

Let's start with some broad, generally acceptable concepts.

Human beings are, among other things, creatures of action. We do things. We perform actions.

Morality is the assigning of values to our actions. A value of good is assigned to some actions, meaning that these actions are those that we should perform, and a value of bad is assigned to some other actions, meaning that these actions are those that we should not perform. We say that performing good actions is the right thing to do, and that performing bad actions is the wrong thing to do.

Not all human actions are necessarily assigned a moral value. Some actions we perform are neither good nor bad, right nor wrong, prescribed nor proscribed. We refer to these actions as amoral, or not in the sphere of morality.

These broad, generally acceptable concepts are our starting place as moral philosophers. Moral philosophy encompasses so many spheres of thinking that the fact that we can come up with any starting place at all is rather surprising. Everything we've said so far would be agreed to by the Pope, Freddy Nietzsche, John Stuart Mill, and Manny Kant, not to mention Moe, Larry and Curly. I would venture that even Shemp, after squinting his eyes a bit and thinking, would concur that, so far, so good.

But, honestly, we haven't gotten very far at all. All we've done is define morality in a most neutral fashion, trying to shake from it any presuppositions about right or wrong per se. At this point our goal is merely knowing what it is we're talking about.

BEHOLD THE PARABLE OF THE PUP:

Dogs are, among other things, creatures of action. They do things. They perform actions.

Dogs that are domesticated have human masters. A human master often assigns a value of good to some of a dog's actions, meaning that these actions are those that the dog should perform, and a value of bad is assigned to some other actions, meaning that these actions are those that the dog should not perform. The human master says that performing good actions is the right thing to do, and that performing bad actions is the wrong thing to do. Good actions for dogs might include sitting, or staying, or pooping outdoors. Bad actions for dogs might include biting, jumping on the bed or pooping indoors. The determination of good and bad is entirely up to the master.

Dogs might be trained to do good things by a system of rewards, and trained not to do bad things by a system of punishments. But not all canine actions are necessarily assigned a moral value by a master. Some actions a dog performs are neither good nor bad, right nor wrong, prescribed nor proscribed. Sleeping in the backyard under a tree, chewing on a bone, drinking water, could be actions a master does not assign a value to.

In the parable of the pup, the dog does not determine right and wrong. This determination is done by the master. The master reveals to the dog which actions are good and which actions are bad.

How the master decides which actions are right and which actions are wrong is not seen in this parable.

THUS ENDS THE PARABLE OF THE PUP

The parable of the pup tells us a number of things, and raises a number of questions. The most important issue is the determination of right and wrong for the dog. The dog in the parable is a mechanical actor, following orders (or not following orders) blindly. The dog has no internal sense of right or wrong, although no doubt the dog does have an internal sense of pleasure and pain. If the dog were to act on its own, most likely it would act in favor of pleasure rather than in favor of pain. The point is that the dog is not a rational actor: the dog acts on instinct alone until its actions are modified by a master. And even obeying a master is, for a dog, instinctual, tied into its roots as a pack animal. Humans have been able to domesticate dogs because domestication is congruent with canine instincts. Humans have not been able to domesticate rattlesnakes because domestication is not congruent with pit viper instincts.

Because a dog is not a rational actor, we do not consider it a moral actor. Since it cannot make a moral determination for itself, its actions are amoral, even though from the perspective of the master certain of those actions are “good” and certain of them are “bad.” Realistically the master knows that these are trained behaviors, guided by the master’s desires, rather than conscious choices on the part of the dog. Because of the dog’s pack animal instincts, it presumably does not even choose to follow or not follow its master’s training: it’s instinct is to perform as a member of the pack, following the rules set down by the leader. When it does what the master wants, it is only following its own instincts.

AN ADDITIONAL FACT ABOUT HUMAN BEINGS

As we have said, human beings are, among other things, creatures of action. We do things. We perform actions.

Among those other things, we are creatures of thought. We are rational beings. We can think about our actions. We can choose to perform, or not to perform, those human actions that are voluntary (as compared to, say, having blood pumped by our heart, which is involuntary as long as we don’t terminate our own existence, which is certainly possible, but merely a footnote to the discussion). This ability to think is referred to as rationality, and the possession of this ability, combined with our nature as creatures of action, means that we are rational actors. Human beings are both creatures of action and creatures of reason.

