[This is the pre-peer-reviewed version of the following article 'Consequentialism with Wrongness Depending on the Difficulty of Doing Better' *Thought* 5 (2): 108–118, 2016.]

Consequentialism with Wrongness Depending on the Difficulty of Doing Better

Johan E. Gustafsson*

Moral wrongness comes in degrees. On a consequentialist view of ethics, the wrongness of an act should depend, I argue, in part on how much worse the act's consequences are compared to those of its alternatives and in part on how difficult it is to perform the alternatives with better consequences. I extend act consequentialism to take this into account, and I defend three conditions on consequentialist theories. The first is consequentialist dominance, which says that, if an act has better consequences than some alternative act, then it is not more wrong than the alternative act. The second is consequentialist supervenience, which says that, if two acts have equally good consequences in a situation, then they have the same deontic status in the situation. And the third is consequentialist continuity, which says that, for every act and for any difference in wrongness δ greater than zero, there is an arbitrarily small improvement of the consequences of the act which would, other things being equal, not change the wrongness of that act or any alternative by more than δ. I defend a proposal that satisfies these conditions.

Traditional consequentialism states that an act is right if and only if its consequences are at least as good as those of every alternative act and, if the act is not right, it is wrong. The wrong acts, however, might differ in ways that could be morally relevant. The consequences of one of them might be almost optimal while those of another are catastrophic in comparison. In addition to value difference, a further intuitively morally relevant difference between different wrong acts is how hard it would have been to avoid them—how hard it would have been to perform a better act instead.[1] In this study, I shall defend an extended version of consequentialism that is capable of taking these kinds of differences into account for a more fine grained grading of the wrongness of acts.

* I would be grateful for any thoughts or comments on this paper, which can be sent to me at *johan.eric.gustafsson@gmail.com*.

[1] If we grant the possibility of being passive and not performing any act in a choice situation, there might be a difference between how hard it is to avoid an act and how hard it is to perform an alternative act. For the purposes of this study, however, I shall count being passive in a choice situation as an act.

It is natural to assume that consequentialism can account for degrees of wrongness by taking into account how much worse the consequences of a wrong act are than those **[p. 109]** of the right acts.[2] Consider a choice situation where there are three available alternatives with consequences valued as follows:

Situation 1

Act	Value
a_1	100
a_2	99
a_3	0

Here, it appears reasonable to say that a_3 is more wrong than a_2, because the loss of value compared to the optimal alternative, a_1, is greater for a_3 than for a_2. Moreover, it seems that, a_3 is not just more wrong than a_2, it is much more wrong than a_2 because the loss of value relative to a_1 is much greater for a_3 than for a_2.[3]

We introduce the following notation:

$W(x)$ is the degree of wrongness of x.

$V(x)$ is the value of the consequences of x.

Ω is the set of available alternatives.

[2] See, for example, Mackie (1977, p. 205), Eriksson (1997, pp. 218–219), and Calder (2005, p. 229). Mill's (1969, p. 210) proportionality criterion of wrongness, which says that acts are wrong 'as they tend to produce the reverse of happiness', might be read as suggesting a similar view.

[3] The main theoretical interest in such cardinal information about how much more wrong an act is compared to another act (rather than merely ordinal information about whether it is more wrong than the other act) is that it affords a neat solution for dealing with the following kind of case—first discussed by Regan (1980, p. 265)—where two states of nature are equally likely:

Act	State 1 (0.5)	State 2 (0.5)
a_1	Very minor wrong	Very minor wrong
a_2	Major wrong	Right
a_3	Right	Major wrong

In this case, it seems that a morally conscientious person would choose a_1, which is known to be slightly wrong, rather than one of a_2 and a_3, each of which might be right but they both have a 0.5 chance of being very wrong. However, if a_1 were instead known to be almost as wrong as a_2 and a_3 are in states 1 and 2 respectively, it appears that a morally conscientious person would instead choose one of a_2 and a_3. Hence it appears that morally conscientious people are not just concerned with avoiding wrong acts; their concern with avoiding wrong acts is proportional to the acts' degrees of wrongness; see Graham (2010, p. 99) and Bykvist (2011, p. 37). A natural way of spelling this out is that morally conscientious people act as to minimize expected moral wrongness, and that requires cardinal rather than merely ordinal degrees of wrongness.

