

Egoism

For the last two classes, we have been discussing the question of whether any actions are really objectively right or wrong, independently of the standards of any person or group, and whether any one is ever really morally responsible for their actions. Today we turn to a different topic; we will concede that some actions are objectively right or wrong, and ask: is it possible for people to ever act morally?

This might seem like a silly question. But most views of ethics imply that, in at least some circumstances, we ought to set aside our own interests in favor of the interests of other people, or other things. But, you might think, any such view is in trouble from the start -- after all, aren't **all** human actions done in self-interest?

The view that all of our actions are done in self-interest is called **psychological egoism**.

Psychological egoism is a universal claim: it is a claim about all human actions; the idea is that something in the nature of human beings, or in the nature of intentional action, implies that genuine **altruistic** action -- action done for the sake of others rather than oneself -- is impossible for human beings.

So it is important to distinguish psychological egoism from the claim that human beings sometimes, or quite often, act out of their own self-interest. These claims are so obviously correct as to be hardly worth discussing. The psychological egoist doesn't just claim that some actions are done for self-interested reasons, but that all are.

It should also be distinguished from the claim that we **ought** to act out of our own self-interest. This is a claim about what our actions ought to be like whereas psychological egoism is a claim about what they always in fact are like.

The view that all of our actions are done in self-interest is called **psychological egoism**.

Psychological egoism is a universal claim: it is a claim about all human actions; the idea is that something in the nature of human beings, or in the nature of intentional action, implies that genuine **altruistic** action -- action done for the sake of others rather than oneself -- is impossible for human beings.

An interesting feature of psychological egoism is that many people think that it is obviously true, whereas others think that it is obviously false. This suggests that we have to do some work in understanding what, exactly, the view says.

One clarification needed to understand the view involves knowing what should count as a genuine **action**. There are, after all, some things which I do which I clearly do not do out of self-interest -- yesterday I tripped on some stairs, and surely I didn't do *that* out of self-interest. But, equally obviously, I didn't do that for the sake of others; rather, I didn't do it for any reason at all, since I did not do it on purpose. For the purposes of the discussion of psychological egoism, we will be setting such actions to the side, and focusing only on *intentional actions* --- actions done on purpose. Psychological egoism is a claim about all intentional actions.

A second question we need to ask to understand the view is: what does it mean for an action to be **self-interested**?

The view that all of our actions are done in self-interest is called **psychological egoism**.

A second question we need to ask to understand the view is: what does it mean for an action to be **self-interested**?

One appealingly simple answer to this question is called **hedonistic psychological egoism**. This is the view stated by Jeremy Bentham in the following passage.

NATURE has placed mankind under the governance of two sovereign masters, *pain* and *pleasure*. It is for them alone to point out what we ought to do, as well as to determine what we shall do. On the one hand the standard of right and wrong, on the other the chain of causes and effects, are fastened to their throne. They govern us in all we do, in all we say, in all we think: every effort we can make to throw off our subjection, will serve but to demonstrate and confirm it. In words a man may pretend to abjure their empire: but in reality he will remain subject to it all the while.



“Present, but not voting.”

The view that all of our actions are done in self-interest is called **psychological egoism**.

A second question we need to ask to understand the view is: what does it mean for an action to be **self-interested**?

One appealingly simple answer to this question is called **hedonistic psychological egoism**. This is the view stated by Jeremy Bentham in the following passage.

NATURE has placed mankind under the governance of two sovereign masters, *pain* and *pleasure*. It is for them alone to point out what we ought to do, as well as to determine what we shall do. On the one hand the standard of right and wrong, on the other the chain of causes and effects, are fastened to their throne. They govern us in all we do, in all we say, in all we think: every effort we can make to throw off our subjection, will serve but to demonstrate and confirm it. In words a man may pretend to abjure their empire: but in reality he will remain subject to it all the while.

According to this view, all human actions are self-interested in the sense that they all aim at the **pleasure** of the person performing the action. So, on this view, all of our actions are done in the interest of causing a certain sensation in us: the sensation of pleasure.

Is this true? Do we always act in the way that we think will bring about the most pleasure?

Consider, for example, eating when you are very hungry. Perhaps sometimes we eat in order to cause pleasurable taste sensations; but surely this is not the case **every** time we eat.