As creatures of reason, we are able to think about our actions.

Let us return to the parable of the pup for a moment. One of the questions that remained unanswered was why the master selected some actions as right and some actions as wrong. Without answering that question, let's look to a comparable situation of the master and the dog in the human model.

THE NUMBER ONE SOURCE OF HUMAN MORALITY

At the moment, we are doing our best to analyze morality from a neutral, objective position. From this position, a little research in the field will bring back to us the realization that the question of what is right and what is wrong is often approached by humans not as a question of rational consideration of various actions but as a subject of spirituality. Human beings, in addition to being rational creatures of action, are often, perhaps always, creatures of faith. By faith we mean the ability to hold as true ideas and concepts that we cannot define rationally; this is not necessarily religious faith, although religion, often by its own admission, does fall into this category. The religious person does not need to "prove" the existence of God; the religious person accepts the existence of God on faith. No proof is needed, or else something other than traditional proofs are accepted as warrant enough for belief. And even the non-religious person holds beliefs in things for which proofs are impossible or at the very least difficult and elusive.

Religion is a complicated subject and the point of this essay is not to question religious beliefs which, by their very nature, are beyond questionability. Nevertheless, religion is important to the moral philosopher because so many people derive their moralities from their religion. Many religions, including those of the Judeo-Christian tradition, have moral laws given by divine revelation. These are literally the laws of God. God has said that certain actions are good, and certain actions are bad. The authority for this version of morality is ultimate. The rightness or wrongness of human actions is made clear, or at least clear enough for the scholars of the religion to establish individual analysis of smaller acts (e.g., self-defense) in the light of broader strictures ("Thou shalt not kill"). The point is, many of us—perhaps most of us—derive our sense of right and wrong, our morality, from our religion.

Which takes us back to the parable of the pup. As we said, one of the questions that remained unanswered was why the master selected some actions as right and some actions as wrong. If you're willing to accept that the model of God and human is comparable to master and dog, then the question can also be posed regarding God. Why has God selected some actions as right and others as wrong? This leads to a classical conundrum: is an action right or wrong because God says it's right or wrong, or are actions inherently right or wrong and God is simply pointing out to us which ones are which? This is probably not answerable, whereas the master and dog metaphor does seem resolvable: the master determines, for whatever reasons, what is right and wrong for the dog, and while many, many masters might agree, there is no reason why they all would. Some people might raise their dogs with radically different rules from other people. It's up to the master. The dog will follow the rules, whatever they are. Humans, on the other hand, even believing that a rule is God's rule, can choose whether or not to obey it. Part

of our inherent rationality, our minds, is the ability to use them. This ability to act on our thoughts is our will. In the religious context, this is often referred to as free will.

Human beings are rational creatures of action, which implies the possession of free will to make choices about which actions to perform. Many of us have faith in a religion that provides us with a moral framework for making those choices. But the faith does not answer the unanswerable question: Is an action right or wrong because God says it's right or wrong, or are actions inherently right or wrong and God is simply pointing out the difference to us? Whichever way religion answers this unanswerable question, the question of what is right and what is wrong can become a subject of analysis outside of a religious context. If an action's morality is inherent, then we can attempt to figure out why. And if an action's morality has been determined by God, then we can attempt to figure out why. A rationalist approach need not conflict with religious belief.

Which brings us back to our moral philosophy square one, that we are rational creatures of action. We are able to analyze our actions and we have demonstrated that even a belief in a religious morality to determine what our actions ought to be does not obviate the use of our reason to do that analysis. I think it is important for the moral philosopher to understand the importance of religion and address it as we have here before addressing pure rationalism. But I think we are now on safe ground going forward, and we are almost ready take that purely rational view of the subject. As soon as we address one other important reality.

THE OTHER NUMBER ONE SOURCES OF HUMAN MORALITY

A problem with our attempting to take a purely rational view of a subject is that our rationality is far from pure.

Our minds do not exist in a vacuum, a raw organic computer full of hard-wired operating instructions but without any data, and cannot operate as such, no matter how much we convince ourselves otherwise. We are, rather, creatures of society. We live within cultures, and on top of that, we have long maturation periods. That culture that we live in, even removed from religion, affects our rationality.

