

University of Richmond

UR Scholarship Repository

Honors Theses

Student Research

5-6-2021

Doing Good Better: The 'Do Good' App

Emily Trumble

University of Richmond

Follow this and additional works at: <https://scholarship.richmond.edu/honors-theses>



Part of the [Leadership Studies Commons](#)

Recommended Citation

Trumble, Emily, "Doing Good Better: The 'Do Good' App" (2021). *Honors Theses*. 1544.
<https://scholarship.richmond.edu/honors-theses/1544>

This Thesis is brought to you for free and open access by the Student Research at UR Scholarship Repository. It has been accepted for inclusion in Honors Theses by an authorized administrator of UR Scholarship Repository. For more information, please contact scholarshipprepository@richmond.edu.

Doing Good Better: The ‘Do Good’ App

by

Emily Trumble

Honors Thesis

in

Leadership Studies

University of Richmond

Richmond, VA

May 6, 2021

Advisor: Dr. Marilie Coetsee

Abstract

Doing Good Better: The 'Do Good' App

Emily Trumble

Committee Members: *Dr. Marilie Coetsee, Dr. Geoff Goddu, Dr. Terry Price*

Although we all want to do the right thing, doing so reliably is often extremely difficult because we do not always know what the right thing to do is. Personal biases undermine our ability to perceive information in a way that does not favor ourselves. Additionally, even without those biases we often do not have enough information to make adequately informed decisions. The solution to this problem is the Do Good App, an app which would assist people in making the correct decisions. Unlike humans, the Do Good app would not be biased and it would have easy access to limitless information. As a result, individuals who use the Do Good app would make more reliable decisions.

While the app would have many benefits, there are also potential objections. First, there is the problem of which moral theory the app should rely on to make its recommendations. The three main first moral principles that I assume the app will allow users to use, and that I will focus on, are utilitarianism, Kantianism, and Virtue Ethics. The second set of objections concern whether each moral theory would support the use of the app. Although each moral theory may have objections to the app, I will argue that with adjustments there are versions of the app that each moral theory would support.

Signature Page for Leadership Studies Honors Thesis

Doing Good Better: The 'Do Good' App

Thesis presented

by

Emily Trumble

This is to certify that the thesis prepared by *Emily Trumble* has been approved by his/her committee as satisfactory completion of the thesis requirement to earn honors in leadership studies.

Approved as to style and content by:



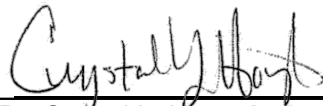
Dr. Marilie Coetsee, *Chair*



Dr. Terry Price, *Member*



Dr. Geoffrey Goddu, *Member*



Dr. Crystal L. Hoyt, *Associate Dean for Academic Affairs*
Jepson School of Leadership Studies

Introduction

Think for a moment about all of the decisions that you've made today. Before even leaving the house, you probably chose when to get up, if you were going to turn your light on, what clothes to wear, if you were going to brush your teeth, what to eat for breakfast, and a million similar things. Then think of all the decisions that you made prior to today that influenced your decisions today. You chose a job that requires you getting up at a certain time. When you bought your house you decided what kind of lighting would be in your room. Last week you went to the grocery store and picked out what kind of toothpaste and breakfast supplies that you were going to want this week. It is clear that our decisions are incredibly casually interrelated. The decisions that you make today impact the decisions that you will be able to make tomorrow.

The decisions you make today also impact the decisions that other people will be able to make tomorrow. Choosing when to get up impacts what time you will be driving on the roads which impacts other people's drives to work. Choosing to turn on your light affects light pollution and therefore other people's ability to see the stars. Buying toothpaste tubes instead of toothpaste tablets adds to the amount of waste that you are producing and contributes to climate change. Buying sausage for your breakfast supports the production of cows and therefore methane pollution which also contributes to climate change. With the increase in technology that we have seen in the last two hundred years the world has gotten very small: our actions have casual consequences that affect not only our immediate communities but also people across the globe. Moreover, those causes often have significant moral weight due to how largely they can affect people's lives. With the current and quickly growing technology, you can so easily buy something on Amazon from across the world and that purchase affects the people who create,

produce, ship, and eventually dispose of that product. This increase in technology has also accelerated the climate crisis through purchases and shipments of this sort. The effect that our decisions have on each other makes every decision an ethical decision. Since every decision has the potential to impact other people, it is incredibly important that we make good decisions.

As I've suggested, so often decisions that we think of as mundane or ordinary do end up having ethical consequences. For instance, we regularly purchase from Amazon, almost without thinking about it and those purchases can affect people whom we'll never see. However, making a good decision in cases like this is incredibly difficult because there are so many unknowns to us. In the case of buying from Amazon, we often do not know who is making or shipping the product we are buying or what their pay or working conditions are. Additionally, even if we did know those facts it is often still difficult to weigh our options. For instance, if you choose to buy from a 'fair trade' producer on Amazon it is going to be incredibly expensive, and--since it is easier to get this expensive 'fair trade' certification in moderately developed countries--profits from 'fair trade' purchases may not benefit those who are most in need.¹ Is it better to buy that 'fair trade' product or to use that money to buy something else? Is it even better to take the time to go to an in person store to buy what you are looking for or is it better to use that time doing something else? Even in less mundane or ordinary situations we still often do not know the exact context and consequences of our decisions are unknown to us. Determining the possible results of our decision requires knowing that context, but gathering that information takes a considerable amount of time. Determining what the roads will look like on your way to work, who else will be driving at the time, and how driving at that time will affect all of those people requires experience driving on the roads at a certain time every day, but it also requires specific hard to

¹ Sylla, Ndongo Samba. "Fairtrade is an unjust movement that serves the rich" *The Guardian*, (2014)

gain information about what other people are planning on doing that day. Deciding the most environmentally friendly light sources in your home takes research as does deciding what kind of toothpaste to buy or what to get for breakfast. Everything we buy or consume has impacts on everyone in that supply chain as well as in the disposal of the product after we use it. Knowing that full chain before and after the product comes in our hands takes a significant amount of research and that kind of research would need to be done for every product we use.