Or consider my changing my daughter's diaper. This is often something which fails to cause sensations of pleasure, and yet it is surely something that I do intentionally quite often.

About any such case, the hedonistic psychological egoist can come up with something to say. For example, he can say that in the case of the diaper changing, that I believe that changing this diaper will lead to her being healthy, and observation of her being healthy will cause feelings of pleasure in me.

Are explanations of this sort plausible? Should they convince someone who is not already a psychological egoist?

The view that all of our actions are done in self-interest is called **psychological egoism**.

Of course, to raise problems with hedonistic psychological egoism is not to show that psychological egoism is false -- it just shows that the psychological egoist should come up with some definition of "self-interested action" other than "action which aims at bringing about sensations of pleasure in the agent."

One intuitively plausible idea about how to do this is to say that self-interested actions are those in which the person performing the action **aims at satisfaction of his or her own desires**. Let's call this **desire satisfaction psychological egoism**.

What does it mean to say that I always aim at the satisfaction of my desires? Here is one interpretation: every time I do something on purpose, I am moved to action by some desire of mine. And, in every such case, the desire which moves me to action is a special kind of desire: it is a desire that some other desire be satisfied.

For example, consider my action of showing up for lecture today. Perhaps it is the case that I was sitting in my office, considering my various desires, and found among them my desire not to be fired. I then came to desire that this desire not to be fired be satisfied, and it was this that caused me to come to lecture.

Maybe some actions are like this, but it is not very plausible that all actions are like this. For one thing, it makes actions way more sophisticated than they plausibly are. It seems that in many cases I am moved to action by a desire like a desire to get out of the cold, or a desire for pleasure --- I am not always moved to action **by a desire about another desire**.

We need to come up with a better interpretation of the idea that all actions are aimed at the satisfaction of our own desires.

The view that all of our actions are done in self-interest is called **psychological egoism**.

One intuitively plausible idea about how to do this is to say that self-interested actions are those in which the person performing the action **aims at satisfaction of his or her own desires**. Let's call this **desire satisfaction psychological egoism**.

We need to come up with a better interpretation of the idea that all actions are aimed at the satisfaction of our own desires.

A natural idea is that our actions are not aimed at the satisfaction of some desire in the sense that the desire which moves us to action is always a desire about some other desire, but rather the simpler fact that **we are always moved to action by our own desires**. Perhaps all of our actions are self-interested in the sense that they are always motivated by our own desires.

A benefit of this sort of 'desire satisfaction psychological egoism' is that we can give a very plausible argument that every intentional action is brought about by one of the agent's own desires: after all, if an action is **not** caused by one of the desires of the agent, then it seems as though it must have been done by accident rather than on purpose, and so is not an intentional action.

To establish the truth of the desire satisfaction version of psychological egoism, then, we need to show only one more thing: that any action caused by *my* desires is done in *my* interests.

The view that all of our actions are done in self-interest is called **psychological egoism**.

One intuitively plausible idea about how to do this is to say that self-interested actions are those in which the person performing the action **aims at satisfaction of his or her own desires**. Let's call this **desire satisfaction psychological egoism**.

To establish the truth of the desire satisfaction version of psychological egoism, then, we need to show only one more thing: that any action caused by *my* desires is done in *my* interests.

However, this thesis faces some apparent counterexamples. Consider this one:

Jane does everything for other people. She desires only that they be happy, and this desire -- that others be happy -- is what moves her to act. She has no cares about herself, her pleasure, her health, or her own well-being.

The following seems to be true: **if** there were people like Jane, then psychological egoism would be false. Jane, as described, is not doing things out of self-interest; so she is the sort of person that psychological egoism claims does not exist.

The example of Jane does not show that psychological egoism is false -- the psychological egoist can just deny that there are, or ever have been, any people like Jane. But it does show that something is wrong with this form of argument:

Everyone's actions are caused by their own desires.

Everyone's actions are selfish.

After all, Jane is a possible example of someone whose actions are caused by their own desires, but who is not motivated by self-interest. So this argument can't be valid.

The view that all of our actions are done in self-interest is called **psychological egoism**.

One intuitively plausible idea about how to do this is to say that self-interested actions are those in which the person performing the action **aims at satisfaction of his or her own desires**. Let's call this **desire satisfaction psychological egoism**.

