

Why be moral?

Second 5-7 page paper

Due: in discussion sections, Friday, December 5 (Friday, the week after Thanksgiving)

Below are four topics for your second paper. You are welcome to come up with your own topic, though you must get the written approval of your TA by email first. If you do this, the question that your TA approves should be on the first page of your essay. The papers should be at most 5-7 pages in length, double-spaced and with reasonable margins and font.

A late penalty of 3 points per day, including weekends, will be assessed for any papers which are handed in late.

It is your responsibility to acquaint yourselves with the University's honor code, as well as with the philosophy department's guidelines regarding plagiarism. Both are linked from the course web site.

1. We discussed several responses to the problem of the apparent incompatibility of free will and determinism. Which of these do you think is the most plausible? Why?
2. David Hume held that all knowledge belongs to one of two categories: knowledge of relations of ideas and knowledge of matters of fact. Explain this distinction and the role it plays in Hume's argument that our use of induction cannot be justified. Is Hume right that all of our knowledge falls into one of these two categories? Why or why not?
3. Explain and defend what you take to be the strongest form of psychological egoism. It may be a variant of one of the views discussed in class, or one that you come up with on your own.
4. What is the problem of moral luck? What is the best solution to it?
5. Explain Frankfurt's example of Black and Jones (given in the notes for both classes on free will). How can the example be used to defend compatibilism? Does it succeed? How should an incompatibilist respond to the example?

Last time I noted that most ethical theories say that you should sometimes act in ways which are contrary to your own pleasure, your own financial prospects, or other sorts of self-interest. The psychological egoist claims that this task is impossible to carry out: we cannot, in the end, act in a way which is contrary to our own self-interest.

Last time I noted that most ethical theories say that you should sometimes act in ways which are contrary to your own pleasure, your own financial prospects, or other sorts of self-interest. The psychological egoist claims that this task is impossible to carry out: we cannot, in the end, act in a way which is contrary to our own self-interest.

Today, we set this view aside, and assume for purposes of argument that it is possible to do something which is morally right even though contrary to at least some sorts of self-interested aims. Our question today is: why should anyone ever do this? Why be moral?

One answer to this question is that one should act morally -- or, as Socrates and Glaucon say in the reading, act with justice -- because this will help you to achieve your ends. Glaucon says that this is the view of “the multitude,” and describes it as the view that acting morally

One answer to this question is that one should act morally -- or, as Socrates and Glaucon say in the reading, act with justice -- because this will help you to achieve your ends. Glaucon says that this is the view of “the multitude,” and describes it as the view that acting morally

“belongs to the toilsome class of things that must be practiced for the sake of rewards and repute due to opinion but that in itself is to be shunned as an affliction.” (358a)

One answer to this question is that one should act morally -- or, as Socrates and Glaucon say in the reading, act with justice -- because this will help you to achieve your ends. Glaucon says that this is the view of “the multitude,” and describes it as the view that acting morally

“belongs to the toilsome class of things that must be practiced for the sake of rewards and repute due to opinion but that in itself is to be shunned as an affliction.” (358a)

This is one general sort of reason for performing an action: that action will be the means to some end, which is valuable, or which we have reason to want. In this case, the ends are “rewards and repute due to opinion.” In such cases we say that the action has *instrumental value*.

One answer to this question is that one should act morally -- or, as Socrates and Glaucon say in the reading, act with justice -- because this will help you to achieve your ends. Glaucon says that this is the view of “the multitude,” and describes it as the view that acting morally

“belongs to the toilsome class of things that must be practiced for the sake of rewards and repute due to opinion but that in itself is to be shunned as an affliction.” (358a)

This is one general sort of reason for performing an action: that action will be the means to some end, which is valuable, or which we have reason to want. In this case, the ends are “rewards and repute due to opinion.” In such cases we say that the action has *instrumental value*.

Can we explain why one should be moral in terms of the instrumental value of moral action?