Another issue that makes this problem even worse is that in addition to gathering the necessary information, it is also necessary that we sort through that information and apply it in a way that is unbiased. Often we have unconscious personal biases that negatively impact our ability to filter and process this information. For instance, we tend to assume that it is more effective to mitigate costs that are closer to us rather than further away. Therefore, worrying about climate change in the future or the ethical standards of sweatshops across the world does not immediately concern us when we weigh that information. Though we could possibly overcome these biases, doing so would take time to address so that they do not affect our ability to apply the information. Therefore, both gathering the information and dealing with the biases take time. Taking the time necessary to do this for every single small decision we make would take a considerable amount of time--time that we simply do not have for every decision. Even if it was possible to make the time, there would be other negative externalities to using all of our time thoroughly examining every decision we make.

The fact that we are very bad at constantly making decisions and those decisions have cumulatively big impacts can easily cause us to become discouraged and simply give up trying to live a fully ethical life. However, it would be really bad if we all decided to give up trying to live ethical lives and so we need a solution. There are many areas of our lives currently where we use

computers and algorithms to help us make decisions and in this thesis I will argue that we should leverage that power to make better ethical decisions. One example of algorithms being used to make decision making processes easier for humans is the way that insurance providers use an algorithm to help clients determine which health care package is best for them. Clients are asked multiple questions about the status and history of their health, and from that information the algorithm is able to place them with the health care package that makes the most sense for that client to choose. The algorithm takes the inputs that the client gives it and applies those to the information it already has about the given packages. In this case, the algorithm is not choosing for the client, rather it is only giving them a suggestion about which package would be best for them, given the clients own health goals. This suggestion then allows the client to make a more informed decision for their health care. There are many other examples of this kind of algorithm such as algorithms used for housing decisions or in applying for credit cards. These algorithms can have their problems, but they start us in a good direction in thinking about how to solve our problem with making ethical decisions. Since the technology is becoming more and more advanced for helping us systematically make decisions such as what credit card to get and what health insurance to buy, perhaps they can also help us more effectively make ethical decisions.

Through this essay I will argue that the solution to our poor decision making problem is the ‘Do Good’ app, an app that would help individuals apply their good moral principles to complex situations in a way that helps them more reliably make the correct moral decisions.² Unlike humans, the ‘Do Good’ app would have easy access to limitless information and not be biased in its filtering and processing of this information. As a result, the ‘Do Good’ app would help individuals make ethical decisions more reliably. There are a lot of ways that the app would

² Clearly this ‘Do Good’ app does not yet exist, but throughout the paper I will be treating it as though it does. Later I will touch on the feasibility and logistics of such an app.

be able to help people act ethically. The most obvious way is that the app is able to process far more information than people are which enables it to be more aware of the full potential consequences of an act. Knowing the full context and effects of the act allows the decision to be more informed. Another way that the app enables people to act more ethically is that it pushes them to determine what their first moral principle is. Before downloading the app, many people will not have thought about ethics in a principled or systematic way. When most people try to do what is right they do so based on their intuitions in one given scenario, and these intuitions may not be consistent between cases. Rather than users simply making decisions based on scattered intuitions about ethical principles, the app will have each user go through a process of answering questions about principles and proposed scenarios (like iterations of the Trolley Problems, Singer's drowning child example, or the Ursula Le Guin's "The Ones Who Walk Away from Omelas"³) to find out what moral principle they are most fundamentally committed to. The app would give people an opportunity and a structure for processing their intuitions and boiling them down to a consistent and logical principle from which they can act in future scenarios.

While the app would have many benefits, there are also potential objections. First, there is the problem of which moral theory the app should rely on to make its recommendations. If the app is designed to help individuals make more ethical decisions, why should it give them a choice as to what fundamental moral principle to use in the app? Shouldn't it simply direct them to the truth? In Chapter 1, I argue that peer disagreement over which moral theory is true means that--within a range--the app does not have a good enough reason to decide for the user which moral theory would be used. The three main first moral principles that I assume the app will allow users to use, and that I will focus on, are utilitarianism, Kantianism, and Virtue Ethics.

³ Le Guin, Ursula K. "The ones who walk away from Omelas." (1975).

The second set of objections concern whether each moral theory would support the use of the app. In Chapter two, I argue that a utilitarian would support the use of the app because the app maximizes utility by minimizing the errors that humans make in their decisions and speeding up the process by which the user makes decisions. I will consider utilitarian objections such as the possibility of the app not truly maximizing utility because of the potential that the app tells the user the wrong thing to do. I will then consider the potential importance of the user being able to question the app and give an argument for the reliability of algorithms for handling these kinds of decisions. I will also consider an objection that the app turns the user into a robot, ridding the user of any of the kinds of “higher” pleasures that John Stuart Mill argued comes from determining things for ourselves. In response to this I will argue that the utility gained in some cases of the user turning into a robot will be higher than the utility lost by them losing that opportunity for higher pleasures. Ultimately I will argue that for the utilitarian version of the app, the app will need to weigh the user’s willingness to be controlled, their likelihood of making the correct decision on their own, the risk of making the wrong decision, and the possibility of gaining higher pleasures from the user making the decision on their own. All four factors will be important to be considered for the app to determine how much it should control the user.

