

# On Going Above and Beyond the Call of Duty

---

Michael Licciardi  
*Rutgers University*

## **Abstract:**

In this paper, I endeavor to discuss and analyze the concept of morally supererogatory acts—of going above and beyond the call of moral duty. There is reason to suggest that the topic needs more clarification, and that a presumed understanding of it is not wholly sufficient. I will present these reasons before pursuing an investigation as to what further reflection yields when we examine the concept more closely. Ultimately, I intend to show that the concept of moral supererogation is not so coherent a concept as we might initially have thought.

In this paper, I endeavor to discuss and analyze the concept of morally supererogatory acts—of going above and beyond the call of moral duty. There is reason to suggest that the topic needs more clarification, and that a presumed understanding of it is not wholly sufficient. I will present these reasons before pursuing an investigation as to what further reflection yields when we examine the concept more closely. Ultimately, I intend to show that the concept of moral supererogation is not so coherent a concept as we might initially have thought.

### Stilts

Ethics, as a sub-discipline of philosophy, similar to most sub-disciplines thereof, must, in a certain sense, stand upon stilts. That is to say, if one wishes to see far, one's natural gifts alone will not be enough, and, invariably, the shoulders of a number of artificial giants will be necessary—which is really just to say that one cannot gain ethical insight, propose substantive ethical theories, or even assign oneself to the task in the first place without taking a number of things for granted, assumptions on which one's theories can rest, stand, or run for that matter.

As such, if one wishes to provide the philosophical community, or at least one's ethics professor, with a theory involving ethics, it would seem most prudent and honest of that individual to provide a brief account of exactly what assumptions he/she is making. This seems only fair to all parties involved. Believing, or at least wishing, myself to be the prudent and honest type, this will be my first task. Before I approach the problem of supererogation directly, I will lay out those prejudices in my ethical theory, which, for lack of space and time, I cannot defend in this essay.

First and foremost, I suppose, I assume that there is some thing<sup>1\*</sup> called “the good,” some thing called the “bad,” possibly even the “evil,” some actions which are “right,” and some which are “wrong”. As such, I assume that there is a moral fact about every morally relevant act or situation, rendering it either good or bad, right or wrong—excluding those cases which we might call morally neutral. The meaning I assume for the word “good” is somewhat abstract, but possibly helpful—by good I mean, that which one should do (or possibly value).<sup>2\*\*</sup> I assume, then, that bad is

---

1       \* I use the word thing here because I do not wish to commit myself to exactly what the good might be—whether Form, property, or otherwise.

2       \*\* While this definition is certainly only provisional, and by no means perfect, it seems, in its current formulation, to bridge the epic gulf between meta and normative ethics—generally phrased in the assertion, “no theory or definition of what the word ‘good’ means will ever give us a directive to, or show us why we should, be good.” If we assume that the good is that which we should do, then it follows analytically that we should be good.

that which we should not do, that right is any action which is, or leads to, good, and that wrong is any action which is, or causes things to become, bad.

Next, I assume that moral situations are of such vast uniqueness and complexity that, on the whole, general normative theories will not be able to encompass them. Thus, normative theories will be able to provide no more than guidance to individual moral agents, rather than definitive accounts of the absolute moral fact regarding a given situation. In this way, I assume a special brand of moral particularism, which, as noted, cannot be defended here.

Given that this sort of particularism requires certain epistemic conditions/qualifications, those conditions/qualifications will comprise my next assumption. I assume a sort of ethical intuitionism, whereby individuals, upon reflection, are presented with the moral fact of their individual situation. These moral facts are disclosed privately, and are, ultimately, unimpeachable by any moral theory or the judgment of other individuals<sup>3</sup>.\*

Lastly, I must assume a certain definition of “supererogation,” or at least an approximation of its meaning. I assume “supererogation” to mean: doing more than one is obligated to do, or going above and beyond the call of duty. It is my hope that this definition/approximation will conform to most everyone’s judgment about how the word is typically used, in order that it not have been constructed solely to serve my own theory—as such a construction would make my theory vacuous and unimportant. I hope also that the simplicity of my definition will not count against it, but rather, will leave room open to fit our general understanding of the term.

