

ENGINEERING ETHICS PART II – ETHICAL FRAMEWORKS AND ETHICAL ANALYSIS

Why intention is not always enough

As discussed in Part I of the EDC Engineering Ethics reading, instilling the desire to be ethical is generally not the object of professional ethics instruction. That basic motivation is integral to our personal moral development, something we should each bring with us already. Although that development is never fully complete, ethics instruction presumes and relies upon the *intention* to be ethical. However, if intention alone were sufficient, we wouldn't need to worry very much about meeting professional ethical challenges. The fact is that ethical problems in the real world can be exceedingly complex; even detecting that there may be ethical implications in a decision often requires substantial thought or experience. This is why Part I of our ethics reading, as well as the lecture, focused on the many non-obvious ethical consequences implicit in design.

When faced with actually having to make an ethical decision, some means to manage this complexity are required. The field of moral philosophy is dedicated to the study of the foundations of ethical thought, and understanding some of these ideas can help us better conceive of and frame complex ethical problems, in turn helping us draw conclusions in complex situations where the right thing is not obvious. This second reading offers some ways to help you think about complex ethical problems.

Moral Systems and Principles

Moral systems are the provenance of philosophy.¹ But that does not mean that only philosophers have use for them. Indeed, they are essential tools in the engineer's toolbox when confronted with ethical challenges – which, as we have seen, may happen more frequently than you might think. So what are ethical systems? Though some philosophers would disagree, for our purposes ethical systems can be thought of as somewhat analogous to scientific theories. That is, they are attempts to systematically explain, account for and extend an already-familiar phenomenon, and put this understanding to use in new situations to make better decisions.

For example, people did not first begin to notice that things fall to earth when Galileo began some of the first scientific work on gravity. Indeed, it was only this plain fact that prompted and made possible the work to systematically understand why and how things fall. Although we knew that all objects fell to earth, Galileo was the first to point out that, once you control for varying drag, they all fall at the same rate. This insight helped to systematize our understanding, and is obviously useful to know when designing something.

Similarly, people generally recognized that lying, cheating and stealing were wrong long before philosophers began conceiving of morality as a system. Moral philosophers did not invent ethics any more than Galileo invented gravity. But moral systems help us to think about ethics more rigorously and productively. For example, theories of justice attempt to articulate this concept by describing what all just arrangements have in common, despite what may be varying appearances among them. Just as

¹ Although some authors draw a semantic distinction between them, I will use the terms 'ethical' and 'moral' interchangeably.

with a physical theory, this helps us better apply the concept to real situations and make better decisions.

This does not mean that ethics itself is a science – it is not, at least not in the same sense that physics is. However, like science (and many other fields) ethics is not formless – it has its own structure, and it can be approached systematically.

There are dozens of kinds of moral systems –ways of putting a structure around ethical problems. But a useful classification begins with whether an ethical system judges morality on the basis of what comes out of an act, or on the basis of what goes into it. The first kind looks at the good or ill consequences of an act, and are collectively referred to as varieties of Consequentialism or Utilitarianism. The second kind is more concerned with motivations for an act, specifically what duties and responsibilities we are acting to fulfill, and these are collectively known as varieties of Deontology (*deon* being Greek for duty).

Utilitarianism and Consequentialism

The broad class of theories of ethics which judge the rightness of an act by its consequences or social impact, broadly defined, are called consequentialist. The most common consequentialist theory, often associated in name with nineteenth-century English philosopher John Stuart Mill, is utilitarianism. Utilitarianism is a particularly “quantitative” variety of consequentialism – the consequence it seeks is the greatest utility possible.

It is important to understand that the criterion used as the basis to evaluate the consequences of an act – that is, the particular “utility” – varies among versions of the theory. In general, the utility can be any outcome which is measurable (in principle, if not actuality) and valued. For example, Mill himself always referred to “happiness” in a general sense. Other thinkers have been more specific, looking to maximize individual freedom or opportunity. For example, Amartya Sen, a Nobel laureate economist, propounds a metric whereby each person has sufficient resources (including food, education and political freedom) to exercise his or her personal capabilities to their fullest extent. Other potential utilities that might be the subject of utilitarian calculation include personal choice or lifespan. In some medical ethics, lifespan is made even more specific by weighting life expectancy into something known as a QALY, or Quality Adjusted Life Year. Under such an approach, the ethical act is that which produces the highest QALY number for the patient or patients.

In all these examples, the option which best provides for the relevant utility, in net terms, is the most ethical choice to make according to the consequentialist approach.

