

INTRODUCTION TO ETHICS

**“BUT MEN MUST KNOW, THAT IN THIS THEATRE OF MAN’S LIFE,
IT IS RESERVED ONLY FOR GOD AND ANGELS TO BE LOOKERS ON.”**

— Francis Bacon (1561-1626),
Advancement of Learning, xx.8.

**“MEN ARE NEITHER ANGELS NOR DEVILS;
THAT MAKES MORALITY BOTH NECESSARY AND POSSIBLE.”**

— H. L. A. Hart (1907-1992)

[40] WHY BE MORAL?

How shall we live? At what shall we aim with our lives? And how shall we get on with those around us? **Ethics** is the science devoted to answering these questions and **morality** is the body of possible answers. The words ‘ethics’ and ‘morality’ are closely related, stemming from Greek and Latin words that mean “custom” (Greek: *ethos*; Latin: *mos*). In common English usage, ‘morality’ usually refers to the way that people believe they ought to behave (which may differ significantly from the way that they do in fact behave!), while ‘ethics’ refers to the discipline that studies these various beliefs; as such, ‘ethics’ is synonymous with ‘moral theory’. Occasionally we use the word ‘ethics’ to refer to a certain morality or body of moral belief (e.g., “Such behavior is inconsistent with Buddhist ethics”) or even to the moral rightness of an action (e.g., “I question the ethics of that decision”); but here I hope to limit the use of ‘ethics’ to refer to the discipline that studies such beliefs.

Ethics is a *normative* science — it concerns itself with how we *should* behave, not how we, in fact, *do* behave. How we *do* behave is studied by descriptive sciences such as sociology or anthropology.

WHAT IS THE SOURCE OF MORALITY?

Where does morality come from? How do we get it? That is, how do we arrive at our moral beliefs? And how does morality get its authority over us — if, indeed, it has any?

Conscience?

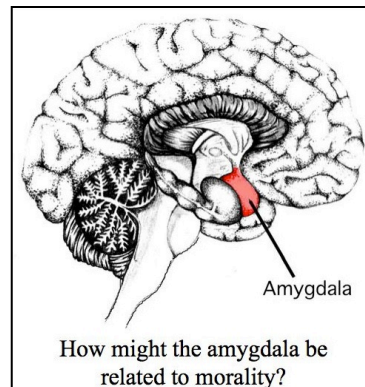
In matters of morality, the discussion sooner or later gets around to one’s conscience.

If we equate morality with “the voice of conscience,” then we need to give an account of the conscience. Let’s assume that one’s conscience is an inner sense of right and wrong, an inner voice that — if only we listen closely and carefully — will speak to us unambiguously. Where does this conscience find its voice? Perhaps it has some objective basis that makes it roughly the same for all human beings (either because God instilled it within us, or because it stems from our

**“It is better to give
than to receive.”**

Why?

And: Better *for what*?



biological makeup).¹ Or perhaps it is the product of the many social forces surrounding us? Or some combination of these? We may find that the human conscience differs too much from person to person, and that it must therefore be a result of our various (and idiosyncratic) life experiences — and if my conscience is wholly subjective in this way, it is unclear what use it could ever be to me. It may *try* to guide me, but why should I follow? Regardless of its source and its scope, I can still meaningfully ask: “Why should I give it any authority?”

First-Order Pleasure: Doing what I want

Could morality be based, ultimately, on pleasure? Everyone seems to seek pleasure, so at least this motivation is universal in its scope. But where will the pursuit of pleasure lead us? And will it take us all in the same direction, or will it set us at odds with each other?

Suppose you could do anything you wanted: What would you do? What gives you the most amount of immediate pleasure or gratification or satisfaction? (This is just for your own reflection; be honest with yourself.)

What this is exactly — and it differs from person to person, as you might imagine — is not always obvious. You might ask yourself: What do I do most often when I have a little free time? Or: What do I *think* about doing most? Or simply: What do I think about in my off moments when there isn’t anything else demanding my attention?


Perhaps what you enjoy doing most of all is not socially acceptable. If that’s the case, then you probably don’t do it as often as you’d like, or you do it only in secret, or perhaps you don’t do it at all, and instead merely think about it.

Second-Order Pleasure: Doing what I want to want

Figure out what it is that you like most, and then ask yourself a second question: In considering the entirety of your life — imagine here that you’ve grown quite old, and are now reflecting back over your many years — how would you most like to have spent it? (There’s also the different, but related question: How would you like to be remembered? What epitaph would you like engraved on your tombstone?) Presumably there are all sorts of things that you will have wanted to do; but is there anything in particular that you would like to have favored, some activity that you would like to have devoted as much of your time as the activity required? Does the activity that you “most enjoy” fit into this picture of your life anywhere?

HEDONISTIC RATS

Brain researchers in the 1950s located the existence of “pleasure centers” in the limbic system of the brain. When they implanted an electrode into the pleasure center of a laboratory rat, and arranged things so that this center could be stimulated by the rat pressing a bar, it was found that the rat would press the bar thousands of times per hour, often for fifteen to twenty hours at a stretch, before finally collapsing from exhaustion — and after reviving it would get right back to pressing the bar, preferring this to every other activity, including eating and drinking.²



¹ Research has suggested, for instance, that the amygdala is essential for learning “care based” forms of morality, e.g., empathizing with the distress of others; cf. R.J.R. Blair, “The amygdala and ventromedial prefrontal cortex in morality and psychopathy” in *Trends in Cognitive Science* 11 (2007): 387-92.

² James Olds and Peter Milner, “Positive reinforcement produced by electrical stimulation of the septal area and other regions of the rat brain,” *Journal of Comparative and Physiological Psychology*, 47: 419-29 (1954), and James Olds, “Pleasure Centers in the Brain,” *Scientific American*, 195: 105-16 (Oct. 1956).

For most of us, these two questions result in two very different answers. Your ruling passion and first-order pleasure might be playing video games, or shopping at malls, or smoking pot, or watching the soaps, or engaging in sex, or just lying about doing nothing. Yet few people, when they reflect on the entirety of their life, would look back and say: “I’m *really glad* that I played so much Nintendo! That was a truly significant thing to do with my life!”

This is perhaps the first step in constructing a personal morality: Living your life as though it really was *a* life — a life that is unified, that has meaning, and that — at its close — will have “made sense.”

If you’re lucky, the social expectations of your society push you towards a life that (for you, at least) makes sense. But we aren’t always that fortunate; and, in any event, can we assume that the larger society knows what is best for us?

How *do* we decide what kind of life we want to live? Is this wholly a matter of taste — and so, in theory, totally different for different people? If the only requirement is that the life “makes sense” (i.e., is somehow consistent) then perhaps the good life will in fact be wholly arbitrary.

On the other hand, perhaps, we might find a few constraints that push these conceptions of the good in the same direction. These constraints could include the need for psychological unity (as mentioned above), or the requirement that value reside in what is permanent, or perhaps those conditions necessary for a community in general to exist.

THE GOOD LIFE VS THE GOOD SOCIETY

A life that “makes sense” is what we could call “The Good Life.” Everyone has some idea of what this is, and in our pluralistic society we quickly discover that there are many different conceptions of this Good Life.

How we arrive at our ideas of the good life probably has more to do with religion or art or psychoanalysis than it does with rational argumentation. Why do some people decide that they should give away their belongings and serve the poor in some back alley of Calcutta? Why do others devote their lives to the amassing of material wealth? Or to writing the definitive novel? Or to raising children? And, of course, many people lack devotion to much of anything, and simply put in time before they die (normally watching a lot of television during the interim).

There may be only one thing that a philosophy professor can do with respect to this question about the good life: To encourage you to think closely about what you really believe to be important, and then to encourage you to pursue these things, so that in your old age you don’t look back on your life only to find it a wasteland of ill-spent years. “Live so as not to regret the past!” — not so much a maxim for living the good life, as a call to take the good life seriously, whatever it turns out to be for you. Another way to ask this question: “What is your road to happiness?”

“What is the Good Society?” might seem to be the same question as “What is the Good Life?”, but it is in fact rather different, since it presupposes a more basic question, namely, “How must I constrain my behavior so that life in a community is possible?” For instance, there exist in our society certain individuals who would like everyone else to adopt their own vision of the good life, with the resulting society being what in their mind is the Good Society — but these social visions rarely strike the rest of us as especially good, for the obvious reason that they often require

OUR TASK

“To have been young, and then to grow older, and finally to die, is a very mediocre form of human existence; this merit belongs to every animal. But the unification of the different stages of life...is the task set for humans.”

— Søren Kierkegaard (1813-55),
Concluding Unscientific Postscript (1846)

GOMPERS’ VISION

Samuel Gompers (1850-1924), the founder and first president of the American Federation of Labor (AFL), offered these thoughts on the good society:

“What does Labor want?

We want more schoolhouses and fewer jails,
more books and less guns,
more learning and less vice,
more leisure and less greed,
more justice and less revenge.
We want more opportunities
to cultivate our better natures.”

So what is *your* vision of the good society?

us to abandon our own conception of the good life. So perhaps it is best to leave off all talk of a “Good Society” and speak instead of “The Merely Possible Society,” or more precisely, “The Possible Society containing Diverse Conceptions of the Good Life.”

While we will likely never reach a consensus as to the nature of the good life, we may be able to reach a consensus on how best to deal with these differences among ourselves. We need to find principles for justifying our behavior in the eyes of all others, principles that will also justify the various laws and policies of the community.

This focuses specifically on how I am to relate to other people, and asks nothing about my vision of happiness as such; it *does* assume that a pre-condition of my happiness is the ability to live in society (which, for most people, is a safe assumption).

How wide is the community? Does the ideal of our social relations consist in those conditions necessary for social relations to exist at all? This might be a kind of “minimal morality,” but a true morality will stretch us — it will make our community not merely *possible*, but *desirable*. As humans we are interested in more than mere survival. Like Nietzsche’s *Übermensch*, we want to stretch, to flourish, to press the limits. To merely survive is to be an object of pity; but to really live is to sing and dance and invent and explore...to be fully human. Surely morality speaks to our flourishing if it speaks to anything. Morality is for life, not for mere survival; it speaks to us of ideals, not of mere necessary conditions.

This, finally, might give us the difference between the good life and the good society, the difference between the private and the public spheres of morality. What we can demand, as members of a pluralistic society, is a “minimal ethic” that cuts roughly along the lines of those necessary conditions for social existence. The pushing, flourishing part extends beyond this and occurs on behalf of the individual and her private conception of the good (or in smaller communities, such as religious or intentional communities, where all the individuals share a common conception of the good).

HOW DO WE DISTINGUISH AN ETHICS-BASED MORALITY FROM ... ?

... Conventional Morality

Unlike conventional morality, moral theory is reflective. It strives to understand the principles underlying our moral intuitions, and then to order our moral beliefs into a coherent system. Conventional morality may or may not be well thought out, but normally it is not, since people tend to accept it unthinkingly from cultural pressures and from the influence of authority figures when they were still unreflective children.

... The Law

Clearly the law and morality are quite different. The law is whatever the ruling authority (e.g., a monarch, a parliament, Congress) says must be done or not done. But there can be much that is considered morally right or wrong

CONFUCIAN RECIPROCITY

Someone said: “What about ‘Repay hostility with kindness’?” The Master said: “How then do you repay kindness? Repay hostility with uprightness and repay kindness with kindness.” [...]

Zigong asked: “Is there a single word such that one could practice it throughout one’s life?” The Master said: “Reciprocity, perhaps? Do not inflict on others what you yourself would not wish done to you.”

— Confucius (6TH-5TH CENTURY BCE)
Analects, 58, 62

JEWISH RECIPROCITY

“What is hateful to you, do not to your neighbor — that is the whole Torah, while the rest is commentary; go and learn it.”

— Hillel (1ST C. BCE-1ST C. CE)

HINDU RECIPROCITY

“This is the sum of all true righteousness: Deal with others as you would yourself be dealt. Do nothing to your neighbor which you would not have them do to you.”

