

ON DESCRIPTIVE ETHICS

Andrew Boucher

August 2009

In its descriptive sense ethical language allows one to make assertions, which like other assertions may be true or not. “One should not torture,” descriptively, makes an assertion about torture - that it is an act that one should not do. While the peculiar force of ethical language comes from its overloading of different types of uses - descriptive, imperative, and emotive -, our concern here will be with the descriptive. Many of our assertions will focus on the English word ‘should,’ although *mutatis mutandi* they hold as well for other ethical terms, such as ‘just.’

Meaning of Ethical Descriptions

I mean ‘should’ in a particular way; because of the way I mean it, “One should not torture” is true. Presumably I learned the meaning of the word ‘should’ by being confronted with various acts while I was growing up, and being instructed by my parents, by others, and by society in general, that such and such was an act that I *should* do, and others that I *should not* (e.g. don’t hit your brother); or by learning principles, rules, or maxims which provide a reason to do or not to do certain types of acts (e.g. do unto others as you would have others do unto you). Sometimes the specific instruction and the principle might be linked (e.g. don’t hit your brother because you wouldn’t like it if he hit you).

‘Should’ would be learned imprecisely even if it were taught only by pointing to specific acts (this act is one you should do, that act is not), because there are aspects of the act, such as its consequences or the intention of the actor, which cannot be seen or sensed or exactly known, which are important to determining its category. The use of principles only complicates the task, because of their abstraction, number, and differences.

It is possible, at different times, that I mean ‘should’ in ways which are conflicting or incoherent. At times I think people should be paid more when they work more, and at times I think people should be paid equally. Perhaps these times even coincide, although one spontaneous instance of meaning is more apt to draw on one aspect. A simple model of how this might arrive: from one parent I am taught about the value of hard work, and from another I am taught the importance of equality. In the case of an abstract term like ‘should,’ incoherence may never be remarked and so does not get sorted out.¹

¹ Henry Sidgwick, *Methods of Ethics*, 7th ed., p. 508: “For the negation of this connexion [between universal and egotistic hedonism] must force us to admit an ultimate and fundamental contradiction in our apparent intuitions of what is Reasonable in conduct.”

Another simple way to arrive at conflict is, by learning that an act is one that one should do because of certain principles, when it turns out that they are not actually supportive. For instance, one can imagine that a person has been taught and so believes that one should do an act A because the Bible says so (call this reason X) - but it actually turns out that the Bible does not, and indeed for the sake of argument actually says the contrary. If this is brought to the person's attention, one response might be to deny or diminish the fact. But if the person accepts that the Bible says the contrary, then, based on the person's psychology, he may believe still that one should do A but no longer accept the reason X, for instance by modifying it in some way; he may believe that one should not do A and that the reason X is still valid; or he may simply stay conflicted, still believing both that one should do A and that reason X is valid. In the latter case, one would expect the conflict to be reflected in the meaning of the person's 'should', and that the meaning itself is conflicting or incoherent when related to A.

And another simple way is that one learns principles which conflict. One can imagine an individual assimilating both rules "One should do unto others as one would have them do unto oneself" and "One should be true to oneself," which may conflict for anyone who is not perfectly saintly (such as myself).

Suppose that one means 'should' such that, superficially, one should do A and one should do B, when A and B are mutually exclusive (doing A means one cannot do B, and *vice versa*). For instance, one may mean 'should' such that one should follow rule X and one should follow rule Y; that following rule X one would do A, and following rule Y one would do B. In this case *neither* one should do A *nor* one should do B, and neither "One should do A" nor "One should do B" are true. That is, "Not one should do A" (meaning, "It is not the case that one should do A") and "not one should do B" are true.

Suppose B is the act of not doing A. Then while "Not one should do A" and "one should do B" are usually equivalent, they are not in cases where 'should' has a conflicting or incoherent meaning. In the case of incoherence, "Not one should do A" is true, because it is a denial of incoherence; but "one should do B" is not true, since it is an assertion of incoherence. It therefore can be a logical error to conclude, from "Not one should do A," that "One should do not-A," since the former may be true but the latter not. And so it is a logical error to assume it is the case that either "One should do A" or "One should do not-A" is true. (Still, of course, "One should do A or not one should do A" is always true; in particular, in the case of incoherence, the latter disjunct is true.)

