

Philosophy 221 Final Exam

John Price, 98039216

December 11, 2002

1 The Categorical Imperative vs. The Hedonistic Calculus

Kant's Categorical Imperative and Bentham's Hedonistic Calculus share the same goal. Ideally, both systems seek to take the uncertainty out of making moral decisions, and provide a simple decision process to solve ethical questions. However, both are fundamentally different in their origins, and their workings.

The Categorical Imperative can be summarized in one test: "If I act on this maxim, would any rational agent refuse to consent?" It is one part of Kant's overall ethical theory which is partially founded on a belief in a quasi-metaphysical Will \ddot{e} . This Will \ddot{e} is a common reason shared between all humans. In other words, if we were able to strip away our differing desires and interests and act solely using our rational capacity, we would all share the same idea of what constitutes good and bad.

On the other hand, Bentham's Hedonistic Calculus, a part of his theory of Utilitarianism, requires no metaphysical claims, but is based on an idea about human nature. Namely, Bentham believes that all human action is based solely on our desire to experience pleasure and avoid pain. Correspondingly, the Calculus decides whether an action is good (or bad) based on the action's capacity to create said pleasure (or pain). Deciding between two courses of action is a simple matter of determining the net utility for the two actions, and choosing that which encourages more happiness in the world.

Obviously, since different people have different ideas about pleasure - one person might find happiness in reading a good book, while someone else may find the same book a downright bore - the Calculus results in a subjective morality. This is in stark contrast to Kant's universalism, which is a direct result of his belief in a universal Reason for all people.

An important distinction between the two decision processes is that while the Categorical Imperative answers the question "is this action morally permissible," the Hedonistic Calculus also seeks to assign relative values to different acts. That is, under utilitarianism, it is possible to say that action A is morally "better" than action B, by virtue of having generated a higher

overall utility.

This leads to a potential criticism of Kant's system - it tells us what we are permitted or not permitted to do, but does not tell us what we *should* do. Given two (possibly mutually exclusive) morally permissible actions, how are we to decide which to pick? The obvious solution would be to look at the expected consequences of each and choose the results we prefer. However, this solution seems to go against both Kant's staunch deontological views, as well as his stipulation that an act is only morally good if we pursue it because it is our duty to do so. It seems that there may be the potential for Kant's system to leave us totally clueless as to what our "duty" is.

Kant's dependence on reason, and specifically on a universal Reason, also leaves some room for criticism. The weakest link is the metaphysical and untestable claim of the existence of the Willë - unless I choose to affirm said existence I have no reason to accept any of Kant's further claims. Secondly, rational people often disagree on what counts as rational behaviour. Even if there *is* some universal Willë, it does not seem that we know what acting in accordance with it would entail. If there is no agreed-upon universal definition of rationality (regardless of whether or not such a rationality actually exists), how can any individual possibly hope to determine if "any rational agent [would] refuse to consent"?

Bentham avoids these problems by avoiding metaphysics. When trying to choose an option out of many, he feels we naturally tend to make a guess as to which action will result in the best outcome for us anyway. Bentham's goal is to formalize the process that we're already familiar with and provide us with a more consistent decision-making process. As such, the resulting system may seem much more "common-sensical" than Kant's, but is not without its own problems.

One common criticism is that the Calculus is too difficult, that while we have some basic capacity to *guess* the utility payoffs of certain actions, we cannot hope to rigorously perform the numerous mathematical exercises Bentham sets before us. I would propose the opposite criticism - that rather than enhancing, Bentham's calculus instead *oversimplifies* the moral decision-making process we already perform. Critics of utilitarianism point out the potential for utility monsters and the call for excessive personal sacrifice - things we see as being morally wrong, yet which classical utilitarianism does not condemn. It would seem that while we do make use of utility calculations to a certain degree, we also use other criteria which allow us to make decisions like "utility monsters are bad."

Consider the mathematician who can calculate the trajectory of a thrown ball relatively easily, but not nearly as quickly or intuitively as a ball-player

can reach out his arm and catch it. Nor can the mathematician use trigonometry alone to convey enough information that a child might learn to play catch. In this particular application mathematics serves to make catching a ball less intuitive, while at the same time providing a vastly oversimplified description of the same process. It could be argued that the Hedonistic Calculus succeeds in its goal of codifying a process for making ethical decisions but, like trigonometry, it oversimplifies the system we already use while at the same time making it less intuitive.

Both systems have their strengths and weaknesses. The Hedonistic Calculus seems to model our decision making processes up to a point but doesn't manage to exclude certain situations we think should be considered morally reprehensible. Similarly, the Categorical Imperative seems to hold up well as a general principle, but its usefulness can be questioned if we try to apply it too rigorously. It may be best to consider both as models to help us understand and improve our underlying moral codes. In physics, scientists often look at light in two different ways - either as being composed of waves, or of particles. Neither perspective seems to be entirely "true" but each is able to explain certain phenomena that the other cannot and so is useful in different circumstances. Similarly, while it may be that neither Kant nor Bentham have stumbled onto a complete and consistent theory of morality, both can be seen as models representing our underlying ethical system.

