Wilson, A. T. (2017). A Dilemma for Neo-Aristotelian Supererogation. *Ethics*, 128(1), 199-211. <https://doi.org/10.1086/692948>

Peer reviewed version

Link to published version (if available):
[10.1086/692948](https://doi.org/10.1086/692948)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the author accepted manuscript (AAM). The final published version (version of record) is available online via the University of Chicago at <http://www.journals.uchicago.edu/doi/abs/10.1086/692948> . Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

This is the penultimate draft of work published in *Ethics*, Vol. 128, No. 1, pp. 199-211 (Oct. 2017). Please do not cite this earlier draft without permission.

The full and final version is available at: <http://www.journals.uchicago.edu/doi/full/10.1086/692948>

A DILEMMA FOR NEO-ARISTOTELIAN SUPEREROGATION*

Alan T. Wilson

ABSTRACT: It has recently been argued that virtue ethics cannot accommodate the possibility of supererogation. In response, Rebecca Stangl proposes a neo-Aristotelian account of supererogation that, she argues, generates plausible verdicts, while also being compatible with the doctrine of the mean. I argue that Stangl's response is unsuccessful. First, I demonstrate that the proposal in its current form is problematically indeterminate, meaning that we cannot know what verdicts would be produced in response to classic examples. Second, I argue that anyone attempting to develop the account faces a dilemma, and that both options for responding to this dilemma generate problematic results.

INTRODUCTION

The concept of supererogation presents a challenge for ethical theorists. Following J. O. Urmson's influential "Saints and Heroes", it has become necessary for theorists to demonstrate that their chosen theory provides a plausible account of supererogation, or else to explain why such an account is unnecessary. Rebecca Stangl (2016) has recently addressed this challenge from the perspective of virtue ethics. Stangl proposes a neo-Aristotelian account of supererogation that, she argues, generates plausible verdicts, while also being compatible with the Aristotelian doctrine of the mean. In this way, Stangl rejects the claim that virtue ethics fails to accommodate supererogation.

In this paper, I argue that Stangl's attempt to reconcile virtue ethics and supererogation is unsuccessful. First, I explain why the level of detail provided by Stangl is not sufficient to demonstrate that the account would generate intuitively plausible results. I then argue that

Stangl faces a dilemma in deciding how best to add the necessary detail. In particular, I argue that Stangl's account will generate unintuitive verdicts regardless of the position taken on the importance of the *internal targets* of virtue. Given these issues, virtue ethicists may be forced to choose between accepting a highly revisionary account of supererogation, or else refusing to accept the possibility of supererogation altogether.

I. SUPEREROGATION

Urmson's "Saints and Heroes" argues for the existence of a category of action that cannot be reduced to the conventional categories of obligatory, permissible, and forbidden acts.¹ Allowing for only these three categories fails to acknowledge the possibility of actions that *go beyond* what is morally required. Such actions would be neither obligatory nor forbidden, but they would also not be *merely* permissible. Actions of this sort have come to be referred to as *supererogatory*, and Urmson argues for their existence by way of examples. Two of these examples can be set out as follows:

Soldier: We may imagine a squad of soldiers practising the throwing of live grenades. A grenade slips from the hand of one of them and rolls on the ground near the squad. One soldier sacrifices their life by throwing themselves on the grenade, thereby shielding their comrades.

Doctor: We may imagine a doctor, no differently situated from countless other doctors, who becomes aware of the great suffering of people in some far away, plague-ridden city. Upon learning of this fact, the doctor volunteers to travel to the city and join the depleted medical forces that are based there.²

Urmson's claim is that the agents in these examples are acting in ways that are neither forbidden nor obligatory. It would be overly demanding to claim that the other soldiers or other doctors had committed a moral wrong by failing to throw themselves on the live grenade, or by failing to move to the plague-ridden city. The actions involved in *Soldier* and *Doctor* also appear to go beyond

what would be *merely* permissible. We, therefore, have intuitively appealing examples of supererogatory actions, of actions that are beyond the call of duty. This presents a challenge for any ethical theory that does not accommodate the existence of such actions.