Given that each alternative might be performed without difficulty, the following appears plausible:[4]

(1) $W(x) = \max\left\{V(y) - V(x) \,\middle|\, y \in \Omega\right\}.$

In other words, one calculates the degree of wrongness of an alternative x by, for each alternative y, calculating the result of subtracting $V(x)$ from $V(y)$. According to equation (1), the greatest of these results is equal to the degree of wrongness of x. Given that all acts in situation 1 can be performed without difficulty, this yields that $W(a_3) = 100$ and $W(a_2) = 1$. These results are, I think, in line with what one would expect from a consequentialist theory capable of handling degrees of wrongness. Another welcome feature is that the acts with optimal consequences—that is, the right acts according to traditional consequentialism—always get a zero degree of wrongness: In situation 1, for instance, equation (1) yields that $W(a_1) = 0$.

In addition to value difference, the wrongness of an act might depend on the difficulty of the acts with better consequences.[5] In the literature, there are various accounts of degrees of easiness and difficulty of acts.[6] I shall try to remain neutral between these accounts. I shall assume, however, that each act among your alternatives is securable in the sense that you cannot fail to perform the act you have chosen to perform. Hence the easiness of an act should not be thought of as the probability of successfully performing the act given that you have chosen to perform it. Non-securable acts—such as, *buying a winning lottery ticket*—are better modelled as securable acts with an uncertain outcome—such as, *buying a lottery ticket* when one does not know whether the ticket is a winner. A heroic act such as *throwing yourself on landmine to save your fellow soldiers* might very well be easy in the sense that you are likely to succeed in performing it if you choose to perform it. But this act is plausibly very difficult in the sense that it is very difficult to bring yourself to choose to

[4] Note that this is not the only formula that could account for the intuitions we have appealed to so far. Taking an act's degree of wrongness to be proportional to the value difference between the act's consequences and those of the right act (or acts) appears, however, to be the simplest, most straightforward way to do so. The aim of this paper is mainly to find one plausible account of degrees of wrongness depending on the difficulty of doing something better and not to show that it is the only plausible account of this kind.

[5] Compare Parfit (1984, p. 33), who suggest that if it is very hard to not act in a certain way, acting that way is only morally bad in a very weak sense. Moreover, compare Berlin (1958, p. 15n) who suggests that the morally relevant sense of freedom depends not just on which possibilities are open to the agent but also on how easy or hard they are to achieve.

[6] For an early account of degrees of ability, see Benson (1987, p. 329). Moreover, see Portmore (2007, p. 10) for a distinction between degree of difficulty and degree of effort, and Cohen (1978, pp. 238–239) for a distinction between difficulty and cost.

perform it, due, among other things, to the sacrifice the act involves. It is easiness and difficulty of acts in this latter sense that I wish to discuss **[p. 110]** in this study. That is, the relevant easiness of an act is the easiness by which one can bring oneself to choose to perform the act.

We will express the degree of easiness of an alternative on a scale from 0, representing impossible, to 1, representing maximally easy, that is, performable without any difficulty.[7] Consider the following situations:[8]

Situation 2

Act	Value	Easiness
a_1	100	0.1
a_2	10	0.5

Situation 3

Act	Value	Easiness
a_1	100	0.9
a_2	10	0.5

The idea is that, even though the value difference is the same in situation 2 as in situation 3, act a_2 is more wrong in situation 3 than in situation 2, because the better alternative is easier in situation 3 than in situation 2. We introduce the following notation for easiness:

$E(x)$ is the degree of easiness of x.