With these assumptions out in the open, I will now proceed to outline my theory of moral supererogation, detailing, first, why I believe the concept to need further

---

<sup>3</sup> \* (A somewhat lengthy point)—It might be objected that this allows one to do what one wants, provided one can tell a convincing story about one’s ethical intuitions. This is not necessarily true, though. For one, one’s ability to cheat a system does not make the ideal of that system inherently false. Second, ultimately, we may only be punished, deemed worthy of censure, or officially blamed for some act via the imperfect laws of the legal system. In that we must all live together, we must all be governed by the same laws, even if these cannot, or could not, mirror moral fact. In so far as morality is concerned, we stand alone as individuals, and, for better or for worse, ultimately, we cannot fully know that our moral judgments of others are correct, in virtue of the fact that we are not them (although it is not uncommon that individuals are presented with similar, if not the same, moral intuitions, making moral judgment, if not infallible, not always unreasonable). In summation, law is rigid, and, while it can be cheated as well, it provides more objective standards for judgment. As far as ethics is concerned, one cannot fully cheat oneself. Thus, this theory does not prevent one from being judged by the law; it only prevents one from being conclusively judged by others. The inability to be judged conclusively by others, however, does not make one’s actions right when one wants them to be.

clarification, moving on to a theory of moral worth and character, before explaining why the concept's presumed coherence does not hold up to scrutiny.

### **Pre-Reflection on Supererogation**

It is often assumed that morally supererogatory actions are, in some sense, morally superior to morally obligatory actions. This assumption seems to come from the generally accepted view that morally neutral supererogatory acts are, in some sense, superior to obligatory actions. For instance, James, Greg's boss, will think more highly of Greg, if Greg stays late once a week—not out of any contractual obligation, but out of a desire to increase productivity at the office—than he will of Steve, who only works the hours that he is required to. This is not a terribly uncommon phenomenon: students are admired by their teachers for reading ahead; our friends greatly appreciate it when we pick them up from the train-station at short notice; and restaurant-goers are generally very grateful—their gratitude often expressed in the form of a higher tip—for exceptional service provided by waiters and waitresses.

The pertinent question is whether this principle extends to the moral realm. In other words, would it be correct to assert that actions which go above and beyond the call of moral duty are morally superior to actions which are simply in accord with moral duty? By way of a sort of analogy, we might think that the answer to this question is simple: of course morally supererogatory acts are superior to morally obligatory acts.

Consider an example, however, that might complicate this idea. Suppose that Mary sees her daughter drowning out at sea, and, quite naturally, swims out and saves her. Suppose, next, that Carol sees her dog drowning out at sea, and, in a moment of courage, swims out and saves it. Each, Mary and Carol, has taken a great risk, but, given that Mary has a greater moral obligation to her daughter than Carol has to her dog, we would say that, while Mary has fulfilled her moral obligation, Carol has gone above and beyond her duty. If morally supererogatory acts are somehow morally superior to morally obligatory acts, then we seem to be pushed into claiming that Carol is more praiseworthy for saving her dog than Mary is for saving her daughter—this result, however, seems unacceptable.

This leads us to the more reserved assumption that morally supererogatory acts are, in some sense, more praiseworthy than morally obligatory ones, but only in certain ways—ways, which we are given the task of describing. Determining exactly what these ways are, however, does not seem to be a promising task. Perhaps this is

due to what would seem to be a problem in our method. It may be argued that we are attempting to discuss certain features of morally supererogatory acts—namely, in what ways they do or do not warrant greater moral praise than morally obligatory acts—before having a real understanding of what a morally supererogatory act is—we are, as Socrates would say, forgetting the priority of *ousia* over *pathos*. Thus, it seems, if we are to make any claims regarding the relation between morally supererogatory acts and morally obligatory acts, we need to have a clearer understanding of the former.

These considerations seem to strongly suggest that our presumed understanding of supererogatory acts is unsatisfactory. It might be wondered why, if we really are so in the dark on the matter, any presumed understanding of the concept has arisen in the first place. In other words, if we don't understand, why do we think we do? While I do not pretend to have a conclusive answer to this question, I would hypothesize that our understanding of merely supererogatory acts—supererogatory acts which are not necessarily moral acts—leads us to believing that no further thought is needed to understand morally supererogatory acts.

Let us ask, then, what we would want from an account of what a morally supererogatory act is. It seems that such an account would need to help us, somehow, to distinguish between morally supererogatory acts and neutral supererogatory acts, as well as between morally supererogatory acts and morally obligatory acts.

