



4-25-2016

Effort and Ethics

Edward J. Lee
Ursinus College, lee.j.edward@gmail.com

Adviser: Kelly Sorensen

Follow this and additional works at: http://digitalcommons.ursinus.edu/phil_hon



Part of the [Ethics and Political Philosophy Commons](#)

Recommended Citation

Lee, Edward J., "Effort and Ethics" (2016). *Philosophy Honors Papers*. Paper 1.

This Paper is brought to you for free and open access by the Student Research at Digital Commons @ Ursinus College. It has been accepted for inclusion in Philosophy Honors Papers by an authorized administrator of Digital Commons @ Ursinus College. For more information, please contact aprock@ursinus.edu.

Effort and Ethics

Edward Lee

04/25/2016

Submitted to the Faculty of Ursinus College in fulfillment of the requirements for Honors in the Philosophy Department.

Abstract

In this paper, I start by describing two categories into which effort can be divided. The first category I call objective effort, which refers to the external factors associated with an action; the second I call subjective effort, which deals with the internal psychological or mental effort required to bring oneself to perform an action. I then track my intuitions about how objective effort and subjective effort are relevant to the moral evaluations of actions and agents, and further develop this descriptive account by explaining these intuitions normatively through the lenses of contractualism and virtue ethics. Having given my account of how effort plays a role in ethics, I show how understanding the role of effort can change the way we look at the distinction between doing harm and allowing harm. I do this by describing how actions and omissions can be understood through effort, giving my account of the distinction between doing harm and allowing harm, and then comparing my account with others' accounts. Finally, I touch on the distinction between intending harm and foreseeing harm to show that thinking about effort is relevant not only to the particular example of the dialogue surrounding the distinction between doing harm and allowing harm, but also to other discussions as well.

It's Always Sunny In Philadelphia is a TV sitcom about “The Gang”, a group of friends who own a bar named Paddy’s Pub in Philadelphia. The friends are incredibly selfish and seem to have no moral values at all, and episodes often revolve around their crooked plans like the time two friends fake being a recovering crack addict and being mentally disabled in order to obtain welfare. But despite the general vice and debauchery the friends are used to, in Season 10, Episode 9 (“Frank Retires”), even these friends are shocked when one of them, Dennis, reveals one way he cheated Frank, another friend and the man who raised Dennis, out of money for years:

“All right, look, guys, it's very simple. Seven years ago, I walked into a Wawa, and I saw this guy and I thought, ‘Holy shit! That guy looks just like Frank, but, you know, Mexican.’ And then I came up with a plan where I could make Frank believe that he had a son with our maid Josefina all those years ago and I could bilk him out of a shitload of money under the guise that the money was all going to his firstborn son, Franquito, when, in fact, the money was going to me.”

This seems wrong.

There are, of course, many, many reasons why this plan seems wrong. We can explain the wrongness of this plan by talking about the financial harm it caused the victim involved, Frank. We can also reference the particular heinous nature of tricking a person about one’s children. (Interestingly, Frank was also previously in the TV series falsely led to believe that Dennis was his biological son.) Others of us might frame the discussion around the rights such as the right to property that were violated. But one factor that seems to me to go ignored often in discussions of ethical behavior is the factor that I want to focus on here (and for the remainder of the paper) – effort. Dennis clearly put in a lot of effort into this plan, and I think this fact is

particularly shocking to us. This is a plan that required deception that could not have been easy to pull off, recruitment of other people (namely Franquito) that would also have taken a lot of energy, and years and years of upkeep of this façade. All of this would have taken an incredible amount of effort and there is something quite deplorable specifically about the effort involved that makes this plan so much worse than your average case of fraud. But despite how much effort Dennis would have had to put in to carry out this plan, Dennis at the same time introduces this plan as “very simple” and tells the story in a quite casual manner. His offhand manner of telling the story of his plan reveals to us that while Dennis was willing to put in a lot of effort to carry out this plan, it was also, in a way, quite *easy* for him to carry out this *hard* plan. He did not have to put in the effort that most of us would have had to expend in order to bring ourselves to do something so horrible.

Effort itself seems to play an important role in the way we look at actions and people. This importance of effort will be my focus throughout this paper. I begin by describing effort on its own and then observing the way in which effort seems to affect the way in which we evaluate an action and an agent. Next, I try to explain why effort seems to matter in the way I observe. Having attempted to give an explanation for why effort matters, I then look at how effort matters in other matters like the difference between an action and an omission and the difference between doing harm and allowing harm.

Objective and Subjective Effort

It seems to me that there are two main categories of effort when we talk about the difficulty involved in doing something. I will call the two categories objective and subjective effort.

Objective effort includes things like money spent, time spent planning, time spent acting, or physical exertion – what we typically mean when we use the word effort conversationally. I call this effort objective because it is the effort involved in doing an action that anyone with certain relevant characteristics would have to expend in order to do that action. I specify “certain relevant characteristics” because, while I use the word objective to point out that objective effort holds constant across a certain group of people, it does not hold constant for absolutely everyone. In other words, if an action required an hour of time to perform, I am calling this hour of time objective effort not because an hour is the same for absolutely everyone including a person who works fourteen hours a day and a person who works on their own flexible schedule; I am calling the hour of time objective effort because if we hold certain relevant characteristics constant, which in this case would be work and commitment schedules, we would see that an hour would be the same for people with similar work and commitment schedules.

Now if we were to ask who spends more objective effort, a person who spends an hour of time to do an action compared to another person with a similarly busy schedule who spends two hours of time to do an action, the answer would be easy because they share the same relevant characteristics: the person who spends two hours of time. But we also may have to compare the amount of objective effort spent between people who do not share the same characteristics. For example, we might ask whether a person who spends an hour of time to do an action expends more objective effort or another person who is much more busy who spends an hour of time to do an action expends more effort. In this case, even though both people spent an hour doing an action, the busier person spends more objective effort because it requires more from the busier person to spend an hour than it does the less busy person. In situations like this when we compare people who do not share the same characteristics, we would have to come up with some

sort of (rough) conversion scale to convert objective effort for one person with certain characteristics to objective effort for another person with other characteristics. For example, \$1,000 is not the same for people who make \$12,000 a year and for Tony Stark, a.k.a. Iron Man, but we might be able to come up with some sort of progressive scale using percentage of income and living expenses where $X\%$ of income for the \$12,000-a-year group and $Y\%$ for the Tony Starks of the world could be considered to have a similar amount of impact for both groups. A similar strategy might be used for the other examples of objective effort I mentioned above like physical exertion or time spent: lifting X pounds would require the same amount of objective effort for people with a certain weight and strength training level as lifting Y pounds would require people of another weight and strength training level, and X hours would require the same amount of objective effort for people with a certain schedule as Y hours would require people with another schedule. Regardless of whatever these X 's and Y 's end up being, however, the main point to keep in mind is that when we say that someone spends more objective effort than another person (or I say in this paper), this sort of conversion should be implicit in that statement.¹

If objective effort is that effort that holds constant across people with certain characteristics relevant to the action, subjective effort is the effort that varies even for people with the same relevant characteristics. Subjective effort is the psychological and mental effort associated with an action, or how hard it is for an individual to carry out an action given a certain amount of objective effort. For example, if my mom asked my sister and me to clean our rooms, the objective effort involved in doing so would be roughly the same for the both of us. The size

¹ Of course, I don't want to pretend that coming up with such a conversion scale would be easy at all, but I am merely suggesting that such a scale is conceptually possible.

of our rooms and the spatial density of objects that need to be thrown out or organized are roughly the same, and my sister's characteristics and my characteristics relevant to cleaning our rooms are the same: we are both strong enough that cleaning our rooms requires about the same physical exertion from both of us and the time spent cleaning our rooms would have the same impact on both of us as we are both students with similar schedules. However, the subjective effort that my sister needed to exert in order to perform the act of cleaning her room is higher than mine, so my room was consistently cleaner than hers, even though in terms of objective effort, it seems like there was no reason why I seemed to clean my room more often than my sister cleaned hers.

Distinguishing these two categories of effort is important because they are important in different ways when it comes to moral evaluation.

Let's first look at the way in which subjective effort plays a role in moral evaluation. Consider two students who vandalize school property. The first, A, is a student who does not think twice before vandalizing school property. B, on the other hand, hesitates and finds it harder to bring himself to do so, but eventually does it anyways. Both students appear to have done an action that is equally bad, and it would be strange for the school to mete out different punishments for the two, assuming that they have similar disciplinary histories with the school. Yet, we are still inclined to evaluate A more negatively than B. In other words, a person who expends less subjective effort to commit a harmful action is evaluated more negatively than a person who needs to expend more subjective effort to commit the same action.

The reverse seems to be true for good acts, which seems intuitive. For example, consider two people of the same economic status who each donates \$10 to a charity. The first, A, finds it

very easy to donate the \$10. On the other hand, the second, B, also donates the \$10, but she finds it difficult to do so and needs to spend additional subjective effort in order to bring herself to actually donate the money. In this situation, the two people perform equally good actions that will presumably result in the same amount of good. But while we do not look at A's action and say that she has performed an action that is better than B's action, we do look at the two people differently. A appears to be the type of person who finds it easy to donate, and a better person who generally finds it easy to help others. B, however, does not seem as good of a person as A because of how much harder it was for B to donate. In other words, the lower the amount of subjective effort it takes for a person to do a good act, the more positively the person is evaluated.

One might point out that there seem to be cases that contradict my observation that we evaluate people more positively when they only need to spend a small amount of subjective effort to do a good action. Consider Nigel², a person who has difficulties donating to charity not only because of normal selfish desires but also because of a near pathological phobia of donating. Despite this, Nigel expends a high amount of mental effort to work through his issues and eventually brings himself to donate \$10. Indeed, there seems to be something that is praiseworthy about Nigel here for spending a high amount of subjective effort to do a good action. However, I do not think this case necessarily has to contradict what I said earlier, which is that an agent who only needs to spend a small amount of subjective effort in order to do a good deed is evaluated more positively. This is because there seem to be two ways to look at (or set up the case of) Nigel. In the first way, the enormous difficulty it takes for Nigel to donate is closely tied to him as a person, to his moral character. So even though he is able to work through his issues to donate, there is still something wrong with Nigel as a person that he had to exert so

² Kelly Sorensen, "Effort and Moral Worth," *Ethic Theory Moral Prac* 13 (2010), pp. 89-109.

much subjective effort just to donate \$10. In this interpretation of Nigel, we in fact don't evaluate Nigel positively, and any such feeling reminiscent of positive evaluation that we have is really only positive in the context and to the standards of Nigel as a bad person who is trying not to be as bad as he is. In the second interpretation, the enormous difficulty it takes for Nigel to donate is truly pathological – a genuine illness outside of and separate from Nigel's person, his moral character. In this case, it seems that we could categorize working through his illness as a factor of objective effort. So just like how we would say that a person with a busy schedule who spends an hour to do an action spends more objective effort than a person with a relaxed schedule who also spends an hour, we should say that Nigel's donating requires more objective effort than person A or person B's donating did. In this case, the \$10 donation by Nigel, a person with this mental illness, does make Nigel a more admirable person because of a difference in objective effort, which is not related to my observation about subjective effort. Thus, I don't think Nigel's case or similar cases are contradictions to the general trend that when a person does not need to expend much subjective effort to do a good action, we evaluate him more positively.

It may be tempting to use objective effort as an agent-evaluating tool as well – after all, a person who volunteers for three hours may seem more praiseworthy than a person with a similar schedule who spends less objective effort and volunteers for one hour. But I believe that this temptation should be resisted. Though one person may have done something better (volunteering for longer), the person may have had to exert a high amount of subjective effort in order to bring himself to do that action and had to really force himself to do something good. The other person who volunteered for one hour, on the other hand, might have hardly expended any subjective effort at all, and perhaps would even have volunteered for three hours without

hesitation and with no more subjective effort if given the chance. Thus, knowing subjective effort provides critical information about the person acting that objective effort does not provide.

Objective effort is useful, however, in helping us evaluate an action on its own, apart from who performs it and without considering how much subjective effort it took, specifically because it involves characteristics of an action that are objective and consistent across different people. Furthermore, its usefulness in act evaluation also seems to make it a useful factor to consider in cases of decision making when one is choosing to do or not do an action because when a person decides to do or not do something, he considers factors having to do with the action itself (like objective effort). Subjective effort, on the other hand, does not seem quite as useful. If a person finds that he must expend a lot of subjective effort in order to do good, and as a result discovers he is less admirable than someone who does not need to spend that much subjective effort, this is no reason for him to not do a good action, which is still good regardless of his lessened admirableness. Or if a person finds that he would have to spend a lot of subjective effort in order to do harm and is thus more admirable than someone who does not spend much subjective effort at all doing harm, this person's increased admirableness would not somehow make doing harm less bad. Such considerations about subjective effort and its implications on one's character seem like they should be made separately from the considerations about the characteristics of an action (such as the objective effort involved in performing the action) that are relevant to deciding on whether to perform an action or not. So objective effort seems relevant in act evaluation and decision making in a way that subjective effort is not.

Objective Effort and Act Evaluation

To see how objective effort affects how we think about the moral evaluation of an action, let us consider cases where the same person carries out two different actions. Compare a case in which a person trips someone walking past them in the hallway by sticking her foot out and a case in which the same person instead sees someone walking in the hallway on the other side of the school, runs over, and sticks her foot out to trip the person. Though both actions result in the tripping of another person, the second action, which requires additional time and physical exertion to complete the action, seems worse than the first. Thus, it seems that the amount of objective effort expended on its own can make the action of harming seem worse and magnify the wrongness of an action.