Keeping with the "raw organic computer" metaphor, culture affects the data in our mental machine, and it could even affect the processing instructions. Culture can be compared to a computer operating system, one that is powerful enough to potentially rewire the motherboard. The raw organic computer can run any of a variety of operating systems, each of which totally takes over the computer, making it a very different sort of machine for doing pretty much the same thing regardless of the operating system, just as with real computers. The operating systems Vista and Leopard and Linux all can connect to the internet or run a word processing application, but they each do it differently, with some aspects being done more easily and efficiently in one than the other, some being easier to learn, more malleable, whatever. Cultures are the same thing. Some are easy and welcoming, some are dense and closed, but they all contain the instruction set for individuals to live within a certain society, and they are in total albeit often subtle control

of much that we consider individual behavior. The play of society and the individual, or culture and personality, is complex. No individual within a culture is without that culture, to torture a phrase. Culture is a part of the individual from the moment of birth, if not the moment of conception.

The first connection of the individual personality to the group culture is through the individual's parents. This is where the long maturation period comes in. It is the parents, who are already enculturated, who are the first agents of the enculturation of the child. The parents themselves are already a part of the culture, transmitting that culture to the child. The child is very much like the dog in the parable of the pup, following whatever distinctions the parents make between good and bad behavior. The first thing the child learns about right and wrong is from the parents, which once again raises the question of where the parents/masters get it from. For the child, the source for the parents doesn't matter, however, because the child, a not yet fully rational actor, will simply follow (or not follow) the instructions of the parents. Doing the right thing is not found in doing the action itself but in following the parents' instructions, as doing the right thing for the dog is not found in doing the thing itself but obeying the master.

So it is easy to find the authority for children's morality, which resides in the parents. It is similarly easy to find the authority for religious morality which resides in the revelations of God. But is there morality attached to culture aside from religion? If not, then the parents are simply channeling the morality of their religious training to their children, and in fact, this is the usual case. But there can be more to morality in culture than the purely religious. This too distracts us from our quest for a rational approach to morality, but can only be ignored at peril to our entire analysis, which is why we're addressing it now.

Culture comprises all the common practices of a large group. For ease of analysis, we can equate a culture with a polity; let's say we're talking about an island nation, with a single state government. This island, separated from the rest of the world, has its own language, its own religion, its own music, its own art, its own history. Self-governed, it has its own laws. All of these, and more, contribute to the island's culture, and that island's culture, in turn contributes to the conceptions of morality of the island's inhabitants. Another island nearby, with any variations on its practices from the first island, might have totally different conceptions of morality. Life is complicated, and rich in detail. Island A might be monogamous while Island B is polyandrous. Each would consider the other's form of marriage immoral, but within the native culture, that form is the norm.

The preceding paragraph posits law as a part of culture, and therefore a determinant of morality. Law can also arguably be viewed as codified morality, that is, the legislated and enforced morality of a society. Personally I find this model inadequate to explain all law, but certainly some laws are exactly that. A law that claims that murder is illegal would be hard to separate from a cultural sense in the polity that murder is wrong. And certainly law is the result of a rational process of analyzing right and wrong in that a legislature of some sort has had to envision and describe that particular illegal act. But the morality of law is another footnote to this discussion, interesting perhaps, but not germane to pure rational analysis of right and wrong. Pursue it on your own time.

So we now have two gorillas in the morality room. The first is religion, but we've already explained how religion and moral philosophy can be compatible. Separating culture and moral philosophy may be more difficult, in that any attempt to rationalize anything by any individual is an attempt made within that individual's culture, and therefore congruent with that culture's "operating system." How can I know if I'm not merely rationalizing my own preexisting cultural norms?

The short answer to that is, I can't. But what I can do, as much as possible, is apply rational thought to morality on a cross-cultural basis. Whatever our conclusions are, they must be as valid in India as Guatemala as Japan as the USA. Since this is not intended as a cross-cultural analysis, and will not compare ideas from culture to culture, we will have to take our neutrality, to some degree, on faith. But as we have already explained, human beings are often, perhaps always, creatures of faith, so we do not ask for more than can be given. If at any point our analysis, because of its cultural bias, becomes too parochial, we should be taken to task on it. The promise is that, to the best of our ability, we will try to avoid that bias. It is not a hundred percent possible, but we can get close. And an understanding of morality that is close to perfect is only slightly less good than one that is completely perfect. We do not seek to prove absolutes in a short essay. We are simply trying to understand a few difficult ideas, to provide a meaningful framework for future inquiries.