There are two ways which we might approach the first of these requirements: we might give an account of neutral supererogatory acts, and then seek to show how we would want morally supererogatory acts to differ from them; or we might begin by giving an account of some necessary condition of an action having moral character, and then proceed to show how neutral supererogatory action lacked this condition, and, subsequently, to show what ways in which morally supererogatory acts would need to satisfy this condition, to be worthy of the title. Each method has its merits, but, for the sake of perspicuity, or so I would like to believe, I am going to choose the second method.

### **Moral Credit and The Kantian Moral Characteristic**

I would like to defend a semi-Kantian view about what makes an action moral. It is important, however, for me to emphasize that I am not making any claims about what makes an action morally good or bad, right or wrong—merely what makes an action moral. Along Kantian lines, then, I would like to assert that, in order for an action to be a moral one, it must be grounded in a desire to do the right thing, because it is

what is right—in other words, to do one’s duty for the sake of doing one’s duty. This account does not presuppose a more detailed ethical theory—one may believe, on this view, that one’s duty is to maximize the good, just as well as one may believe that one’s duty is to follow a set of deontological rules.

My motivation for accepting this view derives from a theory of praiseworthiness, proposed by Holly Smith, which I will be provisionally accepting. Smith’s claim is that an agent’s being worthy of praise for their performing of a certain action is a function, first, of that action being a good action, and second, of that action arising from a commendable configuration of desires and aversions in the agent. Smith stresses that the idea of configuration is an essential element of her account. This is because it is not necessarily the possession of certain desires or aversions which makes our actions good or bad, but rather, their relative strengths. For example, we would not want to assert that the desire to promote one’s own self-interest is inherently an immoral desire, but rather, that it’s being a stronger desire than other, more morally commendable desires is what makes the configuration of desires to which it belongs an immoral one, or at least one worthy of moral criticism.

Smith divides our desires/aversions, with respect to some act, into three categories: cooperative, countervailing, and morally good<sup>4</sup>\*. To illustrate how these operate we can take an example of a moral action, and show how these desires relate to it. Suppose, then, that John decides, based on his relevant desires/aversions, to donate money to charity. His cooperative desires are those which might give him reason to donate to charity, but which are not necessarily moral desires—for instance, the desire to get a tax refund, or the desire to be regarded as a philanthropist. His countervailing desires are those which might give him reason not to donate to charity—for instance, the desire to use the money in question for some other end, say, to buy his wife a nice anniversary present. It is important to understand that countervailing desires need not be immoral ones—in fact, their being morally neutral is somewhat crucial to Smith’s theory, as it stands. Finally, morally good desires are those which are in accordance with generally accepted moral principles—for instance, the desire to give what one can to those in need. It should also be noted that morally good desires have the ability to be directed either for or against the commission of some action. For example, John might have both the moral desire to give what he

4

---

\* *Morally Wrong* can be added to this list, but, for our purposes, it need not be.

can to those in need, as well as the moral desire to make some person—namely, his wife—happy. Each of these conflicting moral desires comes to bear in his reasoning.

According to Smith, then, the relative strengths of these three types of desires will determine one's choice in performing some morally good act, and it is this configuration that enables us to bestow praise upon an agent when he performs a good, or right act (or, conversely, blame, if he performs a bad, or wrong act). More specifically, she claims that, in order for an agent to be praiseworthy for his action there must be a relation of moral fitness between his configuration of desires/aversions and his action. She lists two conditions for moral fitness: (1) the agent must have at least one moral desire that comes to bear in some minimal way on the agent's decision; and (2) this desire must be strong enough—although, not necessary—given the strength of the agent's other desires, to result in his making the morally right decision.

While this account is both insightful and well developed, there is one element which, I believe, it lacks. It seems to me that, as it stands, this account leaves a number of cases underdetermined. For instance, imagine that, given John's configuration of desires/aversions, his final desire—assuming, for the sake of convenience, that we can quantify such matters—has a strength of negative-10—in other words, his overall desire not to donate to charity is ten points stronger than his overall desire to donate. As we have been, thus far, framing it, our account, I believe, leaves an answer to the question of what John will do underdetermined. It does not seem unreasonable to say that, despite, given his other desires, not really wanting to donate to charity, John will, nevertheless, donate anyway. It seems accurate, here, to say that, the final factor, which plays a role in determining John's ultimate decision—allowing us to truly credit a moral action to him—is what I would call the moral desire: the desire to do the right thing. In other words, one can, given his general set of desires and dispositions, want to, or not want to perform some moral act M, but his final decision will—if his decision is a moral one—depend upon the strength of his desire to do the right thing, either in accordance with, or in the face of his other relevant desires.

