for?

In general, there's no good inference from the fact that X is a trait of an evolved organism to the conclusion that X was itself selected for in evolution.

The evolutionary argument against moral beliefs

```
1. My moral beliefs are due to evolutionary
   processes.
2. If a moral belief is due to evolutionary
   processes, it is very unlikely to be true.
------------------
C. My moral beliefs are very unlikely to be true.
```

Let's mention on last, somewhat more tricky, reply to this argument. Consider the following line of argument:

> It is vanishingly unlikely that a long process of genetic mutation and natural selection would lead to an organism who could ride a bicycle. I am an evolved organism. So, it is very unlikely that I can ride a bicycle.

This is clearly a bad line of reasoning: I can definitely ride a bicycle. But it is wildly unlikely that evolution would have led to bicycle-riders, and I am an evolved organism. So what as gone wrong?

Let's mention on last, somewhat more tricky, reply to this argument. Consider the following line of argument:

> It is vanishingly unlikely that a long process of genetic mutation and natural selection would lead to an organism who could ride a bicycle. I am an evolved organism. So, it is very unlikely that I can ride a bicycle.

This is clearly a bad line of reasoning: I can definitely ride a bicycle. But it is wildly unlikely that evolution would have led to bicycle-riders, and I am an evolved organism. So what as gone wrong?

Here's a plausible diagnosis. If I were (somehow) present at the start of the evolution of life on earth, and I were asked how likely it was to lead to bicycle-riders, the right response would be to say that it is very unlikely indeed.

But that is not where I am. I live in a world of bicycle-riders; so, given my current evidence, that evolution would lead to this result is not in the slightest bit unlikely.

Here's a plausible diagnosis. If I were (somehow) present at the start of the evolution of life on earth, and I were asked how likely it was to lead to bicycle-riders, the right response would be to say that it is very unlikely indeed.

But that is not where I am. I live in a world of bicycle-riders; so, given my current evidence, that evolution would lead to this result is not in the slightest bit unlikely.

Could there be a similar mistake in the evolutionary argument? One might say that, if I were (somehow) present at the start of the evolutionary process, I should judge it vanishingly unlikely that it would lead to organisms with true moral beliefs.

But that is not where I am. I am in a world with a bunch of organisms with true moral beliefs. So, given my current evidence, this is not unlikely at all.

the problem of dependency of seemings on belief

Let's turn to our third topic. This is the problem of how we can trust the way things seem to us, if the way things seem to us can be affected by beliefs which may well be false.

Let's look first at how things seem to us in our visual experiences. Some interesting studies have been done which seem to show that our background beliefs, expectations, and desires can have an effect on how things visually appear to us.

In one well-known study, white Americans were first shown a picture of either a white man's face or a Black man's face, and then shown a picture of either a tool or a gun. Under time pressure, they had to categorize what they were shown. Participants primed with a Black man's face mischaracterized tools as guns significantly more than those primed with a white man's face.

The best way to interpret this study is controversial. But what seems reasonably clear is that whether the participant saw a white face or a Black face affected whether it seemed to them that they were being shown a gun or a tool.

In one well-known study, white Americans were first shown a picture of either a white man's face or a Black man's face, and then shown a picture of either a tool or a gun. Under time pressure, they had to categorize what they were shown. Participants primed with a Black man's face mischaracterized tools as guns significantly more than those primed with a white man's face.

The best way to interpret this study is controversial. But what seems reasonably clear is that whether the participant saw a white face or a Black face affected whether it seemed to them that they were being shown a gun or a tool.

Similar results have been obtained in less politically charged contexts. In one case, people are given two beers, one of which has some balsamic vinegar in it, and asked to pick which one they liked better. A majority chose the one with balsamic vinegar in it. The experiment was then repeated with the change that participants were told in advance that one of the beers had some vinegar in it (but not which one). A majority chose the one with ought vinegar in it. Some infer that the expectation of a vinegar taste changed the way the liquid tasted to the subjects.

One reason why these cases are interesting is that they call into question Seems → Belief. If our background beliefs can affect the way things seem to us, then it is tempting to say that we should trust the seemings only if we should have the belief on which the seemings are based.

One reason why these cases are interesting is that they call into question Seems → Belief. If our background beliefs can affect the way things seem to us, then it is tempting to say that we should trust the seemings only if we should have the belief on which the seemings are based.

Here's an interesting case, from the contemporary philosopher Susanna Siegel:

> Jill, for no particular reason, has the belief that Jack is angry. This is a belief which Jill should not have.
>
> When Jill sees Jack, Jill's belief that Jack is angry at her makes Jack look angry to her — it causes it to seem to her that Jack looks angry.
>
> On the basis of the fact that it visually seems to her the Jack is angry, Jill's belief that Jack is angry at her is strengthened.

At the start, Jill should not believe that Jack is angry. If Seems → Belief is true, it looks like at the end she **should** believe that Jack is angry. But can this be right? Does Jill really have a better reason for her belief at the end than at the start?

At the start, Jill should not believe that Jack is angry. If Seems → Belief is true, it looks like at the end she **should** believe that Jack is angry. But can this be right? Does Jill really have a better reason for her belief at the end than at the start?

