

Directed Studies: Philosophy
Professor Gregory Ganssle

By submitting this essay, I attest that it is my own work, completed in accordance with University regulations.—Brian Earp

Telling a Lie to Save a Life: Kant's Moral Failure and Mill's Mere Suitability

by Brian Earp

Suppose your best friend has taken refuge in your house. She is hiding from an evil murderer who intends to kill her and has the means to kill her—only he doesn't know where she is. A knocking at the door. You open it. The murderer, whom you recognize from your friend's description, stands on your porch. He asks abruptly, "Did someone just run into your house? I'm looking for her. It's important. I must know the truth." Your choices here are limited. You may try to avoid the question, but this will only give you away. So in fact you must either reveal the truth—"Yes, she is in my house"—or blatantly lie—"No, there is no such person in my house."

(Now perhaps "in reality" there are some clever alternatives; one can think of a few at least. For our purposes, however, you may only answer Yes or No. A Yes is both honest and *factual*—i.e. your friend cannot "slip out the back door" unbeknownst to you—while a No is both dishonest and counterfactual. Answering Yes, while truthful, results in the murderer's entering your house and killing your friend; answering No, while dishonest, results in the murderer's giving up and your friend living another day.)

Surely common sense warrants a lie in this situation. Someone could argue that in fact you have a *moral responsibility* to lie, in order to protect your friend (who is innocent) from this criminal (who is evil). Or perhaps the murderer has given up his "right" to know the truth, since his own ends are harmful and immoral...

Non-intuitively, however, Immanuel Kant, in response to a scenario exactly analogous to the one just described, claims that one is *always* morally obliged to tell the truth. In *On a supposed Right to Lie*, he writes, "To be truthful (honest) in all declarations is...a sacred and unconditionally commanding law of reason that admits of no expediency [including philanthropic or life-saving expediency] whatsoever" (65). John Stuart Mill's Utilitarianism, on the contrary, would seem actually to require as a matter of moral duty that you *lie* to the murderer: evidently you increase the greater happiness of all by saving the innocent and thwarting the guilty.

But despite this contrast, one must keep in mind an important similarity: at base, both philosophers attach an actual moral duty to your choice. In other words, although Kant's and Mill's respective moral programs result in different (opposite) answers to the same problem, neither is a "to each his own" moral relativist.

With these points in mind, I will attempt the following. First: to spell out each philosopher's moral system with regard to the case at hand. In doing this, I hope to illustrate that Kant's view of his own philosophy clearly requires categorical honesty¹ while Mill's Greatest Happiness Principle (GHP) clearly requires dishonesty. Second: with regard to whether there is a moral duty in the first place (to act one way or the other), I will attempt to argue (a) that Kant's categorical imperative has a stronger and more consistent *theoretical* basis for establishing such a duty, (b) that Kant's system, however, falls apart when one tries to apply it practically (especially since he insists, erroneously, that a lie is *always* immoral), (c) that Mill provides a comparatively weak theoretical basis for his assertion that one has a *moral duty* to act in a Utilitarian fashion (and therefore in our case tell a lie), and (d) that Mill's system, however,

¹ I believe it is possible to read Kant's *Grounding* such that lying to the murderer is morally permissible, even though Kant himself does not allow it. We will address this possibility later.

actually does provide the "right answer" to the problem, and is therefore more successful practically. To clarify this last point: I will argue, ultimately, that it is moral to lie to the murderer. I will also try to show that it is in fact *immoral* to speak truthfully to the murderer (thereby permitting him to kill my friend), and that Kant therefore fails despite his promising theoretical foundation. However, I do not take Mill as my authority for the moral permissibility of lying, but rather David Hume, as we shall see.

But let us start with Kant. He first of all rejects the (Utilitarian) idea that **happiness**—one's own or anyone else's—is the moral end of action, writing in *Grounding for the Metaphysics of Morals* that "If...happiness were the real end of nature in the case of a being having reason and will, then nature would have hit upon a very poor arrangement in having the reason of the creature carry out this purpose" (8). Indeed, an animal, say, a cow—one whose actions are determined much more by instinct than by reason—could be said to be much "happier" (or at least less unhappy—less able to *be* unhappy) than us human beings—we who are tortured by existential awareness, a "free will," and an ability to reason. If happiness were the end, Kant seems to say, we would have been created such that reason played a very small role for us. Like a cow, *instinct* would fulfill our needs and keep us at peace. However obviously we are not cows. So what is the role of reason for us?