The example of Jane does not show that psychological egoism is false -- the psychological egoist can just deny that there are, or ever have been, any people like Jane. But it does show that something is wrong with this form of argument:

Everyone's actions are caused by their own desires.

Everyone's actions are selfish.

After all, Jane is a possible example of someone whose actions are caused by their own desires, but who is not motivated by self-interest. So this argument can't be valid.

How could the psychological egoist repair this defect in the argument? One plausible idea is to add a premise to the argument, as follows:

Everyone's actions are caused by their own desires.

Everyone's desires are all, ultimately, only desires for their own well-being or self-interest.

Everyone's actions are selfish.

Arguably, adding this premise makes the argument valid. Is the argument as stated a convincing one?

The view that all of our actions are done in self-interest is called **psychological egoism**.

One intuitively plausible idea about how to do this is to say that self-interested actions are those in which the person performing the action **aims at satisfaction of his or her own desires**. Let's call this **desire satisfaction psychological egoism**.

We can sum up our discussion of the 'desire satisfaction' version of psychological egoism as follows.

On one interpretation of the theory, it says that we always act out of a desire that some desire of ours be satisfied. This interpretation seems implausible, because we sometimes act out of a desire for something -- like chocolate, or money -- other than one of our own desires.

On a second interpretation of the theory, it says that we always act out of our own desires. This version of the theory seems to be true, but it does not seem to imply that we always act out of self-interest. To get this result, we need to add the assumption that all of our desires are self-interested, which is (at least closely related to) what we were trying to establish.

But the psychological egoist is not out of options. Let's consider a third version of psychological egoism, which defines 'self-interested actions' as 'actions done in order to increase one's own welfare or well-being.' Let's call this **welfare psychological egoism**.

The view that all of our actions are done in self-interest is called **psychological egoism**.

But the psychological egoist is not out of options. Let's consider a third version of psychological egoism, which defines 'self-interested actions' as 'actions done in order to increase one's own welfare or well-being.' Let's call this **welfare psychological egoism**.

This is similar in form to hedonistic psychological egoism, but aims to improve upon that theory by allowing that we sometimes have self-interested motives other than desire for sensations of pleasure. So, for example, sometimes we might eat food to satisfy our hunger, rather than just to cause sensations of pleasure.

Why might we believe this sort of view?

One sort of argument might be based on the idea that humans evolved by natural selection. Oversimplifying, the theory of evolution by natural selection leads us to expect, in general, that evolution will favor those traits which improve the chances of the the bearer of the trait having a relatively large number of viable offspring. So, if this theory is correct, we should expect that the tendency to help others at the expense of one's own reproductive prospects will not be passed on to future generations -- unless that tendency leads, in some other way, to the propagation of your own genes.

This last point might help to explain the example of diaper-changing we discussed above. Perhaps in that sort of case, I am not acting directly in my own self-interest, but am acting in the interest of the propagation of my own genes, by doing what I can to ensure that my daughter will remain healthy until an age at which she too will be able to reproduce, and hence pass along (in part) my genetic material.

One clear and popular example of this sort of argument is the book *The Red Queen*, by Matthew Ridley, in which he claims

“no creature could ever evolve the ability to help its species at the expense of itself.”

The view that all of our actions are done in self-interest is called **psychological egoism**.

But the psychological egoist is not out of options. Let's consider a third version of psychological egoism, which defines 'self-interested actions' as 'actions done in order to increase one's own welfare or well-being.' Let's call this **welfare psychological egoism**.

One clear and popular example of this sort of argument is the book *The Red Queen*, by Matthew Ridley, in which he claims

“no creature could ever evolve the ability to help its species at the expense of itself.”

This sort of claim is fine on its own. But it does not establish any form of welfare psychological egoism, because it does not imply the following:

No evolved creature could ever help its species at the expense of itself.

This is because the fact that a given creature evolved via natural selection does not imply that every trait of that creature is one that it has because that trait was selected for by natural selection. So showing that it is impossible for evolution to select for a trait does not show that no evolved creature could have that trait.

So arguments from evolution to the impossibility of certain kinds of altruistic behavior are, to say the least, problematic.

In our other reading for today, though, Glaucon gives what can be thought of as another sort of argument for psychological egoism.