We have been focusing on perceptual cases. But the moral of the above cases would seem to apply even more strongly to cases of non-perceptual seemings.

Consider the way in which your political beliefs can affect what claims seem true to you. This is an instance of the well-known phenomenon of **confirmation bias**.

Cases of confirmation bias are structurally the same as the Jack/Jill case: one begins with a belief (which might well be a belief one should not have), that belief causes other claims to seem true, and those other claims support the original belief.

If Seems → Belief is true, this kind of thing is perfectly ok. One's belief in P can be justified by one's belief in Q, even if one believes Q because Q seems true and Q seems true because one believes P.

But doesn't this seem like the kind of circular reasoning we would reject in other contexts?

But doesn't this seem like the kind of circular reasoning we would reject in other contexts?

It is worth thinking about how we might modify Seems → Belief in response to cases of this kind. Here's one suggestion:

Restricted Seems → Belief
If it seems to you that P is true, and you have no argument against P, and the seeming is not caused by a belief you should not have, you should believe P.

This rule restricts the seemings you should trust to the ones that are not caused by beliefs you should not have. This would block the result that Jill should believe that Jack is angry.

The problem, though, is that it is hard to know how one could employ this rule. After all, the problem with the cases under discussion is that one can't tell from the inside when a seeming is caused by one of one's beliefs.

This rule restricts the seemings you should trust to the ones that are not caused by beliefs you should not have. This would block the result that Jill should believe that Jack is angry.

The problem, though, is that it is hard to know how one could employ this rule. After all, the problem with the cases under discussion is that one can't tell from the inside when a seeming is caused by one of one's beliefs.

Here's an idea. Perhaps you should not trust seemings when **you have good reason to think that** the seeming is based on an unjustified belief, and hence good reason to think that the seeming is unreliable:

Restricted Seems → Belief 2.0
If it seems to you that P is true, and you have no argument against P, and you have no good reason to think that the seeming is unreliable, you should believe P.

What does this say about the case of Jack and Jill? It says that, if Jill is not aware that her belief played a role in Jack seeming angry, she should form the belief that he is angry. (After all, she had no way of knowing that the seeming was unreliable.)

Restricted Seems → Belief 2.0

If it seems to you that P is true, and you have no argument against P, and you have no good reason to think that the seeming is unreliable, you should believe P.

What does this say about the case of Jack and Jill? It says that, if Jill is not aware that her belief played a role in Jack seeming angry, she should form the belief that he is angry. (After all, she had no way of knowing that the seeming was unreliable.)

But suppose that Jill is told that beliefs about people can play a big role in determining how you perceive their emotions. This information would give her reason to think that Jack's seeming angry to her is unreliable — and in that case she should not reinforce her belief that Jack is angry.

In general, it seems like the best course of action for those who accept Seems → Belief is to restrict it in some way, and to educate themselves about the various situations in which background beliefs (or other mental states) are most likely to affect how things seem to them.

the problem of disagreement

Let's turn to our last main topic of the day: the question of how we should respond to the fact that people (who seem to be our equals in every respect) disagree with us.

Let's start with an example which seems to show that we should modify our views in response to disagreement of this kind.

Let's start with an example which seems to show that we should modify our views in response to disagreement of this kind.

The horse race

Imagine that you are at a horse track with a friend. Two horses, A and B, are competing for the lead down the stretch. At the finish, it is extremely close, but it looks to you that horse A won. You are highly confident that you are correct.

Your friend then turns to you and says "I can't believe that B won."

Should you now be less confident in your initial judgement?

Splitting the bill

You are in a restaurant with some friends, and the bill comes. You've agreed to split the bill equally. You think that everyone owes $19.

Your friend says, "OK, everybody should chip in $18."

Should you now be less confident that everyone owes $19?

These are simple cases of disagreement. Many people have the intuition that, in cases like these, disagreement should lead us to revise our beliefs.

Here is one way to state this view:

The Equal Weight View

In cases of disagreement, you should give equal weight to your own opinion and the opinion of the person with whom you disagree.

There are two (related) ways to understand what exactly this view implies about the above cases.

The Equal Weight View

In cases of disagreement, you should give equal weight to your own opinion and the opinion of the person with whom you disagree.

Here is the first, and simplest:

The judgement suspension rule

If you believe P, and then come across someone who believes not-P, you should respond by suspending judgement over whether P or not-P is true (and so should they).

This seems to explain our intuitive judgements about the horse race and check splitting cases.

The Equal Weight View

In cases of disagreement, you should give equal weight to your own opinion and the opinion of the person with whom you disagree.

The judgement suspension rule

If you believe P, and then come across someone who believes not-P, you should respond by suspending judgement over whether P or not-P is true (and so should they).

But this can't handle all of the cases of disagreement we might want to think about. Suppose that you believe P, and you come across someone who has suspended belief in P. What should you do?

The natural answer to this question introduces the fact that, in ordinary life, we don't just believe or disbelieve things; we also take them to have a certain probability of being true. The probability that you take P to have is called your credence in P. Credence can be expressed as a percentage, or as a number between 0 and 1 (1 means that you are sure that P is true, 0 that you are sure that P is false).