Kant, unlike Hume, believes that reason can (and is in fact designed to) influence the will. But reason is "not competent enough to guide the will safely as regards its objects...[since, as has just been shown] to this end would an implanted natural instinct have led much more certainly" (*Grounding* 9). Therefore, according to Kant, the role (that is to say *the very purpose*)

of reason (and therefore our purpose as rational creatures)² is "**to produce a will which is not merely good as a means to some further end, but is good in itself**" (*Grounding* 9, em³). If we accept this as reason's role—to "produce a will" which is "estimable in itself and good without regard to any further end" (*Grounding* 9)—we must still ask *how*? How does "reason" produce a "good will"? It is uncertain whether Kant effectively answers this question. Immediately after defining the role of reason, Kant shifts gears rather awkwardly and takes up a discussion of **duty**, which, *like* a good will, "holds first place in estimating the total [moral] worth of our actions" (*Grounding* 9). The connection ultimately seems to be that as rational creatures we can use our *reason* to derive *duty*. And when our actions are done *out of* duty (as determined by reason), they promote the development of a good will per se. To clarify: "An action done **from** [as opposed to merely in accordance with] duty must altogether exclude the influence of inclination and therewith every object of the will" (*Grounding* 13, em). In other words, when we act out of duty, we do so not to achieve some object of the will (happiness, for example) but rather to promote the goodness of the will itself. An action may happen to *conform* with duty, but if it is done out of an inclination external to the will itself (an inclination towards happiness, to repeat our example) it cannot have been done out of a pure concern for the goodness of the will in itself. Therefore the action is a misuse of reason and an abuse of our nature as rational creatures.⁴

² A potential problem with Kant: he seems to suggest that, because we are rational creatures, we have a *responsibility* to act in accordance with reason—that is to say, in accordance with *Kant's* notion of reason in particular. Rather, I believe it is more "reasonable" to conclude that (1) as rational creatures we certainly have an *ability* to act in accordance with reason, but that unless it can be proved that we have a moral duty to do so in all cases, other factors (such as emotional inclination) might (legitimately) influence our behavior, and (2) that there may be different types of "reasoning"—i.e. Kant's particular brand of reasoning does not have exclusive rights to the title.

³ The letters "em" indicates "emphasis mine" throughout.

⁴ Here is an example. A man is drowning; I decide to rescue him because I hope he'll pay me some money in gratitude. Now, as it turns out, it is actually *my duty* to save this man from drowning (this can be proven a priori/rationally according to Kant, as we shall see); however, since I did not save him *from* duty my action does not have moral worth, even though it happens to *accord with* duty.

All of this is to set the stage for Kant's moral program. As far as I can discern, he has tried to establish or assert the following points:

1. We are autonomous (free-willed) creatures with an ability to reason.
2. Reason is a "practical faculty," designed to influence the will .
3. If reason were designed to guide the will towards some ultimate end outside of itself, such as happiness, Nature made a big mistake: instinct would have been better suited to the job.
4. Therefore reason must have been designed to promote *the goodness of the will itself*.
5. Reason is able to determine duty, and acting in accordance with duty promotes the goodness of the will itself. Therefore by determining duty a priori (and not *for* some external goal) reason fulfills its very purpose, i.e. producing a good will.
6. An action is moral if it is done from duty and purely from duty. Aristotle wrote that an action is only truly virtuous if it is done *by a virtuous man*—therefore a "virtuous" action done for the wrong reason (i.e. saving a man from drowning so that he'll pay me for my efforts) is not truly virtuous. Kant seems to be borrowing this idea.

So Kant is not a consequentialist. An action is only moral if it is done for the sake of duty; duty cannot be determined by an "a posteriori incentive" (*Grounding* 13) lest it contradict the very purpose of a rational creature's reason (the production of a good will *per se*); duty therefore must be determined a priori.