MAKING MORAL CHOICES

Humans, as rational creatures of action with free will, can both choose what to do and analyze what they are doing.

Morality, as defined earlier, is the assigning of values to our actions. A value of good is assigned to some actions, meaning that these actions are those that we should perform, and a value of bad is assigned to some other actions, meaning that these actions are those that we should not perform. We say that performing good actions is the right thing to do, and that performing bad actions is the wrong thing to do. But how can we tell the difference?

There is a number of possibilities for any given action. It can be good, it can be bad, or it can be indifferent. Since humans are animals, one thing we can say with certainty is that, like the dog in the parable of the pup, while the explanation for the difference between right and wrong may elude us, we are, as literal animals, aware of the difference between pain and pleasure. We can, at an elemental level, equate good with pleasure, bad with pain, and indifferent with neither. This is a good starting point for our analysis.

Moral is equated with pleasure
Immoral is equated with pain
Amoral is equated with indifference

Problem: If we look at actions only insofar as they relate to a single individual, we may not ever be in a realm of even remotely objective morality, at least if our test is going to be pleasure and pain. An individual in a box, able to perform either pleasurable or painful actions, could perform the former and forego the latter without ever venturing into what we would seriously consider as the moral sphere. If our test is going to be pleasure and pain we should look to a number of individuals, where the model is easier to analyze, and see how an action affects not just the performer but how it affects other individuals as well.

Morality measured by pleasure and pain, therefore, is morality analyzed across the members of a group, and not a single individual. Since humans are social animals, this method of analysis appears to be relevant. With luck, our study will also provide a measure for the morality of hermits. We shall see.

So we are dealing with how the actions of one individual affect not only that individual but also other individuals. All the individuals who are a party to the action on either end need to be considered. Taking either the actor or the acted-upon out of the equation makes it a partial equation of no value. We need to analyze how everyone is affected by the action.

Moral is equated with pleasure
Immoral is equated with pain
Amoral is equated with indifference

Question: Is it the action itself that is equated with pleasure or pain, or the result of the action? A painful action can lead to pleasurable results, and vice versa.

Answer: It cannot simply be the action itself that we must test. It must be the total sum of the pleasure and/or pain of both the action and its consequences. An action and its consequences, in this measurement, cannot be separated from one another, or more to the point, although they may be evaluated separately, it is the sum of the two that is the whole test of pleasure/pain.

Since humans are animals with comparable biology, it seems reasonable to assume that they have similar concepts of pain and pleasure. Perhaps a single individual might have switched circuits, but groups of individuals tend to agree. This does force us to consult the culture, unfortunately, but let us make the assumption that reasonable people will choose similarly: that test of a "reasonable person" is a traditional one in some legal circles, and it should be good enough for us.

So, we will test an action in its entirety on the group involved according to a generally accepted schema of pain/pleasure, to determine the morality of an action. If it causes a net increase in pleasure for the group, we will deem it a moral action. If it causes a net increase in pain for the group, we will call it an immoral action. And if it causes no effect, it is amoral.

PAIN/PLEASURE CALCULUS

For argument's sake, let's select some objectively immoral actions and see how they stand the test. Granted that I seem to be skipping past the proof to the conclusion if I say I already have some objectively immoral actions we can test, but realistically there are actions that are universally held to be wrong by all religions and all cultures. That is, certain actions have already been tested by experience and been universally accepted as wrong/bad/immoral. Stealing, murder and lying would fall into this category, or at least they would fall into this category if we specify stealing something we don't need from someone who does need it, murdering an innocent child at random, and lying about an innocent person to cover our own guilt.

Stealing something we don't need from someone who does need it should be a confined enough wording of a commandment against stealing to be fairly unshakeable. If Jean Valjean steals a loaf of bread to feed his sister's starving family, one might be able to condone the action. But a rich guy stealing a loaf of bread from Jean Valjean's sister's starving family is pretty nasty, and that's what we're evaluating. In terms of pain, we are inflicting a lot of pain on a number of people, in return for the rather minor pleasure of the act of stealing by the rich guy plus that rich guy's enjoyment of the bread. His enjoyment of the bread < sister's family's enjoyment of the bread, so his enjoyment of the bread < pain of sister's family not having the bread. No matter how you evaluate the math, the act of stealing the bread seems to result in more pain than pleasure, so by the (admittedly vague) measure of the calculus we're using, it would be an immoral act.