The first to develop an extension of consequentialism that takes degrees of easiness into account was, as far as I know, Björn Eriksson.[9] He proposes that

$$(2) \qquad W(x) = \max\left\{\left(V(y) - V(x)\right)\frac{E(y)}{E(x)} \,\middle|\, y \in \Omega\right\}.$$

This means that you calculate the degree of wrongness of an alternative x by, for each alternative y, first subtracting $V(x)$ from $V(y)$ and then multiplying the result by $E(y)$ divided by $E(x)$. According to equation (2), the greatest of these products is equal to the degree of wrongness of x. The theory yields that $W(a_2) = 18$ in situation 2 and $W(a_2) = 162$ in situation 3. Consequently, we get the desired result that a_2 is more wrong in situation 3 than in situation 2.[10]

---

[7] What I call degrees of easiness, Eriksson (1997, p. 219) calls degrees of difficulty. That is, degree of difficulty 0 represents impossible and degree of difficulty 1 represents something that can be done without difficulty. But this seems backwards. If, for example, a certain math exercise has a higher degree of difficulty than another, it should be more, not less, difficult. Portmore (2007, p. 6) similarly takes a goal of zero difficulty to be a goal that can be achieved without difficulty. Another difference is that Eriksson (1997, p. 219) identifies the least difficult degree with something that cannot be avoided, while I allow that there may be two or more mutually exclusive but maximally easy acts in a situation.

[8] Eriksson (1997, p. 219).

[9] Eriksson (1994, 1997).

[10] The division in equation (2) might raise worries about division by zero if alternatives might have the degree of easiness 0. Yet it seems plausible that, if an act is an alternative, it can be performed. Not just ought, but also the property of being an alternative, implies can. This entails that all alternatives must have a degree of easiness greater than 0. Hence we avoid division by zero.

This line of thought brings us to a further reason for taking degrees of easiness into account. The difficulty of acts is relevant for moral rightness on all plausible moral theories. The acts that are morally relevant are just those that are not so hard that they are impossible to perform. So to let moral rightness and wrongness depend on degrees of easiness is to take into account in a more fine grained way something that all plausible moral theories take into account. One way of thinking about this could be that the morally relevant property of being an alternative comes in degrees or that an act's relevance as an alternative comes in degrees depending on how easy it is to perform.

Folke Tersman offers the following counter-example to equation (2):[11]

Situation 4

Act	Value	Easiness
a_1	101	0.1
a_2	100	0.001
a_3	0	0.899

[p. 111] In this situation, equation (2) yields that $W(a_2) = 100$ and $W(a_3) \approx 11.2$. It is implausible that a_2 would be almost nine times as wrong as a_3. The consequences of a_2 are almost as good as those for the right act, a_1, while the consequences of a_3 are much worse. That a_2 is much harder than a_1 while a_3 is relatively easy should not make a_2 more wrong than a_3. Situation 4 shows that (2) violates the following principle:

The principle of consequentialist dominance
If acts x and y are available in the same situation and x has better consequences than y, then x is not more wrong than y.

If an act has better consequences than another act, it comes closer to what is, according to consequentialism, the ultimate aim of morality at large: that consequences be as good as possible.[12] Hence the act with better consequences should not be judged morally more severely. Should one give up the principle of consequentialist dominance, there would not be much left of traditional consequentialism.

Situation 4 illustrates why it is unreasonable to let an act's degree of wrongness depend on its *own* degree of easiness. We find a perhaps even clearer illustration of this problem in the following situation, where two wrong alternatives have equally good consequences but vary in terms of difficulty:

[11] Tersman (1997, p. 50).
[12] See, for example, Bentham (1970, p. 282) and Parfit (1984, p. 24).

Situation 5

Act	Value	Easiness
a_1	100	0.5
a_2	10	0.9
a_3	10	0.1

Here, (2) yields that $W(a_2) = 50$ and $W(a_3) = 450$. Whether you perform the easier a_2 or the harder a_3 appears irrelevant for the degree of wrongness of what you do. What matters is that, irrespective of whether you performed a_2 or a_3, you have achieved an outcome worth 10 units of value while you could have achieved an outcome worth 90 units more with degree of easiness 0.5. That is, a_2 should have the same degree of wrongness as a_3.