"But what sort of law" he asks, "can that be the thought of which must determine the will without reference to any expected effect, so that the will can be called absolutely good without qualification" (*Grounding* 14)? He answers,

Since I have deprived the will of every impulse that might arise for it from obeying any particular law, there is nothing left to serve the will as principle except the **universal conformity of its actions to law as such, i.e., I should never act except in such a way that I can also will that my maxim should become a universal law [the categorical imperative]**....The ordinary reason of mankind in its practical judgments agrees

completely with this, and always has in view the aforementioned principle. (*Grounding* 14)⁵

I believe Kant makes an error in answering this way. So far he has established that rational creatures are not designed to will on account of external incentives but rather for the sake of the will's own goodness. This much is a strong premise. The question he poses is, what sort of principle will "determine the will without reference to any expected effect?" That is to say, **what principle will determine the will only with reference to/ for the sake of its own goodness?**

The problem arises when Kant rhetorically equates "particular laws" with "a posteriori incentives." He has only shown that *a posteriori incentives* do not promote the highest good of the will. But by calling these external incentives (such as the desire for happiness) "particular laws" he unfairly forces "universality of law" as the reverse and therefore proper solution. I am uncertain whether Kant is justified in saying that the *only way* for an action to be moral is if it could be willed to be universal. I believe he has merely shown that the only way for an action to be moral (assuming "to be moral" means "to act in accordance with the highest good of the will" as Kant suggests) is if it promotes the goodness of the will per se...however, what it *is* that promotes the goodness of the will per se might be any number of things.

But this objection aside, Kant's foundation is actually very strong. He links morality to our very nature as rational creatures, suggesting that *it follows from the fact that we are rational*

⁵ The last part of this statement is particularly problematic. Kant will ultimately argue that his categorical imperative shows conclusively that it is immoral to lie to the murderer (in our opening example). However, I believe the "ordinary reason of man" shows something quite different, and in fact would not agree with any principle which forbids a lie to save an innocent life. "Ordinary reasoning" goes something like this: I have an enormous ethical responsibility towards my innocent friend who is in danger; I harm no one by telling a lie in this instance, and I have a comparatively very limited responsibility to tell the truth to this criminal—by lying I save a life; therefore by telling a lie I fulfill the greater ethical responsibility.

*that we should be moral.*⁶ He thereby raises a strong argument against the idea that one should only behave morally for some benefit, such as perceived future happiness/comfort.

Before turning to Mill, let us briefly examine the specifics of the categorical imperative as it relates to the example with which we started the essay. To put it simply, Kant argues that immorality consists in making oneself an exception to a universal rule. For example, making a promise without intending to keep it is immoral because *it is impossible to will such behavior to be a universal law*. If *everyone* made promises without intending to keep them, "making a promise" would quite literally cease to exist because no one would ever believe you—it would become impossible to make a promise in the first place. By the same token, *lying* for some expediency is immoral because if *everyone* lied for some expediency, the "very source of right" would be "vitiating" according to Kant (*Right* 65).

But there can also be an impossibility in willing. I could never will the maxim "I don't need to save someone's life when I don't feel like it" because if this were universalized, I myself would be in danger of not being saved during some crisis.

Things get complicated, however, when one starts to refine his maxims. What if I were to put it this way: "I am allowed to lie in order to save my best friend's life from an evil criminal." Seemingly, there is no problem here... if everyone lied for *this particular expediency—i.e. the saving of an innocent life*—presumably the world would be better for it and fewer best friends

⁶ Morality defined as acting from duty; duty defined as the principle which determines the will to be good in and of itself; this "principle" defined as the "categorical imperative," i.e. **I should never act except in such a way that I can also will that my maxim should become a universal law**. There is still some question in my mind, however, about whether this "should" is truly categorical, or whether it is only hypothetical. If I *want* to promote the highest good of the will, then perhaps I should act according to the law of universality—i.e. I do not necessarily have a *duty* to promote the highest good as defined by Kant; I have a choice. However, Kant may be saying: *if you are a rational creature, morality (duty) is following the categorical imperative*, in which case there *is* a duty, only it becomes a question of whether I decide to follow it or not. Kant seems to be trying to show that we can derive a priori what principle morality must follow—he is *not* saying, of course, that everyone will choose to be moral, knowing this principle.

would die. There is no contradiction in willing since, if universalized, I myself might indeed benefit from a best friend's lying to save *my* life.

However, there may be a contradiction in possibility, since, if *everyone* lied to save a life, no one would believe you when *you* lied to save a life, and therefore it would be impossible to lie to the murderer in the first place.