Murdering an innocent child at random takes the act of murder and removes any chance of it being justifiable, as in, say, self-defense. The murder of the innocent child causes the pain of that death plus the pain of the loss felt by the child's family, which is easily measurable as less than the pleasure of the presumably psychotic killer performing the action. No matter how you evaluate the math, the act of murdering an innocent child at random seems to result in more pain than pleasure, so by the measure of the calculus we're using, it would be an immoral act.

Lying about an innocent person to cover our own guilt, like the other two examples, is worded in such a way that there can be no doubt, instinctively, that it is an immoral action. But does it satisfy our calculus? Let's assume that a crime was committed by Person A. Person B, for some reason, was arrested for the crime and is now on trial. Guilty Person A is called as an eyewitness to the crime, and fingers innocent Person B as the perpetrator. As a result, Person B is found guilty, and punished for the crime. Person A enjoys a certain amount of pleasure, while Person B suffers a certain amount of pain. But what if the situation is reversed and the guilty Person A is arrested and the innocent Person B testifies, and Person A is found guilty and punished? Isn't the amount of pleasure and pain identical? According to the calculus as we've presented it, wouldn't the lying here give us no net surplus on either side? Sure, you might be able to dig up some extra pain, for instance the guilty conscience bothering Person A after the fact, but that's stretching it. Take it on face, and it's pleasure = pain. Since pleasure and pain are being

churned up, we couldn't call this amoral, but we would, with no other tools to judge by, be unable to perform the required measurement calculus on the action.

You will suggest that our example is faulty, and that is the problem. But I would reply that, even if you can find a flaw here, sooner or later we could come up with a flawless example that would lead to the same conclusion. That pleasure and pain might be undeterminable doesn't even become an issue at the point where, even if they are determinable, we can't see how something we "know" is patently immoral does not register as immoral when we analyze it. Pleasure and pain alone do not seem to be enough to measure the morality of the action.

What is missing in the example of the false witness? Simple enough: culpability. The wrong persons receive both the pain and the pleasure. They have not earned these results. Their pleasure and pain are not warranted.

Can we throw something into our vague calculus to cover this contingency? Can we add that the pleasure and pain must be warranted? It seems that we have no choice. But the problem is that we're getting further and further from a workable formula for determining morality. We've already sailed past the indeterminate nature of some pleasure and pain, and now we're adding that the pleasure and pain must be deserved. That's asking an awful lot of an attempt at a simple formula. Yet we need a simple formula; a complicated, Byzantine calculus would not be particularly useful to the moral practitioner, and would only satisfy the ruminator with lots of time and, perhaps, no horse in the race.

Killing one innocent person to save the lives of more innocent people, the 2008 Sept-Oct resolution, seems to succeed on the overall less pain premise but fail on the warranted pain premise. The one innocent person, by virtue of the inherent innocence, has done nothing to warrant the pain of being killed, aside from fitting the calculus of $1 < \text{more than } 1$. By this logic, killing one innocent person would be immoral on face, the other innocent persons notwithstanding. The example of one innocent and totally healthy person being eviscerated for transplantable body parts for five other innocent albeit sick persons epitomizes this category: there's a net gain in pleasure or a net loss of pain, but the one healthy person has not warranted being eviscerated.

Nevertheless, plenty of complicated situations can be evaluated through the pleasure/pain calculus. Is it moral to sacrifice your own life to save the lives of many other people? If we can assume a base of equal innocence, then this would seem to result in a net gain of pleasure over pain. If I am very old and the others are very young, it almost seems to become a mandate, given the vast amount of potential pleasure $>$ my personal old-guy pain. And plenty of situations seem to resist evaluation through the pleasure/pain calculus. Even if we knew every single aspect, every single variable of the pleasure and pain on all sides, we wouldn't be able to make it work unless we also included warrant in the calculation. This can, perhaps, be done, but by no math easily available to the rational person attempting to make a moral decision on the fly. And let's face it: many moral decisions must be made quickly. Practical philosophers may not have a lifetime to

choreograph all the angels dancing on the head of the proverbial pin. There is no value to tests of morality that are virtually impossible to perform.