In a similar manner, more objective effort seems to magnify the goodness of a helping or benefiting act. For example, a person might donate blood at a local blood drive that is conveniently located at her place of work, or she may instead be in a situation in which she does not have the same convenience so she instead finds a place where she could donate blood, makes her way to that place, and then donates the same amount of blood. Though both actions help people equally, the second act seems to be the better act because of the additional effort involved. There is something about effort itself that plays a role in the way we look at these actions despite holding other aspects of the actions, such as the consequences, constant.

Some might point out certain cases that seem to contradict this magnification effect in which more objective effort makes a harmful action seem worse or a good action better. Consider a person, A, who casually kills others – it seems like there might be something more sinister going on there than that with someone else, B, who plans out a murder. Especially in

movies and TV, people like A seem even scarier than B. Gus Fring in the TV series *Breaking Bad* seems like a perfect example of someone like A. Gus Fring is a man who is polite, professional, calm, and skilled at his job as drug kingpin. Despite his calm demeanor, in one of the most memorable episodes of the series, Fring quietly changes into hazmat gear, picks up a box cutter, and effortlessly slices the throat of his right-hand man for seemingly no good reason. These sorts of actions cement Fring as one of the most villainous characters in the series despite being portrayed as a quiet, calm person. However, I think this sort of contradiction to my observation that high objective effort makes a harmful action seem worse is one that is based on muddling the distinction between objective and subjective effort. When we look at people like A and Gus Fring, we are shocked because this person commits such a harmful act with such little *subjective* effort. If A or Gus Fring then proceeded to commit another similarly horrible murder that was much more planned and effortful, that second murder would be even worse than the first murder because of the additional *objective* effort; this would conform with my observation that more objective effort magnifies the wrongness of a harming action. Furthermore, if B, after carrying a planned out murder then also killed someone casually like A did, of course we would also be equally shocked by B as we were with A because of how little subjective effort B spends to kill. But B was initially only presented as the person who spent more *objective* effort, and as a result, the higher objective effort was muddled with what we assumed was her level of *subjective* effort. In other words, we were initially more shocked by A because we assumed B's planning meant she was not as flippant about murder as A, but in reality, B's planning and objective effort said nothing about B as a person and her subjective effort. So when setting up cases to help us understand objective effort, it is important to hold subjective effort constant by not using two different people (like A and B) and ensure that our evaluation and visceral reaction to subjective

effort does not cloud our understanding of objective effort.³ Taking this into account and looking again at the actions of casual and planned murder with the *same* person, we can better understand how objective effort matters. So instead of comparing a casual murder by A and a planned murder by B, let's compare a casual murder by A to a planned murder by A. It seems to me that the murder that A does think about, plan, and try harder to execute seems like the worse action than the one where she does not, even if the consequence to the victim is held constant. This again follows the trend in which more objective effort magnifies the wrongness of a harming action.

Keeping the person the same in these cases like this is important in order to hold other characteristics of the person constant as well. While my focus is on effort, consequences, and actions, I do not want to pretend that other factors do not play an important role in the way we evaluate agents and actions. These factors might include the reasons for which a person performs an action or the intentions behind an action⁴, which are unique to the individual performing the actions. But because objective effort is my focus here, it would be ideal to hold constant the person performing the action while only changing the objective effort involved in an action. In this way, other potentially conflating factors affecting our intuitions of moral evaluation can be controlled for while focusing on objective effort.

So generally it seems that more objective effort magnifies the goodness or wrongness of an action. However, the way objective effort has this magnification effect is not symmetrical

³ This is why when I set up the subjective cases earlier, I compared two different people (e.g. "Person A" and "Person B") while in setting up the objective cases I used one person (e.g. "the same person").

⁴ Some might say that the *only* factor that matters is the agent's intentions behind an action and cannot see how objective effort matters at all. I do not think addressing this objection fully is possible here. However, I do hope that, first, my cases above comparing actions performed with different levels of objective effort seem at least somewhat plausible even for people who do not see how actions could be considered on their own, and, second, that we might find some common ground in considering how subjective effort plays a role in agent evaluation.

between helping and harming actions. A harming action seems bad whether it requires a low amount of effort or a high amount of effort, though more effort magnifies how bad such an action actually is. Tripping someone next to me requires only a small amount of objective effort (sticking my foot out), and it seems wrong, though not as wrong as running across a school campus to trip someone.

However, helping or benefiting actions do not always seem to be good in the same way that harming actions always seem bad. If I walk through a door and hold the door for the person right behind me who is on crutches, this would not be praiseworthy – it is merely neutral. If I instead noticed someone some distance away who is on crutches and stay at the door for some time so that I could hold the door open, this seems somewhat praiseworthy. I need to expend more effort than just sticking my arm out to hold the door for a second before my helping action of holding the door open can be called good or praiseworthy. This is different from harming actions. Even if I only expended a small amount of effort, the harming action of tripping immediately seems wrong. I don't need to expend a certain amount of effort before it becomes wrong like I would have to in order for holding the door to become good. This difference between the way in which objective effort affects helping and harming actions might be even clearer with a more extreme case. If a child starts to drown in a kiddie pool in someone's lawn as I am walking by on the sidewalk and all I need to do is take one step and upright the child, doing so is not heroic at all. It is neutral and expected. If the child were instead in the middle of an extremely wide river and I jump in, swim the tremendous distance despite my burning, aching muscles, and drag the child back to land, this action seems to me to be praiseworthy. So it seems that a certain amount of objective effort is necessary to expend before even saving a life becomes praiseworthy. Contrast this with how murder is immediately wrong regardless of the amount of

effort. If I pressed a button that would kill someone, even this low amount of effort would immediately be wrong. It might be less wrong than if I instead spent more objective effort and constructed a device that would kill someone with a push of a button and then pressed the button, but it is nonetheless still wrong. Thus, a helping action does not seem to actually become good until the action requires a certain amount of effort; this is different than with harming actions – a low effort harming is not neutral in the same way a low effort helping is (Figure 1).⁵




Figure 1. The asymmetric effect of objective effort on the praiseworthiness or blameworthiness on harming and helping actions.⁶

⁵ Some will share my intuitions here; some may not. Below, I offer some explanations of my intuitions that I believe would satisfy both those who share my intuitions and those who do not because I think these explanations can stand on their own apart from what one's intuitions actually are.

⁶ Of course, I do not want to pretend this graph shows anything exact or that the slopes on the helping effort and harming effort side should be the same. I am only trying to show what I mean by the asymmetry.

A similar asymmetry is seen with cases of not harming and not helping or benefiting (Figure 2). Not harming seems to be mostly neutral, regardless of whether the harm that was not done would have been low effort or high effort. If I don't trip someone right next to me, the omission of my tripping is simply neutral. If I don't trip someone across school campus, again, my omission is neutral. So whether or not the harming action would have required only a little amount of objective effort or a lot does not matter, and the omission of the harming action is neutral regardless. However, effort does seem to matter with omissions of helping or benefiting actions. Not doing a high effort helping or benefiting seems neutral. But not doing a low effort helping or benefiting seems like a blameworthy omission. For example, I might not hold the door for someone on crutches who is somewhat far off. It might be nice if I stayed there and held the door, but not doing so seems neutral. If the person were right behind me, however, and I let the door slam right behind me and I did not even expend the small amount of effort to hold the door for a second, this instead seems wrong. Additionally, it seems to me that the blameworthiness of not expending a small amount of effort to help someone increases "asymptotically" in that the blameworthiness seems to increase by a greater amount than the amount by which the omitted objective effort decreases. For example, the blameworthiness of not waiting one second to hold a door for someone seems more than "five times" worse than the blameworthiness of not waiting five seconds to hold a door for someone. These intuitions about the blameworthiness of not spending a small amount of effort to help someone and the neutrality of not spending a large amount of effort to help someone and of not spending effort to harm are shown in Figure 2.

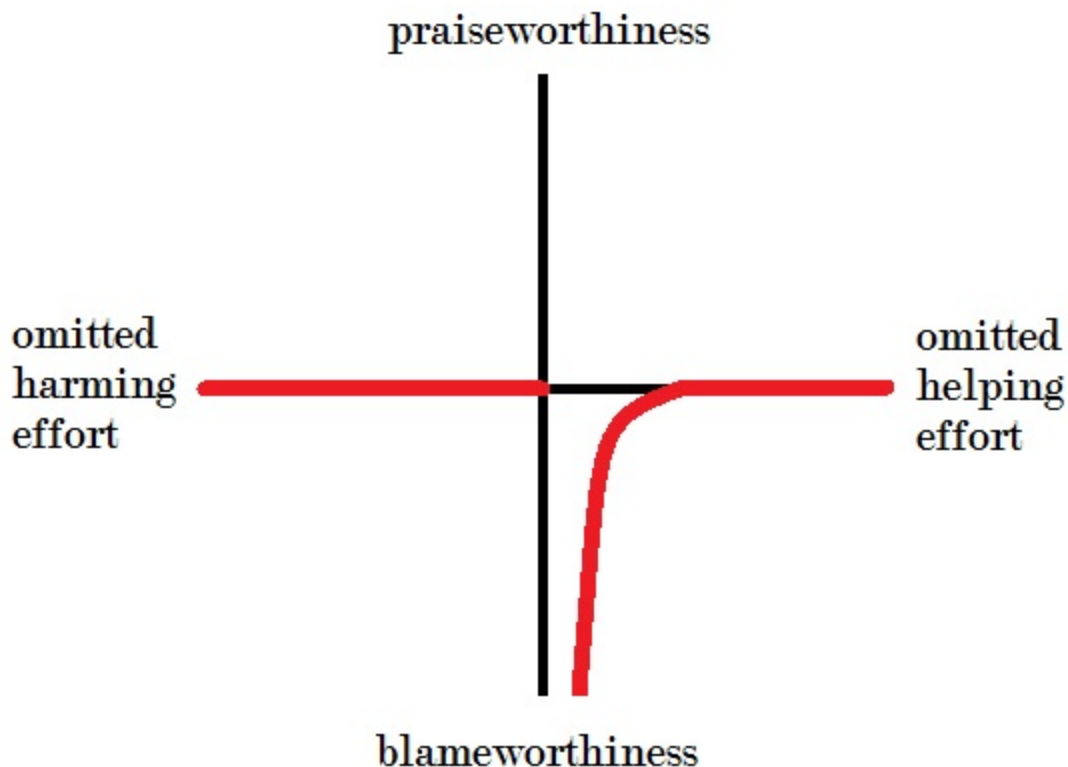


Figure 2. Objective effort and the praiseworthiness or blameworthiness of omitted actions.

Much of what I have been doing until now was describing what our (or at least my) intuitions are about the types of effort there are and how effort matters when it comes to moral evaluation. I will now try to explain *why* I believe effort matters and explain where our intuitions come from.

Why Does Objective Effort Matter?

It seems to me that the importance of effort is better explained in a non-utilitarian context than a utilitarian one. Consider the form of contractualism set forth by T. M. Scanlon. In this form of contractualism, “an act is wrong if its performance under the circumstances would be disallowed by any set of principles for the general regulation of behaviour that no one could

reasonably reject as a basis for informed, unforced, general agreement.”⁷ Contractualism differs from other non-utilitarian theories in several ways. Unlike other related theories like contractarianism which is more focused on how those engaged in the social contract would maximize their self-interest, contractualism asks how individuals who respected each other as moral equals would justify the pursuit of their interests to others. In other words, an agent under contractualism justifies himself to others because others are also equal moral agents, while an agent under contractarianism justifies himself to others because of the strategic value of working with others in order to maximize his own self-interest. Contractualism also differs from the moral theory of Kant and Rawls in that contractualism asks what principles no one would be able to reject, as opposed to principles that all would agree to, which Kant and Rawls focus on.⁸

As a contractualist, one would ask what kind of burden a certain principle places on him and how that burden compares to the burdens imposed on others by that principle. If a burden placed on him is too great or if a burden is distributed unfairly based on morally arbitrary criteria, he might reasonably reject that principle. However, if the burden is great but alternatives to the principle result in even greater burdens on others or give others more pressing reasons to reject a principle, the principle might not be rejected. Using this kind of reasoning and by paying attention to the strengths of the reasons those affected by a principle have to reject it, a contractualist would assess the acceptability of a principle. Here, there seems to be space for effort to pay a role. If a principle requires an excessive level of effort to help someone, it can be reasonably rejected by the person on whom the burden of effort is placed. However, a principle might also demand a reasonable level of effort from someone to help others, and cannot

⁷ T. M. Scanlon, *What We Owe to Each Other* (Cambridge: Cambridge University Press, 1998).

⁸ Ashford, Elizabeth and Mulgan, Tim, "Contractualism", *The Stanford Encyclopedia of Philosophy* (Fall 2012 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/fall2012/entries/contractualism/>>.

reasonably be rejected. Therefore, actions of this level of effort that follow the principle are simply what are required because they abide by this principle, while omitting such actions are a violation of the principle. Thus, the change from required to not required as the effort of a helping action increases can be explained by what levels of effort can be reasonably demanded of someone in a moral framework in which moral equals consider the reasons for which others might reasonably reject a moral principle, and whether the effort exerted for the helping action merely satisfies an acceptable principle or goes above and beyond what can be demanded or expected.