I will not here repeat the arguments that have been used to show that standard virtue ethical accounts cannot accommodate supererogation.³ Instead, I will focus specifically on Stangl's recent attempt to respond to this issue.⁴

II. STANGL'S NEO-ARISTOTELIAN ACCOUNT

Stangl's account of supererogation builds upon the work of Christine Swanton, and is set out as follows:

An action is supererogatory if and only if it is overall virtuous and either (a) the omission of an overall virtuous action in that situation would not be overall vicious or (b) there is some overall virtuous action that is less virtuous than it and whose performance in its place would not be overall vicious.⁵

To evaluate this account, it is necessary to first clarify the key concepts of "overall virtuous" and "overall vicious" action. For this reason, it is important to briefly explain those elements of Swanton's work that are relied upon by Stangl.

Swanton argues that we can helpfully think of virtues as having various items within their characteristic fields of concern.⁶ Different virtues will be concerned with a different range of items. For example, "bodily pleasures" will be included as items within the field of temperance, whereas "dangerous situations" will be of more concern to courage.⁷ Importantly, each virtue can be expected to have a plurality of items within its field, and possessing a given virtue *just is* a matter of being disposed to respond to those items in an appropriate way. When an agent succeeds in responding in an appropriate way to one of these items, we can say that they have

“hit the target” of the corresponding virtue.⁸ Swanton provides the following examples of items within the fields of certain virtues, and of the appropriate modes of response:

The virtue of justice is primarily concerned with the honouring of rules of justice by adhering to those rules oneself, and with respect for the status of individuals... Many virtues, such as that of friendship, exhibit many modes of moral acknowledgement. A good friend does not merely promote the good of her friend: she appreciates her friend, respects, cares for, and even loves her friend.⁹

The virtue of justice has rules and other persons as items within its characteristic field of concern, and the appropriate modes of response are to honour the rules (by following them) and respect the persons. The proposed virtue of friendship (or friendliness) has the friend as an item of concern, and appropriate modes of response include promoting the friend’s well-being, as well as caring for, loving, and respecting the friend.

With this understanding of virtue in place, Swanton then provides her account of the distinction between an *act from virtue* and a *virtuous act*.¹⁰ An act from virtue is one that expresses or exhibits the agent’s actual possession of the relevant dispositions. For example, an *act from justice* would be one that expresses the agent’s underlying dispositions to adhere to the rules of fairness and to respect other people. A virtuous act, on the other hand, is one that hits the target of virtue, regardless of whether the agent in question possesses the relevant dispositions. For example, a *just act* could be one that adheres to the rules of fairness, even if the agent in question is not generally disposed to do so. Stangl follows Swanton when defining virtuous actions and overall virtuous actions:

Virtuous Action: An action is virtuous in respect V (e.g., benevolent, generous) if and only if it hits the target of (realizes the end of) virtue V (e.g., benevolence, generosity).

Overall Virtuous Action: Overall virtuousness is a function of the virtuousness of actions in these more particular respects. If only one virtue is relevant in a situation, then overall virtuousness will consist simply in hitting the target of that virtue. But in those cases in which more than one virtue is relevant, which action counts as overall virtuous is determined by the relative importance of the differing virtues at issue.¹¹

Overall virtuousness, then, is a matter of hitting the target(s) of those virtues that are relevant in the situation. This provides us with a definition of the first key concept that appears in Stangl's account of supererogation. The second key concept is that of overall viciousness. Building on the work of Swanton once more, Stangl defines overall viciousness as follows:

Vicious Action: An action is vicious in respect R (e.g., stingy, unjust) if and only if it constitutes a failure to hit the target of its corresponding virtue V (e.g., generosity, justice) in the manner distinctive of that vice (e.g., giving too little, failing to respect rights).

Overall Vicious Action: Overall viciousness, in turn, is a function of the viciousness of actions in these more particular respects. If only one virtue is relevant in a situation, then overall viciousness will consist simply in a failure to hit the target of that virtue.¹²

With these details in place, it is possible to understand Stangl's neo-Aristotelian account of supererogation. A morally right action is one that is overall virtuous. A morally wrong action is one that is overall vicious. A supererogatory action is one that is overall virtuous and where *either* (a) failing to perform an overall virtuous act would not have been overall vicious, or (b) some other overall virtuous act was available that was less virtuous than the action performed.

Stangl has two aims when defending this account. The first is to demonstrate that it is compatible with the Aristotelian doctrine of the mean. The second is to demonstrate that it generates

intuitively plausible judgements concerning supererogatory acts. I will focus on the second of these aims.