So, pleasure/pain can be used. Sometimes, in some very clear or very simple circumstances. Sometimes, however, it does not appear as if it will work. We need to look somewhere else.

INHERENT RIGHT AND WRONG

Let us go back to the beginning. Human beings are rational creatures of action with free will. Human beings are also animals. Our first attempt at establishing a scheme for determining morality was through our animal feelings of pain and pleasure. We thought that if an action resulted in more pleasure than pain it would therefore be a good action, while one resulting in more pain than pleasure would be a bad action, insofar as that pleasure and pain were measured across a group. Of course, our analysis ultimately led us to a dead end, and a close reading showed many demurrals along the way, and a lot of exceptions lurking at the edges. Maybe good and bad, morality and immorality, are something else altogether. Maybe we should be asking a deeper question: Is there anything such as right and wrong in the first place, or are we just playing a jejune mental game?

Originally we defined morality as the assigning of values to our actions. A value of good is assigned to some actions, meaning that these actions are those that we should perform, and a value of bad is assigned to some other actions, meaning that these actions are those that we should not perform. We say that performing good actions is the right thing to do, and that performing bad actions is the wrong thing to do.

With this definition, we really don't have to have right and wrong as platonic or absolute ideas preceding action. Right or wrong are our assignments of values to action, not the recognizing within the actions of an inherent rightness or wrongness. Since humans are rational creatures of actions, analyzing our actions is a logical outcome of our rationality. Preferring some actions over others is, similarly, a logical outcome of our rationality. Since we are human, we have no choice but to perform actions. A claim that right and wrong precede the action is no different from the core religious question, is an action moral because God says so, or does God say so because it is moral in the first place. If the latter is true, then attempting to understand why is attempting to read the mind of God, a human impossibility. And if an action is inherently right or wrong, God notwithstanding, attempting to understand why it is right or wrong seems about as humanly impossible as reading the mind of God.

So, we have two possibilities. Either we find out what is inherently right and wrong, and assign those values to our actions, or we find out what our actions are, and assign values of right—actions we should perform—and wrong—actions we shouldn't perform—to them. Since the former possibility is probably unattainable, we have no choice but to pursue the latter possibility. But the end result is probably the same. We are not pondering imponderables. We are pondering how to assign values to actions. We are

perfectly capable of doing this. Hence, we are not simply wagging our brains in the wind. We can proceed with our analysis.

We asked ourselves earlier, in analyzing the pleasure/pain of actions, would we be analyzing the act itself, or the results of the act, and we decided that we would analyze the totality of the two. What if we break things down, and simply analyze an action in and of itself? Is this a meaningful attempt at assigning a value of right or wrong to that action?

PICK AN ACTION. ANY ACTION. IS IT RIGHT OR WRONG?

Tough problem. Since we've agreed that a knowledge of the absoluteness of the rightness or wrongness of an action is probably unattainable, if such absoluteness indeed exists, we have to come up with something else. And we've also discounted our physical reactions (pain and pleasure) distributed over the group affected by the action. So what else is there?

Well, using my rationality, I not only have it within my mental power to analyze an action as one that I should perform, or one that I should not perform, but I do so all the time. Every action I perform as an exercise of my free will is subject to my mind deciding that action should be performed. I do not act blindly. I act according to my own mental proscriptions and prescriptions.

But do I actually follow my sense of right and wrong, and only perform actions I think are right? Not always. Sometimes I evaluate an action and I decide that the action is wrong, and do it anyhow, or else I decide that it's right, and I don't do it. My rational evaluation of the action does not determine whether I perform the action; that determination is through my will. I am perfectly capable of deciding to do something I think is wrong. But when I do, I know that it is wrong, at least according to my analysis, and I know that therefore I am performing an immoral act.

The main question is, how do I decide, in my day-to-day life, that an action is right or wrong? Generally I pull a set of criteria from various places and measure the act against those criteria. I know what is culturally thought about a particular action or comparable actions, I know what that action is like because I've performed it in the past or performed similar actions, or at least considered them, or seen others perform or refrain from performing them. In other words, I have a wealth of experience, both personal and cultural, to inform my decision. Which is why we made the point about culture earlier on; the cultural component of morality can not be removed the discussion.