Now, I mentioned that this must be so, in order that John's act M be a moral act, and such a claim seems to require further justification—let me turn to that now. I would like to suggest that Smith's category of morally good desires is an artificial one. This is not to say that the desires contained therein do not have morally good acts as their object, but rather to suggest that these desires are accidental. I mean, by accidental, not that the acts themselves, or the desires directed thereat, are only accidentally, contingently, or conventionally morally good—my argument rests neither on conventionalism nor realism—but that John's possession of these desires is

accidental. In other words, John has a whole set of desires which he cannot help—we can call these first-order desires. His biology, schooling, rearing, and social training have left him the way he is—John could not choose his desires before he became John, and becoming John involved, at least in part, a process of desire, belief, and disposition implantation. Accordingly, by the time John has any say in the matter, John already is who John is—he cannot help this. In that John cannot help what types of first-order desires he has acquired in, and as a result of, his early life-stages, they are accidental—and thus he cannot be morally credited with their moral goodness. In other words, in so far as John has, what I'd like, for now, to call a blind desire to help others in need, he is merely morally lucky that such a desire happens to have morally good actions as its objects. I would like to suggest that it is in his grounding desires—namely, his desire to do the right thing, to do what is morally required of him—that John's actions take on a moral character. In other words, the desire to do good things, which provides the ground for his other desires, the reason upon which he acts on or does not act on those others, is what gives those others their moral character—what makes them, what I would call, fully intentional, and deserving of moral praise.

In this way, I have tried to show why one should accept a semi-Kantian model for what makes a moral action a moral one—and subsequently what makes an action worthy of moral praise. Moral action, worthy of moral praise, must be rooted in a desire to do what is right, because it is right—to do one's duty for the sake of doing one's duty.

### **The Incoherence of Moral Supererogation**

I will now indicate how neutral supererogatory acts do not have the characteristic just described—that of being grounded in the desire to do one's duty for the sake of doing one's duty—and how they are, thus, essentially extra-moral. We need only one brief example to show this. Consider, again, Greg, the employee, who puts in extra hours in the office. We would say that, in so doing, Greg has gone above and beyond the call of duty. But such an act, or set of actions, has no moral character. Greg works more than is required of him to increase productivity at his office—this is true—but his desire to increase productivity does not arise out of any desire to do his duty, for the sake of doing his duty. His desire to increase productivity might arise from a desire to keep his office open—and thus to keep his job—or it might arise from a desire to

have his boss like him more, thus creating a more pleasant working environment, or providing him with a positive professional reference in the future.

Of course, in providing these possible motivations for his desire to increase productivity at his office, I seem to be leaving open the possibility that Greg might have some other motivation that would constitute a moral choice—and thus, perhaps, leaving open the possibility that this supposedly neutral supererogatory act could possibly have been a morally supererogatory act. In this way, the character of a supererogatory act would not play a determining role in whether it was moral or not; rather, its motivation would. This indeed would be true, if it were possible to have such motivations with respect to supererogatory acts. I am now going to argue, however, that this is not possible.

I will give two main reasons that, in the light of what we have been discussing, will show why I believe the concept of a morally supererogatory act is an incoherent one. The first falls out of my acceptance of what, in the first place, makes an act a moral act—this one we can deal with more quickly before moving on to a more substantive line of reasoning. My first argument is this: given that any moral action must be motivated by a desire to do one's duty for the sake of doing one's duty, it is not coherent, then, to claim that doing more than one's duty is motivated by a desire to do one's duty.

Issues of linguistic comfort aside, an objection of a slightly more substantive nature could be made to the proposed claim that would run as follows: practically speaking, one could certainly do more than necessary with the intended outcome of doing what is necessary. Consider the college student, Tom. Tom knows that for his philosophy class he needs to get a 90 on his last exam, in order to get an A in the class. One might object that my theory implies that it is incoherent to speak of Tom trying to get a 100 on his exam, with the intended outcome of getting an A in his class, even though a 90 is all he needs. This might be called the covering-one's-bases objection. The main idea is that one can attempt to do more than one believes is necessary in order to make sure that what is necessary does get done, in virtue of its being a subset of the set of actions which one does take.