Glaucon's speech is, in effect, an argument that we cannot. He gives an argument for the conclusion that "the life of the unjust man is far better than that of the just man."

Glaucón's speech is, in effect, an argument that we cannot. He gives an argument for the conclusion that "the life of the unjust man is far better than that of the just man."

He does this by the comparison of an unjust man who is skilled in his injustice, and a just man who is committed only to acting justly.

Glaucon's speech is, in effect, an argument that we cannot. He gives an argument for the conclusion that "the life of the unjust man is far better than that of the just man."

He does this by the comparison of an unjust man who is skilled in his injustice, and a just man who is committed only to acting justly.

Since the unjust man is skilled in his injustice, and knows that knowledge that he is unjust will have bad consequences for him, we should imagine that his injustice will escape detection; he will give the appearance of justice, and so will be praised as a just/moral person by everyone who knows him.

Glaucón's speech is, in effect, an argument that we cannot. He gives an argument for the conclusion that "the life of the unjust man is far better than that of the just man."

He does this by the comparison of an unjust man who is skilled in his injustice, and a just man who is committed only to acting justly.

Since the unjust man is skilled in his injustice, and knows that knowledge that he is unjust will have bad consequences for him, we should imagine that his injustice will escape detection; he will give the appearance of justice, and so will be praised as a just/moral person by everyone who knows him.

In considering the just man, Glaucon says, we should imagine

Glaucón's speech is, in effect, an argument that we cannot. He gives an argument for the conclusion that "the life of the unjust man is far better than that of the just man."

He does this by the comparison of an unjust man who is skilled in his injustice, and a just man who is committed only to acting justly.

Since the unjust man is skilled in his injustice, and knows that knowledge that he is unjust will have bad consequences for him, we should imagine that his injustice will escape detection; he will give the appearance of justice, and so will be praised as a just/moral person by everyone who knows him.

In considering the just man, Glaucon says, we should imagine

"a simple and noble man who ... does not wish to seem but to be good. Then we must deprive him of the seeming. For if he is going to be thought just he will have honors and gifts because of that esteem. We cannot be sure in that case whether he is just for justice's sake or for the sake of the gifts and honors."

Glaucón's speech is, in effect, an argument that we cannot. He gives an argument for the conclusion that "the life of the unjust man is far better than that of the just man."

He does this by the comparison of an unjust man who is skilled in his injustice, and a just man who is committed only to acting justly.

Since the unjust man is skilled in his injustice, and knows that knowledge that he is unjust will have bad consequences for him, we should imagine that his injustice will escape detection; he will give the appearance of justice, and so will be praised as a just/moral person by everyone who knows him.

In considering the just man, Glaucon says, we should imagine

"a simple and noble man who ... does not wish to seem but to be good. Then we must deprive him of the seeming. For if he is going to be thought just he will have honors and gifts because of that esteem. We cannot be sure in that case whether he is just for justice's sake or for the sake of the gifts and honors."

But when we imagine the life of someone who always acts morally but is thought by others to be immoral and unjust, we see that he

Glaucón's speech is, in effect, an argument that we cannot. He gives an argument for the conclusion that "the life of the unjust man is far better than that of the just man."

He does this by the comparison of an unjust man who is skilled in his injustice, and a just man who is committed only to acting justly.

Since the unjust man is skilled in his injustice, and knows that knowledge that he is unjust will have bad consequences for him, we should imagine that his injustice will escape detection; he will give the appearance of justice, and so will be praised as a just/moral person by everyone who knows him.

In considering the just man, Glaucon says, we should imagine

"a simple and noble man who ... does not wish to seem but to be good. Then we must deprive him of the seeming. For if he is going to be thought just he will have honors and gifts because of that esteem. We cannot be sure in that case whether he is just for justice's sake or for the sake of the gifts and honors."

But when we imagine the life of someone who always acts morally but is thought by others to be immoral and unjust, we see that he

"will have to endure the lash, the rack, chains ... and so will learn his lesson that not to be but to seem just is what we ought to admire."