Kant latches onto this contradiction in possibility (content), and refuses to let go—and here he fails as a moralist. He writes,

Truthfulness in statements that cannot be avoided is the formal duty of man to everyone, however great the disadvantage that may arise therefrom for him or for any other. And even though by telling an untruth I do no wrong to him who unjustly compels me to make a statement, yet by this falsification....I do wrong to duty in general in a most essential point. That is, as far as in me lies I bring it about that statements...in general find no credence, and hence also that all rights based on contracts become void and lose their force, and this is a wrong done to mankind in general. (*Right* 64)

Kant has said something false here. It is not actually true that when I lie to the murderer "statements in general find no credence"—nor is it true that "all rights based on contracts become void and lose their force." Kant is mixing reality with unreality, illegitimately—he tries to string his principle through both worlds at the same time. In the hypothetical world of universalizability of maxims, it is true that I cannot tell a lie without undermining the very possibility of lying in the first place. And if all of his previous reasoning is true, this may serve as an indication that lying is indeed immoral.⁷ That is a complete system unto itself. However, *in reality*, if I tell a lie to the murderer in order to save my friend—*nothing happens to actual statements in general, and nothing to actual rights based on contracts*. In reality, if I tell a lie, I damage truthfulness in

⁷ Because of the absurdity of the outcome, it might be reasonable to conclude that Kant has made one or more mistakes in his previous reasoning. The fact that his system breaks down in reality—a system which he claims "completely" agrees with "the ordinary reasoning" of men whereas in truth it does not—may be an indication that the initial steps are dubious. He writes, "A principle acknowledged as true must never be forsaken, however apparently danger is involved in it" (*Right* 67). **Another way of looking at this same axiom is to say, "If there is a great danger involved in something acknowledged as true, perhaps the question of its truth and how it was established initially should be reevaluated."** I think that is what must be done with Kant.

general only very little: I myself become less trustworthy by a miniscule degree and that is all. Or perhaps someone may hear of my example and decide to lie when his/her own best friend is being hunted down—in which case I will have damaged that person's trustworthiness a little as well. But it is patently false to imply that telling a lie causes *actual* contracts to become void—that would indeed be devastating! An important feature of reality is that real people are still able legitimately to trust other real people even though at times lies are told. It is not an all-or-nothing game of truth: in short, Kant's fears would only be justified if when a person lied (in reality), she *actually* caused the universal adoption of her behavior (in reality). This is impossible; his fears are unjustified.

Kant commits additional errors when he writes,

Wrongdoing would occur if I made the duty of truthfulness, which is wholly unconditional and which constitutes the supreme juridical condition in assertions, into a conditional duty subordinate to other considerations. And although by telling a certain lie I in fact do not wrong anyone, I nevertheless violate the principle of right in regard to all unavoidably necessary statements. **This is much worse that committing an injustice against some individual person...** (*Right* 67, em)

By what standard does Kant claim it is "much worse" to tell a certain lie, wronging the "principle of right," than it is to allow a murderer to kill my friend, thereby "committing an injustice against some individual person"? If we trust his philosophical arithmetic, he is only justified in saying,

1. Morality consists in following the categorical imperative;
2. Lying always fails the universalizability test due to a conflict in possibility (content);
3. Lying under any circumstance is immoral.

But wherefore the value judgment that telling a certain lie is "much worse" than allowing one's friend to be killed? Even if it were so that lying was *always* technically immoral, as points 1-3 (above) potentially attest, one could argue that it is better by far to be technically immoral than to let one's friend be murdered. (Of course, if Kant's standard for betterness/worseness **is** morality/immorality, then at least he is consistent.) However, the apparent "real-life" absurdity

of this consequence of his system leads one to believe that there might be a better way to determine morality than Kant's. And finally, Kant may be inexplicably in conflict with himself. Recall that his categorical imperative demonstrates that it is immoral not to save someone's life if one is able—since universalizing such a maxim puts one's own life in danger (and is therefore conflicted in willing.) However, in the murderer/friend example, the *only way* you can save her life is to tell a lie. So two categorical imperatives come into direct conflict—how does one give priority to one over the other? Why does Kant seem to hold truth-telling in such high esteem that he subjugates life-saving under it? This last question I cannot begin to answer.

Now let us turn to John Stuart Mill and Utilitarianism. We will cover him much more briefly. He writes in *Utilitarianism*, "all action is for the sake of some end, and rules of action, it seems natural to suppose, must take their whole character and color from the end to which they are subservient" (2)—this is diametrically the opposite of Kant's view, which we have just considered. For Mill, the end of action is Happiness, which he defines as pleasure, which he claims is inherently good. When Mill, like Kant, asks, "what principle makes an action moral?" his answer is consequentialist: "an action is moral if that action causes pleasure; it is immoral if it causes pain." Specifically, Utilitarianism requires one to act in such a way that he increases the overall pleasure for the most number of people.