There are situations in which there is no act, including the act of doing nothing, that one should do.

Similarly, when X are reasons or a rule to act, it is a logical error to suppose that "Not one should follow X" and "one should follow not-X" are equivalent.

For instance, one can imagine a situation where, if one were to torture an individual, then the individual would give information which would save millions of

people from unspeakable agony. In this situation there is an incoherence in my meaning of 'should' - torture is bad but saving millions is good - and so "not one should torture in this case" and "not one should not-torture in this case" are both true. There is *nothing* in this situation that one *should* do. It is a logical error to infer, from "not one should not-torture," that one should not-not-torture, i.e. that one should torture. In all situations, one should not torture, in the sense of "it is not the case that one should torture."

The overloading of the descriptive 'should' makes reflection on what I mean by it, and its reflective use (such as in considered judgments or philosophical discussions), where I make claims or draw inferences or create theories based on how I mean the word in some abstract or general way rather than how I mean it in a specific instance, especially problematic.

Removing conflict or contradiction *changes* how a word is meant. The replacement of a nebulous meaning with a clear meaning is a logical error; if one has really succeeded in making a meaning clear, then it cannot be equivalent to the original, obscure meaning precisely because one is clear and one is obscure. One can answer more questions and so forth with a clear meaning, but precisely because one can answer more, shows that one is not answering the original questions with the original meaning.

Perhaps answering a different question helps in showing how to answer the original; but perhaps it doesn't; and perhaps it even hurts, for instance by creating confusion or by making it seem there are similarities when there are not.

Perhaps considering a different, clearer meaning has interest or value in its own right. In order to provide a test whether this is the case, a different word, such as 'plok,' could be used in replacement for the original word. One would write, for instance, *A Theory of Plok* instead of *A Theory of Justice*; in such a book one would first explicate what one means by 'plok' and then show what inferences can be drawn from this meaning. Interest or value will then appear, based on the meaning, rather than based on the word.

Replacing the descriptive content of an ethical term with another, however apparently clearer or better in some way, and keeping the same word, is especially problematic because of the emotive and imperative senses overloaded into the word. There is no reason to suppose that the new descriptive meaning should be bundled with the same emotive and imperative senses, although the use of the same ethical term probably is an indication that the author of the replacement would surreptitiously like to benefit from them.

As with other words individuals may mean ethical terms differently. Indeed, from the indirect evidence that I have - the uses of ethical terms by different people and how people reason about them - individuals *do* mean them differently, although of course there is overlap. As with all words a certain uniformity in meaning of 'should' is

enforced, but unlike more concrete terms such as 'giraffe,' where it is possible to point to an object or a picture, there is no ready manner to enforce precise uniformity, and there are both subtle and important differences in how different people mean the same ethical term.

A community of speakers means ethical terms. The communities in which I belong mean 'should' such that "One should not murder" is true. Given how differently individuals mean 'should,' a community's meaning of 'should' may be particularly nebulous.

Evidently, in order to be understood as one would want to be understood, it would be best that one's own meaning hewed as closely as possible to how the community which one is addressing means it. But in some, if not most cases, this is problematic with a term as indeterminate and disputed as 'should.' Not only is the community's meaning varied and nebulous to begin with, and not only is it difficult if not impracticable to determine its meaning, in any case I cannot change my life or my experiences, on which I draw in my meaning, and to try to change it to suit the community at large could be extremely difficult.

Assertions of the form "One should do A" and "One should not do A" can be analytic. It is so simply because I mean 'should' in a particular way. "One should not torture" would be an instance.

However, not all true ethical assertions are analytic. For example, "One should vote for A rather than B" would depend on features about the world - who the two candidates are, what they propose, and how the world is.