The utilitarian version of the app seems the most intuitive because it is focused on getting the user to simply do the right thing. However, for Kantianism, it is also important that the user do the right thing *for the right reasons*. Therefore in Chapter Three, I then argue that a Kantian could support a version of the app that maintains the user’s autonomy by giving them tools to ensure that they are acting from the right reasons. Objections to the Kantian version of the app will include the objection that the app turns morality into a technocratic game where people stop doing the acts for the right reason and therefore the acts done purely on the basis of the advice of

the app would not have moral worth. In response to this objection I will argue that the app will need to be modified from the utilitarian version in order to give the user the ability to learn what the reasons behind the correct action are.

Whereas utilitarianism argues that the user has to do the right thing and Kantianism argues that the user has to do the right thing *for the right reasons*, Virtue ethics argues that the user has to do the right thing, for the right reason and *with the right feelings*. Therefore in Chapter Four I will argue that a Virtue Ethicist might support a version of the app that helps the user become a virtuous person--who both acts for the right reasons, and with the right feelings--rather than just focusing on doing the right things. I will raise an objection that the app cannot determine what the right thing to do is because that requires a kind of emotional disposition that only moral exemplars possess and that the app is not capable of manifesting or teaching. I will also object that the Virtue Ethic's theory is not codifiable and therefore an app's algorithm could not be programmed to give advice on the basis of it. Although each moral theory may have objections to the app, I will argue that with adjustments there are versions of the app that each moral theory would support.

I will conclude by arguing that although the app may not be a complete replacement for independent moral thinking and though there are additional issues with the practical funding and creation of the app, it would still be worthwhile for someone to pursue the development of the app due to the enormity of our problem with making good decisions in our increasingly interconnected and complex world. Although the app may not make us perfectly good, it can make us better.

Chapter One: The Problem

I. The Ethics of Belief: Sally and the Shipowner

Sally wakes up for work one Monday morning. She works at a big executive firm making 50 thousand a year. She likes her job enough. It pays the bills. Today, she has a big presentation with a new client. She stands to earn an additional two thousand five hundred dollars from this client, so she dresses in her nicest suit and shoes. Before she heads out the door to her presentation she checks her mail and sees a fundraising flyer from Wasti, which is a highly effective organization that pairs patients needing life-changing surgeries with people who can fund them. The flyer shows a picture of a girl named Joyce from Kenya who needs three thousand dollars for a spinal surgery that will save her life. She quickly glances over the flyer. She is briefly tempted by the potential appeal of the flyer but quickly discounts it. “These charities tend to be all for show. The money is hardly ever used effectively,” she thinks. She hesitates, because she realizes that it is possible that this charity is different, but she looks at the clock over her oven and realizes she is running late for her meeting with the new client. She thinks, “I was hoping to buy a new cell phone with that money anyway.” So she leaves the flyer on the counter with the other junk mail and gets in the car.

Sally drives her Jeep to work and on her way she passes a pond that she passes every day. She stops at a traffic light next to the parking lot for a park near her house. As she's looking out her window she sees a child in the water, flailing their arms, clearly drowning.⁴ There are other people around but no one seems to see the child. Sally was a lifeguard as a teenager and has swam in this lake countless times before, so she knows that she will be able to successfully save

⁴ Singer, *The Drowning Child and The Expanding Circle*, 1997

the child if she attempts to. Sally hesitates for a moment. She is already running late and she knows that if she gets out of the car to go save the child she will miss the meeting with the client, lose the two thousand five dollars, and also ruin her suit. Yet, Sally's moral principles tell her that she should save a life when she can. She can see that in this situation her moral principle dictates that she should save the child's life. Although it would be nice for her to gain the client and the two thousand thousand five hundred dollars, she does not need that money to live. She can buy all her necessities with her salary and would not be putting herself in jeopardy by not having the client. So, Sally pulls the car into the parking lot and runs into the water to save the child. Sally saves the child's life. Sally believes she made the correct decision, even despite ruining her five hundred dollar suit and losing the client that she stood to gain a two thousand five hundred dollar contract with.

Sally's two choices are inconsistent. Sally believes that she should save a life when she can. In the second case Sally was correctly acting in line with this moral principle. She was able to save the drowning boy's life and she did. However, Sally also had the capacity to save a life in the first case, and she failed to do so. Therefore she failed to act correctly in accordance with what her moral principle demanded of her.

Part of the reason for the inconsistency in Sally's choices is that Sally lacks knowledge in the first case that she has in the first case. In the second case Sally knows what is true about the situation. She knows that she has the moral principle that she should save lives when she can and she knows that in this instance she can save a life. Sally is able to analyze the situation clearly and act in line with what she knows to be true. However, in the first case, Sally has obstacles in the way of her getting to the truth. She does not have complete knowledge of the situation, since she does not know that the charity is effective. Compounding the problems with Sally's lack of

knowledge, Sally also has personal biases that get in the way of her getting to the truth. She has a personal bias towards helping herself and buying the new phone.

One classic critique of Sally's actions in the first case would come from William Clifford. In "The Ethics of Belief," Clifford describes a shipowner who is about to send his emigrant-ship off to sea.⁵ However, he knows that the ship is old and has needed repairs, which has caused him to doubt if the ship would be safe at sea. Although he has these doubts, he succeeds in stifling his doubts. He decides to trust 'Providence' and comes to fully believe that the ship will be safe. After convincing himself that the ship will be safe, the ship-owner decides to send the ship to sea along with all the emigrant families that were on board. However, the ship sinks in the middle of the ocean killing everyone on it.