— from the *Mahabharat* (C. 150 BCE)

CHRISTIAN RECIPROCITY

So whatever you wish that men would do to you, do so to them; for this is the law and the prophets. (Matthew 7:12; LATE 1ST CENTURY CE)

Give to everyone one who begs from you; and of him who takes away your goods do not ask them again. And as you wish that men would do to you, do so to them. (Luke 6:30-31; LATE 1ST CENTURY CE)

that is not required or forbidden by the law, and similarly there may well be laws that are immoral. For instance, many would say this of the Fugitive Slave Law of 1850, which required citizens to help return runaway slaves to their masters (this law was even upheld by the U.S. Supreme Court in the Dred Scott decision of 1857).

... Obedience to Authority

Moral agents are normally held to be responsible for their actions; consequently, insofar as they are moral agents, they must not blindly obey authority. Rather, they must first assess the command placed on them and make certain that carrying it out is morally permissible. This is true whether the authority is political (in the form of rulers or laws) or religious (in the form of priests, scripture, or the wider congregation).

... Etiquette

That which falls within the purview of advice from Miss Manners and Emily Post is what we call ‘etiquette’, the art of being polite, which would seem to be entirely defined by one’s culture. Whether we eat with a knife and fork, or with our hands — and how we hold our forks, or which hand we use to hold it — will differ from culture to culture, and has little basis in the nature of humans as such.

In matters of etiquette as well as morality, we speak in terms of ‘ought’ and ‘should’, and so it might seem that they are two aspects of the same normative endeavor. If etiquette and morality are part of the same game, however, it is clear that moral matters are more serious. My moral obligations will always override any duties I might have to be polite. More importantly, while it is non-controversial to claim that etiquette is strictly governed by social convention and has a high degree of arbitrariness, it is far from clear that the same can be said of morality.

... Wishful Thinking

Morality is not meant to be some ideal system that is noble in theory but useless in practice. Any moral theory that is impractical or unworkable is a bad theory and needs to be replaced. Morality is for acting and for getting on with our business of living well.

... Prudence

Morality is commonly thought of as different from prudence, insofar as it often involves placing constraints on our selfish behavior; prudence, on the other hand, is just acting in one’s own best interests. Although prudence and morality may often tell us to do the same thing, they often disagree.

This distinction can also be explained in terms of **action-guides**, where an action-guide is of the form: “I ought not to do X” or “I ought to do X.” Prudential action-guides are **self-regarding** (they always appeal to your own well being and desires), whereas moral action-guides are typically **other-regarding**.

With many issues, there might be both prudential and moral action guides (or claims), and it is frequently helpful to sort these out. Consider vegetarianism: there are various reasons why someone might want to eliminate or reduce their meat consumption. Some are merely prudential (such as benefiting one’s health or saving money), others are explicitly moral (such as minimizing the suffering of animals or freeing up scarce resources for others, or reducing

VAMPIRE BAT RECIPROCITY

Vampire bats get their food by drinking the blood of larger mammals, typically domestic livestock. They do this by making a small incision in the skin and lapping the blood with their tongue (their saliva contains an anticoagulant that keeps the blood flowing). In the space of fifteen minutes, a bat can drink up to 40 percent its own weight in blood: an insignificant amount for the donor animal, but a good night’s meal for the bat. Not infrequently, however, a bat will fail to find a suitable mammal to feed from, and thus will return to its colony hungry. Bats cannot fast for very long, and what is noteworthy in their behavior is that those bats with full stomachs will share their food with the less fortunate. But they don’t share with just any hungry bat; rather, they prefer relatives, as well as unrelated bats with whom they often associate. Specifically, they are more likely to share with a bat that has shared with them in the past. Bats who refuse to share their meals are much less likely to be helped the next time they come home hungry.

Now imagine two separate colonies of bats: one that has developed this strategy of sharing, and one that has not. Which colony will be more reproductively successful? [See the research of Gerald Wilkinson, “Reciprocal Food Sharing in the Vampire Bat” in *Nature* 308 (1984): 181-84.]

environmental waste and pollution). In the chapter on “Psychological Egoism,” below, we will consider the claim that ethics is simply a glorified form of prudence.

... Religion

A March 2002 survey conducted by the Pew Charitable Trust found that 47% of U.S. respondents think that “belief in God” is necessary for one to be moral.³ This varies from 56% of those respondents with a high school education or less to 33% for college graduates.

Moral thinking and religious thinking often arise together. Nonetheless, religion and morality are not the same thing, nor does one appear to be a mere subset of the other. We can all, for instance, think of religious leaders and organized religions performing apparently immoral acts. For example, Jim Jones forcing the members of the People’s Temple to drink cyanide laced Kool-Aid at their Guyana temple on November 18th, 1978, which killed 914 men, women, and children, or of the actions of the Roman Catholic Inquisition, or of the man who killed several people on the Staten Island ferry in 1995 because God had told him to do so.

Should scripture guide us? If so, then *which* scripture? And *whose* interpretation? Should the Holy Spirit guide us? The Staten Island killer thought he was listening to the Holy Spirit. Should the organized church guide us? If so, then *which* church?

The question at work here — first raised in Plato’s *Euthyphro* — is whether what is morally correct is determined by whatever God commands us, or whether instead God commands us to do certain things *because* these things are morally required. So, for instance, if God commands you (just as God commanded Abraham) to slay your first born child, would it be morally required of you to do so? Does God’s command determine morality? A slightly different way of approaching this: Is the highest moral principle “obedience to God”?

It could well be that your moral beliefs are closely tied to your religious beliefs, but this connection may not be necessary, and there is some advantage to finding a non-theological basis for morality, insofar as you are then able to use moral persuasion with people who don’t share your particular conception of God or the divine. Much of the divisiveness and violence in the world is grounded in differences of religious opinion. Insofar as we want to be able to use moral reasoning to reach some peaceable accord with others, we can hardly base such reasoning in our religious beliefs, since our only audience will likely be the members of our own religious sect.

HOW DO WE CLASSIFY OUR ACTIONS?

One traditional classification scheme sorts our actions into four groups: required, prohibited, supererogatory, and morally neutral.⁴ With actions that are morally **required**, our failure to do them merits *blame*, while actions that are morally **prohibited** merit blame if we do them. So, for instance, if I take to stealing books from the library for the sole purpose of trying to dam the Kenapocomoco with them, you will likely view my behavior as blameworthy, since I shouldn’t be stealing books and I shouldn’t be throwing them into the river. Similarly, if I fail to read any of your exams and papers, and instead assign grades to them randomly, you would find this sort of behavior blameworthy as well, for part of my duty as a professor is not merely to assign grades to your work, but to do this in a way that reflects certain features of that work.

Praise is rarely appropriate where the behavior in question is either prohibited or required. No one will praise you or pat you on the back if you point out that you haven’t murdered anyone in the last 36 hours. (More likely they will call the police or campus security.) The only time we might praise someone for doing what they are supposed to do is when they are in some sort of “remedial moral training,” as might be the case with young children and the criminally insane.

³ This was a nationwide survey of 2,002 adults, conducted Feb. 25-March 10 by the Pew Research Center and the Pew Forum on Religion and Public Life. The full report is available at: http://pewforum.org/uploadedfiles/Press_Room/Press_Releases/poll2002.pdf

⁴ This discussion draws from distinctions made by Alexius Meinong (1853-1920) and Ernst Schwarz .

Supererogatory actions are morally relevant but are neither required of us nor prohibited. Consequently, doing them merits *praise*, but failure to do them does not merit blame. For instance, suppose while out on a walk you notice occasional pieces of broken glass along the way, and rather than ignoring the glass you take the time to carefully pick it up so that young children and people in bare feet don't accidentally come to harm falling or stepping on it. Doing this sort of thing is surely praiseworthy, and while we may likely blame those who broke and left the glass in the first place, we might not think that those who ignore the glass are doing something morally wrong, and therefore blameworthy. (If, on the other hand, you are unable to imagine any actions that fit this description of supererogatory, then you have made an important moral discovery about yourself.)

Finally, **morally neutral** actions are neither required, prohibited, nor supererogatory, and so doing them (or failing to do them) merits neither praise nor blame. For instance, whether in getting dressed in the morning you pull on your right or your left sock first is a matter of utter irrelevance with respect to morality, and no one would think to blame or praise you for such actions.

It so happens that, for the utilitarian, every action will be either required (if it tends to maximize happiness) or prohibited (if it tends to minimize it) or indifferent (if it tends neither to increase nor to decrease happiness). Consequently, if you find the class of supererogatory actions to be empty, then you might just have utilitarian sympathies. (This will be discussed further below.)

WHAT IS ETHICS GOOD FOR?

This is perhaps the central question of ethics: What good is it? For as much time and effort as humans devote to this subject, surely some good must come of it — but what could that possibly be?

Ethical reflection ultimately pushes us to develop a moral theory. The purpose of any **theory** is to unify information (facts, ideas) that would otherwise remain scattered or disconnected. Theories simplify the world of information by finding or creating connections, and turning a mere heap or *aggregate* of information into a *system* of information. In unifying information, a theory is said to *explain* each particular item of information, and quite often the theory is also able to *predict* new information. This is true of any theory, whether of ethics, physics, or economics. For instance, the greatness of Newton's theory of gravitation was that it provided a single explanation for the motions of the heavenly bodies as well as for the motions of falling bodies here on earth (and when combined with his three laws of motion, he was also able to explain so-called "violent" motions of thrown objects, or projectiles). A famous example of how this theory allows for *predictions* comes from the 1840s, when astronomers used Newton's theory to predict the existence of a new planet. They based this prediction on the observed and occasionally puzzling motions of Uranus — which at that time was thought to be the most distant planet in our solar system; and on September 23, 1846, Johann Galle spotted the new planet in the predicted location: Neptune was officially discovered.

Each of us encounters a great many **facts** in the moral world surrounding us. These facts include our many moral **intuitions** (that is, our spontaneous, unconscious affective responses — our "gut reactions" — to situations we encounter) as well as our consciously-held **beliefs** about how we ought to behave or exist and how we value different things. Add to this the responses and beliefs of *others*, as well. Because these different beliefs often come into conflict, we are naturally driven to discover or develop a set of moral **principles** that will collect together as many of our beliefs as possible into a unified system.

The study of ethics should help me make sense of or explain my moral intuitions, it should guide me when my intuitions fail, it should help me justify my actions in the eyes of my neighbors, and it should help me become a better person than I am, creating a coherent moral existence — my life — and one that I will not regret at the end of my days.

[41] FROM MORAL INTUITION TO MORAL THEORY

All of us already have moral characters formed in varying degrees and directions, and because of this we respond to the world around us with any number of moral **intuitions**, on the basis of which we pass moral **judgments**. Most of us are quite proficient at passing moral judgment, but rather less skilled at discovering the moral **principle** from which such a judgment might plausibly be derived, to say nothing of discovering a moral **theory** that might explain and justify those various moral principles.

In this section I hope to begin a more systematic inquiry into our moral intuitions and judgments, exploring what moral principles might plausibly underlie those judgments, and then briefly introduce three moral theories, any one of which might be used to justify our moral principles.

As noted above, this exploration of moral principles and theories is useful in a number of ways: it will guide us in times when our intuitions are weak, vacillating, or in conflict with one another, it will help us explain our behavior to others, and finally it will help explain our behavior to ourselves. But first, we would do well to consider some recent work in moral psychology.