Our knowledge of ethical facts can advance with time. Revelation of new matters-of-fact do occur, such as the consequences of acts hitherto not or not well known. This happened, for instance, with the discovery of the ill effects of smoking cigarettes, and it was a new fact that one should not smoke.

Suppose again that the way that one means 'should' is such that, superficially, one should do A and one should do B, where A and B are mutually exclusive (doing A means one cannot do B, and *vice versa*). Then, as I have said, neither "One should do A" nor "One should do B" are true. Still it may seem to the individual that "one should do A" (or "...do B" or both) is true and indeed analytically true, and he may have the same intuitive feeling for it as when he is faced with a real analytical truth, such as "All bachelors are unmarried." This makes reflection on ethical assertions and ethical arguments especially problematic, e.g. when an individual considers "one should do A" but fails to consider "one should do B" and asserts mistakenly the former's truth and then draws inferences.

Relativity of Moral Assertions

We can imagine someone learning an entirely different meaning of 'should,' e.g. where one 'should' lie. But this does not imply that one should lie, because in the last 'one should lie,' I mean 'should' in the way *I* do. And according to how *I* mean it, it's not the case that one should lie.

Perhaps some people mean ethical terms so that they are incomplete and need an indexical. In assertions of taste, I do mean using an indexical; for instance, when I say, 'Chocolate chip ice cream tastes good,' I mean, "Chocolate chip ice cream tastes good *to me*." And some people, when they say, 'Lying is bad,' may mean "Lying is bad *to me*." But that is not how *I* mean ethical terms. Mine are complete and do not need an indexical.

Imagine the case where everyone is like myself and means ethical terms completely i.e. not with an indexical. Then everyone means ethical terms in their own way - the meaning is relative to the individual - but each make assertions which are not relative and which are, accordingly, universal. So there could well be the situation where my meaning of 'One should not lie' is true (true for *everyone*), while someone else's meaning of the sentence is false, because we mean 'should' in different ways.

Remark that here I am distinguishing a *sentence*, which is a particular sort of concatenation of words, and its meaning. (I call the meaning of a sentence a *proposition* or a *statement*. And here I am using *assertion* to cover both sentences and statements.) So while the meaning of a particular descriptive sentence with ethical terms may mean different things to different people, there are ethical statements and propositions which are true, and thus true to everyone, and not just true to some and false to others. Again, that, at least, is how *I* mean my ethical language.

An individual may even mean ethical terms differently at different times, so in the morning, in one context, his meaning of 'One should not lie' is a proposition which is true, but in the evening, in another context, his meaning is a proposition which is false (in my case, however, 'should' does not seem to change its meaning in this way).

What the word 'should' means is relative to an individual (and indeed may be relative to an individual at a particular time). Ethical *propositions* are not relative. The proposition "one should not torture" is true - even though one could mean 'should' differently, so that the sentence means a different proposition and even a proposition which is false.

Different people may disagree, correctly, whether 'One should do X,' for a particular X, because they mean 'should' differently. Indeed, it is possible, and seems actually to occur, that 'One should do X' is meant as a true analytical proposition by one individual, while 'One should not do X' is meant as a true analytical proposition by another, for certain X. For instance, one can imagine where a poor man asserts, in a

true, analytical way, 'One should take money from a wealthy man,' while a rich person asserts, similarly, 'One should not take money from a wealthy man.'

Different people may have different values. Values are relative. For example, while I value honesty, some people may not, and indeed do not, value it. Nonetheless, they should be honest, and indeed they should value honesty. Evidently, the last sentence means a true proposition *as I meant it* (and since I am the one who wrote it as an expression of what he thinks, I am the one who counts), and it would also be a true proposition for the way many communities mean 'should;' but it would probably not mean a true proposition when meant by those people who do not value honesty.

There are universal values, in the sense that all people should value certain things, such as life and liberty. Of course, someone might mean 'should' in such a way that it is not the case that 'one should value life or liberty.'