Clifford argues that the shipowner is responsible for the sinking of the ship because he did not have sufficient evidence to think that the ship was safe and yet he sent it out anyway. This would have been one thing if the shipowner had not had a sufficient opportunity to become aware of the possible problems with his ship. For instance, if the shipowner had had no reason to suspect that the ship may need repairs, then the shipowner could not be blamed for acting on the insufficient evidence. However, that was not the case for the shipowner. What really makes the ship-owner's decision appear especially problematic in this case is that he had a suspicion that there could be something wrong, and nevertheless chose to stifle his doubts. Moreover, he presumably chose to do so for selfish reasons. Although Clifford does not go into detail about why the ship owner chooses to stifle his beliefs, the circumstances of the case suggest that he may have been eager to profit from letting the ship sail, or that he may have not wanted to pay

⁵ Clifford, William K. "The Ethics of Belief." *Contemporary Review*, (1877), 1.

the costs of repairing the ship. Whatever the precise reason was, the ship-owner chose to believe the ship was safe because that was what was best for him.

In the decision to allow the ship to sail or not the truth about the safety of the ship really mattered. The truth about the ship had a large impact on other people and so it seems that the ship-owner should have taken a greater interest in the truth. Given the importance of the truth to others' well being, he should have made a greater effort to discover truth rather than to prioritize his own interests. Thus, the shipowner's choice to "stifle his doubts" seems to make the ship-owner morally blameworthy for his choice to set the ship to sail.

It seems pretty clear that the shipowner that Clifford describes is in the wrong. After all, his own selfish biases seem to lead to the death of everyone on his ship. Even though Sally's choice to ignore the flyer for Watsi and not donate seems more ordinary than the shipowner's choice to send the ship to sea, Sally could also be seen to be at fault for similar reasons as the shipowner. Sally could save a life by donating to Joyce's fund for spinal surgery, but Sally is not aware of this truth at least in part because, like the shipowner, she does not want to be. All the information that she needs for reaching the truth is in front of her, but she lets her desire for a new cell phone cloud her judgment and keep her from further investigating into the possibility that she may be able to save a life through Watsi. Like the shipowner, Sally is not sufficiently motivated by the truth.

II. The Challenge of Believing Ethically: The Epistemic Regress Problem

We might think that the way to fix the moral problem with both Sally and the shipowner's choice is to morally require them to take more time to consider the evidence for their choices, rather than rushing into a response that is influenced by personal biases. While this

may seem like a straightforward requirement, however, it is not clear that it is a practical response. The goal of requiring them to take more time to consider their beliefs seems to be to ensure that they get at the truth about their judgments (about the ship and about Watsi, respectively). But one problem that arises with imposing a requirement that is aimed at this goal is the problem that they will inevitably face an epistemic regress. The epistemic regress problem starts with the observation that in order to know any piece of information is true we need to have evidence for its being true. However, in order to verify that that evidence is good evidence, we need further evidence to show that it is good evidence. That further evidence itself needs to be backed by yet further evidence, and so on. This creates a regress problem. Clifford argues that the ship owner is responsible for acting from good evidence, but the shipowner cannot guarantee the validity of every piece of evidence he uses to come to the conclusion. He can't guarantee that his evidence is true or getting him at the truth. For instance, the shipowner has this initial doubt about the safety of the ship, so he hires a ship safety inspector to examine if the ship is safe. However, why can he trust that inspector's opinion? Maybe because they are accredited in some way, but then what evidence does he have to support that that accreditation is to be trusted? It clearly initially seems very wrong that the shipowner does not look for more evidence beyond his personal intuitions, but Clifford is not clear how far one would have to go to be justified in believing the evidence they are using to come to the truth.

There are some cases in which the epistemic regress behind the beliefs that underlie our moral choices can be stopped, but those instances are rare. Sally's pond case is an example of a case where I will argue the epistemic regress problem can be adequately solved. Sally has direct experience with all the evidence that's relevant to her decision to save the drowning boy. Sally has swum in the pond many times and has vast experience as a lifeguard, so she can justifiably

feel confident in her ability to save the boy. She is basing her knowledge of the current situation on what she can see and hear at the time.⁶ However, most decisions are not like that. It is normally much harder to make correct moral choices because we do not have time to verify all the relevant evidence with direct experience.

In Sally's case of not giving to the charity, it does not seem obvious that she is doing something wrong. As mentioned previously, her choice to not donate looks less blameworthy than the shipowner's choice to send the ship to sea, and this is arguably in part because, in Sally's case, the challenges that face the shipowner in overcoming the epistemic regress problem are repeatedly replicated, given the many different moral choices and trade-offs she must constantly make. Once we consider the real world context that she faces-and perhaps also that the shipowner would also face, if we filled out the story more-we see more clearly that people have to make trade offs with their time in investigating all the possible epistemic regressions.

To illustrate that more fully imagine that when Sally sees the charity flyer she has three options: she can throw it away, keep it and look at it later, or stop and read the flyer now. Each of those choices has potential and unknown consequences.