But while I believe that part of the importance of objective effort can be explained by contractualism, I do not think that this explanation is complete on its own. Contractualism explains whether an action is right or wrong based on whether or not it violates a set of principles that no one could reasonably reject. But while it explains right and wrong, it does not seem to say much about praiseworthiness. If a helping action has a level of effort associated with it that goes above and beyond what a non-rejectable principle demands, the action is simply categorized as not wrong. I presume that a contractualist would see this action as one that is praiseworthy as well, but it is hard to guess how a contractualist might explain why this would be without a theory of praiseworthiness and blameworthiness distinct from a contractualist account of right and wrong. Can an action not violate a principle more than another action also does not violate the principle? Such a question seems nonsensical. Thus, as effort increases, rather than transitioning from required to praiseworthy as I have been describing, with praiseworthiness scaling up as effort scales up, contractualism seems to simply lump together lower effort helping actions that barely surpass what a principle demands and higher effort ones under the category of

justifiable. Furthermore, a similar issue arises when considering contractualism and blameworthiness. As the effort of a harming action increases, the blameworthiness seems to increase as well. So running across the school to trip someone seems like a more blameworthy action than tripping someone next to me in the hall. But does it make sense to ask whether I violated the principle that requires me to not trip people more or violated it harder in the first case than in the second? Because these principles deal with whether an action is justifiable or not, it seems easy to give importance to effort in affecting what actions can be demanded or prevented by a principle, but harder to explain the spectrum of blameworthiness and praiseworthiness that different levels of effort seems to create.

Therefore, rather than only looking to contractualism, I will also look to virtue ethics for an explanation of why effort seems to matter. In contractualism, right and wrong is determined by whether an action is allowed or not by principles that cannot reasonably be rejected. In utilitarianism, right and wrong is determined by whether an action produces a net increase or decrease in utility. Virtue ethics takes a different approach. Virtue ethics would look at an action and ask whether it is in line with what a virtuous agent would do. This virtuous agent is one who has and exercises the virtues, which are deep character traits that reveal the agent's grasp of *eudaimonia*. Although *eudaimonia* is difficult to describe, it can be thought of as happiness or flourishing in a deeper sense, happiness that is worth pursuing and flourishing that is unique for rational beings. A person living a *eudaimon* life pursues what is "truly worthwhile, truly important, and thereby truly advantageous in life"⁹.

With this understanding of virtue ethics, I believe that the importance of effort and helping actions can be explained. Even a person who is not very virtuous at all would prevent

⁹ Hursthouse, Rosalind, "Virtue Ethics", *The Stanford Encyclopedia of Philosophy* (Fall 2013 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/fall2013/entries/ethics-virtue/>>.

harm or aid someone if it required only a certain low level of effort. Therefore, a low effort helping action, which is an action that even a person who is just barely virtuous could be compelled to do, merely falls under the category of neutral; actions that are characteristic of a person who is hardly virtuous are not praiseworthy. Furthermore, when this low effort helping action is omitted, the omission is blameworthy – if a person who is hardly virtuous would perform a helping action that requires only a small level of effort, omitting that action would of course be typical of someone who is not virtuous at all, and instead someone who is vicious and lazy. But as the effort expended for the helping action increases, these actions become characteristic of what increasingly virtuous agents would do because a virtuous person pursues what is worthwhile and exerts effort doing that which the virtues compel him to do. Thus, helping actions that require more and more effort become more and more praiseworthy because more fully having the virtues and being able to enjoy *eudaimonia* is praiseworthy. These actions embody in a fuller sense what it means to be charitable, generous, and compassionate. But while such actions are praiseworthy, omissions of such actions would be more neutral because omitting such actions would be typical of an agent who, while clearly not highly virtuous, is not necessarily one that lacks any virtue at all or one that is vicious.

Besides the interplay between effort and helping actions, the interplay between effort and harming actions seem to be explained in the context of virtue ethics as well. Just as helping actions are characteristic of virtuous people, harming actions are typical of vicious people. A vicious person would be cruel, stingy, or selfish. The more vicious a person, the more effort such a person would devote towards harming actions. Therefore, increasing effort put into a harming action makes the action increasingly blameworthy because the action is increasingly characteristic of blameworthy vices.

One might notice, however, that virtue ethics on its own is also incomplete in its ability to explain why effort matters in the opposite way that contractualism was incomplete. Virtue ethics seems powerful in its ability to explain the spectrum of praiseworthiness and blameworthiness that arises from varying levels of effort and its ability to explain why subjective effort matters, aspects in which contractualism was inadequate. However, it is difficult to explain how a switch from impermissible to permissible could occur just using virtue ethics. With contractualism, this was easier to explain – whether an action is disallowed or required by a principle dictated by contractualism explained whether an action was impermissible or permissible. But virtue ethics seems to deal more with how much an action embodies a virtue or a vice without dealing with concepts of duty, requirement, permissibility and impermissibility. Without engagement with these concepts, it is not possible to explain why it seems that omitting a helping action that would have required only a low level of effort is impermissible while omitting a helping action that required a high level of effort is permissible. Consequently, it becomes harder to explain why certain do/allow cases appear to make doing and allowing seem equivalent while others make doing and allowing seem significantly different. Considering this weakness, it seems best to understand the importance of effort from multiple perspectives rather than just virtue ethics or just contractualism.

However, despite the differences in the strengths and weaknesses of the explanations from contractualism and virtue ethics for the moral importance of effort, what I believe to be the important, and perhaps in a way, obvious, common premise is that we do not have an unlimited amount of effort.

Imagine that there are two children drowning at two different points in time. One is the child who we are familiar with, the one who is right next to me drowning in a kiddie pool. The other is drowning in the middle of the Atlantic Ocean. Let's suppose that I do not save either. I do not bend over and exert the approximately 20 lbs. of force to upright the child drowning in the kiddie pool; nor do I make the plans to rent a helicopter, seek out a rescue team who would know how to pilot the helicopter and perform at-sea rescue missions, take off for what would presumably be an all-day mission, dive into the ocean, and then save the child. Here, it seems clear to me that the objective effort required to save the child in the kiddie pool is so ludicrously low that saving the kiddie pool is absolutely required of me, so I act impermissibly by not saving this child. But that I do not save the child in the Atlantic Ocean seems to me to be permissible because the objective effort required to do so seems so high. Of course, it would be nice if I did make the effort to save the child, but doing so would be praiseworthy, but not required. So the difference in level of objective effort causes a distinct change between failing to save one child and failing to save the other that makes one impermissible and the other permissible. We can explain how effort causes this change from impermissible to permissible using contractualism and virtue ethics. Saving the child in the kiddie pool is morally required because it violates the reasonable principle that all can agree upon that people should expend a small amount of effort when it means saving a life (contractualism), especially when even a person who minimally expresses the virtues would be able to find the compassion and charity to save the child in this scenario (virtue ethics).

But note that both parts of this explanation, the part from contractualism and from virtue ethics, depend on the fact that objective effort is not unlimited. What makes a principle that can govern our behavior reasonable? The reasonable-ness of the principle comes from the fact that it

does not demand too much of our limited reserve of effort. If we had an unlimited amount of effort and never grew weary of exerting effort, it would not make sense to say something requires “too much” of this unlimited resource and we would not need to even refer to what is considered reasonable. Principles would merely have to refer to what outcomes we should effect with no reference to how hard it would be to effect those outcomes. But we in fact do have a limited amount of effort, so a principle cannot demand too much effort if it is to not be reasonably rejected. Of course, this also makes it such that doing an action that requires a lot of effort is praiseworthy. What makes spending effort to help others virtuous? That it is virtuous to expend effort to help others, and that spending more and more effort makes one more and more virtuous, comes from the fact that, again, we have a limited amount of effort, so spending more and more of that limited amount towards helping others is characteristic of one who is compassionate, charitable.

To make this point clearer, let’s imagine there exists an omnipotent being. If she does not save either the child in the kiddie pool or the child in the middle of the Atlantic, both seem bad, but neither seems worse than the other. It seems to me that the being’s omitting to save the child in the kiddie pool and her omitting to save the child in the Atlantic Ocean are equally wrong because saving one or the other is the same to this being. Additionally, if she saves both, neither rescue seems praiseworthy. The difference in objective effort that would have been involved with the two rescues that was important for us is not important for this being. In fact, the word “effort” would not have any meaning for such a being – this being does not *try*, she simply does. So referencing contractualism or virtue theory at all would be incoherent for this being.

Of course, unlike this being, we as human beings are not omnipotent. We have limitations when it comes to strength, time, stamina, money, and effecting the outcomes we want in the universe. In other words, we do not have unlimited objective effort to spend. So we allocate objective effort throughout our days, weeks, years, and lives to the particular pursuits we deem to be worthy. And it turns out that we ought to allocate at least some of this valuable commodity for certain things – aiding others. Whether we explain why we ought to allocate this amount of effort through the lens of the principles that can reasonably govern our behavior or the virtues that would cause someone to behave in a certain way, the fact that this positive demand exists seems to me to hold true. Though we might permissibly not expend effort beyond what this positive demand requires because we presumably have other worthy pursuits that are deserving of our effort, when we fail to expend the amount of objective effort on which this positive demand has a claim, we have surely acted in an impermissible manner. Furthermore, if we were to go even further and deem harming others as a pursuit worthy of our limited effort, it seems clear that we have organized our lives in a deficient way and that we have failed to recognize what is worth spending our objective effort on, even if it is a small amount that we spend.¹⁰

The idea that we have a limited amount of effort to spend also says something about whether increasing objective effort spent on a helping action always increases praiseworthiness at a constant rate. Consider a person who agrees with my observation that increasing effort

¹⁰ It might seem strange that I keep referencing the person (“we”) performing the action here even though I previously said that objective effort is important for act evaluation and not agent evaluation. However I am not referencing one specific person or another and trying to evaluate a specific person; instead, I am referencing the general property of people that we cannot exert effort *ad infinitum*. Without referencing this general property, I cannot talk about effort at all. It is only in reference to beings that have limitations that effort matters, so I must refer to the being performing the action, at least at this point in the discussion. Once we establish this, we can move on and focus on the relationship between objective effort and actions because the additional step of relating actions to the general nature of beings with limitations for which morality applies (which I am doing here) is assumed.

generally seems to increase praiseworthiness. This person spends his entire day running around his office building holding doors for people, exerting a great amount of effort in doing so. He spends a great amount of physical energy and time, and even stays late after work to hold more doors for people. Holding the door like this seems to be a somewhat more praiseworthy action than holding the same number of doors just causally when the opportunity arises. But surely not by much, and even if it were more praiseworthy, there seems to be something strange about exerting so much effort to hold doors. So increasing the amount of objective effort spent to help others does not increase the praiseworthiness of an action at very high amounts to the same extent it does at lower amounts (Figure 3). This sort of tapering can be explained by the fact that we do not have unlimited effort to spend. A fully virtuous person would see the good in exerting some level of effort to hold doors for others, but because objective effort is not unlimited, he would distribute his limited effort accordingly across other pursuits, relationships, and helping actions throughout his day characteristic of a moral, *eudaimon* life. Therefore, the action of running around to hold doors increases in praiseworthiness with increasing effort at a fairly consistent rate until the point at which holding doors hampers a person's ability to distribute effort among other important pursuits.¹¹

¹¹ Another factor that could explain why there seems to be something off about the holding doors with a great deal of objective effort is that it seems that a person who does so *aims* at being praiseworthy. There does indeed seem to be something strange about trying so hard to be morally praiseworthy that makes the person trying ironically less praiseworthy. This may be due to the fact that such a person might wrongly be motivated disproportionately more by his own moral worth than the needs of others; or perhaps this person is misguided in setting an aim (high moral praiseworthiness) that is unrealistic or at least not achievable through the act of holding doors. (For more on this paradox, see Kelly Sorensen, "The Paradox of Moral Worth," *Journal of Philosophy* 9 (2004) pp. 465-483.) These and other reasons might even make my asking whether there is an upper limit to praiseworthiness seem strange because perhaps even asking the question is motivated by a wrongful focus on aiming directly at being praiseworthy. However, this factor of trying too hard to be morally praiseworthy seems to deal more closely with a specific type of intent that makes high objective effort door holding wrong. So while this factor is surely important, in my aim to deal more generally with the action of high objective effort door holding, I focus on the limited nature of objective effort and praiseworthiness rather than focusing on a factor that only applies if we know something specific about why a person does something and whether the person is in fact aiming at praiseworthiness in performing an action.

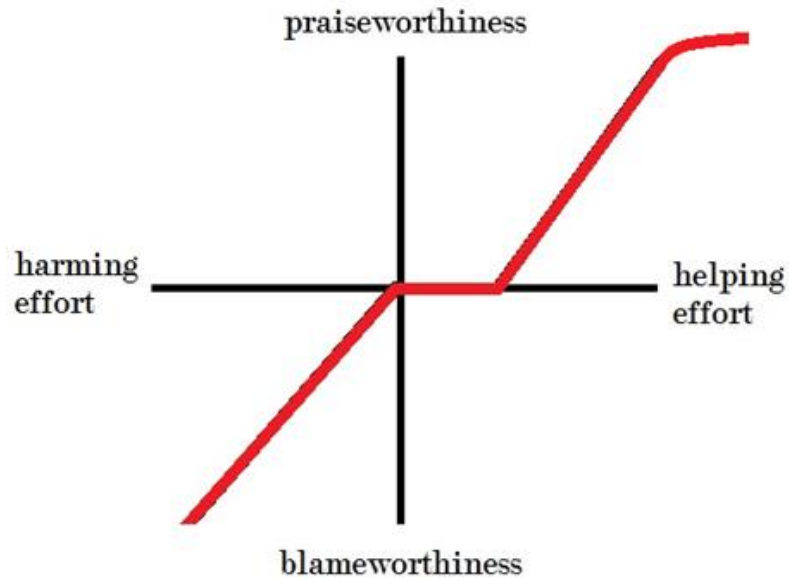


Figure 3. Diminishing returns on praiseworthiness as effort increases for a helping action.