It is said that ancient Romans thought that they should abandon an infant to death if it was deformed. They should not have. No matter of fact has changed; only values and, presumably, the meaning of ethical language.

If someone should provide an ethical discourse which is innovative with respect to current moral language but which, for whatever reason, is convincing or influential, then the meaning of ethical language tends to change towards that of the discourse. After the fact, based on the new meaning of ethical language, a progression in the ethical level of society is sometimes proposed, which is true but trivially so.

While some may think it clearer to say, 'According to how I mean 'should,' one should not lie,' it is clear to me that when people say, 'One should not lie,' they mean 'should' in their own way and not in someone else's. So the explanatory preface is, at least in ordinary conversation, pleonastic.

Axioms, Conditions and Principles

It is possible that some individuals have such simple meanings of 'should' that one can find necessary and sufficient conditions for their meaning of 'one should do A.' For instance, someone raised in a strict environment without contact with other views, might learn to mean 'one should do A' in a way which is equivalent to "God wants A."

The utilitarian equivalence (based on the greater good of the greatest number) does not capture my meaning of 'should,' since, for one thing, according to how I mean 'should,' one should not shoot oneself in the foot in order to save two unknown people from being shot in the foot.

For anyone like myself, who learned his meaning in an environment with many, sometimes conflicting views, it looks highly doubtful that one can state non-trivial conditions which capture precisely how he means 'one should do A.' (A trivial condition

of course is possible: simply, “one should do A.”)² This includes descriptions of the sort ‘an individual of type T would do A,’ where for instance “an individual of type T” might be “a rational, impartial individual.” (Again, there can be trivial types, such as a “perfectly good individual,” since one should do A precisely when a perfectly good individual would do A.)

Part of the attraction of providing conditions and principles for ‘one should do A’ is that, when they are determinate and universal, they may provide a unique answer as to whether one should do any particular act, whereas people may, left to their intuition - their particular meaning -, hesitate or disagree about whether one should do it or not. But having determinate conditions and principles when different people’s intuitions disagree proves that the conditions and principles have not captured at least one person’s intuition (and, of course, probably many peoples’ if not most or even all); the conditions and principles at best have taken a side, or compromised, or done something else entirely. Far from providing an advantage by making differences disappear, the conditions and principles, by being determinant when intuitions are different and nebulous, show they must be different from peoples’ intuition.

An individual who only uses reason, i.e. begins only from logical axioms and uses logically valid methods of reasoning, will not be able to infer any ethical assertion of the form, ‘One should do A’ or ‘One should not do A’ (where ‘A’ refers specifically to some type of act) simply because he has not begun with any non-logical axioms and ‘One should do A’ is not a logical assertion (different interpretations of ‘should’ could make it true or false).

If one allows the individual to begin with axioms which are analytically true, then he may be able to draw inferences of true ethical statements; in my case, one would be able to infer that “One should not torture,” since this is analytically true and it implies itself. Still, it looks highly doubtful that any set of simply stated non-trivial axioms will produce any significant body of ethical theorems which coincide with a person’s, or at least my, meaning of ‘should’.

Reason, when deriving only from logical axioms and using logically valid methods of reasoning, is trivially impartial - it will trivially not favor any individual over any other, simply because it will not be able to draw *any* conclusions along the lines of ‘One should do A.’ When axioms are allowed to include propositions which are analytically true, then moral reasoning may not be impartial and instead may rather favor certain individuals, since rational individuals are by no means necessarily impartial and may in particular favor, to some and even to a large extent, their own well-being, and the well-being of those close to them, over others (e.g. a wealthy person might mean ‘should’ so that no

² Sidgwick, *Methods of Ethics*, p. 32: “What definition can we give of ‘ought,’ ‘right,’ and other terms expressing the same fundamental notion? To this I should answer that the notion which these terms have in common is too elementary to admit of any formal definition.” G.E. Moore, *Principia Ethica*, p. 62: “But far too many philosophers have thought that when they named those other properties they were actually defining good; that these properties, in fact, were simply not “other,” but absolutely and entirely the same with goodness. This view I propose to call the ‘naturalistic fallacy’ ...”

one should take money from a wealthy man). A certain bias which favors my own situation seems to be embedded in my own meaning of 'should.'