2 Kant and Intent

Kant breaks judgments about an action's morality into two steps. The first of these is using the Categorical Imperative to determine if an action is morally *permissible*. However, for an action to be morally *good* it must not only be permissible, but done for the right reason. The only acceptable reason, for Kant, is out of a sense of duty to do the "right" thing - we must do our duty for its own sake. Specifically, an action is *not* morally good if it is done out of any kind of self-interest. As an example, the Categorical Imperative would dictate that telling the truth is one of our moral duties. However, if we only refrain from lying because we wish to avoid the unpleasant consequences of being caught, we are not acting morally, but prudentially. Thus we can see that intention plays a very significant role in Kant's moral philosophy.

This reliance on intention has both its strengths and its weaknesses. It does seem to accord to some extent with our "everyday" understanding of morality. As an example, we generally do not feel that someone who's action results in the accidental death of another should be punished as harshly (if they are to be punished at all) as someone who intentionally murders. It seems that consideration of intent is a necessary part of a moral theory. Additionally, in criticizing consequentialism, Kant points out (quite rightly) that a person has more immediate control over her intention and action than over her action and its consequences. It would seem fair to punish people for something only to the degree that they have some control over it.

On the other hand, at least so far as our justice system is concerned, an action's consequences also bear on the judgment of said action. For example, attempted murder is often judged less harshly than successful murder. In both cases the intent and action may be the same, but the consequences are different, and so we judge the overall situation differently. This can be seen in a couple of different ways: either our "everyday" morality (as manifested in our system of justice) is correct, and Kant's system is leaving out an important factor in making moral judgments, or our morality is incorrect and we should be considering only the action and the intent behind it. In the case of attempted murder, it would seem easy to make the argument that attempted murder *should* be judged as being *as wrong as* actual murder. But consider someone with maximally benevolent intentions whose actions have maximally malevolent effects. Assuming that the action is in accord with the Categorical Imperative, do we really want to say that the action was *good*? On an entirely abstract level, one might concede that the action itself, separated from its consequences, was indeed good. But, back in the

real world, would we be content with giving this person a pat on the back and saying “better luck next time”? If in the process of trying to kill someone an assailant instead saved their life, would we still want to punish them as if they had successfully killed their victim? It doesn’t seem likely. Rather, it seems that while intent is certainly important in making moral judgments, the consequences of an action cannot be entirely ignored, and as such Kant’s system is incomplete.

It also doesn’t seem entirely accurate to say that we have no control over the consequences of our actions. *Occasionally* the unexpected happens, but the majority of the time we can predict the results of a given action with reasonable accuracy, at least into the immediate future. It seems quite possible that a particular action could pass the test of the Categorical Imperative, and be done with the proper motive, but have disastrous consequences that could have been avoided with a little foresight. Do we really want to say that such an action is still “good”? Perhaps it isn’t possible for such a situation to arise - one could argue that a rational agent would forbid it on the grounds of its likely disastrous outcome. On the one hand, foresight seems an integral part of a “rational agent” but at the same time, such an appeal to consequences seems contrary to Kant’s views.

More specifically to Kant’s views on intent, his idea of what constitutes “proper intent” may differ from many of our own ideas. According to Kant, a good person might, for example, give to charity because he has a moral duty to do so. Kant’s system does not preclude such a person from also gaining pleasure (or public accolades, or salvation in an afterlife, or whatever) from helping those in need, but demands that a truly moral individual not do so with such benefits as a specific intent. On the one hand, this ensures that the man who makes hefty donations to charities because of the tax breaks it gets him isn’t judged as being Good. However, it seems quite possible that a man who was by all normal accounts a horrible person, made miserable by performing acts of kindness, could actually be the perfect Kantian moral man if, despite his negative feelings, acted according to his moral duty for its own sake. At the same time, the man who gains joy from helping those in need, and helps them because of this joy, is deemed a less moral person in Kant’s system. Do we really want to say that someone who helps because he feels obligated to is a better person than the person who gains joy from doing so? Kant’s idea of proper intent seems entirely disjoint from what we feel makes a good person. Is someone who, with the intent of being honest, makes children miserable by telling them that Santa doesn’t exist a better person than those who would perpetuate this lie? Most of us would probably say “no”. Similarly, it seems that someone could insist on making

you miserable by (in the interests of truth) constantly reminding you of every one of your faults, yet still qualify as a “good person” within Kant’s standards of intent.

Thus, we have two main criticisms of intent in Kant’s system. First, it seems that while intent is certainly an important ingredient in a moral theory, it cannot entirely remove appeals to consequentialism - we need to consider the intent behind an action, the action itself, and the consequences of the action as a complete set. In different cases, intent may weigh more or less heavily than consequences in our judgment, but neither element can be left out entirely.

Secondly, Kant’s idea of “proper” intent leads to an evaluation of what makes a good person that is inconsistent with our everyday morality. Strict adherence to our Kantian duties could (potentially, but not necessarily) result in a world in which everyone was brutally honest and, as a result, miserable most of the time. Kant wasn’t aiming for happiness and so is not inconsistent on this point, but one would imagine that most people would prefer an “immoral” world where everyone is kind to one another, to a “moral” world in which they are constantly told that they could do to lose a couple pounds.

Overall, while Kant provides us with some interesting insights into the workings of morality, strict adherence to his ethics seems to lead to undesirable results. This discrepancy is largely due to the difference between Kant’s ideas about good intentions and those of the majority of society. It would be interesting to know if Kant would truly prefer to live in a brutally honest society, or one in which the truth is occasionally bent in the interest of kindness.