The Equal Weight View

In cases of disagreement, you should give equal weight to your own opinion and the opinion of the person with whom you disagree.

If we take this fact about credence into account, it is natural for the proponent of the Equal Weight View to adopt the 'probability splitting rule.'

The probability splitting rule

If you assign P credence N, and come across someone who assigns P credence M, then you should assign as P's credence the average of N and M.

Suppose that both you and your friend have credence of 0.9 in your initial views about the winner of the horse race. This rule says that, on learning of your disagreement, you should both adjust your credence to 0.5.

The probability splitting rule

If you assign P credence N, and come across someone who assigns P credence M, then you should assign as P's credence the average of N and M.

Here is a different case which, many think, the Probability Splitting Rule says just the right thing about.

The poll

I put an argument on the screen, and conduct a poll, asking you to say whether the argument is valid or invalid. You confidently answer "Valid." When the poll results show up, you find to your surprise that you are the only student who answered this way.

The poll

I put an argument on the screen, and conduct a poll, asking you to say whether the argument is valid or invalid. You confidently answer "Valid." When the poll results show up, you find to your surprise that you are the only student who answered this way.

What should you say in this case? Why?

We can think of this as a case in which you have many simultaneous disagreements. Supposing for simplicity that everyone initially has credence 1 in her answer, the Probability Splitting Rule would suggest that you should lower your credence in your initial answer to 0.5, then to 0.25, then to 0.125, then to .... a small number.

Here's a problem case for the probability splitting rule:

An argument for astrology?

Astrology is the view that we can predict the events in ordinary people's lives by the time of their birth and the relative locations of the stars and planets. I have the view that astrology is completely unscientific; there's just no evidence to show that it works. But 45% of Americans (62% between the ages of 18 and 24!) think that astrology is either "scientific" or "sort of scientific." So, following the advice of The Equal Weight View, I significantly increase my credence in the scientific status of astrology.

Other, similar examples are easy to come by. 20% of Americans think Obama was born in Kenya; 30% think global warming is a hoax; etc. Should any of these facts lead me to revise my views on these topics?

A reply: we need to restrict the relevant cases of disagreement to disagreement between **epistemic peers**. This was already implicit in our earlier examples; if your friend is drunk, then you will be unlikely to lose confidence in your judgement about how to split the bill at the restaurant.

The probability splitting rule

If you assign P credence N, and come across someone who assigns P credence M, then you should assign as P's credence the average of N and M.

Does the probability splitting rule have any practical consequences?

Consider any religious, moral, or political view you have. There would seem to be plenty of people who have the same evidence as you, have thought about the issues as much as you, and are as smart as you, who have a view opposite to yours.

This suggests an argument with massive consequences for what you believe about these domains.

This suggests an argument with massive consequences for what you believe about these domains.

the disagreement → agnosticism argument

1. For every moral, political, of religious view you have, you have at least roughly as many epistemic peers who disagree with you as you have epistemic peers who agree with you.
2. The probability-splitting rule.

C. You should not have credence >0.5 about any moral, political, or religious view. (1,2)

the disagreement → agnosticism argument

1. For every moral, political, of religious view you have, you have at least roughly as many epistemic peers who disagree with you as you have epistemic peers who agree with you.
2. The probability-splitting rule.

C. You should not have credence >0.5 about any moral, political, or religious view. (1,2)

Is this argument convincing?

It looks hard to deny premise (1), for at least many of our moral, political, and religious views. So it looks like a reply to this argument must involve a rejection of the probability-splitting rule.

It looks hard to deny premise (1), for at least many of our moral, political, and religious views. So it looks like a reply to this argument must involve a rejection of the probability-splitting rule.

The probability splitting rule
If you assign P credence N, and come across someone who assigns P credence M, then you should assign as P's credence the average of N and M.

Is this plausible? Let's look at two arguments against this rule of belief.

The first is that the principle is in a certain way self-refuting. There are plenty of people who have thought about disagreement as much as you have who think that the probability-splitting rule is false.

What, given that, does the probability-splitting rule tell you to think about itself?

So there is a sense in which, given actual beliefs of your epistemic peers, this rule of belief is unstable: it recommends against itself.

The probability splitting rule
If you assign P credence N, and come across someone who assigns P credence M, then you should assign as P's credence the average of N and M.

The second argument is simpler. The main point is that this rule makes the facts about what we ought to believe oddly hostage to the beliefs of others.

It is for that reason a somewhat conservative rule of belief: it argues in favor of thinking what other people think.

Would this make it impossible to be a self-aware radical and to be rational in your beliefs?

The Equal Weight View is not the only view you might take. Here is the opposite view:

The No Weight View

In cases of disagreement, you should give no weight to the opinion of the person with whom you disagree, and should maintain your initial view.

We've already seen the problem for this kind of view: it seems to say very surprising things about the kinds of cases discussed at the outset.

One thing you might want to think about: is there some middle ground between these two rules which would be preferable to both?