Dispositions

There are acts which I have a disposition to commit. I am more disposed to do those acts which I believe one should do, and I am apt to justify those acts which I am disposed to do as ones that should be done. Nonetheless, there are usually many cases in which one is not disposed to do acts that one believes one should do or is disposed to do acts that one believes one should not do. For instance, someone may learn to lie because he hears his parents lie often, even though he learns that he should not. Or he may learn that one should give money to the poor, but he is disposed never to.

I have conflicting dispositions. I may want what is best for myself, and at the same time what is best for my children, even in situations where it is impossible to have both. These conflicts are always resolved, in the sense that at any moment I always do something; in any given situation, I either do A or do not-A, for any act A.

Suppose that someone is taught not to lie but hears his parents lie frequently and so learns to do so himself. If he also learns to be disposed to do acts that are or can be labeled as ones that one 'should' do, then he has conflicting dispositions - on the one hand, to lie, and, on the other, to do what one should do and not lie. In any specific act, these conflicting dispositions are resolved, since he will either lie or not lie.

Human beings have certain natural or innate dispositions. Some of them seem to coincide with or induce acts that the human beings should do - such as self-preservation and love of one's children -, while others - egotism - do not. Of course even these natural dispositions may conflict; e.g. self-preservation and love of one's children will conflict in situations where one has to choose between one's own life and that of one's child.

The acts which I do or would do are not equivalent to the acts which I should do. I might torture an individual if I thought information which would save the lives of those closest and dearest to me could be certainly obtained; but that does not imply that I should torture in such a situation, just that I may be disposed to act in a way that I should not.

Ethical Discourse and Arguments

Descriptive ethical discourse sometimes has the intent of changing how people behave, by relying on the fact that many people are more disposed to do what they believe they should do. So, by making others believe that one should do an act, one is making it more likely that people behave that way.

Ethical arguments can be about matters of fact, for example, whether a government should borrow money or raise taxes. Two people could agree on the meaning of 'should' yet disagree about borrowing versus taxes because they disagree about the consequences of one or both policies. A person can be wrong about the consequences or other particulars of an act, so believe wrongly what one should do.

Ethical arguments can also be semantic, where two people simply mean 'should' in different ways and thus are making assertions which syntactically are the same but which are meant differently by the two.

To determine whether an argument is semantic or not, disputants in an ethical argument might find it useful to ask themselves what, if any, matters-of-fact would change their point of view.

Given a list of facts and the assertion that they are the only relevant facts, one may conclude that one should do an act A; but when an additional fact is added to the list, it is possible that one concludes to the contrary that one should not do A.

For instance, given a list of facts, one might typically conclude that one should not hit a particular person. (Ted enters a room and greets everyone with a smile. If those are the only relevant facts, then one would conclude that one should not hit Ted.) However, if one adds a new one (Ted is also threatening Fred with a knife), and the new list is said to contain all the relevant facts, then one may to the contrary assert that one should hit him.

This instability makes the conclusions of some ethical arguments tentative at best, because they may depend on all relevant facts being presented, and in complex situations - precisely the ones which elicit the most disagreement - this is extremely difficult to do. So someone may well provide a wall of facts to support his point of view, but leave out an important one which would change or diminish the conclusion. Or, in an argument between sides, one side may present facts which support a conclusion, the other side may add a new fact which supports the opposite, and then the first may add yet another fact which supports the original, and so on.

Even ethical arguments which are essentially semantic, can be important, because of the link between what one believes one should do and what one is disposed to do. So if one party can be convinced, even contra his own meaning, that one 'should' do an act, he may become more disposed to doing it.