At this point, I would like to refer back to this rough graph I presented earlier (Figure 4):

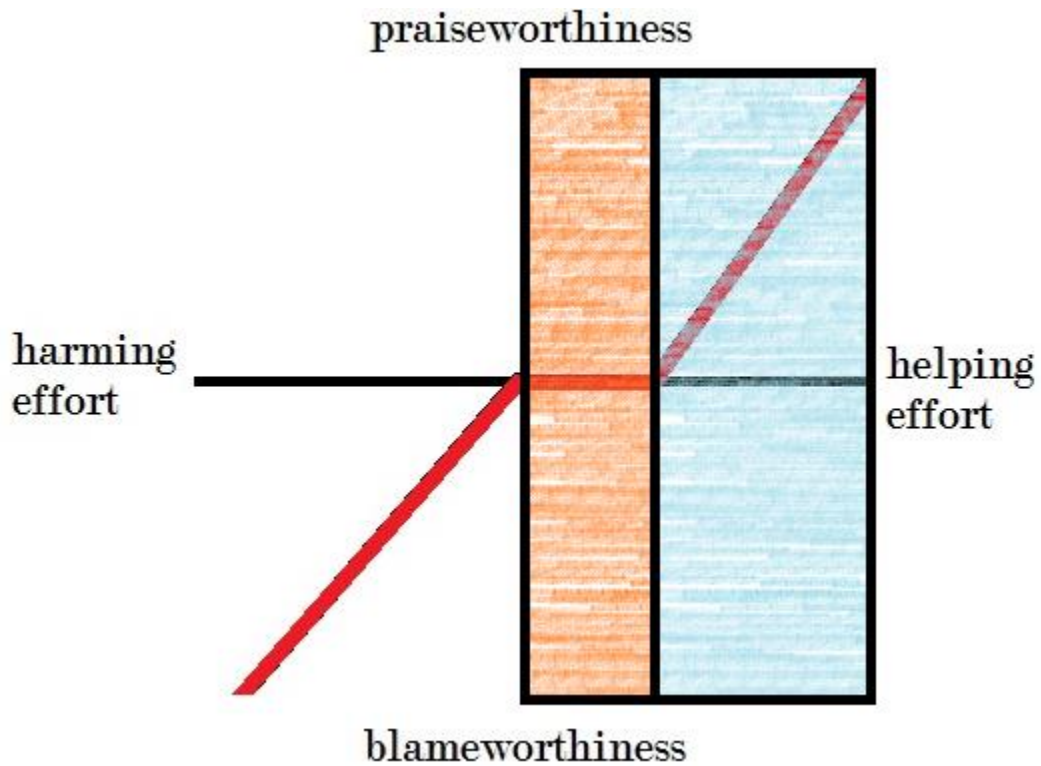


Figure 4. Modified Figure 1 with regions of required levels of effort highlighted in orange and permissible/praiseworthy levels of effort highlighted in blue.

I have boxed here in orange the area that I described earlier as the range of objective effort where effort merely makes an action neutral, and I have boxed here in blue the area where additional effort makes an action increasingly praiseworthy. If we were talking about holding the door, in orange would be holding the door for someone right behind me, and in blue would be waiting that extra time to hold the door for someone somewhat further away. If we were talking about saving drowning children, in orange would be the child in the kiddie pool, and in blue would be jumping into an extremely wide river or going out into the middle of the Atlantic. At the beginning of this paper, I only sought to describe our intuitions, but I hope by now I have given an explanation of why the orange seemed neutral and why the blue seems praiseworthy. The orange seemed neutral because this is the amount of effort that is required of us. The blue seems praiseworthy because this is the amount of effort that is permissible and praiseworthy. In other words, the dividing point between the two depicted boxes is the change from required to permissible that effort itself can cause in the status of an action. Why this shift exists was explained above by contractualism through the principles that govern our behavior morally that we cannot reasonably reject and the effort that such principles require us to expend to help others.

Having pointed out that this explanation depends on the fact that there are limitations to how much we can exert objective effort, I would now like to point out what one formulation of a contractualist principle might look like that would determine how low the effort associated with a helping action needs to be such that the associated action becomes required, instead of merely permissible. One might have noted earlier that when I said that we do not have an unlimited objective effort, I was a bit careless in saying this. Because although I surely cannot always be

jumping into rivers and the Atlantic as I would run out of effort, it seems that when it comes to a few certain actions that require a minimal amount of objective effort and are spaced apart from other actions, spending that level of objective effort actually seems unlimited. For example, wiggling a finger a few millimeters and blinking one's eyes once in a while are actions that require such little objective effort that, if one wanted, one could do indefinitely. So it seems that a certain low amount of effort does not tap into our limited "reserve" of effort but can be spent infinitely. It also seems that I could not reasonably reject a principle that says that I ought to expend part of something of which I have an infinite amount if doing so means helping others. Thus, a way to formulate the specific standard or principle by which a certain amount of objective effort can be required through contractualism would be to say that the amount of objective effort that can be exerted that is low enough and spaced apart far enough from other actions such that one can exert this effort indefinitely without growing tired is the amount of effort that can be required of someone to aid or help another.

Why Does Subjective Effort Matter?

Earlier, I referenced virtue ethics as a way to morally assess actions because understanding the virtues also provides a guide for action. But virtue ethics can assess both actions and agents because a moral theory based on virtues and vices is in a lot of ways, at its core, about morally assessing whether agents are virtuous or vicious. Therefore, virtue ethics seems like the optimal lens through which we can look at the importance of subjective effort, which is primarily useful in agent evaluation. Aristotle describes four categories of people: the virtuous, the continent, the incontinent, and the vicious. The virtuous find it easy to do what is good or even delight in doing what is good. Those who are continent, or those who have

strength of will, are like the virtuous in that they do what is good, but because they are not fully virtuous, they struggle against other temptations in doing so. It is not easy for the continent to always do good, and it takes a certain level of subjective effort for them to do good.¹² The incontinent, on the other hand, do not have the strength of will. While they understand what they should do, they do not have the strength of will to overcome other temptations and instead do wrong. The vicious are altogether bad. They have no problem doing wrong and they do not spend any subjective effort doing wrong because it is easy for them to do so. Using Aristotle's categorization, what I described to be subjective effort seems to arise from the struggle between different parts of the soul as a person is doing a certain action. The more virtuous a person, the less subjective effort she spends doing good because she is naturally inclined to do good and there is less internal struggle in her soul, while the reverse is true when it comes to bad actions. The more vicious a person, the more subjective effort she spends doing good because the part of her soul that wants to do bad struggles harder, while she finds she spends less or no subjective effort doing bad.

Of course, the concept of doing high effort helping actions with low subjective effort as an ideal is not unique to Aristotle. An analog to this type of thought can be seen in the concept of *wu-wei* in mainly Daoist, but also other Chinese philosophy. Just as the virtues are Aristotle and other virtue ethicists' answer to how to live a *eudaimon* life, *wu-wei* is an answer to the

¹² The implication here is that those who are continent are worse than those who are virtuous, and are thus less admirable. One might note that there are cases, however, where a person who struggles to do good nonetheless seems admirable. For example, a person in a seriously difficult financial situation who finds a wallet might struggle to find the rightful owner, yet this person would still be admirable for doing so. However, this does not contradict the implication that the continent are still less admirable than the virtuous. If another person in the same situation found the wallet and had no trouble at all returning the wallet, he seems clearly more admirable. So the person who struggles to give the wallet back is admirable only because most of us are not fully virtuous, so we admire him as a fellow person who is not fully virtuous who was able to overcome his temptations. If we compared the person to the standards of the fully virtuous, the person would come up short. A similar observation was discussed above with the Nigel case, in which Nigel struggles tremendously to give a small amount to charity.

question of how to live a good life. *Wu-wei* can be translated as “doing things in a spontaneous and natural manner”.¹³ This does not mean simply being spontaneous in an action, but spontaneous in a deeper sense – acting without a forced intention or heavy premeditation because of one’s natural flow. Achieving *wu-wei* is difficult, and almost paradoxical as trying to be spontaneous goes against the very definition of spontaneity. Furthermore, the way in which different schools of thought like Confucianism and Daoism advise achieving *wu-wei* are different; the former advises cultivating *wu-wei* through hard work while the latter focuses on more of an attunement of oneself with the cosmic flow. Regardless, efforts like clearing the mind of ulterior motives, cultivating one’s virtues and dispositions, attuning oneself to the cosmic flow, and promoting harmony within oneself and with others seems key. Then, by living with *wu-wei*, a person leads a “naturally rewarding and spontaneously meaningful” life by revealing his “natural self and developing genuine character”. This display of one’s true self also cultivates trust amongst oneself and the community because one is not hiding his feelings or intentions. Therefore, *wu-wei* is not important just for the individual, but for the community as a whole. This vein that runs through both Aristotelian virtue ethics and *wu-wei* of the importance of effortlessness seems like the basis for my intuitions when it comes to subjective effort.

Beyond these moral theories, an additional explanation for the importance of subjective effort is based on the fact that we are limited with respect to the amount of effort we can spend. However, the fact that we are limited with respect the amount of effort we can spend is important here in a way that is different from the way it was important in our discussion of objective effort.

¹³ Edward Slingerland, *Trying Not to Try: Ancient China, Modern Science and the Power of Spontaneity*, (New York: Crown, 2014).

Consider the experiment performed by Baumeister et. al.¹⁴ In this experiment, Baumeister et. al. would introduce a college student to a room that smelled of freshly baked cookies. In front of the student were two plates: one with cookies and the other with radishes. One lucky group of students was asked to eat the cookies while ignoring the radishes, while the other group was asked to eat the radishes while ignoring the cookies. Afterwards, the students were given a difficult questionnaire to complete. It turned out that the students who ate the radishes (and had to exert more subjective effort to both eat the radishes and refrain from eating the cookies) would be more likely to want to quit taking the questionnaire and reported feeling more tired after taking the questionnaire than the other group of students (who had not spent any subjective effort while they indulged in the cookies). They ran out of subjective effort to spend! So given that we are limited not only with the amount of objective effort we can spend but with subjective effort as well, another reason why someone who has to expend subjective effort to do good seems deficient presents itself (besides the reason that a virtuous person would not need to expend that effort). Eventually, a person who needs to constantly spend subjective effort to do good will run out of subjective effort and end up being unable to do good. So, as we observed previously, the less subjective effort a person needs to spend to do good, the more positively we evaluate the person. In other words, a person should see doing good like they see eating cookies, not like they see eating radishes (unless, of course, they happen to love radishes even more than cookies.) This fact that there is a limit to subjective effort, when taken with the previous explanation in terms of the virtues, which are morally important in and of themselves, seem to provide an account of why subjective effort matters morally.

¹⁴ Roy E Baumeister, Ellen Bratslavsky, Mark Muraven, and Dianne M. Tice, "Ego Depletion: Is the Active Self a Limited Resource?" *Journal of Personality and Social Psychology* 74:5 (1998), pp. 1252-65.

Effort and What It Means to Do and Allow

I believe that the importance of the role of effort also extends outside of its importance in morality. It also seems to play a key role in the metaphysical distinction that makes us see acting in one sort of way as an “action” or a “doing” while see acting in another way as an “omission” or an “allowing”.

On the surface, the question seems fairly easy. After all, it seems pretty clear that when I upright the child in the kiddie pool, we would describe this as my doing something (saving the child) and when I do not do so, we would describe this as my omitting to save the child or allowing the child to drown. However, there are some cases where the distinction, at least initially, seems to be cloudy. Consider the case of the mild earthquake¹⁵:

After a mild earthquake the agent finds herself lying on and crushing a tube connected to a life support machine that is keeping someone alive. Unless the agent moves soon, the victim will die from lack of oxygen. She stays put. The victim dies.

Here, it seems on the surface that the agent did not do anything, but we would still say that the agent killed the victim. Or consider the case of a rock rolling down a hill:

Suppose a man is lying asleep on the ground. He is awoken by a crash and notices a large rock rolling down the hill towards him. He can easily move out of its way, but realizes that if he does so the rock will gain momentum and kill a group of small deaf children further down the hill. He tenses his muscles, fights his desire to run away and stands his ground. The rock hits and seriously injures him. But he stops it.

¹⁵ Howard-Snyder, Frances, "Doing vs. Allowing Harm", *The Stanford Encyclopedia of Philosophy* (Winter 2011 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/win2011/entries/doing-allowing/>>.

Here it seems like the man does nothing just as the agent earlier held still, but we would still say that the man did something – he saves the children.

In looking at the distinction between doing/actions and allowing/omissions, Clarke divides them into two categories.¹⁶ Some omissions are nothing – just absences of action. An example of this kind of omission might be my not picking a friend up from the airport after promising to do so. Clarke also sees some omissions as identical to actions. For example, if a person does not move while playing hide-and-seek, this omission (“not moving”) seems identical to the action of holding still. I imagine that Clarke would also look at the mild-earthquake case and the rock-rolling-down-a-hill case and say that the omissions seen in both of the cases are identical to the action of “holding still”. So although we initially said that the actors appear to kill and save by not doing anything, Clarke would say that their omissions are not really omissions at all, but omissions that are equivalent to an action – in other words, they do do something.