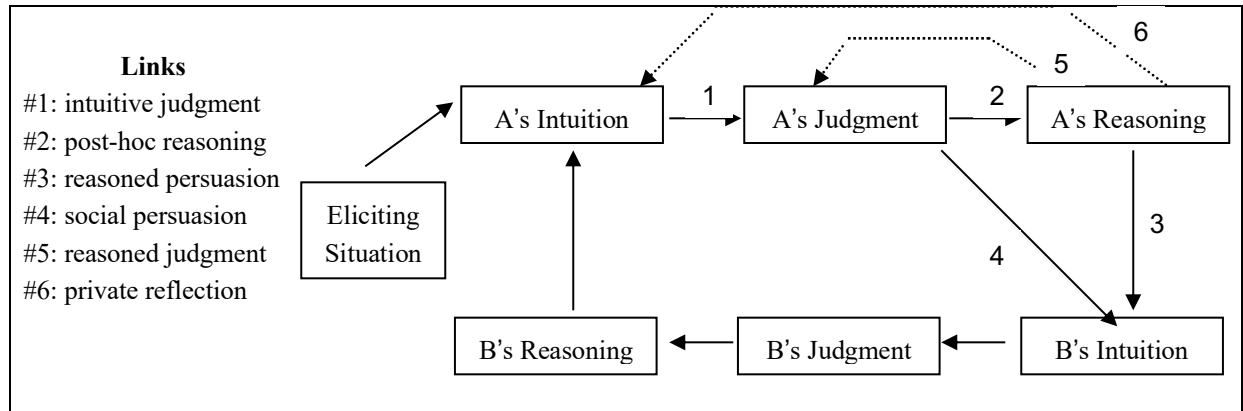
MAKING MORAL JUDGMENTS

Psychologists have amassed considerable evidence for an “intuitionist” model of moral judgment. On this model we arrive at our judgments, *in most instances*, using **intuition** (which happens quickly, effortlessly, and without any awareness of the process for arriving at the judgment) rather than **reason** (which happens slowly, with effort, and involves conscious steps). So for instance, I am confronted by some situation — a man beating a small child — and I intuitively respond with disapproval: I want him to stop. I consider whether I ought to talk to him, I express my disapproval to those around me, and in general I just get “worked up and bothered.” Only if you challenge me will I then try to discover some moral reason to justify this initial response.

Jonathan Haidt, a professor of psychology at the University of Virginia, has advanced a “social intuitionist” model of morality that offers an interesting account not only of how we arrive at moral judgments, but also of how we come to modify our moral beliefs and moral judgments. His diagram of this model is reproduced below (it includes two people — A and B — but is drawn from A’s perspective).⁵

According to Haidt’s research, most of our moral judgments regarding some situation are made intuitively — entirely free of any rational deliberation or conscious reflection. For instance, I see the man beating the child and I say to you, with agitation in my voice, “What that man is doing is terrible; we need to stop him!” This is my “judgment,” which is the product of an unconscious process, immediate and non-reflective (#1, the move from intuition to judgment). My judgment may influence your later judgments by influencing your intuitions (#4, the “social persuasion” link), as there are always strong pressures for one’s own beliefs to conform with the beliefs of others. Having voiced my opinion, however, you might ask: “Why, what’s the problem here? Why do you think what that man is doing is wrong?” It’s at this point that I must justify myself (#2, the “post-hoc reasoning” link), searching for reasons or arguments that will support my intuitive judgment. This is where, as Haidt points out, I typically will behave like a **lawyer**, instead of a **judge**, and look only for evidence that supports my side, rather than trying to discover the truth of the matter. I might then tell you that “It’s wrong to beat children!” — perhaps adding that this is wrong because the beating will only harm and not benefit the child, that nothing ever good comes from such beatings, and so on (#3, the “reasoned persuasion” link).

⁵ This diagram originally appeared in Jonathan Haidt, “The Emotional Dog and its Rational Tail: A Social Intuitionist Approach to Moral Judgment” in *Psychological Review* 108 (2001): 814-834. See also a briefer presentation in: J. Haidt, “The New Synthesis in Moral Psychology” in *Science* 316 (18 May 2007): 998-1002.



Haidt notes that this attempt at using reason rarely influences the other person, and certainly not if you already disagree with my reasons. It might satisfy you if you simply were unclear about the facts of the situation — perhaps you thought the man was beating a rug instead of a child. Otherwise, the most successful approach, if I wish to persuade you, is to frame my reasoning in such a way that it elicits within you new intuitions that bring your judgments into closer conformity with my own.

Haidt emphasizes the primacy of intuition in our moral lives, but he doesn't discount the force of reason altogether. Some people, often of a philosophical bent, are able to reason their way to a judgment (#5, the "reasoned judgment" link); but this is usually in those cases where one's intuition is weak, or where two or more intuitions come into conflict. Finally, one's reasoning might activate a new intuition (#6; the "private reflection" link). This is often what happens in role-play: When you imagine how the situation must feel to the other person — "being in the other person's shoes" — new intuitions will emerge and, with them, new moral judgments.

Several important points emerge from these findings from moral psychology. First, they highlight how *biased* we are in favor of our initial beliefs. Insofar as we truly want to arrive at the best judgment, we need to work actively against these biases. Second, they highlight how difficult it is to change another person's beliefs, especially by using reason. Both of these challenges suggest the same strategy: That we actively seek to expand our own horizons by trying to understand the world from the other person's perspective, and that we encourage others to do the same. What we are attempting here isn't a new way of reasoning about the world (i.e., understanding some new argument), but rather a new way of perceiving the world. Haidt summarizes the situation with the following metaphor:

If moral reasoning is generally a *post hoc* construction intended to justify automatic moral intuitions, then our moral life is plagued by two illusions. The first illusion can be called the "wag-the-dog" illusion: we believe that our own moral judgment (the dog) is driven by our own moral reasoning (the tail). The second illusion can be called the "wag-the-other-dog's tail" illusion: in a moral argument, we expect the successful rebuttal of an opponent's arguments to change the opponent's mind. Such a belief is like thinking that forcing a dog's tail to wag by moving it with your hand should make the dog happy. [Haidt, 2001]

ON THE MORAL CREDIBILITY OF OUR INTUITIONS

We judge our moral theories, in part, by their ability to make sense of our moral intuitions regarding various actions and states of affairs. Sometimes, however, these intuitions must themselves be viewed as morally suspect. Consider a parent trying to save his child. Typically there are strong emotional bonds between a parent and child and, faced with some impending disaster where there is little time to think or deliberate, a parent may well act unjustly in an attempt to save the child. Our deep and emotional attachments to various others can cause us to do the wrong thing.

Would I, for instance, in the course of medical research, consent to having a rhesus monkey tortured and killed so as to save one of my children? Certainly, especially were this posed at the point of crisis, when I must decide promptly. Our first step, however, is to recognize that my actions are *limited* in saving my children. I may not kill

another innocent person to save my child. I may not kill another child and remove its heart so as to save my own child, whose heart is failing. I may not trample another child in an effort to reach and save my own child. These limits I know to exist, because I can sense them — at least now, when my head and heart are cool, and I am not being gripped by the emotion and panic of the moment.

The second step is to discover *where* that limit is. In a moment of crisis or panic, I may well stray past this limit in an effort to save my children. But we need to find those principles that we can live with always — and not just in a moment of panic and strong emotion. Necessarily (almost), these will be principles that others also find acceptable.

One aid in discovering these principles that we implicitly follow is to consider certain stream-lined scenarios, and then observe what we feel is intuitively correct and look for principles that could plausibly be considered the basis of those intuitions. Having once located the principles, we will then have them available for those other occasions when our intuitions are less clear, or when they conflict with the moral intuitions of others.

The Bad Heart

Imagine a pair of one-week old twins, a healthy baby (Hal) and a sick baby (Sal), who are identical in all respects except that Sal has been found to have a bad heart, and will die within a few days if he doesn't receive a heart transplant. No hearts are available except, of course, that of his healthy twin brother, Hal. Being heart surgeons, we're faced with two choices, each of which results in a dead twin and a live twin: either we let Sal die, or we kill Hal and use his heart to save Sal. What should we do?

As it turns out, virtually everyone agrees that we ought to let Sal die. What does this tell us about our moral principles?

Maybe our intuition reflects the moral principle that Hal and Sal have various rights, which include the right not to be killed and the right to be saved. Since we must choose between honoring Hal's *right not to be killed* and Sal's *right to be saved*, we must decide which right is more compelling, and that turns out to be the right not to be killed. This is roughly equivalent to saying that "it is worse to kill than to let die" (but they aren't quite the same).

The Bad Kidney

Let's now change the story slightly. Sal's heart is in fine shape, but he has no functioning kidneys. What is more, he is unable to make use of a hemodialysis machine (for various reasons not worth mentioning here) and the only donor in sight is his dear brother Hal. Sal and Hal are still infants. If we do nothing, Sal will die in a few days. If we give Sal one of Hal's kidneys, then both will live — you only need one kidney to live — although Hal will be burdened with the risk of now having only one kidney (thus, lacking any spare or backup). Should we impose this harm and risk on Hal in order to save Sal?

The Snake Bite

Suppose Sal and Hal are now adults out hiking in the desert. Sal is bitten by a venomous snake and will surely die unless Hal gives Sal some of his snake bite antivenom. Suppose Hal has plenty on hand, and is in no danger of running out. Is he *required* to give Sal some? (If we were a third party with the power to intervene, should we make Hal give-up some of his antivenom?) Some will say that Hal is *required* to help; others will say that he should not be *legally* required, but that he is morally required to help Sal, and that he would be a terrible person not to help.



Kitty Genovese
(Queens, NY, 1935-1964)

BYSTANDER APATHY AND THE CASE OF KITTY GENOVESE

Late one night in Queens, New York, in 1964, a 28-year-old woman by the name of Catherine ("Kitty") Genovese was returning home from work. She had parked her red Fiat and was walking the 100 feet to the door of her apartment building when she was attacked by a man wielding a knife. Her screams eventually woke-up thirty-eight of the neighbors, all of whom peered through their windows to watch. For various reasons, none of these neighbors helped, other than opening a window or turning on a light; no one phoned the police. Twice the assailant was scared away, but twice he returned to continue murdering the young woman, finishing some twenty minutes after her first screams.

Maybe the sick twin/healthy twin scenario is like the case where both Hal and Sal are bitten by snakes, and only Hal has any antivenom, and there is only enough for one person. What is fair here? Should they flip a coin? Is it obvious that Hal should get the antivenom? Would it make a difference if Sal had failed to bring along his own antivenom because he is inexperienced at hiking, or lazy, or forgetful? Or that Hal had some in his pocket because he happened to find some on the ground earlier, or was wearing some borrowed clothes and it was in one of the pockets? In other words, is *merit* an issue here? But that won't help the baby scenario, where presumably merit can't enter into the equation, since neither baby did anything to deserve the kind of heart they have (healthy or otherwise).

The Baby in the Lake

Scenario #1: Watching the Baby Drown. Imagine some business executive who owns a second house on a lake where he goes on the weekends to relax. He's relaxing right now, sitting out on his pier with his favorite scotch. He notices a young child crawling about on an adjacent pier. No other adults seem to be around supervising the child. The executive finds the whole thing curious, and then notices the child fall in the water. It splashes for a minute, and then floats away. He shrugs his shoulders and wonders where the child's parents were. Assume that he could have easily walked over and saved the child by pulling it out of the water, but he doesn't. (This all happens as though the child were nothing more than a balloon that came untethered from the neighbor's porch, and slowly blew away across the lake.)

Scenario #2: Drowning the Baby. Imagine the same scenario, only here the man takes the child from the pier and throws it into the lake and watches it drown. (Suppose the child is a pesky nuisance, always crying and making noise and disturbing his peace. Or suppose the man is simply curious about what it would be like to throw a child into a lake and watch it drown.)

Is the level of moral turpitude the same in each scenario? Or did something worse happen in one of them? There is no conflict of rights here, no question of private property, and many will find that, while the person acts despicably in both scenarios, the level of despicableness is slighter higher in the second scenario, where he actually kills the child (violates its right not to be killed), rather than simply letting it wander into the water and die (violates its right to be saved).

Others will view the action in these two scenarios as equally horrible, and may even view the first scenario as worse, since it portrays the agent as being wholly indifferent to what we generally recognize as a great good ("saving human life"), and as such almost seems inhuman or beyond the pale. Does the moral rightness or wrongness of an action depend upon the *motive* of the actor, or upon the actor's emotional state (e.g., extreme passion vs. cool indifference)?

It seems easier to understand the second scenario, where the evil is actively pursued: maybe the person was intensely angry with the child, or simply angry at the world, or otherwise in the grips of some profound passion. This at least seems *human* to us, whereas the indifference displayed in the first scenario strikes us as a deeply troubling character flaw, a kind of sociopathy. What exactly is this good towards which such indifference is shown in the first scenario? The good of human life as such?

A PROVISIONAL LIST OF "LIFE AND DEATH" PRINCIPLES...