Individuals and groups battle to advance their interests or ideas or justify their dispositions and try to change other individuals' meaning of ethical language. For instance, the music industry over the past decade has tried to establish that downloading music without consent is *stealing*, even though there is no physical thing being taken. Perhaps it will win, and thus establish that one *should not* download music without consent, in which case the meaning of 'should' has been extended in a way which the industry wants. Then this meaning will be taught to the young and become more entrenched so that, in the future, people will be amazed in the year 2000 that

some people thought that one *should* download music without consent (or perhaps not - this is an example, not a prediction).

There are other manners of discourse, besides ethical ones, which can change peoples' behaviour. For instance, one can point out that a particular manner of action is in the person's best interests, e.g. lead to a longer, healthier life, since, in general, people have a disposition to further their own interest. (Remark that reflection on or new information about the consequences of one's acts may change what is in a person's interest. For instance, a person on reflection may realize that an act, otherwise beneficial to him, has consequences which harm others severely. It then may no longer be beneficial to him, since the understanding there are consequences harmful to others detracts from the act, by causing feelings of guilt and so on.)

Rather than asking what an individual should do, one can ask what he should do to achieve a particular objective, for which he has a disposition, such as what he should do to lead a long, healthy life. Or two countries ask what they should do to achieve peace. Or a government may ask what it should do if it is to follow the utilitarian rule.

Evolution and Social Preservation

It may be thought that the particular meaning of 'should' held by a community can be analyzed by using a Darwinian perspective, that is, competition between communities and selective evolution have led to certain communities and their ethical values and ethical language surviving, and those which survive, have survived because they have contributed to the endurance of their societies. Now it is possible that a society has a system, or even lacks one, which contributes to its demise, on its own or in competition with others. For instance, a society which believed that one should murder without restriction is less likely to survive. Or a society which believed that everyone should play but never work, but which needed workers to obtain food, would presumably quickly starve. On the other hand, the idea of a right of an individual to hold property seems to encourage the accumulation of material objects, which produces technological progress and incites other benefits, so that population and the production of belligerent instruments increase, which gives the community an advantage over societies with no or a different notion of property, and thus may endow it with a better chance of survival.

Nonetheless, a society is composed of many features, the moral code being just one part, so the existence or demise of a society may be due to other factors, e.g. access to natural resources, societal organization, and natural disaster or hazard in war, to name only a few. A tidal wave might wipe out an island community. Or, if Lee had won decisively at Antietam, then the Confederacy - and slavery - might have survived and spread. A society with other advantages, might endure despite having a particular moral code, rather than because of it.

Also, what works to help or ensure survival in some circumstances, may hinder it in others; and indeed its very success may backfire and cause its own downfall. So for

instance, a right to private property, by encouraging the exploitation of physical resources, may lead to their eventual exhaustion. Or a strong work ethic may lead to the accumulation of wealth, which reduces the work ethic of succeeding generations. Or the maintenance of an overwhelming military force may protect a country and ensure its survival, or indeed allow it to destroy other societies, but reduce its investment in other areas, and so perhaps eventually lead to a lack of economic health and an eventual reduction in its military standing.

So the only societies which survive, by wiping out the competition, may produce the very conditions which lead to their own downfall and indeed perhaps even the extinction of the species; whereas the competitor societies, or even a state of nature, had they been allowed to continue to exist, would not have succumbed to the same problems.

The example of infanticide in the present day would seem to show the limits of exclusively resorting to evolutionary explanations for morality. Apparently, the killing of babies who are so deformed or unhealthy that they will take up more resources than they can possibly contribute, makes a society stronger without creating compensating weaknesses.

Allowing murder in general may create a weaker situation by installing a climate of fear, since each individual has a higher probability of being killed precisely because another may suspect him of plotting murder and thus act preventively. But a baby does not have the capability to murder.

Allowing the killing of individuals in general (i.e. not just babies) who are extremely deformed may also create a weaker situation, by making healthy individuals worry that they themselves might, by accident or other means, become deformed in the future and thus be subject to murder. But, in order to worry, one must be past the age of being a baby, and so no one can worry about becoming an extremely deformed baby.