Does Stangl's account generate plausible verdicts in response to the classic *Soldier* and *Doctor* examples? Consider first the *Soldier* example. The soldier who dives onto the live grenade in order to save their comrades plausibly hits the targets of relevant virtues. They overcome fear in order to promote something of value in a genuinely dangerous situation, thereby hitting a target of courage. They also successfully protect and promote the well-being of their comrades, thereby hitting a plausible target of benevolence. The action of the soldier, therefore, looks likely to be classed as overall virtuous, and so is eligible for supererogatory status. Whether or not it achieves that status depends on the virtuousness or viciousness of the alternative actions that were available.

Stangl states that it would not have been overall vicious to refrain from jumping on the grenade in this example. We are told that it is not "plausible" to describe the failure to act in such cases as overall vicious, and that "we do not think that the beneficiaries of the soldier's action acted viciously in not jumping on the grenade."¹³ If this is correct, then the soldier has performed an overall virtuous action, the omission of which would not have been overall vicious. In this way, Stangl argues that her account generates the (correct) verdict that *Soldier* is an example of supererogation.

Similar moves are made to suggest that *Doctor* would also be supererogatory on Stangl's account. Stangl asks us to compare the doctor who moves to work in challenging and dangerous circumstances, with a colleague who declines such a move, but who continues to serve the local area in a fair and considerate way.¹⁴ Stangl suggests that both agents act in a way that is overall virtuous, but that choosing to move is slightly more virtuous. Therefore, on the neo-Aristotelian account, *Doctor* will be (correctly) categorised as an example of supererogation. The doctor's actions are overall virtuous, and some alternative overall virtuous action was available that would have been less virtuous. If this is right, Stangl will have successfully demonstrated that her

account generates plausible verdicts in response to Urmson's classic examples. I will now argue that Stangl is not, in fact, successful in this regard.

III. THE NEED FOR DEVELOPMENT

One of Stangl's key aims is to demonstrate that her account will generate intuitively plausible results. However, the level of detail provided by Stangl is not sufficient to allow us to work out *what* verdicts would be generated in response to examples, and so is not sufficient to demonstrate that plausible verdicts would be forthcoming. The difficulty for Stangl stems from the understandings of "overall virtuous" and "overall vicious" that are used in the account.

Regarding the understanding of overall virtuous action, Stangl explains that this concept lends itself to both a strong and a weak interpretation. On the strong interpretation, an overall virtuous action will be one that hits *all* of the targets of *all* relevant virtues, and to the highest degree possible. Stangl instead favours a weak interpretation, which she explains by providing a direct quote from Swanton:

An act is right if and only if it is overall virtuous, and that entails that it is good enough, even if not the (or a) best action. Here it is assumed that there is much latitude in hitting the target of virtues such as generosity. Right acts range from the truly splendid and admirable to acts which are "all right".¹⁵

However, this understanding of the weaker interpretation of overall virtuousness is in conflict with Stangl's account of supererogation. On this understanding, overall virtuous acts will include all of those that are "all right". But, for Swanton, morally "all right" action is a technical term, meant to include all of those actions that are not overall vicious.¹⁶ Therefore, if Stangl follows Swanton in her understanding of the weak interpretation of overall virtuousness, then the first clause in her account of supererogation will be redundant. Overall virtuous acts will include all of those that are "all right", and so will include all of those that are not overall vicious. This means that it could never be the case that failing to perform an overall virtuous action would not be

overall vicious. If Stangl continues to understand overall virtuousness in the way that is borrowed from Swanton, then the first clause of her account of supererogation (clause (a)) will never be satisfied. The only information we are given on how to understand overall virtuousness is presented by way of the direct quote from Swanton. Given that the position stated in that quote conflicts with Stangl's account, it is not clear how overall virtuousness ought to be understood. This means that we cannot work out what verdicts would be generated by the neo-Aristotelian account.

A similar problem arises for the understanding of overall viciousness that Stangl provides. Again, Stangl explains that this concept lends itself to both a strong and a weak interpretation. On the strong interpretation, an act will only be classified as overall vicious if it is "the, or a, worst possible action in the circumstances", where this would involve failing to hit the targets of *all* relevant virtues in the most extreme way possible.¹⁷ Stangl again prefers a weaker interpretation. On the weaker interpretation, overall vicious actions will be those that are "bad enough even if not the (or a) worst action".¹⁸ This definition tells us what is *not* required in order to be overall vicious (the act need not be the worst possible), but it does not tell us what *is* required. This lack of detail leaves us unable to determine what verdicts will be generated in specific cases.