- "It's wrong to harm one person in order to benefit another."
- "It's worse to kill than to let die."
- "The right not to be killed overrides the right to be saved."
- "The right not to be killed overrides the right to be saved, *regardless of the number saved.*"
- "The right to be saved overrides the right to private property."
- "Preserving human life is/is not the highest value."
- "We ought to help others in need."
- "We ought to help others in need (if the personal cost/risk is not too great)."
- "It is worse (or as bad) to be indifferent to evil than to actually bring about evil."
- "The right not to be killed overrides the right to be saved, *regardless of the future of the would be murder-victim.*"

Otto & Donny, Tom & Mary

Here are two outwardly similar scenarios. In the first, we have two sixth grade boys: Otto and Donny. Otto is the neighborhood bully, and he chases Donny home from school one day. Before Donny can reach the safety of his home, Otto catches him, throws him down into the weeds of a vacant lot, and twists his arm until Donny, with tears streaming down his cheek, cries out “Uncle.” Otto then spits in Donny’s face and runs off.

In the second scenario, Tom and Mary are both adults. Tom is waiting for Mary outside the place where she works. It’s dark when Mary leaves to begin her walk home. Tom chases after Mary and eventually catches her, throwing her down in the weeds of a vacant lot where he rapes her. He then spits in her face and runs off.

Something bad happened in each of the above stories: What was it? Was it the same in each? Do these stories differ in kind or only in degree?

Four Sick Children & Hal

Imagine four children, each in need of some vital organ in order to survive, and one healthy child who happens to be the correct tissue type for donating his organs to these sick children. No other donors are available. Assume the children are roughly equivalent as to their likely life expectancies (once the proper organs are in place) vis-à-vis their “contribution to society,” etc. Should I kill the healthy child to save the other four?

Many people believe that it would be morally wrong to kill the innocent person to save the other four children. This suggests that “human life is not the highest value” (since our moral intuitions guide us to that action which *minimizes* those humans which live). Thus, if we feel that our intuitions are grounded in principles of possessing certain rights (whether the right to property, or the right not to be killed), it would seem that these rights are not overridden even when multiple “rights to be saved” are added to the equation. In other words, this thought-experiment might suggest that *numbers are irrelevant*, that regardless of however many people I might save by killing an innocent person, it is better to *let them die* than to *kill* the innocent person.

An alternative case: suppose I am among a group in hiding (a classic example from the last century: Jews hiding in an attic room, with Nazi soldiers searching the rooms below), and a crying child poses an immediate threat to our group’s safety; should I stifle the child — with the risk that the child may suffocate and die — in order to avoid the group’s capture and execution? The Principle of Double Effect (see the box) will likely distinguish between these two cases, possibly permitting us to stifle the child but not permitting us to remove Hal’s various organs.

DOUBLE EFFECT

The doctrine of double-effect places the following restrictions on the permissibility of an action when some of the foreseeable consequences of the action are evil: (1) the act is good in itself or at least indifferent; (2) only the good consequences of the act are intended; (3) the good consequences are not the effect of the evil; and (4) the good consequences are commensurate with the evil consequences.

Yet another alternative case. Suppose you’re hiking out in the backwoods somewhere and stumble upon a children’s wilderness camp, and suppose that this camp has just been taken hostage by a deranged ex-Boy Scout named Don. The ex-Boy Scout sees the opportunity for some fun when you blunder into the camp, and he gives you the following ultimatum: Either you hack to death one of the campers with his scout hatchet, or else Don will hack to death the whole lot of them. Suppose there are exactly five campers, and suppose that the adults leading the camp are not present (perhaps they are off buying groceries, or perhaps Don has already killed them). What should you do? What are the morally relevant features of this unfortunate scenario?

Five Sick Children

Now imagine that one of the sick children has enough good vital organs to save the other four. Imagine the same children as above, only now our healthy child is found to have an inoperable brain tumor that will leave him dead within the month. Since he has less than a month to live, and the other four children will die right away without his organs, is it morally permissible (perhaps even morally required) to kill the child in order to save the other four?

If we decline to kill the child even here, that would seem to suggest that the future of the would-be victim is irrelevant, and that a person's rights are not minimized by having an uncertain or greatly truncated future.⁶

SETTLING OUR DIFFERENCES

We often find ourselves embroiled in moral disagreement. Some consider abortion to be immoral (that is, morally prohibited) while others consider it morally permissible, or in some cases perhaps even morally required. Some consider stealing from large companies to be morally permissible, while others consider it impermissible. How are we to understand — and possibly resolve — these disagreements? It appears that moral disagreement can occur at any one of three levels: disagreement over the facts surrounding the case, disagreement over the moral principles involved, and disagreement over how the principles should be ranked among themselves (such ranking is necessary whenever they come into conflict).

Regarding Facts

Disagreement over the facts is a difference regarding which facts about the world one accepts, but this often leads to differences in behavior which *appear* to be moral differences. For instance, I recall hearing as a child a news report about a group of Hell's Angels bikers. The details are fuzzy, but the general story is hard to forget, and goes something like this (even though it may well be apocryphal, and I was just a child when I heard it): The bikers pulled into a gas station to tank up, the attendant on duty made some sort of rude comment to them, and so they doused him with gasoline and set him on fire with a cigarette. I found such actions horrific, as I suspect nearly everyone would.

Yet such a practice was not always viewed in this light. Consider the life of the 16th century French judge **Nicolas Remy** (1530-1616) — a man of wide-learning, born to a family of jurists, a cultivated scholar, an elegant Latin poet, and the devoted historian of his country — who in the course of his illustrious career as an inquisitor sent from two- to three-thousand alleged witches to the stake, where they were burned alive.

Now, what exactly was Remy doing here? The standard practice of the day was to burn witches so that they might stand some chance of salvation (the fire was to purify their souls). The *children* of witches — because they were the offspring of Satan and the witch — were burned for the same reason; and not burning these children was to foreclose all hope of their escaping the torments of eternal damnation. In writing his memoirs, Remy looked back on his long and fruitful life as a witch-hunter and noted with regret the immature feelings of his youth that had at one time kept him from burning these children of witches. Because of his foolish sentimentality, they were now damned to eternal Hell-fire.

⁶ Whoever still favors letting all the children die rather than killing one of them will, among other things, have all the intuitions necessary for the argument that abortions are *always* wrong (even if the mother will die as a result of the pregnancy coming to term; indeed, even if both the fetus and the mother will likely die, as with ectopic pregnancies).

Despite the horror of killing innocent children (and in such a grisly fashion), it is not difficult to see that the difference between Remy and ourselves is very likely only a difference in factual beliefs, and not a difference in moral principles. Remy wanted the best for those children, and that is why — in his later years — he had them burned at the stake along with their mothers. We also want the best for children, and we do for them what we can; but since we no longer believe in witches and in the burning of witches, our love for children no longer motivates us to set them on fire.

Regarding Moral Principles

Occasionally moral disagreement appears to be a disagreement over moral principles, and this form of disagreement is far more intractable. If we disagree over moral principles, the best we can do is try to find a common criterion by which we might judge those principles. The examples of this sort of disagreement that most readily come to mind are so-called “vice crimes,” such as adultery, pre-marital sex, prostitution, bestiality, the use of certain drugs, pornography, gambling, dancing, chewing tobacco, or spitting in public. Here everyone may be in perfect agreement as to the facts, but one party will say the action is impermissible, while the other will claim it is permissible.

Haidt’s social intuitionist model suggests that this disagreement is occurring at the level of intuition. Insofar as we hope to understand each other, however, we need to attempt to articulate our intuitions in terms of rational principles, and the general strategy here is to keep pushing towards more basic principles, until we finally hit common ground. So in trying to understand why, for instance, someone is opposed to pre-marital sex, one might explore *why* they think pre-marital sex is wrong — is this claim justified by some more basic principle? Or is it wrong “just because it is” — i.e., that the prohibition of pre-marital sex stands as a basic moral principle, and so cannot be justified (such behavior is “axiomatically wrong”)? If they are at all reflective, they will generally have some deeper reason and so can move beyond this superficial level. They might argue on grounds of human well being, or respecting others, or avoiding various risks, and in moving to a deeper layer of moral thinking, the parties in disagreement may well discover a common ground. They may still, in the end, disagree with each other’s reasoning from those principles, but they will at least have found a set of principles that they both accept, and they will better understand why the other party believes as it does.

Regarding the Ranking of Principles

The reason why disagreement might still be possible, even when the two parties agree on the principles, is that they now find that they disagree over how the various moral principles should be ranked. There are competing moral goods, and some choice must be made between them. Such situations typically form some sort of **moral dilemma**, situations where we cannot avoid doing some wrong because two or more of our accepted moral principles collide, and one or more must be neglected in favor of the other. In such instances the principles and facts are all

SCRAPING GENITALIA AND CASTRATING PIGS

Consider the moral principle: “We ought not to cause gratuitous suffering.” This principle claims that any suffering that we cause is permissible *only if* it makes possible some greater good. This is a principle most people endorse. Yet we also endorse many practices that involve immense suffering by animals. So we presumably believe that the good arising from these practices outweighs the suffering caused by them. One such practice is how we harvest musk oil from civet cats, a species of wild cat living in Ethiopia. These cats are trapped for the musk that they produce, which is obtained in the following way: They are kept in cages in dark rooms heated to about 110° F with wood fires. The darkness calms the cats and the heat promotes the production of musk, which is secreted from glands near their genitalia, where it is scraped away daily. This routine continues for the life of the cat. Human beings use musk in perfume as a fixative, where it has the property of making the scent longer lasting. So, presumably, the good that perfume wearers enjoy outweighs the suffering of these cats, or else we wouldn’t support this practice through our purchase of the perfume.

Another example comes from the meat industry. It is standard practice in the US to castrate all bulls and boars raised for meat, as male hormones occasionally give the meat an off-flavor, and because the animals fatten more quickly and are more docile if they lack testicles. The US method of castration is to hold the animal down, slit-open the scrotum with a knife, and pull each testicle out of the scrotum in turn, breaking the cord attaching each testicle with a quick yank. In England, they first administer a local anesthetic to the animal; but this adds an additional expense that the US meat industry or the US consumer finds unacceptable.

agreed upon, but we may disagree as to the ranking: When principles conflict, which principle should be followed and which neglected?

Consider the moral principles: “You ought to care for your children” and “You ought not to steal.” (Or alternatively: “The right to sustain oneself” vs. “The right to private property.”) These are principles that almost everyone endorses; the world would be arguably worse if parents felt no obligation towards their children or if people did not respect the property rights of others. Yet if a father lacks the means to properly care for his children, and his single alternative is to steal food in order to feed them, then his obligation to feed his children stands in conflict with his obligation not to steal. Now, if you are a hard-nosed libertarian, you might stand your ground and insist that the father should not steal the food, even though his children are dying. Here the conflict is not over the facts of the matter, nor over the legitimacy of certain principles, but rather over which principles take precedence when they conflict.

What method do we use to rank our moral principles when confronted by these dilemmas? Is there a higher principle upon which our other moral principles are based? If such a principle can be found, then it will help us resolve such moral dilemmas, and it might also help us argue for the legitimacy of lesser principles that might be thrown into question. We will examine the possibility of resolving this conflict when we turn to specific moral theories.

Negotiating these points of moral disagreement can offer a key opportunity for our own moral development. Even if it is true that most moral judgments occur at an intuitive, pre-conscious level, we still, in the end, have to make sense of them, and making sense of them requires capturing them under principles and higher moral theories. They are still, in the end, my own actions and beliefs, and so I need to be able to see them as expressions of my life. This brings us back to the central task of ethics, which is the **justification** of our actions, both to ourselves and to others. It is important that I am able to justify my actions to others, because I have to live and get along with these other people. It is also important that I justify my actions to myself, because I have to live with myself. Even if the causes of my various moral intuitions are as scattered and unrelated as the many events in my life, I still need to unify these moral intuitions into something like a coherent whole — a fiction, perhaps, but a necessary fiction all the same, a coherent narrative moral thread that strings all my moral responses into a single life.

I justify my actions to others by appealing to principles that I hope the others will accept. If they accept the principles, then I need merely draw a line from the principle to my action. If they don’t accept the principle, then I may need to explore ways to elicit the sort of intuitions in the other that will make my principle appear more plausible. If moral development is to occur and if some sort of amicable reconciliation is to be possible, then this conversation has to happen.