If Sally chooses to read the flyer now then she will be late for work. She loses the potential new client and the money that comes with that. At this point Sally does not know what future decisions reading the flyer may lead to. She suspects that charities are not effective and so she thinks the odds that donating to the charity will help anyone are low. If Sally does choose to do research into whether or not the charity is effective, she will have to determine what evidence she is going to look for to know if the charity is effective or not. She may decide to google 'effective charities'. A thousand or more sites show up. She chooses to click on the first one, but

⁶ I am not going to doubt that Sally can trust her past experience and her senses, so we can assume that in that decision, Sally was justified to believe that she could save the boy.

now she has to determine the validity of this site. Again she has to determine how she is going to tell if this site is reliable. She could use another website, but then, yet again, she is going to have to determine the reliability of that site. This is similar to the epistemic regress problem that the ship-owner faces in determining if the ship is safe.

In deciding whether she is going to do more research or get to work on time, Sally has to rely on the information she already has and she already thinks that the charity is not effective. On the other hand, she has very good reason to think that she will lose the client, and if she does she will be being somewhat irresponsible with respect to her duties to her employer. So in her circumstance, choosing to trade off not reading the flyer in order to not be late for work is reasonable.

If Sally chooses to keep the flyer but look at it later, she has to make new trade-offs with her time. Say that tonight, Sally's friend is having a birthday party and Sally was invited. Once Sally gets home from work, she could read the flyer and do more research about the effectiveness of this charity or she could go to the birthday party. But as far as Sally understands, it's unlikely that that research would show that the charity is effective. On the other hand, she knows that missing the birthday would make her friend incredibly sad, and she could lose her friendship. So suppose she puts it off another night. Tomorrow she has to go grocery shopping. She could go home and do research on the charity, but she thinks it's important to go to the grocery store with ethically produced food across town, and if she does that, she won't also have time to do the research.

One could argue that if Sally's judgment was not at all clouded by her personal bias for wanting to spend her extra money on the phone, that she could and would eventually find time to do the research during her weekend. However, Sally has been working hard all week, and so if

she does not take a day off, she may be exhausted during the next work week. If she is exhausted during her work week, then Sally will not be able to do an effective job at her work. Again, the trade off is the same. Sally can be fairly certain that if she takes her day off to do research then she will not be effective at work, but she is not certain that the research will lead to anything productive. So, Sally is again not unreasonable for deciding to take the day off rather than do research into the charity.

Sally's decision may not be fully justified in the sense that it may not be able to be fully defended on the basis of good moral principles and reasons. Nevertheless, given all the complexities Sally faces, we can understand why Sally might be tempted by the convenience of just throwing the flyer away. Although Sally has good moral principles, the complexity of her choices makes it very difficult for her to apply those moral principles correctly. So, because the situation is so complicated and there are so many trade-offs involved, it seems that she may be morally excused for her decision. She is excused for making the wrong decision because she does not have enough time to get enough information and determine her personal biases without it negatively impacting her ability to make other morally justifiable actions.⁷

In Sally's case there are multiple different biases that could be at play, but I will break it down into two separate ones. First, Sally prioritizes herself over other people. She is inclined to spend money on buying a phone for herself than on someone else. The second bias that affects Sally is a kind of 'availability heuristic'. Availability heuristic is the bias towards believing that information that is more salient is truer than information that is not salient.⁸ Because the drowning boy is right in front of her, Sally is able to experience more of the emotions needed to

⁷ Walzer, Michael. "Political action: The problem of dirty hands." *Philosophy & public affairs* (1973): 160-180.

⁸ Esgate, Anthony; Groome, David . "An Introduction to Applied Cognitive Psychology". *Psychology Press*. (2005), 201.

motivate her to help the drowning boy right in front of her rather than someone she cannot see. She believes that there is more of a need to save the drowning boy because he is in front of her. Both of these biases are things that are a part of Sally. They are mechanisms that humans have developed for survival. However, in Sally's case, they are obstacles to the truth.

The overwhelmingness of information makes her especially susceptible to the effects of these biases. We have so many deeply ingrained personal biases, many of which we are not even fully aware of. These biases get in the way of us reasoning properly to determine what to do. They affect what we see, what we think, and therefore what we do. Having such deeply ingrained personal biases means that it will take time to unravel them. There is already a limited amount of time which is why we've developed certain biases, but it takes even more time to unravel them.

III. Doing Good Better: The Do Good App

The factors that are keeping Sally back and the shipowner back are similar, and these same factors also influence our everyday decisions. The ship owner and Sally have two particular things in common: a lack of knowledge and personal biases. The ship owner does not know that the ship is not fit for sailing and Sally does not know that the charity is effective at using the donations. The ship owner is biased by his ownership over the ship and Sally has a personal bias towards helping herself rather than donating. It is clear that the ship owner stifled his beliefs. He had the thought, "this ship may not be safe for sailing" but then chose to not believe that because it was not what was best for him. The same could be said for Sally. She saw the flyer and had the thought "this charity may be worth donating to" but then she chose to not believe that she should research the charity because that was best for her.

We make decisions like Sally all the time. We often move emails advertising ways to help people to the junk folder before fully investigating the opportunity, or throw away hundreds of disposable water bottles instead of using a reusable, or buy products with horribly unethical supply chains. In all of these instances, the same things that are keeping Sally (as well as the shipowner) from making a good moral decision also keep us from making good moral decisions. We lack knowledge and have a limited amount of time to try to get it. Because there's limited time and an overwhelming amount of decisions to be made, we inevitably have to rely on simplified strategies for making our decisions, and in light of this it is understandable that we sometimes (wrongly) become biased by self-interest.