The view that all of our actions are done in self-interest is called **psychological egoism**.

But the psychological egoist is not out of options. Let's consider a third version of psychological egoism, which defines 'self-interested actions' as 'actions done in order to increase one's own welfare or well-being.' Let's call this **welfare psychological egoism**.

In our other reading for today, though, Glaucon gives what can be thought of as another sort of argument for psychological egoism.

This is based on 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.

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 adamantine 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.

The view that all of our actions are done in self-interest is called **psychological egoism**.

But the psychological egoist is not out of options. Let's consider a third version of psychological egoism, which defines 'self-interested actions' as 'actions done in order to increase one's own welfare or well-being.' Let's call this **welfare psychological egoism**.

In our other reading for today, though, Glaucon gives what can be thought of as another sort of argument for psychological egoism.

This is based on 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.

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.

Suppose that Glaucon is right, and that everyone would act this way if they had the ring of Gyges in their possession. What would this show about our actual behavior? Would it show that it was self-interested in some way? If so, why?

The view that all of our actions are done in self-interest is called **psychological egoism**.

But the psychological egoist is not out of options. Let's consider a third version of psychological egoism, which defines 'self-interested actions' as 'actions done in order to increase one's own welfare or well-being.' Let's call this **welfare psychological egoism**.

Setting this example to the side, it must be conceded there are many examples of actions which seem problematic from the point of view of welfare psychological egoism.

For example, we talked about the welfare psychological egoist's explanation of me changing my daughter's diaper in terms of transmission of my genes. But what about adoption? Many adoptive parents do enormous amounts for their children, with no chance that their genes will be transmitted to that child's offspring.

And there are many other cases, perhaps the clearest of which are cases of self-sacrifice. These can be extreme -- giving your life to save others in a war, for example -- or minimal, like a case in which you hold a door for someone in bad weather even though you are cold and don't feel like it, and are in a strange city and never expect to see them again.

As with our discussion of hedonistic psychological egoism, we can concoct stories about how these actions could be aimed at increasing one's welfare: for example, the soldier could be doing this to avoid feelings of shame in the future, and the door could be held in the expectation that that person will do similar things for you in the future.

The problem is that, as with the examples discussed in connection with hedonistic psychological egoism, these explanations seem a bit forced. One way to see this is by looking at the implausible ascriptions of beliefs they require: does the soldier really think that the shame would be so great that the best way to maximize his own well-being is immediate death which avoids the shame? Does the person holding the door really hold the wildly implausible belief that they will see this person again, and that this person will be in a position to help them, or improve their reputation?

The view that all of our actions are done in self-interest is called **psychological egoism**.

To sum up, we have considered three attempts to explain what, exactly, the psychological egoist's thesis amounts to.

The first, **hedonistic psychological egoism**, seems wildly implausible; we perform many actions which we do not expect to lead to our own pleasure.

The second, **desire satisfaction psychological egoism**, is much more plausible, but is not really a version of psychological egoism, since an action can be driven by one of my desires without being self-interested, so long as the desire is not self-interested.

The third, **welfare psychological egoism**, also faces some difficult counterexamples, and the attempt to establish this thesis on the basis of the fact that human beings evolved by natural selection seems to be a failure.

So, for now, let's assume that psychological egoism is false, and that it is **possible** to do something which is morally right even though contrary to at least some sorts of self-interest. We might still ask: why should anyone ever do this? Why be moral?

So, for now, let's assume that psychological egoism is false, and that it is **possible** to do something which is morally right even though contrary to at least some sorts of self-interest. We might still ask: why should anyone ever do this? Why be moral?

This is the question taken up by Socrates and Glaucon in the reading for today. Or, more precisely, they are considering a specific answer to this question, which might be put like this:

One should be moral because being moral has **instrumental value**: in the end, being moral is the best way to serve your interests.

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)

Glaucon's argument against this view seems convincing. What has instrumental value is not **being moral**, but rather **seeming to be moral**. One has reason to be moral only if this is the best way to seem to be moral; this will be true sometimes, but certainly not always.

One complication here involves the existence of an omniscient God.

But let's grant that we can't explain why we should be moral in terms of the instrumental value of acting morally. How else might we justify morality?