Since (2) yields that a_2 and a_3 differ in degree of wrongness even though they have equally good consequences, it violates the following principle, which Krister Bykvist has dubbed as part of the 'spirit of consequentialism':

> *The principle of consequentialist supervenience*
> If two acts that are available in the same situation have equally good consequences, then the acts have the same deontic status.[13]

If one rejects the principle of consequentialist supervenience, one rejects a fairly basic tenet of consequentialism. This, by itself, might not be so worrisome, but as long as our ambition is to achieve a moderately conservative extension of consequentialism, we should avoid this.[14]

One might still be unconvinced. Bykvist, for example, does not claim that the principle of consequentialist supervenience holds if the deontic statuses of acts depend in part **[p. 112]** on how hard they are.[15] But note that the principle of consequentialist supervenience cannot be denied if one accepts the principle of consequentialist dominance along with the following principle, which any plausible consequentialism with degrees of wrongness should satisfy:

[13] Bykvist (2002, p. 52).

[14] Eriksson has told me that he gladly gives up the principle of consequentialist supervenience since a suboptimal act which involves that *one goes to extra trouble* to act wrongly—that is, a wrong act that is more difficult than it would be to act rightly—is more wrong than an axiologically equivalent act that is easier than acting rightly. Still, the basis of this view is hard to discern. It seems like Eriksson relies on the wrong kind of intuitions, that is, non-consequentialist intuitions. Perhaps the idea is that difficult performances or that one goes to extra trouble is something that is morally problematic in addition to any effect on the value of the consequences. Or it might be that there is something especially tragic with a person who goes to extra trouble for something that in the end still leads to a bad result. Furthermore, it might be evil to go to extra trouble in order to, or with an aim to, do wrong. The problem is that none of this seems to have any basis in consequentialism.

[15] Bykvist (2003, p. 34n).

> *The principle of consequentialist continuity*
> If act x is available in a situation, then, for any difference in wrongness δ greater than zero, there is an arbitrarily small improvement of the consequences of x which would, other things being equal, neither increase nor decrease the degree of wrongness of any available act in the situation by more than δ.

This principle is plausible given consequentialism and graded moral assessment. If moral assessment comes in degrees, gradual changes in the morally relevant factors should yield gradual changes in the moral assessment. Moreover, the principle of consequentialist continuity is satisfied by (2). Hence this principle should, at least in this context, be fairly uncontroversial.

For the argument, assume the principle of consequentialist continuity and the principle of consequentialist dominance. If the principle of consequentialist supervenience is violated, there is a situation where two acts have equally good consequences but the acts are not equally wrong. Consider the act that is more wrong than the other act. Given the principle of consequentialist dominance, any improvement of the consequences of this act would, other things being equal, make it less wrong than the other act. Hence any improvement of this act would, other things being equal, change the degree of wrongness of an act by at least half of the previous difference in wrongness between the two acts. This, however, violates the principle of consequentialist continuity. So, if we accept the principle of consequentialist continuity and the principle of consequentialist dominance, we should also accept the principle of consequentialist supervenience.

Whether you perform an easy act or a hard act with equally good consequences should be irrelevant given consequentialism. This does not, however, rule out that there might still be room for degrees of easiness even in a moderately conservative extension of consequentialism that retains our three consequentialist principles. According to consequentialism, the deontic status of an act depends not only on the value of its consequences. The value of other acts' consequences are also relevant given that they are not impossible to perform in the situation, that is, they have a degree of easiness greater than zero. Hence the relevance or weight of the value of the consequences of other acts depends to some extent on the degrees of easiness of these acts. This opens up for an extension of consequentialism which in a more fine grained way takes degrees of easiness into account without introducing anything too alien, that is, something that is not to some extent already there in the traditional version.