Deontology

Deontology refers to a theory of ethics credited primarily to eighteenth-century German philosopher Immanuel Kant, and is sometimes simply called Kantianism. It stipulates that adherence to moral duty is the soul of ethics. This moral duty is not to be confused with ordinary notions of duty that might accompany one’s social role or profession – for example, a soldier’s duty. The duty in deontology is our duty to the rational moral law as revealed by reason; as such, this duty is universal, and incumbent on all rational beings. Kant gives us several formulations of this universal moral duty, which he calls the Categorical Imperative. One of the simplest and most important formulations is to *never treat people as merely means to an end, but always as ends in themselves*.

Note that this does not mean that people can never serve as means to one another at all. Society is in many ways constituted by the means we serve for each other (that's what having a profession is all about). What it does mean, however, is that we never treat people *solely* as means to an end – that we never *use* people. People may be both a means and an end in themselves at once. This occurs whenever people autonomously choose their profession, for example, as opposed to being forced into it – whether that occurs through outright slavery or having no other social option.

There are many other ways we fulfill this basic obligation to respect personal autonomy. For example, we must never lie to anyone, because to do so disrespects and compromises their autonomy, and so violates the moral law.

Rights

Rights are the inseparable inverse of duties. They always accompany each other, for each one fulfills the logical implication of the other. For example, the duty not to lie logically entails the claim that everyone has the right to be told the truth and vice-versa. They are two sides of the same coin and they only work together.

For consequentialist/utilitarian theories, the ends can – potentially – justify the means. For example, if we sacrifice the interests of a small segment of the population to benefit the larger society, that may be justifiable *if* the circumstances are such that there will be a greater net benefit overall by doing so. Note, however, that this benefit need not apply to the group whose interests are sacrificed – this apparent injustice is one of the difficulties with applying any utilitarian theory. Utilitarian theory is generally sensitive only to the *amount* of benefit achieved – not to whom it accrues or fails to accrue. In such situations, deontological theories are often invoked to defend the rights of the infringed minority. In Kantian/deontological theories, the means, and respect for individual rights, must *always* take precedence over the ends. Nothing good can come from an act that is not in compliance with moral precepts.

Ethical Principles

Moral theories are attempts to account for the nature of ethics at a very fundamental, low level. Yet a more pragmatic approach is also valuable – one which does not try to account for the nature of ethics itself, but merely to account for the most salient, concrete concepts at stake in ethical problems. Philosophically, these ideas do not replace moral theories, which are more basic – but they can be indispensable tools for thinking about ethical problems. You might think of the difference, very roughly, as that between a principle of engineering and a theory in physics.

As ethical problems in medicine and research proliferated over the past 60 years or so, frameworks and precedents for thinking about them began to evolve. These frameworks are distinct from and less abstract than philosophical ethical theories such as utilitarianism and deontology (though they often draw on them). Some of these frameworks, in particular the Belmont Report, have even served as the basis for regulations.

The Belmont Report was the first major American bioethics document, formulated in response to the revelation of the scandalous Tuskegee Syphilis experiment in 1972. The experiment itself began in the 1930s, when there was no effective treatment for the disease, and it essentially consisted in monitoring the untreated progress of the disease in unwitting African-American subjects. The study,

questionable but perhaps defensible in the very beginning, accumulated such inertia over time that even when penicillin had become a routine and effective cure for the disease, Tuskegee subjects were actively denied care in order to preserve the research protocol. At that time, there was essentially no federal oversight of clinical research, and no formal codes about their ethical conduct. Only when the Tuskegee study came to public light was it finally ended among an outcry for binding laws on the conduct of medical research. The Belmont Report was the first product of this effort, and it has guided much policy and thinking in biomedical ethics in the US since its creation in 1979.

While the report is explicitly concerned with the ethical conduct of clinical trials, its application is in fact much broader than that. Unlike most regulatory documents, the Belmont Report outlines a series of principles, rather than a concrete set of rules. These four principles are worth looking at more closely because they offer a practical alternative to using moral theory for analyzing ethical cases.²

Beneficence

Beneficence is the obligation to act in a way that promotes the welfare of others – the obligation to help, aid or do good for others. Beneficence may be both a professional fiduciary duty, as particularly with a physician, or an implicit moral duty that applies to all of us. As a general moral duty, the degree to which it applies to any of us may be debatable (e.g., how much obligation do we have to help others?), but in professional situations, including engineering, the duty to some form of beneficence is fairly clear. Primarily, the principle of beneficence relies on the well-informed intention of the agent, and not necessarily on the resulting outcome.

Nonmaleficence

The original conception of nonmaleficence is expressed in the Hippocratic oath: “first, do no harm.” Nonmaleficence is refraining from causing harm to anyone as much as possible. That does not mean that no harm will necessarily be done; there is after all always risk in any kind of intervention. But nonmaleficence means that these risks should be anticipated, weighed against the potential good, and minimized as much as possible. The harm at stake is also a relative issue. Even certain harm may be tolerable if a net gain is expected – for example, even a successful surgery entails harm, but could save a life, so the net result is a gain.