Consider again Urmson's *Soldier* example. Stangl tells us that her account will generate the intuitively correct verdict here. The soldier's action plausibly hits the targets of relevant virtues, such as courage and benevolence, and so looks likely to be classed as overall virtuous. The soldier's action is therefore eligible for supererogatory status. In order for that status to be awarded, however, the actions of those who act differently must not be overall vicious.

Stangl tells us that it would not be "plausible" to say that those who did not jump on the grenade acted viciously, and that "we do not think" that they would be acting viciously. But this is not the same as demonstrating that her account generates the correct result. Indeed, it does appear that the other soldiers are failing to hit the targets of relevant virtues. They fail to overcome fear in order to protect something of value (the lives of their fellow comrades) in a dangerous situation.

And they fail to act so as to protect and promote the well-being of others. In this way, the other soldiers fail to hit the targets of courage and benevolence. It is not clear, therefore, why failing to jump on the grenade will not be classed as overall vicious on Stangl's account, especially given a weak interpretation of overall viciousness. While it may not be *plausible* to categorise these actions as overall vicious, it is not clear how Stangl's account avoids the implausible verdict.¹⁹ The lack of detail provided when explaining the account means that we cannot confirm that the desired verdicts would be generated, and so Stangl has not succeeded in one of her key aims when defending the neo-Aristotelian account.

This is a significant issue given Stangl's strategy of defending her account by claiming that plausible verdicts would be forthcoming. If we cannot work out what verdicts would be generated, then Stangl's strategy cannot succeed. However, it might be thought that this is merely a temporary set-back. All that has been argued so far is that the neo-Aristotelian account is problematically indeterminate. This is consistent with there being some straightforward way of developing the account so as to satisfy the stated aims (generating plausible verdicts, and being compatible with the doctrine of the mean). And yet, there is reason to believe that successfully developing Stangl's account will be far from straightforward. In the next section, I will argue that anyone attempting to develop the neo-Aristotelian account faces a dilemma, and that both options for responding to this dilemma generate problematic results.

IV. A DILEMMA

To develop the neo-Aristotelian account, it is necessary to clarify the requirements for "overall virtuous" and "overall vicious" action. Stangl's desire to avoid strong interpretations of these concepts is justified. To require that right (overall virtuous) actions must hit *all* of the targets of *all* relevant virtues, and to the greatest extent possible, would be overly demanding.²⁰ Similarly, requiring that wrong (overall vicious) actions must fail to hit *all* of the targets of *all* relevant virtues, and in the most extreme way possible, would result in many wrong actions being classed as morally "all right". We ought, therefore, to agree with Stangl, and reject strong interpretations

of overall virtuous and overall vicious action. It is necessary to develop weaker interpretations of these concepts that can then be deployed within the neo-Aristotelian account of supererogation.

Rival interpretations of the key concepts will differ in terms of what can be referred to as the *accuracy requirement* and the *range requirement*. The range requirement sets out how many of the virtuous targets need to be hit (or missed) in order for an action to count as overall virtuous (or overall vicious). For example, one option is to require that only the main target of the most relevant virtue needs to be satisfied. If the most relevant virtue is friendship, then perhaps an overall virtuous action need only hit the target of promoting the well-being of the friend, even if the agent fails to hit the additional targets of actually caring for and respecting the friend. Alternatively, we might demand that every target of the most relevant virtue be satisfied. Overall virtuousness in the friendship example would then require that the agent care for and respect the friend, as well as promoting the friend's well-being. Different understandings of the range requirement will result in different verdicts concerning when an action is overall virtuous (or overall vicious). The more demanding we make the requirement, the closer we will come to strong interpretations of these concepts. To defend the neo-Aristotelian account, it is necessary to demonstrate that there is *some* understanding of the range requirement that is compatible with the production of plausible verdicts concerning supererogation.