An Example

Suppose you borrow your roommate’s car without first asking permission. You might have done this just because you felt like going somewhere, and you really don’t care what anyone thinks, and your roommate can go jump in a lake if he wants. If this really was your attitude, and you really are such an unpleasant person, then we can easily imagine the sort of moral condemnation others might bring down on you:

1. You must not take what is not yours.

or more generally:

2. You must respect the private property of others.

Are there any moral principles that might make your actions morally permissible? Yes, but they normally can’t be invoked without first learning more about what you did and why. There will be many actions that cannot be morally justified, no matter how many additional facts are considered; but many actions that at first seem wrong turn out to be

A FEW PRINCIPLES...

Non-Maleficence: Do no harm to others.

Beneficence: Help others.

Utility: Bring about the greatest benefit with the least harm.

Autonomy: Rational individuals should be permitted to be self-determining.

Principles of Distributive Justice ...

Equality: Distribute resources equally.

Need: Distribute according to need.

Contribution: Distribute according to contribution.

Effort: Distribute according to effort.

morally permissible, and sometimes even morally required, **once we learn more about the facts**. For instance, what if a neighbor suddenly grows quite ill and needs to get to a hospital right away, and your roommate's car is the only way of getting her there? We might feel justified using the car because of a third principle:

3. You must help others when you can.

Notice that this doesn't make the other principles false or useless; we simply decided in this case, where the two principles collided with each other, that the third principle is more important. It **overrides** the other principle.

But what if the only help someone needs is a ride to the mall for a new pair of socks? Many would feel this need was too insignificant to outweigh the principle requiring us to respect the property of others. There's clearly some sort of proportionality principle at work:

4. When harming one to help another, the benefit must outweigh the harm,

or something like that.

And we might want to tighten up principle (3) to justify using the car without permission, maybe (5):

5. Saving a human life is more important than anything thing else, and therefore justifies every action.

It won't take long to discover the shortcomings of this principle, however. For instance, it would allow killing one person to save the life of another (for instance, harvesting one person's organs to save one or more others in need of the organs), and most would agree that ...

6. It is worse to kill than to let die.

In other words, when faced with the opportunity to kill one person in order to save another, we should let that opportunity pass — unless the person we could save is having his life threatened by the other person. If A is about to kill B, then we might feel justified in saving B by killing A (at least if A is not justified in killing B — suppose he is the public executioner). Here we are faced with considering certain moral features of individual humans, and not just our actions towards them. Perhaps we could make (6) more complicated, but more acceptable, by replacing it with (7):

7. It is worse to kill an innocent person than to let an innocent person die.

Taking the life of another is normally considered worse than failing to save the life of another, but what if the harm is lessened? Are we justified in simply removing (by force) one of your kidneys in order to save your neighbor, who has none? Are we justified in requiring you to donate blood so that your neighbor might live? Perhaps your neighbor needs neither a kidney nor blood, but simply some of your money in order to survive — are we justified in stealing your money on your neighbor's behalf? Or is your neighbor justified in stealing the money to save himself?

Many people feel conflicted by these claims. On the one hand, a human life is surely worth more than, say, \$20. On the other hand, it isn't clear that another person is justified in taking that money from me so as to save his own life or the life of another. Is it true that...

8. The right to life overrides the right to private property?

Whoever fails to save another person's life simply because he's too cheap to part with his money begins to look a bit like a moral monster. But when a third party takes away that money from him by force, in order to save the other's life, he begins to look like a victim as well. How do we resolve these conflicting moral intuitions?

The multitude of principles, and the need to find some way of resolving conflicts between principles, causes us to keep looking for ever more basic foundations for our moral lives. If these principles all arise from some basic **theory**, then the theory should also tell us how to rank the principles, so that we can more readily discern when one principle should override another.

THE VARIETY OF MORAL THEORIES


Ethics is a study of three different aspects of the moral universe, as reflected in three different sets of terms. We can speak of the moral worth or value of an **act** (calling it 'right' or 'wrong'), of a **situation** (calling it 'good' or

‘bad’), and of the **character of a person** (calling it ‘virtuous’ or ‘vicious’). These distinctions help sort-out some of the complexity of our moral lives. For instance, we want to be able to speak of someone as “virtuous” even if they unintentionally bring about bad states of affairs (such as causing more misery in the world) — such people we call “moral bumblers.” Similarly, we want to be able to say of an act that it is morally correct even where it might bring about more bad (through some unforeseen or unintended consequence of the act). And we want to be able to speak of a virtuous person acting wrongly (in a moment of weakness or confusion), or of a vicious person sometimes acting rightly.

These three dimensions of the moral universe correspond to three traditional moral theories: utilitarian, deontological, and virtue-based theories. We will be exploring these three moral theories in much greater detail in later chapters.

Utilitarian Theories

Utilitarian theories emphasize the good or bad consequences of acts rather than the acts themselves. The rightness or wrongness of an action depends wholly on the action’s consequences. One of its earliest and most eloquent proponents was the 19th century British philosopher and social activist **John Stuart Mill** (1806-73). An action is *right* if its good consequences outweigh the bad; specifically for Mill, an action is right if it maximizes the net happiness.



Deontological Theories

A deontological theory was first formulated by the 18th century German philosopher **Immanuel Kant** (1724-1804). Deontological theories emphasize the rightness of acts apart from whatever consequences they might have, and often speak of one’s obligation (or duty) to perform some action — an action is *right* if done from the proper motive. More specifically, for Kant: an action is right if it is done out of respect for a proper rule of action (where the rule is proper if it conforms to what Kant calls the “Categorical Imperative”).

Virtue-based Theories

Virtue-based theories emphasize the character of the person, taking this as primary in understanding morality (e.g., the right act is whatever the virtuous person does when acting deliberately). Here the emphasis is on the ideal human life, and what makes a person virtuous or vicious. This fits with many of our pre-theoretical moral intuitions, for we sometimes feel uneasy about someone’s action not because the action violates anyone’s rights or otherwise has bad consequences, but because of what that action suggests about the person’s character: Imagine some obsequious grandson who always acts kindly to his loving grandmother, but once she dies and leaves him his inheritance, he spits on her grave and laughs. This causes no harm to the dead, nor does it seem to violate anyone’s rights, but it points to a flaw in the boy’s character, that he would behave so disingenuously toward his grandmother while she was alive, and so malignantly once she has died.

As we will see, these three moral theories are not concerned strictly with their one aspect of morality; rather, they emphasize one aspect, while including the other two as well. Mill’s utilitarianism, for instance, while claiming that the good is *happiness*, also informs us of our obligation, viz., *to maximize the good*. Mill also writes about nurturing a moral character in children, that is, encouraging virtue, insofar as this will further maximize the good.

Motives, Actions, and their Consequences

These three dimensions of the moral universe can also be considered in the context of the **self**, her **action**, and those action’s **consequences**. The actor’s *motive* (for performing some action) will often be the desire for some foreseeable consequence of the action. In the example given in the box, suppose that the intended consequence of the action was the saving of a human life. It might be merely foreseen, but not intended, that one might also receive

praise and various rewards for the action. A wholly unforeseen consequence of the action might result from the fact that the drowning person is a deranged serial killer and, once he regains his senses (thanks to your expert ministrations), he proceeds to torture and kill you, along with dozens of your neighbors. Or consider a few other possible arrangements to see how your moral intuitions play out: Suppose you know the person drowning, and you save him because you owe him some favor (perhaps he once had saved your life); or suppose that the person owes you a large sum of money, which you will never recover if he drowns; or suppose you save the man because you think it will make a good bargaining chip when standing before St. Peter; or suppose you are an ambitious pre-med student and want to get a jump on your colleagues by dissecting a human cadaver during the summer (here you drag the body from the lake with no intention of trying to resuscitate it).

To decide whether an action is right or wrong, Mill considers only the *actual consequences*, while Kant considers only the *motive*; and for Kant this motive must not arise from a desire for any possible consequence of the action at all, but instead from a certain logical (or formal) feature of the action itself. For Aristotle, right actions are those that a virtuous person performs when he is acting as a virtuous person; more specifically, the rightness of an action depends on its context: it has to be done “at the right times, with reference to the right objects, towards the right people, with the right motive, and in the right way” (*Nic. Ethics*, Bk. II, ch. 6).

[42] METAETHICS: SUBJECTIVISM AND OBJECTIVISM

NORMATIVE ETHICS AND METAETHICS

Normative ethics is what we have been calling ‘ethics’ thus far, and includes the study of goodness (What is the best distribution of benefits and burdens in the world, and what counts as a benefit?), virtue (What traits of character do we value in people?), and right action (What rules or principles should we follow?). Each of these areas typically involves criteria for deciding when a situation is good or bad, when a person’s character is virtuous or vicious, and when an action is right or wrong.

Metaethics, on the other hand, considers the nature of these criteria and the meaning of the moral concepts used. What is the scope or status of moral values? What does it mean to say that a situation is good, a person virtuous, or an action right? Here I’m not asking how I identify the good, the virtuous, and the right; rather, I’m asking what it is that I’m doing when I make such an identification. Am I recognizing some property, like being red, that makes something good (the property of goodness)? Does moral value exist in the world independently of human beings and their interests, or does it depend on our existence? Do we *create* moral value or do we *discover* it? Is it simply a reflection of our interests, or of our emotional states and feelings? The study of these questions constitutes the field of metaethics.

Deciding between ethical subjectivism and ethical objectivism is an important metaethical topic that concerns the very nature of these moral properties. **Ethical subjectivism** claims that all moral judgment is based ultimately on some arbitrary will, whereas **ethical objectivism** claims that moral judgments are based on some objective (non-willful) state of affairs. The majority of our time will be spent looking at three different forms of ethical objectivism (Aristotle’s virtue theory, Mill’s utilitarianism, and Kant’s deontology), but first we need to examine three varieties of ethical subjectivism: simple subjectivism (where what is moral is decided by the individual will), divine command moral theory (where what is moral is decided by the will of God), and ethical relativism (where what is moral is decided by one’s society). Most of our attention will be focused on ethical relativism, since it appears to enjoy the most adherents.

HAMLET THE METAETHICIST

“There is nothing either good or bad, but thinking makes it so.”

— Shakespeare, *Hamlet*, Act 2, Scene 2

SIMPLE SUBJECTIVISM

Simple subjectivism is the view that moral claims are simply about our feelings or emotions, and not about the way the world is or how people have behaved. On this view, there is no such thing as an “objective” right or wrong. Another term for this position is **non-cognitivism**: moral claims, on this view, lack a truth value. They are neither true nor false, since they are nothing more than expressions of feelings, or commands to act in various ways. On this view, for instance, it is a fact about the world that Hitler and the Nazis murdered in cold blood some ten million Jews, communists, homosexuals, and Roma, but it is not a fact about

the world that they did anything *wrong*. It is a fact about the world that Lt. Calley and his men massacred somewhere between 300 and 500 Vietnamese civilians during a few hours on March 16, 1968, but it is not a fact that they did anything wrong. The claim that “Shooting young children in the face is wrong” is neither true nor false, because it isn’t a claim about how the world is; instead, it is merely a reflection of the speaker’s attitudes or feelings or emotions. **Ethical objectivism**, on the other hand, claims that moral value exists as an objective fact in the world, independently of human feelings, and thus that moral claims *do* have truth-values.

Simple subjectivism was first described by **David Hume** (1711-76) in his *Treatise of Human Nature* (1739-40):⁷

Take any action allow’d to be vicious: Wilful murder, for instance. Examine it in all lights, and see if you can find that matter of fact which you call *vice*. In which-ever way you take it, you find only certain passions, motives, volitions and thought. There is no other matter of fact in the case. The vice entirely escapes you, as long as you consider the object. You never can find it, till you turn your reflexion into your own breast, and find a sentiment of disapprobation, which arises in you, toward this action. Here is a matter of fact; but ‘tis the object of feeling, not of reason. It lies in yourself, not in the object. So that when you pronounce any action or character to be vicious, you mean nothing, but that from the constitution of your nature you have a feeling or sentiment of blame from the contemplation of it. [*Treatise*, Bk. 3, pt. 1, §1]



David Hume
(Scotland, 1711-1776)

For Hume, moral judgments are reducible to judgments about our own feelings and attitudes, for example:

“x is right” = “I approve of x”

“x is wrong” = “I disapprove of x”

“Respecting the elderly is right” = “I approve of respecting the elderly”

“Cannibalism is wrong” = “I disapprove of cannibalism”

Simple subjectivism, as an account of our moral lives, is able to explain certain well-known features. For instance, moral disagreements are often emotional and bitter and intractable, which is what one would expect if the only “real” difference is that you approve of X and I don’t, and there’s no compelling “objective reason” to force either of us to change our minds. Simple subjectivism also accords with the common-sense belief that the world consists of **facts**, while preferences reflect our **values**, and that there is no right or wrong answer in the domain of values.