However, I think that Clarke’s account could be a bit fuller. Just saying that some omissions are identical to actions does not seem to help us think about when exactly we can equate an omission with an action. For example, it seems clear that all cases of the omission of not moving cannot simply be equated to the action of holding still. Consider another case of a rock rolling down a hill:

Suppose a man is lying asleep on the ground. There is a crash, causing a large rock to roll down the hill towards him, but the man remains asleep. Below the man is a group of small deaf children further down the hill. The rock hits and seriously injures the man, but he stops it.

¹⁶ Randolph Clarke, *Omissions: Agency, Metaphysics, and Responsibility* (Oxford: Oxford University Press, 2014).

We would not look at this man and say that his not moving is the same as holding still. He does not do anything to save the children in any meaningful way, except only in the sense that had he not been sleeping there the children would have been hurt. Another case where it seems difficult to answer whether an omission could be equated with an action is seen in the example earlier of not picking a friend up from the airport after promising to do so. Although this was example was given earlier as a case of an omission that is simply an absence of action, it is not clear why we cannot simply equate this omission to the action of staying home.

I believe that instead of using the somewhat vague criterion of whether an omission can be equated to an action, we should consider whether the choice taken by the actor required effort or not. Not moving in hide-and-seek takes effort. It is difficult to not move, not make a sound, and to stay hidden. It does not seem correct to call something that requires effort an omission, and so this “omission” in fact is equivalent to the action of staying still. The rock-rolling-down-a-hill case is similar. The man who wakes up to see a rock rolling down towards him tenses his muscles, fights the natural instinct of self-preservation, and holds still. This is a high-effort doing, not an omission. By contrast, had the man remained asleep, his not moving would have taken no effort at all, and his accidental saving of the children can be called an omission.

One might say here that it seems like I am confusing effort for intent (because intentional actions normally require effort and unintentional actions do not) and am simply describing intentional things like holding still while awake as actions and unintentional things like not moving while asleep as omissions. However, there are clear cases where intent and effort come apart. I might not go pick up a friend from the airport and be fully aware of what I am (not) doing. And despite the fact that I am intentionally not going to the airport, my not going to the airport is an omission because it does not require any effort. Thus, it would not be right to use

intent as a distinguishing characteristic between doings and omissions. Otherwise, only things that I do unintentionally would be classified as omissions.

Rejecting intent as a distinguishing characteristic also means rejecting the level of effort required in decision-making as a distinguishing characteristic between doings and omissions. If I can intentionally omit an action, I am making a decision to omit it. And indeed, there is an amount of effort associated with making a decision. Psychologist David Kahneman describes two systems that our mind uses.¹⁷ System 1 operates automatically with no effort. It is the system of thinking that is used when visually processing the relative distances of two objects or solving $2 + 2$. System 2, on the other hand, is more deliberate and effortful. So effortful, in fact, that continuous System 2 thinking appears to tire out a person. Filling out tax forms, calculating 17×24 , or writing an essay mobilize System 2 thinking. When someone makes a conscious decision like deciding to omit an action, they also mobilize System 2 thinking.¹⁸ However, though the simple action of making a decision on its own seems to require some effort based on Kahneman's research, I do not think this type of effort is sufficient to distinguish between doings and omissions for the same reason I do not think intent is an appropriate criterion. I can surely make a decision to omit to do something, so whatever effort that was associated with making that decision should not be relevant. So a more refined version of my earlier criterion of whether a choice "requires effort or not" would be whether a choice requires more than (action) or less than/equal to (omission) the amount of effort associated with System 2 decision making.

Given this criterion, it seems clear that the hide-and-seek and boulder cases are examples of actions that require more effort than that associated with decision making. A bit trickier is the

¹⁷ Daniel Kahneman, *Thinking, Fast and Slow*, (New York: Farrar, Straus and Giroux, 2011).

¹⁸ Using System 2 thinking seems to require what I have been calling subjective effort and also seems to be what tired out the participants in the cookie-radish experiment mentioned above.

mild-earthquake case. Here, at first glance, it may seem like the effort involved is only that associated with decision making. Lying on the life support tube physically probably requires no effort as one could simply rest on it. Thus, the difficulty with which one needs to hold still in hide-and-go-seek does not even exist, nor does the difficulty involved with tensing one's muscles and fighting the instinct to avoid physical harm in the rock-rolling-down-a-hill case. Thus, according to my criterion, the mild-earthquake case would be an example of an omission. However, I do think that there is effort involved beyond just physical effort in the case of the mild earthquake case. I think the agent who stays on the life support tube exerts "going-out-of-your-way" effort or trajectory alteration effort. The agent's normal course of action after an earthquake would be to get up, to tend to one's injuries, and to do other post-earthquake actions. Had the victim not been connected to the life support tube, the agent would not have stayed on the tube and would instead have gone about this course of action. But because the victim is there, the agent changes the trajectory of actions relative to what he would have done normally and stays resting on the tube. The "going-out-of-your-way" effort associated with this trajectory alteration seems to me the minimal amount of effort needed to distinguish between a doing and an omission, and the reason for which the mild-earthquake case is an example of a doing, not an omission. Cases without this minimal level of effort are cases of omissions. For example, consider again my not picking up a friend at the airport. There may be some decision-making effort associated with my not picking my friend up. However, in deciding to continue on with my day and not alter it in order to pick up my friend, I do not exert the "going-out-of-your-way" or trajectory alteration effort. Therefore, this is a case of an omission. A final version of my criterion for distinguishing a doing and an omission then becomes that a doing involves a level of

effort greater than or equal to the trajectory alteration effort while an omission involves a level of effort less than or equal to the effort involved in decision making.¹⁹

While using this kind of criterion and thinking in this manner generally leads to the same conclusions that Clarke arrives at concerning whether something is an omission or an action (or as Clarke calls it, an omission that is equivalent to an action), I feel that it clarifies why exactly Clarke comes to his conclusions without the messiness of introducing the semantics of whether an omission can be construed and reworded to sound like an action. Furthermore, if we define more clearly the distinctions between doing and omitting, it also makes talking about other types of cases concerning doing and allowing more clearly. For example, to say that someone kills is to say that someone does something to bring about the death of another while letting die implies that someone does not do anything, omits an action, and allows another to die. However, the word “kill” has a negative connotation to it that makes it tempting to say that someone kills simply because we view him as responsible for another’s death. Yet responsibility and whether a person did or allowed something often do not correlate with each other. If a friend asks me to take care of a pet while she goes on vacation and I then proceed to ignore the pet and let it die, my friend might say that I am responsible for her pet’s death and even go so far as to say that I killed her pet, even though I did not actually *do* anything. My response to my friend would then be, “I didn’t kill your pet, I only let it die – I omitted to take care of the pet, the option that does not take effort. But because I seem so responsible for the death, you are mistakenly accusing me

¹⁹ I think an interesting question to ask here is whether the omnipotent being I mentioned earlier, for whom effort has no meaning, can be said to do or allow anything if my account of what it means to do something and what it means to allow something turns on whether a person (or being) exerts effort or not. It seems that I would be committed to saying that it does not make sense to say that an omnipotent being, for which effort has no meaning, can do or allow anything. This commitment actually seems correct to me. For an omnipotent being who controls the very fabric of the universe, does it make sense to say that she allows a child to drown (as I somewhat haphazardly said before)? Or would it make sense to say that she causes the child to drown? It seems to me that what the being does (?) in this case is actually some sort of hybrid between the two for which I’m not sure I have a term.

of killing your pet!” Of course, my friend might still continue to say I killed the pet in everyday conversation, but she would be missing my point that we should be clear with our usage of the terms omissions/letting die and doing/killing without the additional baggage of perceived responsibility when trying to understand what kind of distinctions there are between doing and allowing.

Effort and the Distinction between Doing Harm and Allowing Harm

Given that we have now looked at how effort plays a role morally and in the distinction between what we mean when someone does something or allows something, I now turn my attention to the important question of whether doing harm is worse than allowing harm.

It seems to me that the participants in the discussion of whether doing harm is worse than allowing harm often argue in a way such that one side argues that doing harm is always worse than allowing harm while the other argues that there is never a difference between doing harm and allowing harm. In order to prove or disprove that there exists a moral difference between doing and allowing harm, each side constructs two cases in which everything is held constant besides this whether the person in the case does harm or allows harm. These cases then appeal to our intuitions about the moral permissibility of doing harm in one case and the moral permissibility of allowing harm in the other; if our intuitions seem to say that there is no difference, the side arguing that there is no difference claims they have succeeded, and if our intuitions seem to say that there is a difference, the side arguing that there is a difference claims they have succeeded.

An example of this kind of method is seen in the Smith-Jones case given by Rachels.²⁰ Smith is a man who would gain a large inheritance if his six-year-old cousin were to die. So one evening while the cousin is taking his bath, Smith enters the bathroom and drowns the child in order to gain the inheritance. Jones is another man who also stands to gain a large inheritance if his six-year-old cousin were to die. One night, as Jones enters the bathroom in which the cousin is taking a bath, he sees his cousin slip and hit his head, falling into the water. So Jones stands by and lets the cousin drown. In this case, it seems that Jones's omission is just as heinous as Smith's killing. Therefore, doing harm is no worse than allowing harm, and there is no moral difference between doing harm and allowing harm. But while the intuition that Jones is just as reprehensible as Smith seems correct, this case alone does not seem sufficient to show that there really is no difference between doing and allowing harm. First, the acts or omissions that Smith and Jones perpetrate are so heinous that, if there was indeed a difference between doing and allowing harm, our strong visceral reactions in both cases might really be drowning out whatever intuitions we might have about the potentially subtle difference between doing and allowing harm. Second, if simply showing a case where our intuitions happen to line up with one's claim were enough, it seems that the other side can just as easily come up with cases in which our intuition tells us that there is in fact a difference between doing and allowing harm.

For example, let's create a case comparing a school bully and an older bystander. The bully picks on a victim and steals the victim's lunch money. Meanwhile, a bystander, who is older than the bully and would be at no personal risk if he were to intervene, watches on and allows the bullying to happen. While it would be nice if the bystander did something, our intuition is that the bystander's omission is still clearly better than if he had actually done the

²⁰ James Rachels, "Active and Passive Euthanasia," *New England Journal of Medicine*, 292 (1975), pp. 78–86.

bullying himself. Thus, in this case, it seems that doing harm is indeed worse than allowing harm.

Considering that cases can be constructed in this manner easily to make doing harm out to appear worse than allowing harm or instead make doing harm seem no worse than allowing harm, this method of constructing cases is not a very effective way to discuss the difference between doing and allowing. Because these sets of cases come out very differently despite the only factor being changed is doing vs. allowing, there must be some other feature of these cases that is more relevant beyond simply whether a case is a doing of harm or an allowing of harm. Therefore, a discussion of doing and allowing harm must delve more deeply than simply constructing cases. But not only that, we must look for other factors that are playing a role in our intuitions in these cases besides the factor of whether a person does harm or allows harm. In other words, we should not only look to answering the question of what factors are playing a role in our intuitions in these cases by answering, “This is why doing is always worse than allowing,” or, “This is why doing and allowing are always equivalent”. Answering the question in this manner would simply lead us back wondering why one set of cases of doing and allowing harm turn out one way while another set of cases turns out another way. Instead, an explanation for these cases must be given that is not at its core based simply on whether a case is a doing or an allowing, just as our intuition does not align cleanly and simply with doing and allowing.

However, it seems like often the discussion of doing and allowing does not go the way it should. Quinn, for example, stakes out a position such that doing is always worse than allowing by saying that killing is always worse than letting someone die.²¹ His explanation is based on the difference between negative rights, which protect someone from being actively harmed by

²¹ Warren Quinn, “Actions, Intentions, and Consequences: The Doctrine of Doing and Allowing,” *Philosophical Review* 98:3 (1989), pp. 287-312.

others, and positive rights, which put a claim on others to help someone. Killing would be a case of violating a negative right, while letting die is violating a positive right. Quinn believes that negative rights are more serious because if a person's negative right can be violated by being killed to meet the positive rights of others, then the victim does not own his body in an important moral way – rather, it belongs to the “human community”. Because negative rights are more serious, violating a negative right is also more serious. Therefore, killing is worse than allowing someone to die. But again, this type of explanation is one that ultimately reduces the distinction between cases like the one mentioned above into saying doing harm is always worse than allowing harm, albeit with an added explanation. It does not allow for the type of variation that we see across cases and cannot explain why we do not see any difference between doing and allowing harm in the case of Smith and Jones. On the other hand, Bennett, for example, finds that there is no moral significance to the doing harm and allowing harm distinction.²² Bennett starts by claiming that we say that an agent does something when the behavior that results in the consequence of the doing would not have occurred from most of the other choices that the agent could have made, while we say that an agent allows something when the behavior that results in the consequence of the allowing would have occurred from most of the other choices that the agent could have made. (In other words, his account of what it means to do something and what it means to allow something is different from my account based on effort.) Then, Bennett's argument goes, because whether a consequence could have resulted from many ways that the agent could have behaved or from a few ways that the agent could have behaved is morally insignificant, the distinction between doing harm and allowing harm is also ultimately insignificant. This seems problematic to me on the opposite end from Quinn's argument. While there are cases like the Smith-Jones case where the distinction does seem insignificant, there are

²² Jonathan Bennett, “Acting and Refraining,” *Analysis*, 28 (1967) pp. 30–31.

surely cases where the distinction matters like in the bully-and-older-bystander case. Instead of making doing harm consistently worse than allowing harm like Quinn's argument does, Bennett's seems to consistently make the distinction between doing harm and allowing harm insignificant.

I believe that looking at effort can elucidate this discussion.