Traditional consequentialism yields that an alternative is wrong if and only if there is an act with better consequences which is not impossible to perform. This can be expressed as that an alternative x is wrong if and only if there is an act y such that $E(y)$ multiplied by the **[p. 113]** result of $V(y)$ subtracted by $V(x)$ is greater than zero. A natural extension would be to let the extent to which an alternative is wrong depend on how much such a product of value difference and the better alternative's degree of easiness differs from zero. This suggests that one removes the division by the act's own degree of easiness from equation (2), which yields the following revision:[16]

$$(3) \qquad W(x) = \max\left\{\left(V(y) - V(x)\right) E(y) \,\middle|\, y \in \Omega\right\}.$$

In other words, you calculate the degree of wrongness of an alternative x by, for each alternative y, first subtracting $V(x)$ from $V(y)$ and then multiplying the result with y's degree of easiness. According to equation (3), the degree of wrongness of x is equal to the greatest of these products. This revised proposal satisfies the principle of consequentialist dominance. In situation 4, for example, it gives the more plausible result that $W(a_2) = 0.1$ and $W(a_3) = 10.1$. The theory also satisfies the principle of consequentialist supervenience. In situation 5, it yields that $W(a_2) = W(a_3) = 45$.

Nevertheless, another counter-example demands more thorough revisions. Compare the following situations, where situation 7 is just like situation 6 except for the addition of a_3:[17]

Situation 6

Act	Value	Easiness
a_1	1,000	0.5
a_2	0	0.5

Situation 7

Act	Value	Easiness
a_1	1,000	0.5
a_2	0	0.5
a_3	500	0.99

Each one of equations (2) and (3) yields that a_2 is equally wrong in situation 6 as in situation 7. Intuitively, a_2 seems more wrong in situation 7 than in situation 6. In situation 7, there is in addition to the right alternative, a_1, a further alternative, a_3, that has much better consequences than a_2 and which is also easier than a_1. Even though it is equally hard to do the right act in situation 6 as in situation 7, a_2 seems even more wrong in situation 7, because in that situation one could so easily have done something with better consequences than a_2.

In order to account for this, we need some slightly more complicated mathematics. My proposal can, however, be explained without mathematical notation. It has, as we shall see, a simple graphical interpretation. I propose that

[16] Like before, the claim here is just that this formula accounts for the intuitions we have appealed to so far and not that it is the only formula that does so.

[17] Tersman (1997, p. 52).

$$(4) \qquad W(x) = \int_{V(x)}^{\max\{V(u)|u\in\Omega\}} \max\left\{E(y) \,\middle|\, V(y) \geq z \wedge y \in \Omega\right\} dz.$$

To get an intuitive feel for the integral in equation (4) and also see how it handles situations 6 and 7, it will help to consider the following kind of graph. On the horizontal axis, we have increasingly greater values starting from the value of the consequences of the act whose wrongness is measured—in the graph below, a_2. And, on the vertical axis, we have the maximum degree of easiness with which one can perform an act whose consequences **[p. 114]** have a value at least as great as that on the horizontal axis in the situation—in the graph below, situation 6 (dashed) and situation 7 (solid).