Clifford would respond to what I'm saying by saying that we, and people like Sally, should simply not believe anything. Clifford writes, "it is wrong always, everywhere, and for anyone, to believe anything upon insufficient evidence" and that if you don't have time to check things out you simply shouldn't believe⁹. However, this is an extremely high burden and Clifford knows that people will find this burden unreasonably high. Clifford says that one might complain, "I am a busy man; I have no time for the long course of study which would be necessary to make me in any degree a competent judge of certain questions, or even able to understand the nature of the arguments."¹⁰ Clifford's response to that complaint is that "Then he should have no time to believe."

However, this is not a practical response because if we simply stop holding any beliefs that we are not certain of, then we will have no basis for action, and not acting can cause just as much harm as acting on an uncertain belief. Sally's case illustrates that. In all of her decisions Sally risks doing harm or not acting optimally. If Sally tries to meet that standard of evidence for

⁹ Clifford, William K. "The Ethics of Belief." *Contemporary Review*, (1877), 5.

¹⁰ *Ibid.* 6.

every decision she makes she is going to spend all her time just analyzing the decisions. If she sits there trying to choose between her three options (research the flyer now, research the flyer later, or throw it away) on the basis of sufficient evidence, then she would do much more harm than any of the choices. She will not donate, not give to the charity, and lose her job. So if she chooses to not act she's still risking harm and not acting optimally. However, if we really listened to Clifford, then we would be morally obligated to choose to not act. But she would still also be wrong! She still would be doing the act that was not sufficiently based on evidence. According to Clifford, in Sally's case there would be no right thing for her to do.

We are constantly in situations where we must be making decisions. Like Sally, our ability to get to the truth cannot keep up the demand for action. On the other hand, choosing not to act has its own moral harms. So if we're interested in avoiding moral harms, we can't just fall back into the option of not acting.

Given the complexity of our choices, how can we possibly be counted on to make good decisions? The solution is the 'Do Good' app, an app that would help individuals apply their good moral principles to complex situations in a way that helps them more reliably make the correct moral decisions¹¹. The shipowner and Sally's decisions are bad because they are not made with necessary information and they are influenced by personal biases. However, the app would be capable of "reasoning" without any personal biases and with almost perfectly complete knowledge, through artificial intelligence and access to the internet. The app would also be fast, getting rid of the overarching constraint of time. When you were faced with an ethical decision, you would put the variables of that decision into the app and it would tell you what the correct thing to do is based on your core moral principles. When faced with a decision like Sally's

¹¹ Clearly this 'Do Good' app does not yet exist, but throughout the paper I will be treating it as though it does. Later I will touch on the feasibility and logistics of such an app.

decision to donate to the charity, you could input the variables of the situation into the app: how much money you make, what the charity is, etc. The app would then tell you what choice you should make based on the evidence. Throughout this thesis I will break down the components of the app and what makes it a solution to our problem.

IV. Managing the Effects of Moral Disagreement

Philosophers like Jason Stanley (2005) and John Hawthorn (2004) believe in the knowledge norm of action meaning you need to know something before you act (see also Hawthorn and Stanley 2008).¹² Since knowing something requires having a true belief about it, one might think that when you ask the app what you should do in a particular situation, the app should simply give you the truth about you what to do, regardless of how your own moral principles happen to apply to the situation. However, the problem with this version of the app is that even with no bias and perfect information about empirical facts, the truth about how one should act is still hard to determine.

In the example of the ship-owner, whether or not the ship was safe to sail was an empirical issue. For the sake of this paper I will assume that for all empirical questions there is either agreement amongst experts or with enough time and the filtering out of personal biases experts would eventually agree. However, that is not the case for normative questions. The Kantianism and utilitarianism debate is one example of that kind of disagreement in ethics amongst experts.

The disagreement between ethics philosophers over Kantianism and utilitarianism is considered is disagreement between epistemic peers: on both sides of the disagreement are

¹² Stanley, Jason. *Knowledge and practical interests*. Clarendon Press, 2005., Hawthorne, John, and John P. Hawthorne. *Knowledge and lotteries*. Oxford University Press, 2004., Hawthorne, John and Jason Stanley. 2008. "Knowledge and Action." *Journal of Philosophy* 105: 571-590.

philosophers who are equally educated, equally dedicated, equally intelligent, etc.¹³ In these kinds of disagreements, both sides are equally qualified to have an opinion on the topic, and so there is no obvious reason why one should choose to side with one moral expert than another. In some cases of peer disagreement, we can settle the disagreement by seeking out an expert who has superior knowledge to the peers. However, in the case of basic moral disagreement between Kantians and utilitarians, this is not a possibility. This is a disagreement not only between peers, but also between peers who are experts. This kind of disagreement amongst moral experts is different from most other kinds of disagreement, because the truth about the disagreement cannot be determined based on which of the arguers is more qualified.¹⁴ Disagreement of these sorts can lead us to believe that there is something inherent about the nature of the problem that they are disagreeing over that is causing the disagreement, rather than that one person has a better reason to believe their side than the other.

In light of peer disagreement between utilitarians and Kantians, we cannot be certain whose basic moral principles are true, and so who knows the right thing to do in a given situation. However, even without certainty we may still be epistemically justified in believing and this justification may give us adequate reason to act. We get epistemic justification through experience, evidence, and other reliable indicators of the truth. Thus, if you have epistemic justification for your belief, you hold your belief for good reasons, rather than by luck. So long as you are epistemically justified in your beliefs, you are justified in the choices that you make on the basis of those beliefs. For instance, in Clifford's shipowner case, if the ship owner had done a thorough check of the ship and had good reason to believe that the ship would be safe,

¹³ For an overview of the literature on peer disagreement, see Christensen (2009). For more on what it is to be an epistemic peer, see Gutting (1982) and Lackey (2010).