It is hard not to agree with Glaucon here: as described, the life of the unjust man who seems just will be much more filled with pleasure and honors than that of the just man who seems unjust. This seems to show that what has instrumental value is, as he says, not *being* just, but *seeming* just.

It is hard not to agree with Glaucon here: as described, the life of the unjust man who seems just will be much more filled with pleasure and honors than that of the just man who seems unjust. This seems to show that what has instrumental value is, as he says, not *being* just, but *seeming* just.

This does not show that being just never has instrumental value: sometimes the best way to seem just might be to actually do the right thing. But it does seem to show that if we are motivated only by instrumental considerations, we ought to aim to seem moral rather than to be moral.

It is hard not to agree with Glaucon here: as described, the life of the unjust man who seems just will be much more filled with pleasure and honors than that of the just man who seems unjust. This seems to show that what has instrumental value is, as he says, not *being* just, but *seeming* just.

This does not show that being just never has instrumental value: sometimes the best way to seem just might be to actually do the right thing. But it does seem to show that if we are motivated only by instrumental considerations, we ought to aim to seem moral rather than to be moral.

It is worth noting a limitations on Glaucon's argument: it won't show much in cases in which there is no gap between being and seeming just. Here there are two important cases: an omniscient God, and yourself. In neither case is the sort of deception on which Glaucon's argument trades readily available.

There is a further aspect to Glaucon's argument. He argues not just that seeming just is more valuable than being just, but *we all believe this to be true*. He argues for this using the example of the ring of Gyges: a ring which renders its wearer invisible, and hence makes his actions free from any consequences which might result from the opinions of others.

Glaucon imagines expanding on this example as follows:

If now there should be two such rings, and the just man should put on one and the unjust the other, no one could be found, it would seem, of such adamant temper as to persevere in justice and endure to refrain his hands from the possessions of others and not touch them, though he might with impunity take what he wished even from the market place, and enter into houses and lie with whom he pleased, and slay and loose from bonds whomsoever he would, and in all other things conduct himself among mankind as the equal of a god. And in so acting he would do no differently from the other man, but both would pursue the same course. And yet this is a great proof, one might argue, that no one is just of his own will but only from constraint, in the belief that justice is not his personal good, inasmuch as every man, when he supposes himself to have the power to do wrong, does wrong. For that there is far more profit for him personally in injustice than in justice is what every man believes, and believes truly, as the proponent of this theory will maintain. For if anyone who had got such a license within his grasp should refuse to do any wrong or lay his hands on others' possessions, he would be regarded as most pitiable and a great fool by all who took note of it, though they would praise him before one another's faces, deceiving one another because of their fear of suffering injustice. So much for this point.

Let's grant Glaucon at least this: it is not the case that we should be moral because of the instrumental value of being moral. How else might we try to answer the question, "Why be moral?"

Let's grant Glaucon at least this: it is not the case that we should be moral because of the instrumental value of being moral. How else might we try to answer the question, "Why be moral?"

One attempt to answer this question begins with the idea that it is based on a confusion. On this view, when we ask whether we *ought* to do or *should* do this or that, we are asking a question about morality. But then if we ask "Should I be moral?" we are just asking whether it is moral to be moral -- which, of course, it is. So perhaps we should respond to our question by pointing out that it is just a trivial and obvious truth that we should be moral.

Let's grant Glaucon at least this: it is not the case that we should be moral because of the instrumental value of being moral. How else might we try to answer the question, "Why be moral?"

One attempt to answer this question begins with the idea that it is based on a confusion. On this view, when we ask whether we *ought* to do or *should* do this or that, we are asking a question about morality. But then if we ask "Should I be moral?" we are just asking whether it is moral to be moral -- which, of course, it is. So perhaps we should respond to our question by pointing out that it is just a trivial and obvious truth that we should be moral.