Autonomy

“Auto-nomy” literally means ‘self-law’. Autonomy is the right of a competent person to decide what will happen to him- or herself. The obligation to respect the rights of others to determine what will happen to them means, for example, that everyone has the right to refuse treatment or the right to refuse to take a risk by being part of a clinical trial. Autonomy may also mean that we must allow people to undertake risks so long as they do so knowingly. For example, rather than outlaw smoking, we label the packaging to ensure that smokers are aware of the risks they are engaging in.

² In the original Belmont Report text, only three principles are outlined, as Beneficence and Nonmaleficence were combined; in later writings by the report’s authors and others, these principles are distinguished.

Justice

A central pillar of the principle of justice is the idea that burdens and benefits should be fairly distributed among groups or individuals, so that one person or group does not bear risks or costs from which others benefit most. For example, that might mean assuring that medical test subjects will be able to obtain any treatments that may result from the research they are subjected to, so that they do not bear risks and harms of research from which they might not also reap benefits. Another basic formulation of the idea is *treating like cases alike* – that is, ignoring ethically irrelevant facts in making decisions – for example, discriminating on the basis of race or gender. Of course, the question of what is and is not morally relevant may also be ethically contentious – for example, is nationality or citizenship relevant? What about a person’s economic status or ability to pay?

A related aspect of applying the principle of justice is that it may also imply a fair political arrangement as a context. In a context of deeply rooted social injustice – for example, under Apartheid, it will not do to treat all cases alike, as that would be to ignore the essential and unjust inequality at work in the society. Once again, this requires that we be sensitive to what constitutes “like” cases. While simple, the principle of justice is thus one of the most thorny to apply.

Using Principles

All four principles described in the Belmont Report are *equally important* – and for an act to be ethical, all four principles must be satisfied to their maximum mutual extent. This is the essential challenge of ethical problems. When cases are simple, it is obvious what should be done because perhaps only one ethical principle is at stake. When cases grow more complex, the satisfaction of different ethical principle come into conflict with one another, or their application to the different stakeholders involved conflicts. It is just these sorts of challenging, complex ethical problems with which engineering ethics is concerned (and this is again why ethical intentions alone will not always be enough to know what to do).

The application of the principles approach to evaluating ethical problems reflects these complexities. To use ethical principles to analyze a case and reach a decision, we must decide how each principle applies in the case, how they apply to each stakeholder, and ultimately how they should be balanced against each other. It is here that the principles approach begins to parallel some aspects of engineering design itself.

Ethics and design

Engineering design confronts similar problems of balance: products should be strong but light; powerful but affordable; full-featured (complex) but reliable (simple). All of these conflicting goals must be achieved simultaneously to the greatest degree possible. These multiple conflicting goals are partly why design uses an iterative approach: ideas are generated, prototypes built and evaluated, and new ideas generated, repeating the cycle.

Reflective ethical decision-making has much in common with this “refine and design” process. Like design, ethics also has its paradigmatic historical successes, failures and lessons. And like difficult design challenges, ethical problems may have more than one defensible solution, without ruling out the possibility of a single optimal solution. We may even conceive of the ethical principles as design goals, and the ethical act as the design which maximizes the conflicting goals. Thinking of complex ethical

problems as design problems may enable you to get a better handle on their many dimensions and prompt you to conceive of more creative solutions to them.

Theory and principles summary

In the end, both ethical theories (e.g., utility, deontology) and principles approaches (e.g., autonomy, nonmaleficence) are *lenses* through which we can examine the complex ethical problems we will encounter. They can help us think through ethical problems, but they require careful reflection, judgment and practice to apply effectively. One thing they are most certainly *not* are formulas which will tell us what should be done when we “plug in” our case, like an equation in math or physics. They are sophisticated tools – the distilled result of centuries of moral thought – and they require sophistication, and, ultimately, practice, on your part to put them to good use.

Ethical decision-making process

Whatever concepts you use to frame your ethical analysis, there are a few steps to follow to ensure that you account for the many complex dimensions of an ethical problem and ultimately arrive at an ethical solution that reflects and does justice to them. They include:

1. Gather as much information as possible
 - a. Get all the facts of the case at hand – details can matter
 - b. Know the applicable laws, regulations and company policies
 - c. Realize what the technical, factual or scientific unknowns are and what difference they would make
2. Discern all the stakeholders involved
 - a. Consider the rights and interests at stake for each of them
3. Generate and evaluate different ethical alternatives
 - a. Brainstorm multiple options – find creative solutions
 - b. Consider/predict consequences of each option
 - c. Balance rights or ethical principles
 - d. Ask for advice