Bertrand Russell on

SUBJECTIVISM

“Ethics contains no statements, whether true or false, but consists of desires. If two men differ about values, there is not a disagreement as to any kind of truth, but a difference of taste. If one man says ‘oysters are good’ and another says ‘I think they are bad’, we recognize that there is nothing to argue about. The theory in question holds that all differences as to values are of this sort, although we do not naturally think them so when we are dealing with matters that seem to us more exalted than oysters.”

— Bertrand Russell, *Religion and Science* (1935)

⁷ A good summary of Hume’s moral philosophy is found in his *Enquiry Concerning the Principles of Morals* (1751).

Despite these advantages, critics of simple subjectivism have raised several problems. The first concerns **moral fallibility**: We often admit that we might be wrong about our moral claims (“maybe *x* is wrong”), and yet we are infallible when we report our feelings (if I feel a certain way, then I can’t be wrong about that feeling). If subjectivism were true, then we should be infallible with respect to our moral claims, as well.

A second problem concerns our understanding of **the location of moral disagreement**. If Mary claims that testing the toxicity of drain cleaners on the eyes of rabbits is good (or at least morally permissible), and Edward claims that it is bad, they still will both *agree* as to each others attitudes (namely, Mary will agree with Edward that his attitude is one of disfavor, and *vice versa*). What they disagree on, we all want to say, is whether doing this to rabbits is *in fact* good or bad. Moral disagreement clearly centers on how the world is, while subjectivism would seem to misdirect our attention to our *feelings* about the world.

A third problem is that it is **too inclusive**, for it seems to turn non-moral statements, such as aesthetic claims, into moral statements: “The *Mona Lisa* is a good painting” begins to look a lot like “Gandhi was a good person.”

Finally, it **divorces morality from our everyday moral experience**. Making morality wholly a matter of feeling removes it from the sphere of rational discourse, in which case the commonsensical activities of moral deliberation, moral persuasion, moral justification, and asking for moral advice become meaningless gestures. The Latin motto rightly notes that *de gustibus non est disputandum* (“there is no arguing in matters of taste”); and reducing moral judgments to matters of taste is to ignore an important aspect of our moral experience.


DIVINE COMMAND

A second form of ethical subjectivism is the divine command moral theory, a position that was first explored in Plato’s *Euthyphro*, where Socrates asked whether a thing or action was pious because it was loved by the gods, or whether it was loved by the gods because it was pious. Which, he asked, is the cause of which? According to the divine command theorist, an action is morally right insofar as God has commanded us to perform that action, and wrong insofar as God has prohibited it. On this view, morality is determined by the will of God.

While often not well articulated, the divine command theory has a rather large following. Many people find a natural connection between their religious beliefs and their moral beliefs and the manner in which ethical precepts appear in their particular religious text as God’s commands will lead many to believe that it is the commandment itself that makes a thing right or wrong.

Divine command also fits nicely with certain needs and observations. For instance, unlike simple subjectivism, it actually provides some **moral guidance**, while at the same time sharing with subjectivism a ready explanation for why there is so much **disagreement** — viz., humans are either ignorant of, or else perversely disobedient towards, God’s will.

Before endorsing this position, however, we need to gain more clarity as to its implications. First, it implies that various actions (like setting fire to cats as a kind of Sunday afternoon amusement) are neither right nor wrong prior to the relevant divine command: Until God speaks, there is no right or wrong. Second, if we are to take God’s will seriously, then we need to see these moral commandments as intrinsically arbitrary: God could have commanded otherwise, making wholly possible our being morally required to disembowel our neighbors, along with their children, and to eat them for lunch. In making these commandments, God isn’t consulting some independent set of moral principles, or some independent moral reality; rather, God’s choices create those principles. Third, this view makes moral knowledge dependent on our knowledge of God’s will. In itself, this is not a proper criticism, but one might well believe that it is far easier to come to agreement on many ethical issues than it is to agree on God’s will (insofar as deciding the latter requires appeals to special revelations, which themselves can be contested and countered with other special revelations). Finally, the German philosopher **Gottfried Wilhelm von Leibniz** (1646-1716) argued against the divine command theory, for such a view of morality would ...



“What God wills is not therefore right because he ought or even was bound so to will; but on the contrary, what takes place is therefore right because he so wills.”

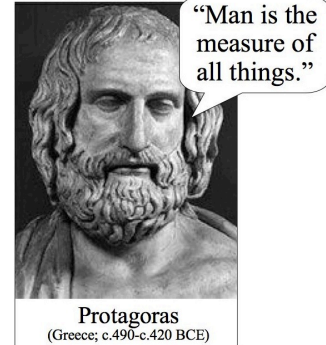


Gottfried Leibniz
(Hannover, 1646-1716)

... destroy all of God's love and all his glory. Why praise him for what he has done if he would be equally praiseworthy in doing the exact contrary? Where will his justice and wisdom reside if there remains only a certain despotic power, if will holds the place of reason, and if, according to the definition of tyrants, justice consists in whatever pleases the most powerful? (*Discourse on Metaphysics*, §2)

ETHICAL RELATIVISM

The etymology of 'ethics' and 'morality' might suggest that they mean nothing more than the study of human habits or customs, since they come from the Greek [*ethos*] and Latin [*moralis*, *mos*] words for habit or custom. If that's all ethics were, however, then it would be nothing more than a branch of anthropology (a descriptive, rather than a normative, science). Ethics deals not only with the way things are, but primarily with the way things *should* be, or the way we *should* behave, and the question facing us here is whether this 'should' is the same for all people, times, and places, or whether it changes relative to these variables. There are customs — normative systems of principles — found in every culture, and individuals who violate those customs are typically chastised or punished. Because these moral principles often differ considerably between cultures, it is natural to ask whether all these principles might be correct in the context of their own culture, or if there might exist a set of principles that are applicable to every culture, regardless of whether or not these principles are recognized or accepted in that culture.



Ethical Universalism & Ethical Relativism

Ethical universalism holds that certain moral principles are true for all people, regardless of who the speaker is or what the context or cultural background might be. This is a natural and widely accepted view, and it would be a natural (although not a necessary) consequence of ethical objectivism. The ubiquity of moral disagreement, however, suggests that perhaps none of our moral beliefs are right or wrong in a "universal" sense — that is, that there is no one correct moral theory or moral principle or set of beliefs, but that they differ from culture to culture or group to group. This position is called **ethical relativism**. According to ethical relativism, moral standards differ between cultures and people, so that the truth-value of a moral claim will depend upon the speaker and her context. Let's borrow an example offered by the contemporary British philosopher Mary Midgley:

There is, it seems, a verb in classical Japanese which means 'to try out one's new sword on a chance wayfarer'. (The word is *tsujigiri*, literally 'crossroads-cut'.) A samurai sword had to be tried out because, if it was to work properly, it had to slice through someone at a single blow, from the shoulder to the opposite flank. Otherwise, the warrior bungled his stroke. This could injure his honor, offend his ancestors, and even let down his emperor. So tests were needed, and wayfarers had to be expended. Any wayfarer would do — provided, of course, that he was not another Samurai. Scientists will recognize a familiar problem about the rights of experimental subjects. [Mary Midgley, "Trying Out One's New Sword"]

Insofar as 20th century Westerners find this method of testing swords immoral, the ethical universalist would argue that *someone* is morally mistaken. Either the people in the samurai culture are wrong, or the 20th century Westerners are wrong, or maybe both are getting morality wrong — the point is that there is somewhere a right answer that holds for all situations. The relativist, on the other hand, would say that such behavior with swords is wrong in our culture, but that it is right in the samurai culture, and that moral truths simply differ from culture to culture.

William Graham Sumner on

ETHICAL RELATIVISM

"The 'right' way is the way which the ancestors used and which has been handed down. The tradition is its own warrant. The notion of right is in the folkways. It is not outside of them, of independent origin, and brought to test them. In the folkways, whatever is, is right. This is because they are traditional, and therefore contain in themselves the authority of ancestral ghosts. When we come to the folkways we are at the end of our analysis."

— W. G. Sumner, *Folkways* (1906), p. 28.

Another way of making this distinction: for any moral claim — for example, “It’s wrong for Mary to try out her new sword on John like that” — the universalist will say that it has a truth-value (although we might not know what it is, if we aren’t sure about the correct moral principles), while the relativist will say that it has no truth-value until its cultural “frame of reference” is provided (for example: The practice of “trying out one’s new sword” is morally right in 17th century Japan, wrong in 20th century Japan). This would be analogous to how we treat claims of fiction. “Sherlock Holmes smokes opium” is neither true nor false, until we provide the frame of reference:

1. Sherlock Holmes smokes opium. [no truth-value]
2. In the stories of Arthur Conan Doyle, Sherlock Holmes smokes opium. [True]
3. In the stories of Arthur Conan Doyle, Sherlock Holmes lives in Brooklyn. [False]

Here the “moral frame of reference” is viewed as a kind of fictional world, within which moral claims will be either true or false. Another analogy might be spatial frames of reference. To say that “Chicago is to the west” is to say something incomplete, and certainly without a truth-value, until you add a frame of reference, e.g., “Chicago is to the west of Minneapolis,” which renders it false, or “Chicago is to the west of Cleveland,” which renders it true. Perhaps moral claims are like this; without their cultural frame of reference, they lack truth-values, and their truth-values might vary from one frame of reference to the next.

Cultural Relativism & Ethical Relativism

This claim of ethical relativism is not to be confused with cultural relativism. **Cultural (or social) relativism** is an empirical claim of anthropology, and not a philosophical or moral theory. It is the view that different cultures accept different moral standards: the same kind of action *thought* to be right in one culture is often *thought* to be wrong in another culture.⁸ **Ethical relativism**, on the other hand, is a moral theory having to do with the truth-value of moral claims. It is the view that what is actually right and wrong (not merely what is *thought* right and wrong) differs between cultures: the same kind of action that *is* right in one culture *is* wrong in another.⁹



Cultural relativism is an obvious fact of human history and anthropology; but ethical relativism, which involves an additional claim, is not at all obviously true. We now need to examine the arguments for and against ethical relativism.

Ethical Relativism: For and Against

Somewhat surprisingly, there’s only one good argument for ethical relativism, namely, that cultural relativism is easier to explain if ethical relativism is true. If there are indeed no ethical universals to guide us, then we would expect to find a wide variety of moral beliefs and practices. Of course, it could also turn out that a set of moral beliefs are shared by all cultures. If so, such a universality of belief might suggest that *those* beliefs have some kind of objective grounding. One might also attempt to support ethical relativism by arguing that there is no convincing account for ethical universalism (namely, what is the theoretical foundation supposed to be?), which speaks *somewhat* in favor of relativism — although such “arguments from ignorance” are inherently weak, as they depend on us never arriving at an adequate moral theory (which, of course, we may — and perhaps already have).

On the other hand, there are several arguments *against* relativism, which meet with varying degrees of success. The first is what I call the “**That’s Simply Horrible!**” argument. A typical response to ethical relativism is that

⁸ See Edward Westermarck’s landmark *The Origin and Development of the Moral Ideas* (1906-8; 2 vols.), in which he studied the differences between societies on a wide host of moral issues.

⁹ Would this claim also hold for such beliefs as: “The right action is one that maximizes the good” or “One ought to treat others as one would like to be treated”?

relativism can't be true because then it would be morally permissible (perhaps even required) to bisect people with your sword, so long as you do it in the appropriate culture — and bisecting people is “simply horrible!”