It is natural to think that this can't be quite this easy -- one way to defend this view is to say that words like *ought and should* are ambiguous between at least two meanings: sometimes we are asking whether we *morally* ought to do this or that, while other times we are asking whether we *rationally* ought to do this or that. When we ask what we rationally ought to do, we are asking what we have most reason to do, all things considered.

Let's grant Glaucon at least this: it is not the case that we should be moral because of the instrumental value of being moral. How else might we try to answer the question, "Why be moral?"

One attempt to answer this question begins with the idea that it is based on a confusion. On this view, when we ask whether we *ought* to do or *should* do this or that, we are asking a question about morality. But then if we ask "Should I be moral?" we are just asking whether it is moral to be moral -- which, of course, it is. So perhaps we should respond to our question by pointing out that it is just a trivial and obvious truth that we should be moral.

It is natural to think that this can't be quite this easy -- one way to defend this view is to say that words like *ought and should* are ambiguous between at least two meanings: sometimes we are asking whether we *morally* ought to do this or that, while other times we are asking whether we *rationally* ought to do this or that. When we ask what we rationally ought to do, we are asking what we have most reason to do, all things considered.

To this point, one might give the following reply:

Let's grant Glaucon at least this: it is not the case that we should be moral because of the instrumental value of being moral. How else might we try to answer the question, "Why be moral?"

One attempt to answer this question begins with the idea that it is based on a confusion. On this view, when we ask whether we *ought* to do or *should* do this or that, we are asking a question about morality. But then if we ask "Should I be moral?" we are just asking whether it is moral to be moral -- which, of course, it is. So perhaps we should respond to our question by pointing out that it is just a trivial and obvious truth that we should be moral.

It is natural to think that this can't be quite this easy -- one way to defend this view is to say that words like *ought and should* are ambiguous between at least two meanings: sometimes we are asking whether we *morally* ought to do this or that, while other times we are asking whether we *rationally* ought to do this or that. When we ask what we rationally ought to do, we are asking what we have most reason to do, all things considered.

To this point, one might give the following reply:

Just as it is a truth of morality that it is moral to be moral, so it is a truth of morality that one always has most reason to do what is moral: moral reasons are always more important than non-moral reasons for action. So, all things considered, we have most reason to be moral, since this is the moral thing to do.

Let's grant Glaucon at least this: it is not the case that we should be moral because of the instrumental value of being moral. How else might we try to answer the question, "Why be moral?"

One attempt to answer this question begins with the idea that it is based on a confusion. On this view, when we ask whether we *ought* to do or *should* do this or that, we are asking a question about morality. But then if we ask "Should I be moral?" we are just asking whether it is moral to be moral -- which, of course, it is. So perhaps we should respond to our question by pointing out that it is just a trivial and obvious truth that we should be moral.

It is natural to think that this can't be quite this easy -- one way to defend this view is to say that words like *ought and should* are ambiguous between at least two meanings: sometimes we are asking whether we *morally* ought to do this or that, while other times we are asking whether we *rationally* ought to do this or that. When we ask what we rationally ought to do, we are asking what we have most reason to do, all things considered.

To this point, one might give the following reply:

Just as it is a truth of morality that it is moral to be moral, so it is a truth of morality that one always has most reason to do what is moral: moral reasons are always more important than non-moral reasons for action. So, all things considered, we have most reason to be moral, since this is the moral thing to do.

Is this a convincing argument? Why or why not?

Let's grant Glaucon at least this: it is not the case that we should be moral because of the instrumental value of being moral. How else might we try to answer the question, "Why be moral?"

One attempt to answer this question begins with the idea that it is based on a confusion. On this view, when we ask whether we *ought* to do or *should* do this or that, we are asking a question about morality. But then if we ask "Should I be moral?" we are just asking whether it is moral to be moral -- which, of course, it is. So perhaps we should respond to our question by pointing out that it is just a trivial and obvious truth that we should be moral.