¹⁴ Rowland, Richard. "The epistemology of moral disagreement." *Philosophy Compass*. 2017;12: e12398. <https://doi.org/10.1111/phc3.12398>

then he would be justified in sending the ship to sea; even if it turned out, against the odds, that the ship had a problem, he could not be blamed if the ship did sink. Because in the case of peer disagreement it is not possible to tell with certainty who is in line with the truth, the standard that we should use is being epistemically justified. The Do Good App should give people moral directions that are based on claims that they can be epistemically justified in believing, and so may give directions that are based on principles of ethics that may ultimately turn out to be false--just so long as they are principles that, like Kantian and utilitarian principles, people are justified in believing.

Because there is epistemic disagreement over moral principles, the app needs to give you a choice of what epistemically justifiable fundamental moral principle you choose to follow. However, the app would not allow you to enter a moral principle that is not epistemically justifiable. The app works by giving you choices in cases where no matter how much time you had it would be difficult to get it right. However, there are other things that it is just difficult to get right because of our own difficulties (the lack of information, lack of perfect reasoning capacities, and lack of time). The app gives you choices for the first and no choices for the second.

If all that is required of the Do Good App is that it give you epistemically justified moral directions, then potentially the app could simply randomly pick a justifiable first moral principle for you to use and then use that principle consistently throughout your decisions. One could argue that so long as your first moral principle is justifiable, there is no important distinction between a moral principle that you have chosen and one that the app randomly chooses for you.

However, this argument is missing an important component of morality: moral understanding. Moral understanding is the ability to explain the 'why' and the reasons behind a

particular moral belief.¹⁵ Though users of the app may not in every instance have a full understanding of why they make a moral choice--an issue I will discuss in greater detail later--the app can preserve some of the value of moral understanding by going through a learning process with a user in which it asks them questions to determine what fundamental moral principle they believe in. Most people do not clearly know what their first moral principle is. They have moral instincts and opinions about what is right to do in a certain situation, but when asked what their first moral principle is, most people would not be able to articulate it. This learning portion of the app gives people the opportunity to distill their values down to their core and then have them clearly articulated. That process leads to greater moral understanding than the user had before, and also gives the user a baseline understanding of what motivates the future moral directions that the app gives the user

Allison Hills, in her article “Moral Testimony”, offers two reasons why moral understanding matters. First, moral understanding allows us to make more correct decisions in the future. Second, moral understanding ensures that the person is acting for ‘the right reasons’.¹⁶ These two reasons align with utilitarianism and Kantianism respectively.

A Utilitarian might argue that moral understanding is important because it allows you to make more correct decisions in the future. For a utilitarian, an action is good if it maximizes utility. By understanding why something is the correct decision, they can apply those reasons again in the future and continue to make the correct decisions. However, this advantage of moral understanding is not relevant to the app. The app is already able to continue to determine the correct decision for you. There is no utility gained by the ability to make future decisions because the app is already doing that. One could respond to this however by arguing that you

¹⁵ Hills, Alison. "Moral testimony." *Philosophy Compass* 8.6 (2013): 554

¹⁶ Hills, Alison. "Moral testimony." *Philosophy Compass* 8.6 (2013): 555

learning to make the decisions for yourself is faster and therefore more efficient than the app doing it for you. However, that is not the case. The app is almost always going to be faster at coming to justified decisions that a human would be for all the reasons we have already discussed. It has the knowledge, lack of biases, and faster processing capabilities. Therefore, this ‘future decisions’ utilitarian argument for moral understanding does not apply to the app.

However, there is another reason why a utilitarian would be in favor of moral understanding: moral understanding provides motivation. If the user of the app understands why they are doing what they are doing they will be more motivated to do it. By doing what they personally understand and believe is right, the user will be more likely to make the decision that the app tells them to do. Even if the user is a Kantian, a utilitarian could be in favor of them using the app as a Kantian because more utility will come from them being effectively motivated acting on principles they understand rather than from principles they do not understand.

A Kantian might argue that moral understanding is important because it ensures that the person is acting for ‘the right reasons’. For a Kantian, morally correct actions are only right if they are done for the right reasons. In order for actions to be done for the right reasons they have to be done with moral understanding. Ensuring that the user understands why they are doing what they are doing is necessary for the action to have moral worth for a Kantian. Thus, it would be important for a Kantian that the app preserve as much room as possible for individuals to have an understanding of the basic moral principles that motivate their action.

V. Conclusion

The truth is difficult to find because we are limited by knowledge, biases, and time. In order to be justified in our beliefs we need to have truth or at least sufficient evidence. Not being

justified in our beliefs can lead to harm, but taking the necessary time to be justified also can lead to harm. Therefore, we have a serious problem. The solution to that problem is the ‘Do Good’ app which would not have the same limitations as we do for finding the truth. The ‘Do Good’ app would be able to tell people what the correct thing is to do in a given situation depending on what their first moral principle is. Throughout the remainder of this thesis I will discuss the logistics of the app and the different objections to this solution to the problem. Specific objections include objections related to the importance of moral understanding and whether or not the app would really give solutions that are epistemically justified. I will also examine the way Kantians and Utilitarians would look at the ‘Do Good’ app.