I cannot deny that this is a common phenomenon, but I contend that there is a crucial difference between covering-one's-bases in an extra-moral context—specifically, this type of context—and covering-one's-bases in a moral context. There seems to be an epistemic condition present in this example that allows Tom's actions to make sense. Tom does not know which questions he will answer correctly on the

aforementioned exam. Moreover, Tom knows that he does not know which questions he will answer correctly. In this way, Tom's attempt to get a 100 on his exam is really an attempt to answer every question correctly, and this attempt is a function of his knowledge of his own fallibility. In other words, in an ideal situation, where Tom knows when he answers a question correctly, and where his only relevant goal is to earn an A in his philosophy class<sup>5\*</sup>, it becomes far more difficult to make sense of his actions.

We can call the characteristic that Tom manifests, in his attempts to get an A in his philosophy class, the A-student characteristic—this is in opposition to, let's say, the impressing-the-teacher characteristic, or the learning-the-most-material characteristic. If Tom is to display only the A-student characteristic, then Tom must be motivated solely by his desire to earn an A in his class—he cannot be motivated by the desire to learn all the information he can, as well as he can, nor by the desire to impress his teacher, or others. Casting the situation in this light allows us to see its relation to moral action. A moral action must be motivated by a desire to do one's duty for the sake of doing one's duty, and not by a desire to appear philanthropic to others, or to receive a tax refund at the end of the fiscal year. Thus, Tom, in exhibiting only the A-student characteristic is like John exhibiting the moral characteristic.

There seems to be an impediment to reasoning analogically here. While it seems true that if Tom were motivated solely by the desire to earn an A in his philosophy course, and if getting a 90 on his exam would earn him an A, and if Tom could know when he answers a question correctly, then we could not make sense of his attempt to get a 100 on his exam, if John is motivated by the desire to do his duty for the sake of doing his duty—i.e. if he is to be motivated by moral desire—and if John believes that his duty, given his situation, is to donate 90 dollars to charity, then we might still be able to make sense of John donating 100 dollars instead of 90.

Being able to make sense of this situation might lead us to accept that morally supererogatory acts themselves make sense, but to reason in such a way would be to move too quickly; for there is still a crucial difference between these two accounts, which, once understood, will allow us to see more clearly why, if John were to donate 90 dollars—instead of 100—to charity, his act would not be a morally supererogatory one. The difference is that, in the case of Tom, we know, as well as Tom, who does himself know, what is required of him, if he wishes to earn an A in his course. John, on the other hand, presents us with a different problem, because not only can we not

---

\* We can, thus, leave aside possible motives of learning the material as best he can, and so forth.

be sure of John's moral duty, John's own understanding of his duty might change or be reevaluated. Hence, if John decides to donate 100 dollars—rather than 90—to a charity, and if such a decision is a moral one, then his decision to do so must arise from his realization that, while he might initially have thought that donating 90 dollars to charity was his moral duty, donating 100 dollars to charity is his actual moral duty.

This leads to our second, more substantive argument—the disjunctive argument against moral supererogatory action. The reasoning is as follows: suppose John decides to donate 100 dollars—rather than 90—to charity. His decision to do so can be motivated by one of two types of desire: moral or extra-moral. We do not want to say, if we are defending the concept of a morally supererogatory act, that his desire is extra-moral. Thus, if we have any hope of attributing to this supererogatory act a moral character, needless to say, his motivation must itself be moral—but first let's take a step back. We can, to simplify, say that John donates 100 rather than 90 dollars to charity because: a.) Doing so would benefit him in some way; or b.) He believes donating 100 dollars is his duty. If John's motivation is of the first kind, then it is not moral—if it is of the second kind, it is not supererogatory, but merely obligatory. My argument, then, is that moral actions must be obligatory to be moral—and hence, an action can never be both moral and supererogatory, given the definition stipulated and defended of moral action, and its relation to human motivation.

I understand that my account entails that one can realize and come to further understand his moral obligations over time, and that such a process, along with the final answer cannot always be fully known or understood by others. These shifts in realization can shift according to circumstances both internal and external to the agent in question, and only he can be the final arbiter of the outcome of his ultimate recognition of duty. In this way, we cannot say what John's actual moral duty is—whether it be to donate 90 dollars to charity, 100, 700, or none at all. As stated above, I realize that I commit myself to a Moral Particularist and Intuitionist standpoint—the defense of such a standpoint cannot adequately or appropriately be addressed here. I will briefly reiterate, now, how these commitment/assumptions, along with my proposal of a sort of Kantianism regarding the moral character of action/motivation, apply to my argument: 1.) An agent has the capacity to recognize his moral duty; 2.) Such a recognition can reorient itself as circumstances, both internal and external to the agent, change; 3.) Once an agent has recognized his moral duty, he is obligated to perform such a duty, for the sake of performing his duty, if his action is to be a moral one; and lastly 4.) Other agents may not always, or ever, be able to fully learn or understand the exact moral obligation that a given agent recognizes as his own. Therefore, what most people take to be morally supererogatory acts are

actually morally obligatory acts—given an agent’s own understanding of his moral obligations—or are simply neutrally supererogatory acts, disguised as moral ones.