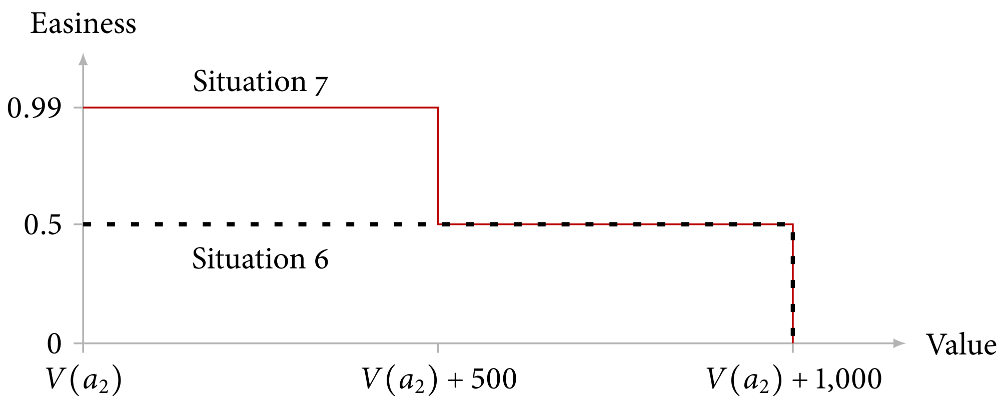

Given this kind of graph, my proposal can be simply stated as that the wrongness of an act is equal to the area under the curve. Hence my proposal yields that a_2 is more wrong in situation 7 than in situation 6. And with the above graph, it is easy to see why. For all z such that $z > 500$, it is equally hard in situation 6 as in situation 7 to perform an act whose consequences have a value at least z units higher than the value of those of a_2. But, for all z such that $0 > z \geq 500$, it is easier in situation 7 than in situation 6 to perform an act whose consequences have a value that is at least z units higher than the value of those of a_2. In each of situations 6 and 7, one has an obligation to perform an act whose consequences have a value of at least 1,000, and this obligation is equally hard to fulfil in either situation. But, in these situations, one also has an obligation to perform an act whose consequences have a value of at least 500, and this obligation is easier to fulfil in situation 7 than in situation 6. It is the violation of this latter obligation that makes a_2 more wrong in situation 7 than in situation 6.

To calculate the degree of wrongness of an alternative x according to equation (4), check first if there is an alternative with better consequences than x. If there is no alternative with better consequences, x's degree of wrongness is 0. Otherwise, let y_0 be an alternative with better consequences than x such that there is no alternative easier than y_0 with better consequences than x. Then make a note of the product of $E(y_0)$ and the difference between $V(y_0)$ and $V(x)$. Then starting with $n = 0$, you (*) check if there is an alternative with better consequences than y_n. If there is no alternative with better consequences than y_n, then x's degree of wrongness is equal to the last noted amount. Otherwise, let y_{n+1} be an alternative with better consequences than y_n such that there is no alternative easier than y_{n+1} with better consequences than y_n. Then calculate the product of $E(y_{n+1})$ and the difference between $V(y_{n+1})$ and $V(y_n)$. Make a note of the sum of this product and the last noted amount. Increase n by one, and go back to step (*).

Like equation (3), my proposal satisfies all three of our consequentialist principles.[18] And it yields the same results as (3) in situations 1–5. It does in fact yield the same results as (3) in all situations where there are just two acts. And, as with (3), it is easily seen that the addition of a dominated act—that is, dominated in the sense that an already available act is at least as easy with at least as good consequences—does not change the wrongness of the already available acts. But, unlike (3), my proposal yields conversely **[p. 115]** that the addition of a non-dominated act will always make acts with worse consequences more wrong—which is exemplified by situations 6 and 7.

[18] To see that (4) satisfies the principle of consequentialist supervenience, note that the integral in (4) is the same for the wrongness of all acts in the same situation except for the lower bound of the domain of integration, which is equal to the value of the act's consequences. Hence acts with equally good consequences in the same situation are equally wrong.

To see that (4) satisfies the principle of consequentialist dominance, note that, if one act has better consequences than another act in the same situation, the domain of integration in (4) for the wrongness of the act with the better consequences is a proper subset of the one for the wrongness of the other act. And note also that the integral in (4) is an integral of a non-negative function, since degrees of easiness are non-negative. It follows that the integral for the wrongness of the act with the worse consequences must be at least as great as the one for the wrongness of the other act. Hence the act with the better consequences is not more wrong than the other act.

Finally, to see that (4) satisfies the principle of consequentialist continuity, note that the integral in (4) only depends on the easiness and the value of the consequences of acts with better consequences than x. Hence an improvement of the consequences of x will only change the lower bound of the domain of integration of the integral for the wrongness of x. Since the integrated function is bounded, a sufficiently small improvement of the consequences of x will thus result in an arbitrarily small change in the wrongness of x, and a smaller improvement cannot result in a greater change in the wrongness of x. And note also that, given (4), improving the consequences of an act can increase the wrongness of another act in the same situation by at most the product of the size of the improvement and the improved act's degree of easiness. It follows that, for any difference in wrongness δ greater than zero, there is an arbitrarily small improvement of the consequences of x which would not change the wrongness of any act in the situation by more than δ.