Recall that allowing something to happen, on my account, means that one does not exert effort for some outcome to occur. Not exerting effort is neutral when exerting it means harming someone or when exerting it means helping someone but the level of exerted effort is high. However, not exerting effort is impermissible in the case in which exerting that level of effort is morally required, either because that level of effort is required by a principle that cannot be reasonably rejected as a principle that governs moral behavior or because that level of effort is minimal and will not be exerted frequently enough that it would deplete a person's effort. Therefore, allowing harm (which is equivalent to not exerting effort to help and prevent the harm) is wrong when it the level of effort that would need to have been exerted to help is required. The blameworthiness of allowing harm increases when as this level of effort decreases even beyond that level which is required.²³ What I believe to be of key importance here is that the nature of what it means to allow something (not exerting effort) and the subsequent explanation of when and why not exerting effort is wrong is the primary explanation for why allowing harm is wrong. In a way, the exact nature of the harm that is prevented is secondary.

²³ As one can see in Figure 2, I believe that this blameworthiness increases asymptotically. Not exerting the pound of force to wiggle my finger to save the child in the kiddie pool seems to me to be much, much more than "twenty times" worse (whatever that might mean) than exerting twenty pounds of force to lift the child directly. But regardless of whether this is in fact true, I hope it is at least plausible that the blameworthiness does indeed increase, asymptotically or not.

Another way to put this might be this: assume you are a parent with your child at a playground and notice that your child stood idly by and watched while another child punched someone. At another point in time, you notice that your child stood idly by while another child stole another child's toy. If you were to lecture your child about why standing idly by in these situations is wrong, the lectures would not be about why punching is wrong and therefore why he should prevent it or why stealing is wrong and why he should prevent it. In both cases, it would probably instead be something about why he should help others in general and go out of his way to help people in settings like the playground where it is easy to prevent these harms. In other words, the way your child acted was not wrong because of the harms of punching and the harms of stealing toys but because of his failure to exert effort to help others. Of course, this is not to say that the nature of the harm does not matter at all. Saving a drowning child and preventing a toy from being stolen surely modifies what level of effort is required. But this is still more of a secondary *modification* rather than *determining* exactly why allowing harm is wrong, which was based on the fact that morality can require one to exert effort to help others.

Notice how this is fundamentally different from the reason why doing harm is wrong. Doing means exerting effort. In this situation of the morality of exerting effort, effort only has a magnifying effect (except when the effort is that effort which is required.) Here, the reason why doing something is wrong or right is primarily based on the ramifications of the action itself. Effort then plays a role in magnifying this wrongness or rightness, but it is not playing the central role of determining right and wrong like it did in the case of allowing harm. For example, consider why it is wrong to trip someone. Tripping someone is wrong because of the nature of the harm being caused, such as the physical pain involved in being tripped. It is not wrong *because* it required effort to trip someone. More effort being involved in tripping someone

makes the action of tripping worse because an action that requires more effort is more characteristic of someone who is vicious (virtue theory) and because it uses more of the limited amount of effort one has for harm (the idea that we have a limited amount of effort). But this is a secondary modification of the blameworthiness of the action. Thus, when we consider why doing harm is wrong, we look towards the exact nature of the harm that is being done, which is different from the way we look at allowing harm, where we look first at the fact that effort that was not exerted to aid another and then at the more specific details what exactly the aid that was not provided was.

I believe that this way of explaining why doing harm is wrong and allowing harm is wrong not only ends up being better at explaining our intuitions when it comes to cases of doing and allowing harm, but I also believe that it reframes the discussion of doing harm and allowing harm. Considering that the primary reasons for which doing harm and allowing harm are different, it does not seem right to directly compare the two and ask whether one is better than the other, just as it would be strange to ask whether a burger is better than a cheesecake – they just fall into different domains. Burgers are good for certain reasons while cheesecakes are good for other reasons. Perhaps we might be able to compare general tastiness levels of one to another, but it would be mistaken to then try to compare how tasty one burger is to how tasty one cheesecake is and conclude that burgers are worse than cheesecakes or conclude that burgers are not better than cheesecakes. In a way, my explanation may be underwhelming for those who wanted to find a way to make it so that a moral difference turns purely on the distinction between doing and allowing harm. Rather than attempting to do this, my explanation instead is more of an explanation why the question of asking whether doing harm is worse than allowing harm is a somewhat flawed question from the very outset of the discussion.

It's also interesting to note that nobody asks whether there is a moral difference between helping someone and allowing a benefit to befall someone. Of course, the answer is obvious here, and it would be pretty crazy to think that helping someone could be equivalent to simply allowing a benefit to happen to someone. Yet, it seems to me that those who argue that there is no moral difference between doing harm and allowing harm would also be committed to arguing that there is no moral difference between helping someone and allow benefits to someone, if they are to truly commit to their conclusion that there really is no difference between doing and allowing. But my point here is not that the proponents of this view are wrong (because I believe both sides are misled in the question that they are asking or in the way they answer the question) but that focusing solely on the metaphysical distinction between doing and allowing as if this alone makes a difference should also mean that this distinction is important (or irrelevant) in any context, including doing good and allowing good.²⁴

But though my account of effort points out the flaws in this question of whether doing harm is worse than allowing harm, I also believe that my account is also superior in explaining why our intuitions line up in certain ways in doing/allowing harm cases.

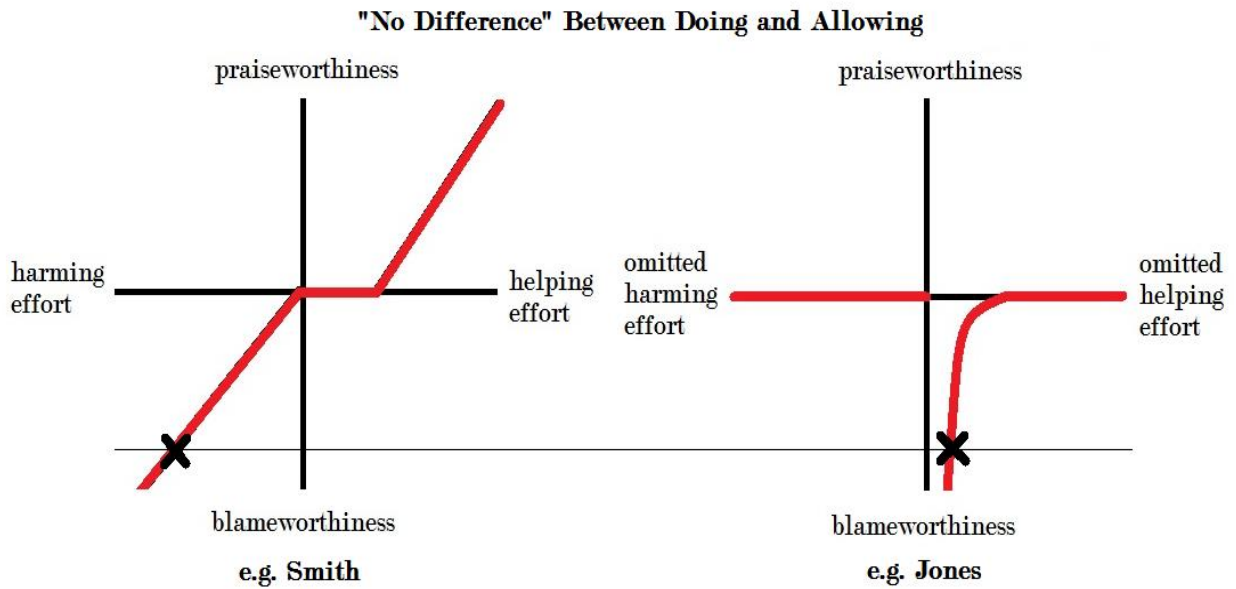
In Smith-Jones, Smith's doing harm was particularly heinous in that it caused the death of a child. Jones's omission was of an action that required an absolutely miniscule amount of effort (pulling a child's head out of the water) when the victim was in a great amount of need, and so Jones's action also seems heinous. Thus, in the Smith and Jones case, doing and allowing harm seem no different (Figure 3A). In bully and older bystander, the bully's doing harm was

²⁴ My explanation of why doing good is better than allowing good would be as follows. Allowing good simply means not exerting effort to do harm, about which morality has nothing to say (unlike with the effort that can be required to do good) and is thus neutral. Doing good, on the other hand, is positive, with how positive the action is being modified by effort exerted. (If the good that was done was merely required, then it would also be neutral and thus equivalent to not doing harm, which seems to align with my intuitions.)

wrong in that he stole his victim's money. The bystander's omission was of an action that would be quite effortful – he would have to confront the bully and take the time to try and reason with him. Because this action that was omitted was high effort, the omission does not seem very bad. Thus, in the bully-and-older-bystander case, there does seem to be a difference between doing and allowing harm (Figure 3B). Furthermore, notice that if the bystander was not just another older student, but an intimidating adult, the omission would be quite bad. An intimidating adult would not have to exert much effort at all to stop a bully besides just saying “stop”, so the omission of the low effort would be comparably bad to (or perhaps even worse than) the bully's bullying.

So while other ways of looking at whether doing harm is worse than allowing harm would have to deal with cases that contradicted their stance either by dismissing them without explanation or coming up with some other factor, my account can deal with the different conclusions we come to quite simply. Cases that make allowing harm out to be the same as doing harm are set up such that the effort that would have been necessary to prevent the harm in the allowing harm case is quite low. This would make allowing harm impermissible, and especially blameworthy if the effort was extremely low or the harm being done was quite great which would make the effort required higher (or both as in the Smith-Jones case). Cases that make doing harm worse than allowing harm are set up such that the effort that would have been necessary to prevent the harm in the allowing harm case is not negligible. This would make allowing harm either permissible or impermissible, but still not as blameworthy as the harm done in the doing harm case. Thus, my account of effort shows how the way in which effort modulates the blameworthiness of allowing harm affects our intuitions regarding these sorts of cases comparing doing and allowing harm.

A.



B.

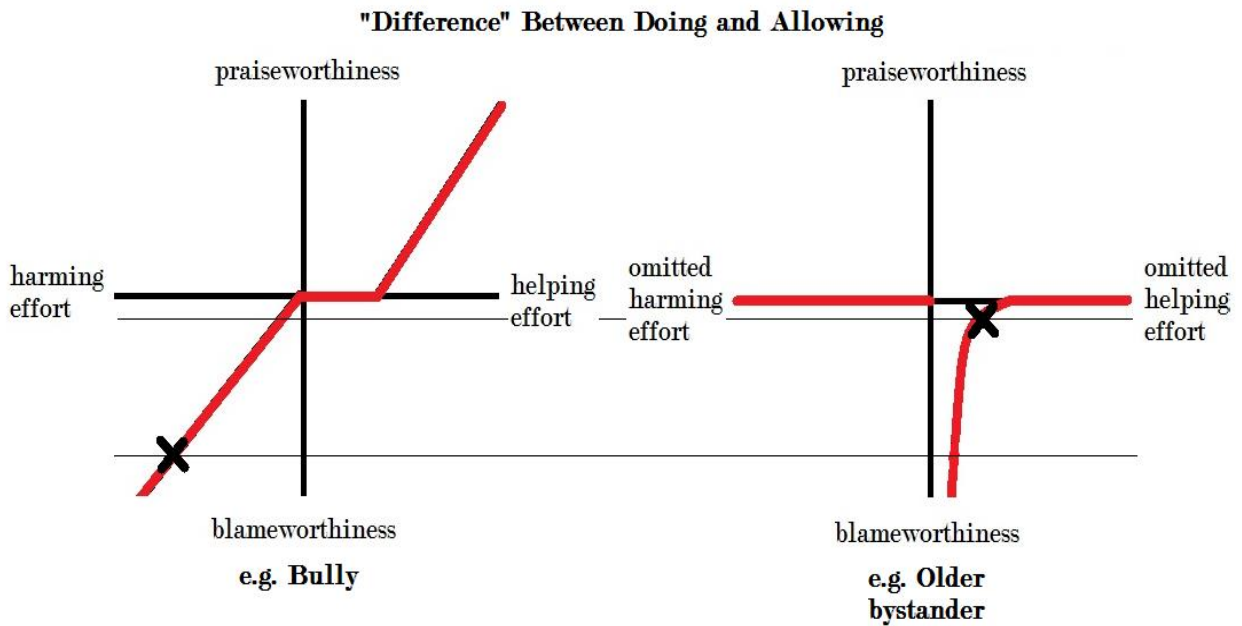


Figure 3. Comparisons of doing and allowing in which there appears to be no difference between doing and allowing like in the Smith-Jones case (A) and in which there does appear to be a difference like in the bully-older bystander case (B).

Other Accounts of the Do/Allow Distinction

Hopefully I have now given a plausible account of the broad importance of effort in morality and in the metaphysical distinction between doing and allowing, as well as the relevance of effort to the distinction between doing and allowing harm.

In explaining my views on the distinction between doing and allowing harm, I briefly mentioned Quinn and Bennett as other examples of ways people have looked at the do/allow distinction. I would now like to take a closer look at such accounts (of which I was admittedly a bit overly dismissive earlier) while also pointing out what I believe to be a flavor of my account playing a role in the background of others' view of doing and allowing.

Let's reconsider Quinn's view on doing and allowing. Quinn believes that doing is worse than allowing because the former involves the violation of negative rights while the latter the violation of positive rights. And because negative rights take precedence over positive rights, the violation of negative rights is a more serious offense, making doing harm worse than allowing harm. The question then, is why should negative rights take precedence over positive rights?

Quinn starts by plausibly claiming that it would not make sense for positive rights to have precedence of negative rights. If positive rights were more important, then we might get the absurd result that we should kill, say, two people to one because the positive rights of the one that tell us to help him can outweigh the negative rights of the two that prevent us from harming them.