Admittedly, the creation of a simpler rule - all human life is inviolate, including that of extremely deformed babies - rather than one based on subtler points, might be said to have helped render life safer for everyone and produced other benefits; but the case of abortion in our own day shows that society is able to allow exceptions to a blanket rule.

Yet our present society developed away from, and not towards, infanticide. Even were infanticide reaccepted in the future, it would seem that exclusively evolutionary justifications would have difficulty explaining its rejection in the present day, after its earlier acceptance.

(I am not, to be clear, arguing that one should practice infanticide; one *should not*.)

Why Bad People Make Good Money

Suppose there are two types of employment available, one where a person will do good and one where a person does not, according to a society's moral code. In general, following their moral leanings with which they have been inculcated, people would tend to want to do the good job, thereby increasing the supply of labor and lowering the remuneration needed to attract a candidate. And fewer people would want to do the bad job, thereby decreasing the supply of labor and increasing the

remuneration needed. So, on average, jobs which allow one to do good will pay less than ones which do not. A good person - that is, someone who follows his society's moral code - is more likely to find the good job attractive in spite of the lower pay, while a bad person will be indifferent to its moral aspects. Hence a good person is more likely to be in the good, lower-paying job, and the bad person in the bad, higher-paying one.³

A Few Comments on Responsibility and Punishment

Distinguish between the persisting person and the actual or present person (the persisting person considered at a particular moment or short interval in time). The present person is never responsible for who he is, since it is the past not the present which influences the present. On the other hand, the persisting person may be responsible for who he is and what he does, to various degrees depending on his history. Of course everyone is influenced externally, by parents, siblings, friends, teachers, and others, in their past, but commonly only a persisting person who has suffered abnormal external influence (such as excessive mistreatment while a child) is held less responsible for who she is.

There is a more prosaic sense of responsibility where a person is responsible for an event or a state of affairs if he caused it in whole or in part. Retributive punishment for crime generally focuses on this notion of responsibility, and one is punished if one has caused a crime. There may, of course, be extenuating circumstances to his responsibility, for instance if intent was lacking, or if he was temporarily unlike his normal self (e.g. under intoxicants or temporarily insane).

Retributive punishment is meted out for past acts, but the current person, not being responsible for past actions, is never responsible for past crimes. This does not, however, imply that punishment should be considered as being meted out to the persisting person; rather, it may be thought of as given to the current person for his likeness to the previous person who committed the crime. This, I think, matches the intuition that a criminal who has changed significantly, for the better, since the crime, is less worthy of punishment.

If retribution is for the current person for his likeness to the previous person who committed the crime, then abnormal external influences, however abominable, which may have influenced the previous person to transgress, are not an extenuating factor. Those external influences are still part of the current person and do not diminish the likeness to the previous one. For instance, in this framework, if a man has committed a crime because he was mistreated as a child, he is not entitled to extenuation because of the mistreatment, because that does not, by itself, diminish the similarity between the current person and the person at the time of the crime.

³ Adam Smith, *The Wealth of Nations*, p. 89 [Everyman's Library]: "The most detestable of all employments, that of public executioner, is, in proportion to the quantity of work done, better paid than any common trade whatever."

Still external influences can be important in determining punishment of a non-retributive sort, e.g. for punishment which is designed to protect society from further wrongdoing. For instance, if a person committed a crime because of external influences which are not likely to be reproduced, or which can be easily counterworked, then punishment with an eye to protection is less needed.

Conclusion

Descriptive ethics does not compel. It compels morally, of course, but it does not *compel* - the fact that I should do an act does not imply that I will do it. And, even though it is true that my belief that I should do an act makes it more likely that I will do it, this is a psychological fact about myself - a disposition that I have -, and there is no logical implication between should doing an act and being more likely to do it. Some people, apparently, do not share this disposition, including many in important and powerful positions. So not only is descriptive ethics often nebulous and tentative and varying, it also is merely descriptive. There are facts in descriptive ethics, but there are *only* facts; and facts are facts, and facts are *just* facts. And yet, and yet - ethics is important to me, and I commend its study to others.