So, drawing on what I can think, I determine if something is right or wrong. But how do I know if I'm correct in my judgment? Well, I can't be certain that I'm correct. But if I am convinced that something is right, I would naturally expect that others think it is right too. This is akin to Kant and the prescription that we should act according to the idea that we would want that action to become moral law. If it's right for anyone, it's right for everyone. There is an interplay that is hard to pin down between my drawing from the

culture and my contributing to the culture, but I would clearly expect that if I think something is right, I think it is right not just for me but for everyone. Moral rules, whatever they are, should apply equally. We would imagine that anyone in our position, using their rationality as we are, would come to the same conclusion of what is right and what is wrong.

By the way, in this line of thinking, and in my rejection of utility, I am implying and accepting a universal sense of humanity and human value, a very Kantian thing to do. But I'm not really basing my logic on universal humanity, so we won't go into it in any detail. Nevertheless, it is a worthwhile area of further analysis, for another time, perhaps.

Anyway, while we are not putting right and wrong to a majority-rule vote, we are drawing on cultural/personal experience to conclude a general rightness or wrongness of an action. We are concluding that moral laws are universal, even if we determine them as individuals, because we posit a norm to the culture and expect that norm to be, well, normal. We believe ourselves to be average moral determiners within the culture, our rationality no better or worse than anyone else's. Regardless of our belief in a relativistic universe, we act on the basis of a relatively objective cultural norm. (This is a set of ideas disputed by Nietzsche, by the way, if you find this logic unacceptable and wish to see a counterargument.)

And what about the differences between cultures? Are we limited on a purely relativistic cultural level? Can Americans only think like Americans, or Chinese only think like Chinese? The data suggests that, while there are certainly specific cultural mores, many ideas about right and wrong either transcend culture or at the very least are a-cultural. Some actions seem to be so right, or so wrong, that all people within their cultures, drawing on their own and their culture's experience, conclude the same as every other culture. It could be that the human animal has some moral instincts. Or, perhaps, human rationality simply always leads to certain constant decisions about the performing of certain actions. It doesn't matter, because the end result is the same. Moral law, which we begin by determining as individuals, ultimately transcends culture in some core instances (murder, theft, incest). The question of the source of cultural morality vis-à-vis instinctual morality is another which-came-first-the-chicken-or-the-tuna-fish conundrum that need not concern us in this essay.

The combination of our rationality with our experience of ourselves and our cultural experience of others, therefore, allows us to consider an action in and of itself and make a moral determination about that action, which we are free to perform or not perform regardless of its morality. Whatever it is we bring to our thinking, we bring it. We can use our rationality to say, this action is good or this action is bad. We can do it, and we do do it. If we want, we can even use a Kantian make-this-moral-law model to enhance the underlying rationality of the process. We can, it would seem, evaluate the intrinsic morality of an action. This would allow us even to make moral determinations about ourselves alone in a box. Would we want to apply our determination of right and wrong to everyone else who is alone in a box? Or for that matter, if it's moral law, would the box even matter? We could decide that suicide is wrong, for instance, based solely on the

fact that it is an act of murder, if we decided that all acts of murder except in self-defense are immoral. (Unless, I guess, you could contend suicide a version of self-defense, which would be quite a rationality-challenging twist on both concepts, if you ask me.)

Are our moral determinations really workable from this model? Have we established a test for the rightness and wrongness of actions—something curiously akin to a popularity measure when we apply the moral law idea—based entirely on the actions themselves? Can we live with the results of that test?

Unfortunately, the answer is, often, yes, and occasionally, no. Stealing would be the sort of action that would never stand up to a test of inherent morality in and of itself. Taking what does not belong to one is wrong on face. But look back at Jean Valjean's sister's starving family. We would probably all agree that it is morally acceptable, or at least not morally unacceptable, for Valjean to steal bread for them. Do we have to set up a paradigm that is so fluid that every single action is subject to a unique test of morality?

CONSEQUENTIALISM

(If we were being academic, I could point you to all sorts of resources covering bits and pieces of what I'm discussing here, but we're simply relaying the results of one particular meditation on the subject of morality. You can do your own research if you want other ideas on the subject. This essay is entirely creative, and not an outline of the subject from a teacher/course perspective. Feel free to disagree on any of it. That's what it's here for. That, and, perhaps, to get you to think about the subject for real, rather than just as the parroting of ideas you've picked up along the way. This is my own trying to think about the subject for real. The result, good or bad, is what it is.)