Then Quinn says that negative rights are more serious because if a person's negative right can be violated by being harmed to meet the positive rights of others, then the victim does not own his body in an important moral way – rather, it belongs to the “human community”.

Furthermore, he states that the extent to which we believe that negative rights are weightier than positive rights will “depend on how important the relevant forms of legitimate control are to us – the extent to which we wish to belong... to ourselves.” So Quinn sees this as an issue of what kind of control the “human community” can have over an individual to meet the needs of others. I think Quinn’s way of explaining why negative rights are more important than positive ones is mistaken.

First, he says that negative rights are more weighty than positive rights because of this issue of control, but positive rights surely also result in a demand on the bodies of others that they aid someone. So while it seems true that positive rights should not take precedence over negative rights because of the strange conclusions to which doing so leads, but this issue of control does not seem like a reason to then place negative rights over positive rights considering that control is an issue for both positive and negative rights.

But let’s assume that Quinn meant to defend a position where control is less of an issue with positive rights than it is with negative rights (though he never really does this). Even then, I think control is not the issue that Quinn makes it out to be, and it seems like control is only a symptom of what is probably actually more fundamentally important to Quinn. If morality correctly determines that one person’s sacrifice is necessary to meet the positive rights of another, then it should also compel that one person to recognize his sacrifice as necessary. That is not a matter of control. Consider the case in which a person must die so that a highly populated city is not bombed with a nuclear weapon. I (like many others) would say that it is acceptable to kill that person. If this is the case, however, I must also accept an implication of my saying so – if it turned out that I was the one who must die so that a city is not bombed, I should accept that I can rightfully be killed so that others may live. In accepting this implication,

I would not feel “owned” by morality or by the human community. I myself recognize that morality or the human community rightfully demands this sacrifice of me, and the recognition that this sacrifice is something I ought to do should make me *want* to give the sacrifice. The only way in which I would see this as control in the sense Quinn uses it would be if I actually thought that one person cannot be sacrificed in order to save a city and the human community is in fact mistaken in believing that I should be sacrificed. In other words, I would feel controlled only if I believed that the human community is mistaken and demands more from me than can reasonably be expected. So it is not the control itself I would be objecting to; fundamentally, it is that more is being demanded of me by the human community than morality actually demands (which incidentally results in the feeling of being controlled). When Quinn makes rights an issue of control, he should dig deeper and ask why it is that it feels like one is being controlled instead of merely objecting to the control itself.

I believe that effort on its own can provide a strong explanation for why negative rights take precedence over positive ones. To violate a negative right, one not only brings harm upon someone, but has to put in some effort to do so. On the other hand, to violate a positive right, one is omitting a helping action or allowing harm; one is not expending a certain amount of effort to help someone. To ask others to not expend effort to harm me (negative rights) is less demanding than to ask that they expend effort to help me (positive rights), even if the effort required by positive rights is a small amount because it is still ultimately easier to not expend effort than to expend effort. Thus, negative rights take precedence over positive rights.

Now Quinn’s focus on control makes me believe that Quinn, like me, is asking the question of what can be reasonably demanded of someone to rights because the feeling of being owned and controlled arises when more is demanded than a right actually does demand. But

Quinn forgets that rights do not exist in a vacuum. Rights are not some sort of special property that a human being would have even if there were only a single human being in the universe (though we often talk about rights as if this were the case). It only makes sense to talk about rights if there are also people who exist who can respect those rights. This relational nature of rights makes it such that it is not only the people who have the rights who should be considered when thinking about rights, but also the other side of the relationship, the people who are being asked to respect the people with rights. Consequently, control is not only an issue when negative rights are violated, which was Quinn's focus, though I certainly believe that the sacrifice that is forced upon the victim when his negative rights are violated is indeed an important explanation for the weightiness of negative rights, a sacrifice that is usually not limited to effort (and is thus not explainable solely by my account of effort). Control is also an issue when we ask others to respect our rights. My suspicion (or at least hope) is that if Quinn were to then ask the question of what kinds of things might be demanded in general in order to satisfy rights, both positive and negative, things like effort would play a similar role for Quinn as it does for me.

Unlike Quinn, Bennett finds that there is no difference between doing and allowing. Bennett starts off by describing what it means to be positively relevant to an outcome and what it means to be negatively relevant.²⁵ A person is positively relevant to an outcome if most of the things that the person could have done would not have resulted in the outcome. A person is negatively relevant to an outcome if most of the things that the person could have done would

²⁵ Bennett uses the terminology of positive/negative instead of doing/allowing or doing/letting in order to talk specifically about whether a conclusion happens because a person did X or did not do X. This is important for Bennett because there are cases where a person does some X (i.e. the person is positively relevant to an outcome) that allows something to happen. Thus it is clearer to talk about positive and negative rather doing/allowing or doing/letting.

have resulted in the outcome. He uses two cases of a vehicle rolling down a hill to show what he means by this²⁶:

A vehicle is rolling down a hill. There is a rock in its path which would stop it. John kicks away the rock. The vehicle crashes at the bottom of the hill and is destroyed.

A vehicle is rolling down a hill. There is a rock near its path which would, if interposed, stop it. John does not interpose the rock. The vehicle crashes at the bottom of the hill and is destroyed.

In the first scenario, John seems positively relevant to the vehicle being destroyed to us because most of the actions he could have taken would not have resulted in the vehicle being destroyed. Only a few ways of moving or types of behavior would result in the removal of the rock from the path of the car, and John chooses one of those actions that fit into this small category rather than choosing any of the vast number of other actions he could have taken that would not have resulted in the removal of the rock. In the second, John seems negatively relevant because most of the actions he could have taken would have resulted in the vehicle being destroyed. There are only a small number of kinds of actions or behaviors that would have resulted in John's interposing the rock to stop the car, but a huge number of actions that would not. So when John chooses one of the many ways of acting that do not interpose the rock, John seems negatively relevant to the outcome of not placing the rock, which is the destruction of the car.

This way of looking can be applied to special cases like the following one that Bennett constructs:

“Henry is in a sealed room where there is fine metallic dust suspended in the air.

If Henry keeps utterly still for two minutes, some of the dust will settle; and if it

²⁶ Jonathan Bennett, “Morality and Consequences,” *Tanner Lectures*, (1980) pp. 47–116.

does, some is bound to fall in such a position as to close a tiny electric circuit which . . . well, finish the story to suit your taste, but make it something big; and let's call its occurrence S. Thus, any movement from Henry, and S will not obtain; perfect immobility, and we shall get S."

Bennett's intuition, which also happens to be mine, is that Henry would be positively instrumental if he were to keep completely still, which would result in the dust settling and closing the circuit. For Bennett, this is because the ways that Henry could have acted such that S obtains (staying perfectly immobile) are fewer than the ways that Henry could have acted such that S does not obtain (any movement at all). Thus, Henry's immobility is a doing.

The reason I believe Henry does something in this situation is different. My explanation is quite similar to my earlier explanation of the hide-and-seek example. Hiding during hide-and-seek is not easy – one must stay still and remain hidden. So in order for a person to hide, he needs to expend effort in order to do so. This effortful behavior, then, should be classified as a doing, not an omission. Similarly, for Henry to remain immobile would be difficult and effortful, so when he remains still so that the metal dust can settle, Henry is doing something rather than omitting something. He is staying still rather than not moving.

But though our intuitions lead us to the same conclusion about whether Henry does something or does not do something in this case, our different explanations of what it means to do something and what it means to allow something result in differences in the way we look at the distinction between doing and allowing. Having defined what being positively relevant and negatively relevant is, Bennett proposes that there is no moral distinction between being positively and negatively relevant to an outcome because the fact that there are many ways one could have acted for a certain outcome to obtain or only a few ways has no moral significance.

As a result, he finds that there is no important moral distinction between doing and allowing. Given this difference, it is worth deciding which way of understanding what it means to do something or to not do something is superior, which I believe can be done with a small modification to the Henry case:

Henry is in a sealed room where there is a heavy metallic dust suspended in the air. Because this dust is heavy, it will still settle even if Henry moves a small amount consistent with how people normally move when they are resting. Such movement might include things like jiggling of the leg, switching positions, shifting his weight, and more. If Henry does not move more vigorously than this for two minutes, some of the dust will settle; and if it does, some is bound to fall in such a position as to close a tiny electric circuit which . . . well, finish the story to suit your taste, but make it something big; and let's call its occurrence S.

In this modified version, Henry does not have to stay immobile. He can still move a small amount, which means he can simply rest comfortably and the dust will settle. For Bennett, Henry's resting would still be a positive action – though this modified version adds to the range of actions that Henry could take while still having the dust settle, this range is still surely smaller than the vast range of actions and movements that Henry could still take that would span from movements that are slightly more active than the movements of a normally resting person to movements like wild, crazy flailing of his entire body and throwing himself around the room. The issue for Bennett's account should now be clear. Given that the relative range of actions that would result in the dust settling is smaller than the range of actions that would result in the dust staying suspended in the air, Bennett would have to be committed to saying that Henry's simply resting is a positive action, a doing. Even more obviously wrong would be the corollary to this

commitment, which is that if Henry, being aware that all he has to do is rest and knowing that he still has the freedom to move to a certain extent, then decides to put extra energy in to make movements that are more active than those he normally would do while he is resting, he would be performing a negative action, or an omission. Even flailing his limbs and running around the sealed room should count as Henry not doing something because it is one of the modes of action that fall into the majority of actions that do not lead to S. It seems clear that Bennett's account leads him to conclusions opposite to what we would expect. My account, on the other hand, would lead to the conclusions that we would normally expect. If Henry were to simply rest, he would not be exerting effort in order to do so. So his resting should be seen as an omission. If he instead exerted effort to move around more actively than he would have had he just rested, he is exerting effort to do so, so his movement should be seen as an action.

Now although I don't want to assume too much about what was going on in the background of Bennett's intuition, I think that something like effort might be in the background of his argument. Bennett, though he thinks that the distinction between doing and allowing does not in and of itself mean anything, seems to have an explanation for why he believes we so often see a distinction there. When we say that someone does something (or as Bennett would say, is positively relevant to an outcome), the person is choosing one of the few ways he could have acted and he "puts his behaviour within a quite small part of the range of all the modes of conduct which are open to him at the relevant time". So to do this positively relevant action would be to "squeeze your behaviour into a designated small area". In doing so, there is likely to be some sort of cost associated with forcing oneself to act within this small range of actions. To allow something to happen, on the other hand, requires only that a person not do something, which is to act in a way that does not fall into some small category of acting and to instead act in

one of the many other ways one could act. Because the number of possibilities that are then available is quite broad, it is more likely that one could act within this range without incurring some sort of cost. “Compare being banished to Liechtenstein with being banished from Liechtenstein.” In this way, Bennett addresses objectors who claim that people like Bennett who see no distinction between doing and allowing are committed to a strenuous commitment like saying that the bully’s behavior and the older bystander’s behavior is morally equivalent. Rather than looking at the distinction between doing and allowing, Bennett says, one should consider other special features like cost that are involved in cases of doing and allowing. For example, the fact that the older bystander has some type of cost associated with intervening in the bullying is what distinguishes his behavior from the bully’s behavior, not the fact that the older bystander allows harm rather than does harm, which is a morally neutral distinction based on the different numbers of ways of acting that lead to a certain outcome.

Given that Bennett acknowledges how cost and effort plays into being positively and negatively relevant to an action, it seems to me that his understanding of doing and allowing is actually somewhat similar to mine – or at least his conclusions might tend to align with mine. But again, the key difference is that Bennett fails to see that the way we understand doing and allowing is inherently tied to effort, not to a metaphysical tally of the number of ways one could act and have some outcome obtain, which I believe the modified Henry case shows. This inherent connection to effort is what naturally leads us to intuitively consider how effort magnifies the good or bad of an action and to how effort affects the demand on someone to perform a low effort helping action. This is an important difference, especially when we consider cases that are set up such that there seems to be no difference between doing and allowing. For example, in the Smith-Jones case, there seems to be no difference between Smith

who drowns the child and Jones who does not save the child from drowning. But the reason why there appears to be no difference is not because of Bennett's reason: that whether or not Smith acted in one of a few modes of acting to drown the child or whether or not Jones acted in one of the many modes of acting does not matter. Rather, the reason is that Smith and Jones do things that are similarly wrong, but wrong for *different* reasons. Smith does something very wrong as he expends effort to drown a child. Jones does something very wrong because he does not expend the minimal effort to aid someone in need when he does not save the drowning child. It just happens to be that Smith and Jones do something similarly wrong, but the reasons why the two did something wrong are significantly different and important to consider separately. So even if Bennett happens to get to the same conclusions that I do often enough by considering cost on top of his account of account of what doing and allowing means, I think there is sufficient reason to reject both his account of what doing and allowing means and his belief that cost and effort is somehow a secondary symptom of his account of doing and allowing (which I believe to be mistaken).