So, for now, let's assume that psychological egoism is false, and that it is **possible** to do something which is morally right even though contrary to at least some sorts of self-interest. We might still ask: why should anyone ever do this? Why be moral?

But let's grant that we can't explain why we should be moral in terms of the instrumental value of acting morally. How else might we justify morality?

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. So what we are asking is really: is it rational to be moral?

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?

But one might still wonder whether we can 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.

So, for now, let's assume that psychological egoism is false, and that it is **possible** to do something which is morally right even though contrary to at least some sorts of self-interest. We might still ask: why should anyone ever do this? Why be moral?

But one might still wonder whether we can 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 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 should be moral?

One possibility is that morality is, even if not in our individual self-interest, at least in our collective self-interest. This is sometimes illustrated by consideration of the following sort of situation:

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 2 years 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 should you reason about this situation? Is it rational to keep silent, or turn State's evidence?

Suppose that we think that it is rational to not confess - to stay silent. How might this be used to support the view that it is rational to be moral?

One might, however, question this defense of staying silent. For one thing, one might doubt whether it is rational to stay silent; the choice, after all, can be represented in the following chart:

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 should be moral?

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 2 years 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 should you reason about this situation? Is it rational to keep silent, or turn State's evidence?

Suppose that we think that it is rational to not confess - to stay silent. How might this be used to support the view that it is rational to be moral?

One might, however, question this defense of staying silent. For one thing, one might doubt whether it is rational to stay silent; the choice, after all, can be represented in the following chart:

Courses of action	Possibility 1: Your partner confesses	Possibility 2: Your partner stays silent
Confess	5 years in jail	go free
Stay silent	10 years in jail	2 years in jail

It looks like whatever your partner does, you are better off confessing. So how could it be rational to stay silent?

One might say that confessing is in your collective self-interest; but this raises a close relative of the question we were trying to answer in the first place: why would it ever be rational to act against my self-interest, in favor of the interests of some group of which I am a part?

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 2 years 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?

It looks like whatever your partner does, you are better off confessing. So how could it be rational to stay silent?

One might say that confessing is in your collective self-interest; but this raises a close relative of the question we were trying to answer in the first place: why would it ever be rational to act against my self-interest, in favor of the interests of some group of which I am a part?

We might reply by pointing out that, even if confessing is rational in a one-turn prisoner's dilemma, in a multi-turn prisoner's dilemma - in which your partner's decision might be based on your past decisions - staying silent might well be rational. Perhaps decisions about whether to be moral are, in important respects, like multi-turn prisoner's dilemma games.

However, this sort of explanation of why one ought to be moral has its limitations. First, it seems plausible that there are some cases of acting morally in which doing so is **not** in our collective self-interest. (We'll discuss some examples of this sort when we turn to discussion of consequentialism.)

Second, many moral decisions are not equivalent to multi-turn prisoner's dilemma games of the relevant sort. Remember the seeming counterexamples to psychological egoism discussed above: the examples of self-sacrifice, and of holding the door for someone in a strange city. The attempt to justify morality on the basis of an analogy with multi-turn prisoner's dilemma games is a close relative of the attempt to justify acting morally on the basis of its instrumental value - hence it is no great surprise that it is open to much the same objections.

However, this sort of explanation of why one ought to be moral has its limitations. First, it seems plausible that there are some cases of acting morally in which doing so is **not** in our collective self-interest. (We'll discuss some examples of this sort when we turn to discussion of consequentialism.)

Second, many moral decisions are not equivalent to multi-turn prisoner's dilemma games of the relevant sort. Remember the seeming counterexamples to psychological egoism discussed above: the examples of self-sacrifice, and of holding the door for someone in a strange city. The attempt to justify morality on the basis of an analogy with multi-turn prisoner's dilemma games is a close relative of the attempt to justify acting morally on the basis of its instrumental value - hence it is no great surprise that it is open to much the same objections.

So it is not clear that we have succeeded in providing a satisfactory answer to the problem posed by the rational egoist.

Here is one attempt to do better. 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?

This is not a direct justification of morality, but rather a critique of an opposed view. But to the extent that one takes rational egoism and morality to be the two most plausible ways of guiding one's behavior, this sort of argument against rational egoism might be a sort of indirect justification of morality.