Consequentialism is the evaluation of actions by their results. If the results are good, then the action is good, and if the results are bad, then the action is bad. It doesn't matter what the action is in and of itself. We probably won't be able to derive universal moral laws from a consequentialist perspective unless a particular action always has an identical result. Each action will have to be tested each time, if there is a likelihood of different results. The consequentialist is, by definition, a busy thinker.

There are many forms of consequentialist ethics, and our earlier discussion of pain and pleasure could certainly be construed as a discussion of consequentialism, although limited to certain specific principles and, as we said, including an evaluation of the action itself into the calculus. What we want to do now, different from that earlier analysis, is look at results for the test of morality as juxtaposed to looking at the actions themselves for the test of morality, which is what we did in the previous section. Removing pain and pleasure and the classic utilitarian view of morality from the picture, can we just make our judgments based on outcomes?

In evaluating the morality of actions by their results, we would have to make the judgment the same way we would make the judgment about the act in and of itself. We would draw on our own instincts and reason, combined with our own and others'

experience culturally, and make a determination. If I do X—and it doesn't matter what X is—will the results Y be good or bad? In that case, X is good.. When evaluating the act itself, we said, If I do X, is that good or bad, regardless of the results Y. Results didn't matter, only the act mattered. Now the act doesn't matter, only results matter.

To be honest, evaluating the results is certainly going to be close to what we said earlier about pleasure/pain, with the same benefits and flaws. At this point we're just stripping it down to the essentials. And the big benefit to this is that we don't have to throw in a lot of qualifiers. If the end result of an action is good, as we perceive good (with our already problematic preconception), that is all we need to know.

Let's look at an example, one that is relatively easy: murder. We would not evaluate murder as murder in this consequential model. We would evaluate a particular murder or type of murder. For instance, is it a good thing (or not a bad thing) to prevent someone from killing us by killing them? With a consequential model, we would probably easily answer yes. The result of our action is saving an innocent party at the cost of killing a guilty party. On the other hand, in evaluating the same act with a deontological model (which is the fancy word for evaluation of the thing itself) we could determine that all murders are wrong, which means that if all murders are wrong then so too murder in self-defense is wrong, and in this case our innocent party would die and the guilty party would live, which doesn't jibe intuitively. But if we have to hem and haw and say some murders are right, and some murders are wrong, this does not give us much of an answer to the particular question, much less clear moral rules for the question in general. By the way, religious conscientious objectors certainly have asserted that all murders are morally wrong, so it's not as if we're being extreme in our choice of example. More people might think otherwise, but conscientious objection certainly seems reasonable to its practitioners (and, for that matter, is supported by US law). At least in this example, a consequentialist approach seems more workable as a test of morality than a deontological approach. But is it?

What we seem to get from a purely consequentialist view of morality is a more portable model, freely movable from action to action, allowing some leeway in our determination of morality. To the deontologist this portability might be considered moral turpitude, but it ties in to the idea above of setting up a paradigm that is so fluid that every single action is subject to a unique test of morality. The more we've talked about it in this essay, the more that fluid paradigm seems to be the only solution to solving the problem of morality, yet it is a solution that is not particularly helpful. We seem to be saying that no unique model of morality works in all situations, and that the subject of right or wrong, and how we perform our actions, must be subject to a mix of tests rather than a simple moral litmus paper turning blue or pink with a single dip into the test tube of action. Morality is so complicated, in other words, that there is no simple solution for it.

And that would seem to be true.

Our morality is informed by a complex web of external ideas combined with our perhaps instinctual animal natures combined with our own unique yet universal abilities to reason.

The conclusion that right and wrong exist is not hard for us to accept, but the conclusion that knowing the one from the other can be tricky should also not be hard for us to accept. Unless we are willing to follow moral precepts laid down for us by others, then we are forced to take on one of the impossible challenges of rationality.

How, then, should we act?

We should think about our actions. Do they seem right or wrong? And will they result in good or bad? More often than not, right actions will result in good and wrong actions will result in bad, and we won't have any trouble evaluating them. But every now and then there will be actions that seem to result in contradictory moral outcomes. What do we do then?

We attempt to do the best we can.

At least, in that case, we are trying to be good. And trying to be good may be the most moral action any human being can attempt to perform.