The Intend/Foresee Distinction

Thus far, I have focused mainly on the distinction between doing and allowing harm as an example of a discussion in which understanding the role of effort in morality can be of help, mainly because of the way this discussion follows naturally from the question of what we mean by doing and allowing in the first place. However, I think that effort can also play an important explanatory role in other discussions as well. As an example, I will now look at the distinction between intending and foreseeing a consequence, or the doctrine of double effect. The doctrine of double effect, which Thomas Aquinas is credited with first introducing, notes that an action

can have two effects, “only one of which is intended, while the other is beside the intention”.²⁷

An action, then, can have an unintended side effect of a serious harm but still be permissible given that the intended effect of the action is good and proportional in a way that the harm can be justified. Another way to put the doctrine is to say that it is wrong to aim at a harm as an end in itself or as a means to an end; therefore, it is permissible to perform an action if a harm that follows is merely an unintended or foreseen side effect that results from pursuing a good, intended effect that follows immediately from the action. This doctrine seems useful in explaining a series of difficult cases. In the famous “Trolley” case, a trolley is barreling down a track along which it will kill five people. If diverted onto another sidetrack, it will only kill one person. Many would say that it is permissible to divert the trolley. In the “Bridge” case, another trolley is also barreling down towards five people. Above the track is a bridge, where a person is standing; by pushing this person off of the bridge and onto the track, the trolley will be stopped and the five people saved. Using the doctrine of double effect, one might say that in “Trolley”, pulling the lever is not an act intended to kill the one person, but rather a death that is merely foreseen as a consequence of an act aiming to save the five people. By contrast, in “Bridge”, pushing the person is an act intending the death of the person, and is thus impermissible despite the fact that five may be saved. Of course, the distinction is relevant in cases that were not initially created by philosophers. For example, if a limited amount of resources at a hospital can only be allocated in a way such that either five people are saved or one is saved, it seems permissible to choose to save the five (the intended effect) while the one dies (the unintended side effect). Or consider the difference between terror bombing during wartime and aiming to kill innocent civilians as a way to demoralize the enemy (seemingly impermissible because it aims at directly at the deaths of civilians) and tactical bombing of a military target that would

²⁷ Thomas Aquinas, *Summa Theologica*.

result in the deaths of an equal number of innocent citizens (seemingly permissible because the deaths are merely foreseen and not intended).

I do not think that the intend/foresee distinction does the moral work that its proponents claim it does. In defending this position, I will look at the various ways in which we might use the word “intent”. It seems to me that there are two categories of ways we use this word. First, “intent” can refer to the mental state of the person performing an action that produces harm. Second, it can refer to the instrumentality of the harm produced by an action.

One mental state that “intent” can refer to is a broad sense of the word intent – that a person intends a consequence so long as he was aware of a consequence of his action. The use of this broad mental sense of the word intention, which I don’t believe many proponents of double effect in fact defend, seems clearly irrelevant in moral considerations. In this sense of the word intention, it is hard to see how intending a consequence and foreseeing a consequence can be different, and thus, how this could affect permissibility. If someone is aware of a consequence that will occur as a result of the action that he takes, and knowing this consequence, decides to take that action, clearly this is intention in the broader sense. To knowingly cause a consequence to occur, whether or not the actor is happy about the consequence or not, and to claim that that this consequence was somehow not something that the actor chose and intended in the broader sense is incorrect, and perhaps almost cowardly. In the “Trolley” case, for example, the two choices are to divert the trolley or to not divert the trolley. We know full well that diverting the trolley will kill one person. Yet many of us still believe that it is the right thing to do to divert the trolley, knowing that to do so would be to kill the person. Somehow trying to

squeeze out of the responsibility of causing the person's death by saying that it was merely foreseen and not intended, not a direct consequence of an action that was chosen knowing that it can be reasonably expected to lead to death, is incorrect and mere wordplay. We know just as well in "Trolley" as in "Bridge" that we are causing a death – to make a distinction between foreseeing and intending in the broader sense in these cases merely describes what we wish to be the case considering our intuition, rather than describing the actual reality. So this sense of the word "intend" seems irrelevant morally.

Another mental state that "intent" can refer to is in how much a person actually delights in the harm that will result from an action. In this sense of the word, a person intends a consequence if he enjoys the fact that the consequence will occur. For example, let's assume that it is permissible to bomb a military target despite the fact that civilians will die as well. If we say that this action is permissible, then it seems that we are committed to say that the action itself is permissible even if the pilot performing the air raid might take pleasure in the fact that civilians will die. It seems strange to even think in the first place that the permissibility of an action can change based on what is going on inside the actor's head. The pilot's intention could not bomb the military target any more, nor could it kill civilians any more (assuming that the pilot's delight does not lead him to kill more civilians than those who would die unavoidably) – the difference in what is inside his head is not manifested in any way in the world. To me, it seems to be overly abstract and philosophical to say that intention can affect the nature of our action to the point where we should ignore the real benefits of the action. Scanlon's way of addressing this case seems quite insightful here. "Suppose you were prime minister, and the commander of the air force described to you a planned air raid that would be expected to destroy a munitions plant and also kill a certain number of civilians, thereby probably undermining

public support for the war”.²⁸ As prime minister, it would be strange to ask what the intention of the people flying the mission is and whether it is to destroy the plant while seeing the civilians’ death as a side effect or whether it would be to kill the civilians.

Besides these mental states, it seems that we also say that a person intends harm based on the structure of the action. In this sense, a person intends a consequence when the consequence is instrumental to the person’s objective in some way. But this instrumentality sense of the word seems morally irrelevant to me as well. A person who defends that the narrower sense of the word “intention” is important morally might say something like this:

In “Trolley”, diverting the trolley, and not the death of the one person, is what is instrumental for saving the five – if the person could somehow have superhuman reaction speed (if the case is presented with the person simply standing on the track and reacting to the diverted trolley) or if the person could escape his ropes somehow (if it is presented with the person tied to the track), the five would still be saved. In “Bridge”, however, the person who is pushed needs to be killed, i.e., his death is instrumental and intended in the narrower sense. This difference, then, makes the difference between permissibility and impermissibility.

The situations in which the one person may survive in “Trolley” are presented in order to show that the person’s death is not directly instrumental to saving the five people. But of course, the survival would depend on some freak circumstances, so it only seems fair to provide a freak circumstance for “Bridge” as well. In “Bridge”, what if the trolley hits the person and stops before hitting the five people, but the pushed person somehow survives in a freak incident

²⁸ T. M. Scanlon, *Moral Dimensions: Permissibility, Meaning, Blame*. (Cambridge, MA: Harvard University Press, 2008).

similar to the stories of people who fall from airplanes and survive? If this happened, clearly the death of the person is not instrumental to saving the five. Rather, what was necessary was merely that something stops the trolley. In other words, the death was a foreseen consequence of the action of pushing something (well, in this case, a person) in front of the trolley to stop it, rather than an intended one. This mirrors the argument for “Trolley” saying that the death is merely a foreseen consequence of the action of diverting a trolley onto a track with a person on it rather than an intended one.

Of course, I am not laying out this example argument for the permissibility of pushing the person in “Bridge” because I think it is a good argument to demonstrate the permissibility of pushing the person. Rather, I am laying it out to attempt to show that an argument of this form can be made for any scenario by focusing down enough onto one aspect of an action such that the death is no longer instrumental, be it “Trolley” or “Bridge”, and thus that these arguments are not useful. Such arguments stray away too far from reality. It makes more sense to consider what can be reasonably expected as a consequence of an action rather than argue about whether that consequence was actually intended or not. In other words, it is the consideration of to what extent a consequence is expected or how much risk of harm is acceptable to assume in order to bring about some good effect that is the factor explaining the differences in the way we feel about these cases, rather than some shift between intending and foreseeing a consequence.

One might make the point here that if intent is really about risk and certainty of harm as I am suggesting here, intent should still be relevant morally because risk and certainty of harm are important morally. But while I agree that risk and certainty of harm are morally relevant, it would not make sense to take into account risk and certainty of harm as relevant in the cases that we are talking about here. Once the certainty of harm reaches a certain threshold, small

differences hardly matter. In other words, though we might say that “Bridge” is a case of intended harm because the certainty of death is 99.99% which is more than the 99.98% certainty of death in the foreseen harm of “Trolley”, to claim that this small difference matters seems absurd. Therefore, intent in the instrumentality sense seems irrelevant morally as well.

So if intention and foreseeing do not affect the permissibility of an action, why does it seem to matter to us? One reason I believe the doctrine of double effect initially seems to hold some truth is because of its real or imagined connection to subjective effort. This connection can be observed when we go through each meaning of the word “intend” and compare someone who “intends” harm and someone who “foresees” harm.

In the first mental sense in which intent seems to mean something like taking ownership of a consequence that one knows will follow from his action while foresight seems to mean something like distancing oneself from the known consequence, we see that a person who foresees harm is someone who wants to see some distance between the harm and his action, and therefore has a harder time doing an action if he feels connected to the harm and would presumably have to use more subjective effort to do harm. I described earlier that that having to use more subjective effort to do harm is praiseworthy, so this person may feel that intent matters because of this sort of effect. Of course, to me, such a person ironically seems cowardly because he falsely tries to distance himself from a harm that he knows will follow from his action. Regardless, if this person’s false idea that he really could distance himself more from his “foreseen” harm than another person who intends harm could, this sort of justification using subjective effort seems to be what is floating around in this person’s ethical subconscious.

In the second mental sense in which intent means enjoying a harmful consequence and taking the harm to be an additional motivator to perform an action while foresight means having an attitude of regret towards the harm, the person who enjoys harm does not have to expend any subjective effort at all. This, along with the additional enjoyment that this person derives from the harm makes him as a person more blameworthy regardless of the actual permissibility of his action. The person who foresees the harm, on the other hand, does not want the harm to occur. Therefore, such a person would have to expend subjective effort and bring himself to cause this harm. This additional subjective effort he expends makes him less blameworthy for the harm he causes than the person who intends harm who spends little to no subjective effort. Unlike the above mental sense of the word “intent”, this mental sense seems to be one that does in fact matter morally in the way we look at the person performing an action; it just does not matter in determining the permissibility of an action.

Lastly, in the instrumentality sense of the word “intent”, a person who intends harm sees the harm as instrumental in some way to the person’s objective while a person who foresees it does not. We saw earlier, however, that “intent” here really meant a perceived greater certainty of harm, so we might also say that the person who intends harm does a harm that is more certain than a person who foresees it. Assuming that both people spend an equal amount of subjective effort in order to perform their respective actions, the person who intends harm seems more blameworthy than the person who foresees harm because the former spends the same amount of subjective effort to do a more certain harm, implying that he would spend less effort to do a harm that is less certain like the one the latter causes. If the difference in certainty of harm between these cases did indeed matter morally, then the proponent of the difference between intending and foreseeing harm in this sense could hold this plausible position that would give reason for us

to evaluate the two actors differently based on whether they intended or foresaw harm. But of course, as I argued above, the certainty of harm between the intending harm case and the foreseeing harm case seem to cross a threshold point where any small difference between the two is irrelevant. So any reason to evaluate the intender and the foreseer in this instrumental sense of the word differently based on subjective effort would be illusory.

Conclusion

I hope that I have been able to shed some light on the importance of effort, a topic which seems to me too often ignored. Though I do not pretend that effort is the only important moral factor, I believe that the way in which it plays a role in the way we understand what it means to do something and what it means to allow something; the way in which subjective effort affects how we evaluate people's character; and the way in which objective effort plays a role in magnifying the goodness or wrongness of an action and the moral impermissibility of allowing harm are just some examples the widespread relevance of effort. As a demonstration of this relevance, I have shown two example discussions, the doing/allowing harm distinction and the doctrine of double effect, in which it seems to me that applying an understanding of the moral importance of effort can be of great help. Additionally, I also imagine that if we put in the effort, we would see how this widespread relevance would cause effort to show its face not only in these examples, but also in other important philosophical discussions, in our day-to-day lives, and even in TV shows like *It's Always Sunny in Philadelphia*.

Bibliography

Thomas Aquinas, *Summa Theologica*.

Elizabeth Ashford and Tim Mulgan, "Contractualism", *The Stanford Encyclopedia of Philosophy* (Fall 2012 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/fall2012/entries/contractualism/>.

Roy E. Baumeister, Ellen Bratslavsky, Mark Muraven, and Dianne M. Tice, "Ego Depletion: Is the Active Self a Limited Resource?" *Journal of Personality and Social Psychology* 74:5 (1998), pp. 1252-65.

Jonathan Bennett, "Acting and Refraining," *Analysis*, 28 (1967) pp. 30–31.

Jonathan Bennett, "Morality and Consequences," *Tanner Lectures*, (1980) pp. 47–116.

Randolph Clarke, *Omissions: Agency, Metaphysics, and Responsibility* (Oxford: Oxford University Press, 2014).

Frances Howard-Snyder, "Doing vs. Allowing Harm", *The Stanford Encyclopedia of Philosophy* (Winter 2011 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/win2011/entries/doing-allowing/>.

Rosalind Hursthouse, "Virtue Ethics", *The Stanford Encyclopedia of Philosophy* (Fall 2013 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/fall2013/entries/ethics-virtue/>.

Daniel Kahneman, *Thinking, Fast and Slow*, (New York: Farrar, Straus and Giroux, 2011).

Warren Quinn, "Actions, Intentions, and Consequences: The Doctrine of Doing and Allowing," *Philosophical Review* 98:3 (1989), pp. 287-312.

James Rachels, "Active and Passive Euthanasia," *New England Journal of Medicine*, 292 (1975), pp. 78–86.

T. M. Scanlon, *What We Owe to Each Other* (Cambridge: Cambridge University Press, 1998).

T. M. Scanlon, *Moral Dimensions: Permissibility, Meaning, Blame*. (Cambridge, MA: Harvard University Press, 2008).

Edward Slingerland, *Trying Not to Try: Ancient China, Modern Science and the Power of Spontaneity*, (New York: Crown, 2014).

Kelly Sorensen, "Effort and Moral Worth," *Ethic Theory Moral Prac* 13 (2010), pp. 89-109.

Kelly Sorensen, "The Paradox of Moral Worth," *Journal of Philosophy* 9 (2004) pp. 465-483.