This is not a very good argument. Moral views differing from our own are despised in proportion to the firmness that we hold our own particular view. But this in itself does not undermine the truth of relativism, for even if relativism is *true* we would expect that people would hold tightly to at least some of their learned moral beliefs, and that they would consequently find repulsive many of the moral beliefs found in other cultures. Indeed, this is not an argument against relativism so much as a mere uninformed reaction to it.

A second, and more promising, argument against ethical relativism is the **corrigibility argument**. The strongest support of ethical relativism is that it can explain the prevalence of cultural relativism, and the corrigibility argument seeks to undermine this support by noting that there have been widely diverging views about certain empirical matters — for instance, whether the earth is flat, or whether the earth revolves around the sun, or whether disease is spread by witches and spirits. The diversity of opinion regarding empirical matters does not support the relativity of their truth; no one would claim that it was once true that the earth is flat, but that now it is false. A diversity of belief clearly does not require a diversity of truth, and we need to keep separate the claims about moral belief and the claims about moral truths.

All the same, it would seem that the cultural diversity of moral norms is best explained by the *absence of* an objective standard of morality, rather than by our *ignorance of* it. The former claims that there is no objective standard, and moral norms reflect the way of life and local needs of the people. The latter claims that there is an objective standard, but most or all people badly misperceive it or are ignorant of it. It is here that relativism finds its strength, for the actual variations in moral codes are more readily explained by the hypothesis that they reflect ways of life, etc., than by the hypothesis that they express badly distorted perceptions of objective moral truths.

A third argument against ethical relativism is that it would seem to make all moral criticism impossible (the “**moral criticism argument**”). If ethical relativism is true, then criticism between different cultural groups becomes incoherent: If other cultures have other moral standards, then we can neither criticize nor praise their practices. Their practices are morally correct, by definition.

This spirit of tolerance and open-mindedness towards the practices of others can seem commendable, of course, but is anyone willing to follow this all the way to the end? For instance, in the spring and early summer of 1994, the Hutu majority of Rwanda — a small country in the heart of Africa — began massacring members of the Tutsi minority. Over a million Tutsis — one-tenth of the entire population of Rwanda — were hacked to death with machetes.¹⁰ Because there were more Hutus than Tutsis (especially after the massacre), and because the Hutus thought that what they were doing was right, then what they were doing *was* right, and it was silly and misguided of the United Nations to condemn this slaughter.

In the Nazi culture of 1930's Germany, it was believed to be morally right to dispossess Jews and others of their property (this also was within the boundaries of the law), and it was morally right to pursue their policy of purifying the country by systematically murdering Jews, Roma, Sinti, and other undesirables. Do we really believe that these actions, while immoral in our own cultural context, are in fact morally correct in the context in which they occurred?

And within one's own cultural group, if the norms of the group define what is right for that group, then an individual cannot properly criticize those norms — so the social reformer is *always* wrong, by definition. Further, given the problems of defining the relevant group, it isn't clear that moral comparisons can be made between individuals, either. Was Hitler a better or worse man than Albert Schweitzer? Who can say, if they belonged to different normative groups?



¹⁰ See Philip Gourevitch's useful account of the massacre — “After the Genocide” — in *The New Yorker* (December 18, 1995), pp. 78-95.

This is a powerful argument against relativism once we realize how differently things appear when relativism is consistently adopted. Insofar as moral criticism *makes sense* as a practice and has importance and meaning to our lives, then ethical relativism is to be rejected.

Finally, it might turn out that, upon closer analysis, the apparent moral difference between cultures is rather a disagreement over certain facts about the world. For example, some cultures believe in witches, in animal reincarnation, in an evil spirit or power that can possess people, in an unseen God who commands certain actions, etc., and these different beliefs often result in widely different actions — recall the story of Nicolas Remy discussed above. But despite the widely diverging moral practice, we might discover a core of moral beliefs shared by all people that motivate these actions. This core group of beliefs might serve as the necessary conditions for the survival of any human community whatever — such as prohibitions against taking human life, lying, and neglecting the young. The ethicist Peter Singer (b. 1946) suggests three groups of universal principles:

There seems to be a popular belief that the taboo on incest is the only moral rule that holds everywhere. The reality is that some much more significant ethical principles carry weight in virtually every human community. These include: obligations on members of a family to support their kin; obligations of reciprocity, to return favors done and gifts received; and constraints on sexual relationships.... The precise form of the obligations or constraints varies from one society to another, but the significance of these universals lies in the fact that obligations of kinship, reciprocity, and sexual relationships form the core of all human ethical systems — and they also guide the behavior of our close non-human relatives.¹¹

Having said all this, we need to remember that even if these arguments against ethical relativism succeed, it does not follow that *bisecting people* is morally wrong. The samurai culture might have gotten its morality right, and it is *our* culture that stands to be corrected. And with the handful of universal moral principles concerning kinship and reciprocity, it is likely that these were simply favored by natural selection in the deep past of human existence; it would be a fallacy to argue that whatever natural selection favors is itself morally right. So even if we are successful in arguing against relativism, we still will need some criterion for deciding *which* moral belief is the right one (in other words, we need some more theory, such as that developed by Aristotle, John Stuart Mill, or Immanuel Kant).

Moral Humility and Moral Fallibilism

Perhaps one of the more pervasive *causes of* (as opposed to *reasons for*) believing in ethical relativism is the desire to challenge the attitude that we possess some privileged moral position from which we can judge others: “*Who isn’t weary of the moral arrogance dripping from these self-righteous hate-mongering bigots?*”

But while adopting a more humble attitude is surely laudable, does it require that we also adopt ethical relativism? Does ethical universalism lead to the sort of ugly, intolerant moral arrogance that typifies a significant portion of public discourse these days?

Of course not. Humility does not require that I adopt ethical relativism; it requires only that I accept my own moral fallibility. It requires that I recognize the possibility of my own moral judgments being wrong, not that I give up making moral judgments altogether. It should not stop me from judging my own practices and those of others, but it *should* change the spirit of my criticism. We ought also to practice humility with respect to our various truth-claims about the natural world, and yet that surely doesn’t mean we can’t believe that there is a truth to the matter; rather, it simply recognizes that we might be mistaken.

There is also a rather unfortunate side of relativism that needs to be confronted head on. Ethical relativism is sometimes the result of a kind of intellectual laziness; it allows us to disagree with one another, but in appearance only. We “agree to disagree,” and so are able to set the problem aside and stop thinking about it altogether. True disagreement, on the other hand, requires more effort, more engagement with the arguments made by others, while relativism suggests that all such disputes, all efforts to understand the other person’s position, is just so much wasted time and effort. Relativism allows us to ignore the other, cloaked under what we think as the civil virtue of toler-

¹¹ Peter Singer, *Ethics* (Oxford University Press, 1994), p. 57.

ance; but such tolerance is really just a refusal to engage with those who are different from us, coupled with a stubborn refusal to grow from our encounter with them. We must be tolerant of differences, of course, but we should not use this as a means to avoid taking those differences seriously and not trying to understand them and perhaps even to become persuaded that they are correct, and so to adopt them as one's own. And thus to grow as moral beings.

[43] PSYCHOLOGICAL EGOISM

WHY BE MORAL?

What motivates us to be moral? Is “doing the right thing” something that is good in itself, that is worth pursuing for its own sake? Or is the moral life just a compromise, the lesser of two evils, where we “do the right thing” not because we want to, but because we fear disapproval and punishment by others? Psychological egoism is the view that self-interest is the only possible motivation of our actions. In other words, that human beings are entirely selfish in all that they do, and that morality is merely a social arrangement whereby it is in our own best interest to act morally.

Plato's most famous work, the *Republic*, is a wide-ranging, ten book examination of the nature of morality, and of justice in particular. Book One is dominated by a discussion between Socrates and the sophist Thrasymachus, who argues that “might makes right” — that the best men are unjust, since injustice is what we all *wish* to pursue and what we *do* pursue, if we think we can get away with it. Although Thrasymachus is eventually caught up in contradictions under Socrates' close examination, it wasn't clear to the onlookers — and certainly not to Thrasymachus — that Socrates has actually defeated his position.

Consequently, Book Two of the *Republic* opens with one of Socrates' own disciples, Glaucon (the older brother to Plato) challenging Socrates on this point. While Glaucon hopes that Socrates' claim is right that justice is a good to be pursued for its own sake, he doesn't think Socrates has shown this yet, so Glaucon develops Thrasymachus' position as strongly as he can. Glaucon asks us to imagine possessing a magic ring like that of the mythical Gyges, a ring with power to make its owner invisible. With such a ring, Glaucon claims, no one would hesitate doing whatever he or she wanted — “moral” or “immoral” — and they would care nothing for justice:

Suppose now that there were two such magic rings, and the just put on one of them and the unjust the other; no man can be imagined to be of such an iron nature that he would stand fast in justice. No man would keep his hands off what was not his own when he could safely take what he liked out of the market, or go into houses and lie with any one at his pleasure, or kill or release from prison whom he would, and in all respects be like a God among men. Then the actions of the just would be as the actions of the unjust; they would both come at last to the same point. And this we may truly affirm to be a great proof that a man is just, not willingly or because he thinks that justice is any good to him individually, but of necessity, for wherever any one thinks that he can safely be unjust, there he is unjust. (*Republic*, 362c)



Never are we motivated to do “the right thing” from love for the just; rather, fear is the prime mover of our souls. This is Glaucon’s challenge.

Altruism

We might reach a better understanding of the nature of selfishness by considering what it takes for an action to be called *altruistic*. In biology, an organism’s action is altruistic when it benefits a second organism at a cost to the first organism. With humans, however, altruism would seem to require only the correct *intention* of the actor. It appears to have nothing to do with whether another person actually benefits from the action, for benefiting another individual is neither a *necessary* nor a *sufficient* condition of it being altruistic. Consider the following three scenarios:

- (1) A crazed senior accounting major with no job lined-up for after graduation is about to kidnap a child outside the building where I’m lecturing in a second floor classroom about the merits and demerits of psychological egoism. My endless droning on eventually goads a disgruntled student into grabbing hold of me and hurling me out the window, whereupon I fall to the ground and squash the would-be kidnapper.
— Now, this action clearly benefits the child, but it isn’t an altruistic action on my part since it isn’t even *my* action; it’s the action of the disgruntled student. As for myself, I was simply assuming the part of an inert bit of matter in my acceleration towards the kidnapper. Nor is the action of the disgruntled student altruistic, since he simply desired to hurl me out the window; he had no idea that doing so would help the child.
- (2) I leap from this same window in order to escape a fire, and inadvertently squash the would-be kidnapper flat.
— Again, this benefits the child and, what is more, it is clearly my action. But it still isn’t altruistic, because the action I intended was simply to leap from the window to escape from the fire, and not to benefit the child.
- (3) I leap from this same window with the sole intention of saving the child from the kidnapper, but miscalculate the distance and miss the kidnapper, squashing the child instead and breaking both my legs in the process.
— Here my action *does* seem to be altruistic, since its intention was to benefit another (even though it failed to do this, and indeed even harmed the one to be benefited).

Scenarios (1) and (2) indicate that benefiting another is not a *sufficient condition* of an act being altruistic, and scenario (3) indicates that it isn’t a *necessary condition*, either.

THREE UNSUCCESSFUL ARGUMENTS FOR PSYCHOLOGICAL EGOISM

Psychological egoism — the claim that all of our actions are motivated by self-interest or are “selfish” — can be understood in any of three ways:

- (1) I always do what I want,
- (2) I always do what I want, which is whatever maximizes my own pleasure (i.e., “my ultimate goal is my own pleasure”), and
- (3) I always do what I want, which is whatever maximizes benefits to myself (i.e., “my ultimate goal is whatever benefits me”).

The second position helps address problems encountered with the first, and the third does the same for the second. But all three meanings of psychological egoism are, in the end, found wanting.

“People do only what they want to do.”

This position holds that all human actions are, at their root, selfish. Even so-called altruistic actions are just those actions that the person, on balance, most wants to do — so they aren’t really altruistic at all. If someone gives money to famine relief rather than buying a new TV, that only indicates that they wanted to contribute to famine relief more than they wanted a new TV. This position assumes something like the following argument, where ‘S’ is the subject or actor:

- (1) If an action is motivated by S wanting to do it, then that action is selfish. [definition of 'selfishness']
 (2) Every action is motivated by S wanting to do it.
 (3) ∴ Every action is selfish. [1,2-MP]

Both premises of this argument can be questioned. Against premise two, one might note that many of our actions are *not* motivated by our wants or desires. Rather, we often do what we *don't* want to do, such as when we do something in itself undesirable to reach a desired end (such as visiting the dentist to get relief from a toothache), or when we simply feel that we ought to do something (such as fulfilling a promise we find onerous).