Interpretations of overall virtuousness and overall viciousness will also differ in terms of the accuracy requirement. This requirement sets out how narrowly the targets of the relevant virtues ought to be construed. On the most demanding version of the accuracy requirement, the target of a virtue will be construed such that there is only one way of successfully hitting it.²¹ If a target of benevolence is to promote the well-being of others, then a demanding version of the accuracy requirement will state that hitting this target requires promoting well-being to the highest degree possible. On a less demanding version, hitting the target might only require promoting the well-being of others to some lesser extent. Again, the more stringent this requirement, the closer we will be to the strong interpretations of overall virtuous and overall vicious action. To fully develop

the neo-Aristotelian account, we need to explain how both the *range requirement* and the *accuracy requirement* are to be understood.

The problem is that it is not clear how we can understand these two requirements so as to provide interpretations of overall virtuous and overall vicious action that will ultimately generate plausible results. In order to demonstrate this problem, I will focus on one way in which we could understand the range requirement so as to generate the desired verdict in *Soldier*. My aim is to show that unintuitive results will be generated regardless of whether we accept this version of the requirement, and so supporters of the neo-Aristotelian account face a dilemma in deciding how to proceed.

When deciding how to understand the range requirement, it is necessary to take a stand on the importance of the *internal targets* of relevant virtues. These are targets that involve feeling or being motivated in a particular way. For example, Swanton suggests that the virtue of friendship involves not only promoting the well-being of your friend (an external target), but also actually caring for and respecting your friend (internal targets). Similarly, the virtue of justice might involve acting in accordance with the rules of fairness, such as ensuring a fair distribution of a particular resource (an external target), as well as feeling the proper respect for others while doing so (an internal target). Focusing on internal targets provides a way of generating the desired verdict in examples like *Soldier*.

On Stangl's account, an action can only be categorised as supererogatory in cases where some alternative action was available that was not overall vicious. This means that the act of jumping on the grenade in *Soldier* will only be categorised as an example of supererogation if the alternative action of failing to jump on the grenade is not overall vicious. The problem is that this alternative action *does* appear to miss the targets of relevant virtues, such as courage and benevolence. However, the problematic verdict can be avoided if the range requirement is understood in a way that acknowledges the importance of internal targets. That is, we need to

understand the range requirement as allowing for the moral status of an action to be improved when the internal targets of relevant virtues are satisfied.

In *Soldier*, the soldiers who fail to jump on the grenade do appear to miss the external targets of courage and benevolence. However, those soldiers might well be hitting the internal targets of these virtues. They might care about the well-being of their comrades, thereby hitting an internal target of benevolence. And they might generally be resolute in the face of danger, thereby hitting an internal target of courage. If we understand the range requirement as acknowledging the importance of internal targets, then we can say that the other soldiers are not acting in a way that is *overall* vicious. And this is what Stangl needs in order for the neo-Aristotelian account to correctly judge the act of jumping on the grenade in *Soldier* as supererogatory.

Understanding the range requirement in this way would also allow the neo-Aristotelian account to generate plausible verdicts in additional cases. There is good reason to accept that internal states or motivations can determine whether or not an action should be classified as supererogatory.²² Consider the following variation on Urmson's classic example:

Modified Soldier: Angela is among a squad of soldiers practicing the throwing of live grenades. In an attempt to impress her comrades and deceive them into awarding her esteem, Angela plans to drop a dummy grenade on to which she will throw herself. However, Angela unknowingly drops a live grenade. Thinking there to be no real danger, Angela throws herself onto the live grenade.²³

Angela's act of jumping on the grenade once it has been dropped hits the same external targets as in the original *Soldier* example (protecting the well-being of her comrades). However, *Modified Soldier* does not deserve to be classed as supererogatory. This can be explained by the fact that Angela's motivations and concerns are not focused on the well-being of her comrades, and by the fact that her incorrect assessment of the situation means that she is not overcoming fear while protecting something of value. Angela therefore fails to hit the internal targets of relevant

virtues such as courage and benevolence towards her comrades.²⁴ Understanding the range requirement in a way that acknowledges the importance of internal targets would allow a neo-Aristotelian to rule out *Modified Soldier* as an example of supererogation, as well as generating the correct verdict in response to the original case.

However, accepting that the moral status of an action is improved by hitting the internal targets of relevant virtues has other, more problematic, implications. The problem is that this will result in granting supererogatory status to actions that are more intuitively viewed as obligatory. Consider the following example:

Debt: Brenda receives a significant loan from an acquaintance and promises to repay the money by a particular date. Feeling a deep respect for the rules of justice, Brenda carefully plans to ensure that she will make the repayments. When the due date arrives, Brenda successfully repays the loan.