One might object that there is no moral excuse for doing a_2 in either of situations 6 and 7 because one could have done something with better consequences with *no more* difficulty. But, while it might appear intuitive that there is especially little moral excuse for doing a_2 as it is no more difficult than doing a_1, we cannot plausibly take the relative easiness of acts into account when we assess their wrongness. To see this, consider again situation 5. In this situation, it might appear that what makes a_3 especially wrong is that one could instead do a_1 with no more difficulty. And, if that is what makes a_3 especially wrong, it appears that a_3 should be more wrong than a_2, since a_2 has equally bad consequences as a_3 but is at least easier than a_1. Yet a_3 cannot be more wrong than a_2 according to the principle of consequentialist supervenience. And that principle is, as I have argued, supported by the principles of consequentialist dominance and consequentialist continuity, both of which are hard to deny.

I wish to thank Gustaf Arrhenius, Richard Yetter Chappell, Björn Eriksson, Daniel Ramöller, Nicolas Espinoza, Marc Fleurbaey, Christopher Jay, Jesper Jerkert, Martin Peterson, Christian Piller, Mozaffar Qizilbash, Tor Sandqvist, Folke Tersman, Fredrik Viklund, an anonymous referee, and the audiences at the Philosophy Research seminar, Royal Institute of Technology, 27 August 2013, and at the Practical-Philosophy-Group Seminar, University of York, 28 October 2015.

References

Benson, Paul (1987) 'Ordinary Ability and Free Action', *Canadian Journal of Philosophy* 17 (2): 307–335.

Bentham, Jeremy (1970) *An Introduction to the Principles of Morals and Legislation*, The Collected Works of Jeremy Bentham, London: Athlone.

Berlin, Isaiah (1958) *Two Concepts of Liberty*, Oxford: Clarendon Press.

Bykvist, Krister (2002) 'Alternative Actions and the Spirit of Consequentialism', *Philosophical Studies* 107 (1): 45–68.

— (2003) 'Normative Supervenience and Consequentialism', *Utilitas* 15 (1): 27–49.

— (2011) 'How to Do Wrong Knowingly and Get Away with It', in Rysiek Sliwinski and Frans Svensson, eds., *Neither/Nor: Philosophical Papers Dedicated to Erik Carlson on the Occasion of His Fiftieth Birthday*, pp. 31–47, Uppsala: Uppsala University.

Calder, Todd (2005) 'Kant and Degrees of Wrongness', *The Journal of Value Inquiry* 39 (2): 229–244.

Cohen, G. A. (1978) *Karl Marx's Theory of History*, Princeton: Princeton University Press.

Eriksson, Björn (1994) *Heavy Duty*, Stockholm: Almqvist & Wiksell.

— (1997) 'Utilitarianism for Sinners', *American Philosophical Quarterly* 34 (2): 213–228.

Graham, Peter A. (2010) 'In Defense of Objectivism about Moral Obligation', *Ethics* 121 (1): 88–115.

Mackie, J. L. (1977) *Ethics: Inventing Right and Wrong*, Harmondsworth: Penguin.

Mill, John Stuart (1969) *Essays on Ethics, Religion and Society*, vol. X of *Collected Works*, Toronto: University of Toronto Press.

Parfit, Derek (1984) *Reasons and Persons*, Oxford: Clarendon Press.

Portmore, Douglas W. (2007) 'Welfare, Achivement, and Self-Sacrifice', *Journal of Ethics and Social Philosophy* 2 (2): 1–28.

Regan, Donald (1980) *Utilitarianism and Co-Operation*, Oxford: Clarendon Press.

Tersman, Folke (1997) 'Om Erikssons grader av synd', *Filosofisk tidskrift* 18 (3): 50–53.