One might think that this second instance could be interpreted as belonging to the first class, where the desired end is "to be moral" or "to keep one's promises" or "to be liked/trusted by others." This fails, however, to take seriously the conflict we feel when performing certain actions (the sort of conflict we don't feel when performing a pleasant action). It seems quite consistent to say: "He didn't want to do it, but did it anyway because he thought it was the right thing to do."

Yet this objection to premise one loses its strength once we distinguish between **direct** and **indirect wants**. It's surely right that many actions are not wanted or desired directly, but it might still be true to say of all actions that they are wanted or desired as a means to some end which we *do* desire directly. Even martyrdom might be a means to a desired end (achieving a more just world, going to heaven, doing God's will, etc.).

The weakness in this argument is actually the definition of 'selfishness' given in the first premise. This definition *makes egoism trivially true*. If the goal of our actions is always the satisfaction of some want, then we can never act altruistically; but this misunderstands the nature of selfish actions. Selfish actions are determined not by their being desired by the actor, but rather by the *object* of the desire, namely, whether it is the good of another or of oneself. This is a definitional problem, but not an unimportant one, for we just mean by 'unselfish action' an action that is done out of concern or regard for another's well-being. Consequently, to say that all our actions are selfish is false, since we just do perform so-called "unselfish acts." Much of what we desire to do is altruistic, because the object of our desire is often the well-being of another.

"People do only what gives them pleasure."

This second way of understanding egoism makes the ultimate goal of one's desires always some subjective state, specifically, the feeling of pleasure. This corrects the above understanding of egoism, for now the argument requires that a selfish act is one motivated by desiring one's own pleasure. Putatively altruistic actions are performed only because they make the actor *feel good* (or else they prevent the self from having unpleasant feelings). This equates psychological egoism with **psychological hedonism**, the claim that our only possible motivation to act is the pursuit of pleasure. This basic argument is well-expressed in the anecdote about Abraham Lincoln (see the adjacent box).

- (1) If an action is motivated by S wanting the feeling of pleasure, then it is selfish. [definition of 'selfishness']
 (2) Every action is motivated by S wanting the feeling of pleasure. [psychological hedonism]
 (3) ∴ Every action is selfish. [1,2-MP]

The point of the argument is not just that we do only what we want to do, but rather that the *goal* of our actions is ultimately the feeling of pleasure.

Against premise two one might argue that *sensations are not enough*. The pleasant sensation we generally have when helping others is seldom (or at least not always) the goal of our actions. We normally don't want the sensa-

ABE LINCOLN ON EGOISM

Abraham Lincoln was arguing for psychological egoism one day while traveling with a friend in a coach down a country road. During the argument, the coach came upon a mudslide where a mother pig was squealing over her piglets, which were drowning in the mud. Lincoln ordered the coach to be stopped and got out to save the piglets. After he returned them to the sow and got back in the coach, his friend asked him whether that wasn't a clear case of altruism. Lincoln replied, "Why that was the very *essence* of selfishness. I should have had no peace of mind all day had I gone on and left that suffering old sow worrying over those piglets. I did it to get peace of mind, don't you see?"

tion; rather, we want to change something in the world. We want actually to *win* the game, not simply have the sensation of having won. In general: our motivation is not simply to alter our own consciousness, but rather to alter the world. To put the same point differently: we must first desire some result of our actions before we can derive pleasure from having attained it. Thus we do not seem to be desiring the pleasure, but rather the various changes in the world toward which our action is directed; the pleasure is just a by-product of bringing about these changes.¹²

We also might still find the definition of ‘selfishness’ in premise one objectionable. As above, this still misconstrues the meaning of selfish and unselfish acts or people. Isn’t the unselfish person just that person who feels pleasure in performing certain actions that benefit others? Even if it’s true that, in some conscious or unconscious way, the motivation for every action is the anticipated pleasure experienced from doing the action, there remains the distinction between those who feel pleasure when acting for others, and those who do not.

“People do only what benefits them.”

Psychological egoism can also be viewed as the claim that the sole motivation for any action is the benefit of the agent. This seems to fortify premise one against the above objections. It also thickens the rather thin notion that we seek only pleasure, as well as specifying the want so that premise one isn’t trivially true. The claim here is that we tend to calculate the general costs and benefits of every action, and on that basis decide how to act. The motivation to act is based on our desire to benefit ourselves — not merely in terms of pleasure (whether in the short-run or the long-run), but rather in some objective bettering of myself or my situation.

- | | |
|---|---------------------------|
| (1) If an action is motivated by S wanting some self-benefit, then it is selfish. | [definition of ‘selfish’] |
| (2) Every action is motivated by S wanting some self-benefit. | |
| (3) ∴ Every action is selfish. | [1,2-MP] |

Unfortunately, we find now that premise two is false, and false for two reasons: We often act in ways that harm us, and we often lack time to calculate harms and benefits. Much of what we do is in the pursuit of pleasure, and this often results in our injury (e.g., using tobacco, drinking to excess, driving fast cars, engaging in unprotected sex). So we would at least need to modify premise two to also include these (viz., combine premise two of this third argument with premise two from the second argument: Every action is motivated by S either wanting some self-benefit or wanting a feeling of pleasure). Further, we often have as our goal the benefit of others; and while it might be thought this is done in pursuit of pleasure, we still have to answer the objections to premise two in the second argument, viz., that we are seeking a good in the world and not just in our minds (as a sensation of pleasure).

Furthermore, *desire is often required* in order for us to act. Much of what benefits us we do out of desire for an immediate pleasure, where this desire seems unrelated to the conscious process of seeing that A benefits me and therefore desiring A. For example, with eating: I often eat from desire, and not from the conscious realization that eating certain things will benefit me, thus inspiring in me a desire to eat. Indeed, we generally have problems doing things that merely benefit us unless we also have some other desire to perform the action.

Finally, there generally is *no time for a cost/benefit analysis*, and so we tend to do either what we *have* to do or what we *want* to do — and much of what we *want* to do often directly benefits others with no obvious gain to ourselves.

In summary, the definition of ‘selfishness’ in the first argument makes egoism trivially true, while the second premises of the second and third arguments are true only if changed to: “*Some* A are motivated...” which, of course, would result in a much weaker conclusion, viz., “Some of our actions are selfish” — a conclusion nearly everyone accepts as obvious, but a far cry from the universality of psychological egoism.

¹² Related here is what Viktor Frankl called “the paradox of hedonism”: those who aim at happiness generally do not find it, while those whose lives have meaning and purpose apart from their own happiness, normally find happiness as well (*The Will to Meaning*, 1969).

[44] RIGHTS

A right is always a right *to* something and *against* someone. So with any right, there are two questions to ask: “To what do I have the right?” and “Against whom do I have the right?” Rights to the action (as opposed to the omission) of another are called *positive rights*. These are rights that some X be done to me (either “rights of contract” against an individual, or “rights of beneficence” against anyone). Negative rights are those rights that another not interfere with me, either with my doing X (active rights), or by doing Y to me (passive rights). As to the question “against whom?”, rights against specific individuals (special rights) are *in personam rights*, while *in rem* rights are rights against all individuals (general rights).

EXAMPLES OF RIGHTS

Positive *in personam*

These are generally characterized as “rights of contract.” They normally arise through a contract or promise, either explicit or implicit. Social Contract theorists might contend that

all rights reduce to rights of contract (e.g., I agree not to interfere with the liberty or security of others in the society by contracting with them that they likewise refrain; and/or I agree to aid others in distress insofar as they agree to do likewise for me).

Some positive *in personam* rights are more accurately called “rights of reparation.” For instance, if you injure my body or property in some manner, then I normally will have a right against you to make good the injury.

Positive *in rem*

These are characterized as “rights of beneficence.” It is common to attempt to base these rights in “natural law,” that is, to claim that we have by nature certain rights to the actions of others. We might include here the right to be given one’s basic requirements for survival, the right to be aided in exercising one’s active rights to life when incapable of doing so, and the right to be informed when such information is important to our well-being and/or the exercising of our other rights — e.g., for someone to inform me if they know that I am (or my property is) in imminent peril.

With positive *in rem* rights, it is important to know who is obliged to act. Does distance make a difference? Does it matter if there are several people in a position to act? Is it mere chance that decides which individual must discharge her duty? It might be noted that political libertarians typically reject all rights of beneficence.

Active *in rem*

These are characterized as “rights of liberty” or the right of autonomy, the right to act in any way I see fit. This right received its classic defense in John Stuart Mill’s *On Liberty* (1859), where he described the principle of liberty (or freedom):

The sole end for which mankind are warranted, individually or collectively, in interfering with the liberty of action of any of their number is self-protection. [...] The only freedom which deserves the name is that of pursuing our own good in our own way, so long as we do not attempt to deprive others of theirs or impede their efforts to obtain it. Each is the proper guardian of his own health, whether bodily or mental and spiritual. Mankind are greater gainers by suffering each other to live as seems good to themselves than by compelling each to live as seems good to the rest.

Rights	Positive	Negative	
		active	passive
<i>in personam</i>	rights of contract		
<i>in rem</i>	rights of beneficence	rights of liberty	rights of security

FAIRNESS

Where would we place our “right to be treated fairly”? It seemingly could be any of the *in rem* rights. E.g., that we be taxed fairly (dispossessed of property), limited in speech or action fairly, given benefits or punishments fairly, and so on. Fairness seems to transcend these distinctions.

Although considered a basic right, it is nonetheless restricted by the right of security, as well as by one's own previous agreements or actions (contracts made or damages incurred).

One might include here the "right to acquire one's basic requirements for survival." In general, the "natural right" of self-preservation has aspects that are both active (acquiring the basic material goods for survival) and passive (defending oneself against others).

Passive *in rem*

These are "rights of security," and amount to our basic right "not to be touched." It may include rights against trespass, seizure of property, damage to property or body, torture, and killing.

Passive *in personam*

These involve rights of security that are normally *waived* (e.g., I allow others to come to my door, write me letters, phone me at home, etc.), but these rights may be enforced against certain individuals, where the "touching" might count as harassment (e.g., court injunctions to prevent individuals from calling or visiting me).

Note that I might *waive* this right not to be touched in some way, but I do not *forfeit* it: the latter would amount to saying that you can touch me not only now, but anytime in the future as well, whether I want you to or not. Generally we feel that we always have the right to make good on our rights of security, although we may waive them temporarily; e.g., individuals within physically-intimate relationships always reserve the right to enforce their right not to be touched.

Active *in personam*

This is the right that some individual not interfere with your performing some licit action, but it isn't clear that there are any rights in this class that do not reduce to either rights of liberty or rights of contract.

EXCEPTIONS TO THE ASCRIPTION OF RIGHTS

The paternal care of others often involves denying certain rights to individuals (which we do, presumably, for their own good). Occasionally we also deny rights to individuals deemed "dangerous to society." There are two questions to ask here:

Whose rights can be waived?

How old must a child be before acquiring a full set of rights? Can rights be lost if one is mentally debilitated or insane, or senile? Does mild mental retardation allow for the waiving of one's rights? (How retarded must one be...?) Are the rights of criminals waived? All criminals? And all rights? How is this decided?

Which rights are to be waived?

Various rights are thought to be waived for different groups. Some examples are **corporal punishment** (waiving a right of security with respect to children, criminals, others?), and **confinement** (waiving a right of liberty with respect to children, criminals, the insane, the contagiously ill, others?). Sir William Blackstone (1723-80), an influential legal theorist, argued that convicted murderers forfeit their right to life (and thus can be executed).

Much of the discussion of what we've been calling the "right of liberty" has proceeded under the heading of "autonomy" and its limiting principles, of which the most widely discussed are the "harm principle" (limiting the autonomy of a person to prevent that person causing harm to others) and the "paternalism principle" (limiting the autonomy of a person in order either to prevent harm to that person or to help improve that person).