If we understand the range requirement as acknowledging the importance of internal targets, then Brenda's repaying of the loan will be classed as supererogatory. An agent could have performed a morally right (overall virtuous) act simply by repaying the money. By repaying the money *and doing so out of her respect for the rules of justice*, Brenda hits both the external and the internal targets of a relevant virtue (justice). Brenda therefore performs an overall virtuous act (repaying the loan from an appropriate motive) that is more virtuous than an alternative overall virtuous act (repaying the loan without the appropriate motivation). Brenda's act will therefore be classed as supererogatory by the neo-Aristotelian account once we acknowledge the importance of internal targets.

This result will generalise. Morally right actions will become supererogatory whenever the agent performs the action while feeling the right way. This means that many plausibly obligatory actions will be categorised as supererogatory. Consider actions such as meeting a friend at an agreed upon time; completing your fair share of the marking load in your department; or ensuring that

the basic needs of your children are satisfied. If hitting the internal target of a relevant virtue can improve the moral status of an action, then all of these actions will be counted as supererogatory whenever the agent performs them while having the appropriate feelings. The actions will be overall virtuous, and will be more virtuous than some alternative action that is also overall virtuous (performing the action without the correct feeling). Those working on supererogation have not typically thought that actions such as paying back a loan or satisfying the basic needs of your children should be included as examples of supererogation. When focusing on the internal targets of virtue, the neo-Aristotelian account leads to unintuitive verdicts in these cases.

A defender of Stangl's neo-Aristotelian account therefore faces a dilemma when developing the account. In order to ensure that the desired verdict is generated in response to cases like *Soldier* and *Modified Soldier*, it will be necessary to allow that hitting (or missing) the internal targets of relevant virtues can affect the moral status of an action. However, a neo-Aristotelian account that allows for this will then produce unintuitive verdicts in response to everyday examples like *Debt*. Regardless of the position taken on the issue of internal targets, the neo-Aristotelian account will produce verdicts that many will find unintuitive.

Given this issue, what options remain for a virtue ethical response to the challenge of supererogation? There appear to be four main options. If we agree with Stangl that the production of intuitive verdicts is a requirement for any successful account, then we might (i) reject the neo-Aristotelian account, and defend some alternative approach. Rival virtue ethical accounts of supererogation have been proposed elsewhere.²⁵ Alternatively, we might (ii) reject the neo-Aristotelian account, and then deny that an account of supererogation is even necessary. Denying the possibility of supererogation has also been proposed elsewhere.²⁶ These look to be the main options if we accept that a successful account of supererogation must be in-keeping with standard intuitions.

A more interesting response from a virtue ethicist would be to deny that their account needs to satisfy standard intuitions, perhaps by arguing that those intuitions ought not to be trusted. In

taking this approach, it would be possible to accept the neo-Aristotelian account, and either (iii) deny that *Solider* is an example of supererogation, or (iv) argue that *Debt* really *is* an example of supererogation. In either case, some argument would then be required as to why we ought not to trust any contrary intuitions.

I will not here discuss the merits of these four options. My aim in this paper has simply been to argue that the neo-Aristotelian account fails to satisfy Stangl's aim of generating intuitively plausible verdicts in response to classic examples. If the neo-Aristotelian account is the best available, then virtue ethicists will be forced to choose between either continuing to accept that account (and so defending a revisionary understanding of supererogation), or else abandoning the attempt to accommodate supererogation altogether.

CONCLUSION

Rebecca Stangl has attempted to reconcile virtue ethics and supererogation by proposing a neo-Aristotelian account of supererogatory acts. I have argued against that proposal in two ways. Firstly, I have argued that the proposal in its current form is problematically indeterminate, meaning that we cannot know what verdicts would be produced in response to classic examples. Secondly, I have argued that anyone attempting to develop the account faces a dilemma concerning the importance of the internal targets of virtue. Regardless of the position taken on this issue, unintuitive verdicts will be generated. Stangl has failed to demonstrate that virtue ethics can accommodate supererogation in a way that is intuitively appealing.