All these things Mill seems to take as essentially self-evident—for how can there be any question that pleasure is good and pain is bad? If pleasure is good, then more pleasure is better, and people ought to act in such a way that they maximize the pleasure for all. With regard to our example (the murderer/friend), it seems obvious to me that more people would be happy if an innocent woman were saved on account of an isolated lie than if she were killed by an evil

murderer on account of a Kantian commitment to the truth. I will quote Mill at length now, since he deals almost majestically with this exact subject,

It would often be expedient, for the purpose of getting over some momentary embarrassment, or attaining some object immediately useful to ourselves or others, to tell a lie. But inasmuch as the cultivation in ourselves of a sensitive feeling on the subject of veracity is one of the most useful, and the enfeeblement of that feeling one of the most hurtful, things to which our conduct can be instrumental; and inasmuch as any, even unintentional deviation from truth does that much toward weakening the trustworthiness of human assertion, which is not only the principle support of present social well-being, but the insufficiency of which does more than any one thing that can be named to keep back civilization, virtue, everything on which human happiness on the largest scale depends—we feel that the violation for a present advantage of a rule of such transcendent expediency is not expedient, and that he who, for the sake of convenience to himself or to some other individual does what depends on him to deprive mankind of the good, and inflict upon them the evil, involved in the greater or less reliance which they can place in each other's words, acts the part of one of their worst enemies. Yet that even this rule, sacred as it is, admits of possible exceptions is acknowledged by all moralists; the chief of which is when the withholding of some fact (as of information from a malefactor)...would save an individual (especially an individual other than oneself) from great and unmerited evil and when the withholding can only be effected by denial. (*Utilitarianism* 23).

Mill gets it right. But I also feel that Utilitarianism is somewhat pointless. Mill's philosophy is so successful in the practical realm because its theoretical underpinnings are practically nonexistent: I claim that Mill has in effect taken information about the way human beings *already happen to act* and then told us, *you have a moral duty to act that way*. Before I conclude, I will try to show that Mill's attempts to show that we have a *duty* to act in a Utilitarian fashion are feeble at best and frivolous at worst.

So why *should* I increase my own happiness, let alone everyone else's? Mill's argument seems to me to go something like this:

1. You *should* work towards ends which are "desirable" because "desirability" includes in its very definition a "should"—this is self evident
2. How do I know if something is desirable and that therefore I *should* want it?

- a. "The only proof that a given object is visible is that people actually see it" (35)
 - b. "The only proof that a sound is audible is that people hear it" (35)
 - c. Therefore "the sole evidence it is possible to produce that anything is desirable is that people do actually desire it" (35)
3. People do actually desire happiness; all other things are evidently desirable only as a means to that end; therefore happiness is in itself desirable and is the greatest end
 4. "Each person's happiness is a good to that person, and the general happiness, therefore, a good to the aggregate of all persons" (35). Therefore you "should" work towards general happiness
 5. But what *is* happiness? "By happiness is intended pleasure and the absence of pain; by unhappiness, pain and the privation of pleasure...pleasure and freedom from pain are the only things desirable as ends" (7)
 - a. "What proof is it possible to give that pleasure is good?" (4)
 - b. It is self-evidently so

My objections are as follows:

1. You *should* work towards ends which are "desirable" because "desirability" includes in its very definition a "should"—this is self evident
Desirable means both "able to be desired" and "should be desired." Just because something is able to be desired does not mean that it "should" be desired. This point becomes relevant in a moment.
2. How do I know if something is desirable and that therefore I *should* want it?
 - a. "The only proof that a given object is visible is that people actually see it" (35)
i.e. "visible" means "able to be seen"
 - b. "The only proof that a sound is audible is that people hear it" (35)
i.e. "audible" means "able to be heard"
 - c. Therefore "the sole evidence it is possible to produce that anything is desirable is that people do actually desire it" (35)
i.e. "desirable" means "able to be desired" —except that Mill trades on its double meaning to establish a "should" here; he suggests that you *should*

desire happiness, because it *is* desired. This is making an "is" out of an ought."⁸ He might have said, "People seem to desire happiness" —but this much we already knew. To suggest that we have a "duty" to increase happiness is unfounded.