## Conclusions

Having given reason to believe that the concept of the morally supererogatory act is an incoherent one, a couple of questions still present themselves. First, what can we say, now, about acts which we initially took to be morally supererogatory—such as Carol saving her dog from drowning—and their relation to morally obligatory ones—such as Mary saving her daughter from drowning? Most immediately, we can state that Carol’s action is not morally supererogatory; it is morally obligatory for her. Thus, Carol’s action holds the same moral-metaphysical status, so to speak, as Mary’s—each is a moral action, and thus a morally obligatory action. The complex and difficult relation that we initially perceived has fallen away.

We see then that acts of the sort that Carol performs cannot be rightly called superior to acts of the sort Mary performs, on the basis of some distinction between obligation and supererogation, because such a distinction is empty—as morally supererogatory acts do not exist—but on what grounds can we justify our remaining intuition that Mary is more praiseworthy for her action than Carol is for hers? The only standard, it seems, by which we could evaluate the praiseworthiness of agents in situations like these, might be to compare the values of the various obligations which they are fulfilling.

The question of whether such evaluation is possible is the second question which our account leaves open. For better or for worse, the answer seems to be no. If our moral obligations become apparent to us on an individual and situational basis, and if our obligations cannot always be fully known to others, then we cannot compare moral obligations across people, for a lack of an adequate epistemological perspective.

The only criterion, on this model, that we might use is force and relevance of obligation. For instance, if Peter knows that he has an obligation to cook dinner for his mother one night, and an obligation to volunteer in a soup kitchen that night, and supposing that his obligation, by his own understanding, to volunteer is more of a pressing and morally necessary nature, then we can assign his level of praiseworthiness based on which obligation he chooses to fulfill—in relation to whether he chooses the more pressing obligation.

This form of evaluation does not seem possible between people, epistemological factors aside, given both that the strength of obligations will differ from person to person, and the somewhat brute fact that an obligation binds a given individual,

regardless of its value in relation to the obligations of others. We are morally required to fulfill our own moral obligations, not the moral obligations of others. We can neither be praised nor faulted for not fulfilling the obligations of others, if those obligations are not our own. Mary's strongest obligation is to save her daughter and Carol's is to save her dog—each is acting in the most morally praiseworthy way, given their situation—as such, we are back to square one, and adjudication of praiseworthiness seems impossible.

Without a theory of what types of acts precisely are objectively morally best<sup>6</sup>—which would amount—answering the question of whose morally obligatory act is more praiseworthy seems profoundly difficult, and is not a task we have the ability to undertake here<sup>7</sup>. This issue aside, we have at least eliminated the problem of comparing morally obligatory acts with morally supererogatory acts, by showing that, given a number of reasonable assumptions, moral supererogation is incoherent.

### Works Mentioned

Smith, Holly, "The Varieties of Moral Worth and Moral Credit," *Ethics* v.101 No. 2 (January, 1991), pp. 279-303

---

6 We might use such a knowledge of which obligations are best to fulfill to claim that the strength and desert of moral praiseworthiness is a combination of our own fulfilling of our moral obligations, and the luck of the draw as to which obligations we happen to be faced with in a given situation—this type of analysis, however, does not seem promising. It does not seem like we would want to claim that two people who each fulfill different obligations, which each involve actions of different relative moral value—say, volunteering at a soup kitchen versus jumping in front of a bullet meant for someone else—can have different levels of praiseworthiness based on these relative differences in these values, when the two individuals do not, and could not, choose their obligations.

7 Adjudication might be feasible by a counterfactual analysis of the strength of one's moral desire versus one's other desires, in comparison to someone else's—i.e. theoretically determining how much more countervailing desire one's moral desire could withstand in comparison to another's. This solution, however, has its problems as well. Unfortunately we cannot discuss them here. For a further discussion of moral praise and blame see Smith *The Varieties of Moral Praise and Moral Worth* (1999).