* I am grateful to Christian Miller, members of the work in progress group at Wake Forest University, and to anonymous reviewers for feedback on earlier drafts of this paper. I am also grateful to Alfred Archer for extremely helpful conversations regarding supererogation. Work on this paper was supported by a grant from the Templeton Religion Trust. The opinions expressed here are those of the author and do not necessarily reflect the views of the Templeton Religion Trust.

¹ J. O. Urmson, "Saints and Heroes," in *Essays in Moral Philosophy*, ed. A. I. Melden (Seattle: University of Washington Press, 1958): 198-216.

² Paraphrased from *Ibid.*, 201-202.

³ See Jason Kwall, "Virtue Theory, Ideal Observers, and the Supererogatory," *Philosophical Studies* 146 (2009): 179-196; and Johan Brännmark, "From Virtue to Decency," *Metaphilosophy* 37 (2006): 589-604.

⁴ While my focus here is on the attempt to account for supererogation, one alternative response would be to reject the suggestion that all ethical theories need to accommodate supererogation. Roger Crisp, "Supererogation and Virtue," in *Oxford Studies in Normative Ethics, Volume 3*, ed. Mark Timmons (Oxford: Oxford University Press, 2013): 13-33 argues that Aristotelian ethics (and pre-Christian ethics more generally) should not be expected to accommodate this concept. For a related discussion, see Julia Annas, "Virtue and Duty: Negotiating Between Different Ethical Traditions," *The Journal of Value Inquiry* 49 (2015): 605-618.

⁵ Rebecca Stangl, "Neo-Aristotelian Supererogation," *Ethics* 126 (2016): 339-365, at 340-341.

⁶ Christine Swanton, *Virtue Ethics: A Pluralistic View* (Oxford: Oxford University Press, 2003), Ch.1.

⁷ *Ibid.*, 20.

⁸ *Ibid.*, Ch.11.

⁹ *Ibid.*, 22.

¹⁰ *Ibid.*, Ch. 11.

¹¹ Stangl, "Neo-Aristotelian Supererogation," 342-343.

¹² *Ibid.*, 345-346.

¹³ *Ibid.*, 356.

¹⁴ *Ibid.*

¹⁵ *Ibid.*, 343 (Quoting Swanton, *Virtue Ethics*, 240).

¹⁶ Swanton, *Virtue Ethics*, 240.

¹⁷ Stangl, "Neo-Aristotelian Supererogation," 346.

¹⁸ *Ibid.*

¹⁹ Indeed, it may be that Aristotle would have endorsed what Stangl takes to be the less plausible view. Crisp, "Supererogation and Virtue," 23-24, argues that Aristotle's verdict would have been that "not sacrificing yourself in

such a case is vicious (it is cowardly, and probably violates obligations of friendship to one's comrades)." That position is at odds with Stangl's own aims in accommodating the concept of supererogation.

²⁰ At the very least, this will be too demanding given the aim of accommodating supererogation. Defenders of a neo-Aristotelian ethical theory who *reject* the necessity of accommodating supererogation might also deny that this interpretation of overall virtuousness is overly demanding.

²¹ This would resemble Aristotle's claim that "it is possible to fail in many ways... while to succeed is possible in only one way", Aristotle (trans. Ross), *Nicomachean Ethics* (Oxford: Oxford University Press, 1998), 38 [1106^b28-31].

²² This issue has been discussed elsewhere, including Alfred Archer, "Supererogation and the Intentions of the Agent," *Philosophia* 41 (2013): 447-462; Phillip Montague, "Acts, Agents, and Supererogation," *American Philosophical Quarterly* 26 (1989): 101-111; and David Heyd, *Supererogation: Its Status in Ethical Theory*, (Cambridge: Cambridge University Press, 1982).

²³ A similar variation of Urmson's example is mentioned in Crisp, "Supererogation and Virtue," 17.

²⁴ There is, of course, the issue of determining which virtues are most relevant in each situation, and therefore which internal and external targets need to be hit in order to improve the moral status of an action. This is an issue that affects virtue ethical accounts of right action more generally, and not just Stangl's approach. I here assume that the virtues of courage and benevolence towards one's comrades are relevant in this example.

²⁵ See Kawall, "Virtue Theory, Ideal Observers, and the Supererogatory" and Brännmark, "From Virtue to Decency".

²⁶ See Crisp, "Supererogation and Virtue" and Susan Hale, "Against Supererogation," *American Philosophical Quarterly* 28 (1991): 273-285.