It is natural to think that this can't be quite this easy -- one way to defend this view is to say that words like *ought* and *should* are ambiguous between at least two meanings: sometimes we are asking whether we *morally* ought to do this or that, while other times we are asking whether we *rationaly* ought to do this or that. When we ask what we rationally ought to do, we are asking what we have most reason to do, all things considered.

To this point, one might give the following reply:

Just as it is a truth of morality that it is moral to be moral, so it is a truth of morality that one always has most reason to do what is moral: moral reasons are always more important than non-moral reasons for action. So, all things considered, we have most reason to be moral, since this is the moral thing to do.

Is this a convincing argument? Why or why not?

An analogy between this argument and our replies to skepticism about the external world, and inductive justifications of inductive reasoning.

But can we give any explanation of why it is rational to be moral which does *not* build in the assumption that, all things considered, we always have most reason to do what is morally right?

But can we give any explanation of why it is rational to be moral which does *not* build in the assumption that, all things considered, we always have most reason to do what is morally right?

One possibility is that morality is, even if not in our individual self-interest, at least in our collective self-interest. The decision about whether or not to be moral might be analogous in this sense to decision one faces in the *prisoner's dilemma*:

But can we give any explanation of why it is rational to be moral which does *not* build in the assumption that, all things considered, we always have most reason to do what is morally right?

One possibility is that morality is, even if not in our individual self-interest, at least in our collective self-interest. The decision about whether or not to be moral might be analogous in this sense to decision one faces in the *prisoner's dilemma*:

You are one of two prisoners arrested for a crime. You, and the other prisoner, are each rational, and you each know that if you both stay silent, and don't confess, you will each be convicted of a fairly minor crime, and get 1 year in jail each. If you turn State's evidence and the other prisoner stays silent, then you will get off with nothing, and the other prisoner will get 10 years; exactly the opposite will happen if the other prisoner turns State's evidence, and you stay silent. If you both confess, you both get 5 years. Is it rational for you to confess, or stay silent?

But can we give any explanation of why it is rational to be moral which does *not* build in the assumption that, all things considered, we always have most reason to do what is morally right?

One possibility is that morality is, even if not in our individual self-interest, at least in our collective self-interest. The decision about whether or not to be moral might be analogous in this sense to decision one faces in the *prisoner's dilemma*:

You are one of two prisoners arrested for a crime. You, and the other prisoner, are each rational, and you each know that if you both stay silent, and don't confess, you will each be convicted of a fairly minor crime, and get 1 year in jail each. If you turn State's evidence and the other prisoner stays silent, then you will get off with nothing, and the other prisoner will get 10 years; exactly the opposite will happen if the other prisoner turns State's evidence, and you stay silent. If you both confess, you both get 5 years. Is it rational for you to confess, or stay silent?

But can we give any explanation of why it is rational to be moral which does *not* build in the assumption that, all things considered, we always have most reason to do what is morally right?

One possibility is that morality is, even if not in our individual self-interest, at least in our collective self-interest. The decision about whether or not to be moral might be analogous in this sense to decision one faces in the *prisoner's dilemma*:

You are one of two prisoners arrested for a crime. You, and the other prisoner, are each rational, and you each know that if you both stay silent, and don't confess, you will each be convicted of a fairly minor crime, and get 1 year in jail each. If you turn State's evidence and the other prisoner stays silent, then you will get off with nothing, and the other prisoner will get 10 years; exactly the opposite will happen if the other prisoner turns State's evidence, and you stay silent. If you both confess, you both get 5 years. Is it rational for you to confess, or stay silent?

How might morality be similar to this situation? The example of the "tragedy of the commons."

But can we give any explanation of why it is rational to be moral which does *not* build in the assumption that, all things considered, we always have most reason to do what is morally right?