3. People do actually desire happiness; all other things are evidently desirable only as a means to that end; therefore happiness is in itself desirable and is the greatest end
Aristotle made an identical "argument." As an empirical observation, it does seem that in perhaps most cases, people desire to be happy. People also seem to try to organize their lives and act in such a way that they will be happy. However, there may certainly be specific instances where happiness is subjugated to other desires. My desire to live outweighs my desire to be happy when I unhappily choose to eat bugs rather than starve to death in a desert. In that instance, did I choose to live so that I could be happy...?
4. "Each person's happiness is a good to that person, and the general happiness, therefore, a good to the aggregate of all persons" (35). Therefore you "should" work towards general happiness
It may be true that I should increase my own happiness: happiness is good/desirable, therefore I "should" try to increase my happiness. But the notion that I "should" try to increase everyone else's happiness, even to the detriment of my own, is relatively unsupported by Mill.
5. But what *is* happiness? "By happiness is intended pleasure and the absence of pain; by unhappiness, pain and the privation of pleasure...pleasure and freedom from pain are the only things desirable as ends" (7)
 - a. "What proof is it possible to give that pleasure is good?" (4)
 - b. It is self-evidently so

With regard to this last point: I believe it is erroneous to say that pleasure is good and pain is bad. This is the crux of Mill's failure, because he rests so heavily on this assumption, saying, in

⁸ David Hume, in his *Treatise of Human Nature*, warned against attempts to derive an "ought" from an "is." That is to say, it is one thing to observe the way things *are*—the way people really feel (or tend to feel) about moral situations—and quite another to talk about the way things *ought* to be—to prescribe some duty independent of mere moral sentiment.

essence, "Well, isn't it obvious?"—and builds his whole theory therefrom. No: it is not obvious. I would say merely that pain hurts and pleasure pleases. I would say that people, in reality, tend to avoid pain. That is all. So a "moral philosophy" which claims I have a "moral duty" to avoid pain and to help others avoid pain is in essence saying to me, "you have a duty to do what you tend to do naturally." Also, this simplistic view of morality runs into problems with itself. I know from experience that if I hurt someone deeply (emotionally) it may actually strengthen their character and lead to greater overall happiness in the future. Do I therefore have a utilitarian moral duty to go around hurting people emotionally? How atomistic does one have to get with regard to "pain" in order to decide whether it really is bad? When I put my hand on a hot stove, presumably the pain that I experience is *good* because it causes me to remove my hand from the stove and thereby avoid getting a nasty burn...Mill's "isn't it obvious?" approach begins to fall apart...

So although Mill's system (as opposed to Kant's) in effect provides the "right answer" to the problem we are considering—by permitting a lie in this special circumstance—he is guilty of making a moral "ought" out of an empirical "is" and therefore does not properly answer to David Hume, who wrote that morality was something "more properly felt than judg'd of." Kant's moral theory has a much stronger theoretical base, but even so, it is unclear whether he proves that I have a *duty* to promote the highest good of the will per se. And in "the real world" his methods totally fall apart.

To conclude, I do not arbitrarily presuppose that it is moral to lie to the murderer, and then congratulate Mill because he agrees with me. Rather I hope precisely to have shown that, unless it can be proved that there is a definite moral duty—an "ought"—attached to our decision, we are unjustified in deriving the morality of that decision from anything other than our "moral

sense"⁹—an "is." Since I hold that neither Kant nor Mill's "ought" is perfectly legitimate, I also claim that it is moral to lie to the murderer simply because *it feels like the moral thing to do*.¹⁰ Mill's system happens to conform with this moral "matter of fact"—to use another of Hume's terms—and therefore I say that Utilitarianism "gives the right answer."

Works Cited

Hume, David. *A Treatise of Human Nature: Being an Attempt to Introduce the Experimental Method of Reasoning into Moral Subjects*. Ed. Ernest C. Mossner. New York: Penguin Classics Reprint, 1986.

Mill, John Stuart. *Utilitarianism and the 1868 Speech on Capital Punishment*. Ed. George Sher. Indianapolis: Hackett, 2001.

Kant, Immanuel. *Grounding for the Metaphysics of Morals and On a supposed Right to Lie because of Philanthropic Concerns*. Trans. James W. Ellington. Indianapolis: Hackett, 1993.

⁹ Hume's term.

¹⁰ That is to say, it feels like the moral thing to do *to me*. I sense that in fact *most* people would choose to lie to the murderer, believing they have a "higher" ethical responsibility to their friend. However, it must be pointed out that Hume's moral system seems to suggest that all moral opinions are equally valid. Therefore, a strict Kantian who chooses to tell the truth to the murderer is on the same (Humean) moral footing as someone like me, who would tell a lie.