One possibility is that morality is, even if not in our individual self-interest, at least in our collective self-interest. The decision about whether or not to be moral might be analogous in this sense to decision one faces in the *prisoner's dilemma*:

You are one of two prisoners arrested for a crime. You, and the other prisoner, are each rational, and you each know that if you both stay silent, and don't confess, you will each be convicted of a fairly minor crime, and get 1 year in jail each. If you turn State's evidence and the other prisoner stays silent, then you will get off with nothing, and the other prisoner will get 10 years; exactly the opposite will happen if the other prisoner turns State's evidence, and you stay silent. If you both confess, you both get 5 years. Is it rational for you to confess, or stay silent?

How might morality be similar to this situation? The example of the "tragedy of the commons."

But can we give any explanation of why it is rational to be moral which does *not* build in the assumption that, all things considered, we always have most reason to do what is morally right?

One possibility is that morality is, even if not in our individual self-interest, at least in our collective self-interest. The decision about whether or not to be moral might be analogous in this sense to decision one faces in the *prisoner's dilemma*:

You are one of two prisoners arrested for a crime. You, and the other prisoner, are each rational, and you each know that if you both stay silent, and don't confess, you will each be convicted of a fairly minor crime, and get 1 year in jail each. If you turn State's evidence and the other prisoner stays silent, then you will get off with nothing, and the other prisoner will get 10 years; exactly the opposite will happen if the other prisoner turns State's evidence, and you stay silent. If you both confess, you both get 5 years. Is it rational for you to confess, or stay silent?

How might morality be similar to this situation? The example of the "tragedy of the commons."

How can dominance reasoning be used to show that it is always in your interest to confess? What does this show about collective self-interest and the question of whether it is rational to be moral?

One way to dramatize this question of whether we can give any non-circular argument in favor of acting morally is by considering the position of a *rational egoist*: someone who holds that they always have most reason to serve their own interests. Is there any argument which could show the rational egoist that he is irrational?

One way to dramatize this question of whether we can give any non-circular argument in favor of acting morally is by considering the position of a *rational egoist*: someone who holds that they always have most reason to serve their own interests. Is there any argument which could show the rational egoist that he is irrational?

Here is one attempt to provide such an argument. It seems that one sort of irrationality results from treating like cases differently: for example, if one responds to an identical bet differently on different days, and can provide no reason for this different response, this seems to be a kind of practical irrationality. Arbitrariness of this sort is irrational.

One way to dramatize this question of whether we can give any non-circular argument in favor of acting morally is by considering the position of a *rational egoist*: someone who holds that they always have most reason to serve their own interests. Is there any argument which could show the rational egoist that he is irrational?

Here is one attempt to provide such an argument. It seems that one sort of irrationality results from treating like cases differently: for example, if one responds to an identical bet differently on different days, and can provide no reason for this different response, this seems to be a kind of practical irrationality. Arbitrariness of this sort is irrational.

But isn't the rational egoist arbitrary in just this way? The rational egoist should agree that he is a being of the same sort as other people, and hence that his interests are things of the same sort as the interests of others. But in deciding what to do, he takes his own interests into account, but not the interests of others. Why isn't this just the sort of arbitrariness that seems to be a mark of irrationality?

One way to dramatize this question of whether we can give any non-circular argument in favor of acting morally is by considering the position of a *rational egoist*: someone who holds that they always have most reason to serve their own interests. Is there any argument which could show the rational egoist that he is irrational?

Here is one attempt to provide such an argument. It seems that one sort of irrationality results from treating like cases differently: for example, if one responds to an identical bet differently on different days, and can provide no reason for this different response, this seems to be a kind of practical irrationality. Arbitrariness of this sort is irrational.

But isn't the rational egoist arbitrary in just this way? The rational egoist should agree that he is a being of the same sort as other people, and hence that his interests are things of the same sort as the interests of others. But in deciding what to do, he takes his own interests into account, but not the interests of others. Why isn't this just the sort of arbitrariness that seems to be a mark of irrationality?

But even if this line of reasoning shows that there is something wrong with rational egoism, it doesn't tell us whether or not we are, deep down, really psychological egoists -- that is, it does not answer Glaucon's argument that the story of the ring of Gyges is a "great proof that no one is just of his own will but only from constraint."