Chapter Two: The Utilitarian Perspective of the ‘Do Good’ App

I. Introduction to Utilitarianism

At first, it seems very likely that the utilitarian would support the use of the app. For utilitarians, the right thing to do is the thing that maximizes utility. This moral principle leads to the Principle of Equal Consideration which states that equal amounts of utility need to be treated equally no matter who they are coming from. There can be no first person exceptionalism where someone favors their own pleasures over the pleasures of someone else.¹⁷ This principle goes against a lot of human instinct. It is very difficult for us to enact this kind of principle where we cannot favor ourselves over other people. The app would be great at fixing this problem. The app would easily be able to see utility as objective and not favor giving it to one person rather than

¹⁷ Singer, Peter. “Practical ethics”. *Cambridge university press*, (2011), 11-12.

another in ways that do not actually maximize utility. However, there are still objections a utilitarian would have to the use of the app. In this chapter I will focus on the utilitarian objections to the use of the app. I will first explain the objections and then argue that utilitarians could still support the use of the app by a utilitarian.

II. Objection One: Imperfect Experts

One objection a utilitarian may have to the use of moral experts in general is that the expert could be wrong and when the expert is wrong utility is not maximized. In the same way that the presence of personal bias and a lack of information can negatively affect average individuals in their decision making, experts could also miss the truth due to bias and lack of information. Though experts may make fewer mistakes than ‘average’ individuals making decisions, those mistakes may matter more to the extent that their mistakes are being replicated amongst all those to whom they are giving advice.

First, the same problem that any individual would encounter with having a bias also exists with experts. Turning our attention back to the example of Sally and the Watsi flyer, imagine if instead of going straight to work Sally had decided to consult with an expert. Even though the expert may not suffer from the same biases that Sally does, they will still be subject to other biases. Suppose the expert in Sally’s case is an expert in effective altruism who also happens to be an American. Though they may not be biased to Sally’s particular interests, they may nevertheless still (for instance) still have a proximity bias where they are placing a disproportionate value on effects on people that they can see and identify with. As a result, they might subconsciously value utility generated in America higher than utility generated in other countries. Therefore, they may tell Sally that the most effective use of her time and money would

be to go straight to work and then when she got home she should donate to the Veterans of Foreign Wars of the U.S. That is not actually the decision that would maximize utility.¹⁸ The expert has the ability to propagate this biased advice not only to Sally but also to anyone else that they are advising. Therefore this bias of the expert may be worse than simply allowing different biases to operate on an individual level.

Another potential barrier towards the truth is a lack of information. Experts are not going to know everything about an individual's particular scenario. In Sally's case, imagine that she decides to consult with the expert, but this time the expert does not know how much money she stands to gain from getting to work on time. The expert in this case may tell Sally to stay home, skip the meeting, and donate the money. That is again not actually the decision that would maximize utility. It would be better for Sally to go to work, make the extra money, and then donate even more than she would have been able to had she not made the additional two thousand five hundred dollars. Although the expert was not biased, they did not have enough information relevant to Sally's particular situation to make the correct decision. A particular problem in this area concerns knowledge about individuals' preferences and motivations. If the expert told Sally to do something that they thought would maximize utility, but Sally was so unmotivated by that thing that she did not do it, then utility would not be maximized.

Additionally, there may be some things that generate more pleasure for Sally personally, that deviate from standard norms that the app uses in its algorithms. For instance, imagine that an expert is helping Sally plan what she wants to eat in a given week. The expert tells Sally that eating dessert on Friday night will maximize utility. The app then tells her to buy an apple pie. However, Sally does not like apple pie; she prefers pumpkin. Because the expert does not know

¹⁸ Barrett, William. "In Pictures: America's 10 Least Efficient Large Charities" *Forbes* (2010) https://www.forbes.com/2010/11/17/charity-10-personal-finance-philanthropy-10-least-efficient-charities_slide.html?sh=37ab74c4556e

all of Sally's preferences, it cannot accurately maximize utility. In this case, Sally would be better at making the decision than the expert because she knows more about her own preferences. Whether the expert does not know about the individual's motivations or their preferences, utility may not be maximized because the expert is not privy to information relevant to the individual.

This argument about moral experts is similar to J.S. Mill's argument for respecting rights and liberties. In *On Liberty*, Mill argues that governments can make mistakes and therefore they should not make excessive laws because those laws could be wrong and have more pervasive and systematic impacts than the variety of 'smaller' mistakes that individuals might make.¹⁹ In the same way that governments could be wrong about the laws and thus could make citizens do the wrong thing, experts could be wrong about decisions and thus could make people listening to the expert do the wrong thing. These dangers would be pertinent to the app because whatever mistakes were made would continue to be made systematically. Whereas individuals may make mistakes of this kind almost randomly and in various ways, the app would make the same mistakes over and over for every user because the mistakes would be imbedded into the algorithm.

Because experts can be wrong, one could argue that it is important to be able to question them. Questioning creates an opportunity for the expert to correct those inaccurate beliefs so that actions are not being taken based on incorrect beliefs. Therefore, actions can be taken with knowledge of the truth and therefore utility can be maximized. Mill gives an argument in *On Liberty* in favor of freedom of expression which also illustrates why it is important to question experts. Mill argues that people should be able to say what they believe because people can be wrong and this right to free speech allows for false beliefs not to persist for very long. Thus, even

¹⁹ Mill, John Stuart, "Liberty". *Yale University Press*, (2003), 3
<https://www.earlymoderntexts.com/assets/pdfs/mill1859.pdf>

if there is extremely good reason to believe that something is true, people should still have the ability to contradict that opinion.²⁰

Mill's argument applies to the case of the user questioning the app. Even when there is a very good reason to believe that the app is generally right, one might argue that there would be utility gained by the user questioning the app because the user can correct the kind of biases that the app might be prone to having. Additionally, the act of questioning allows the app to learn from its past mistakes so that in the future it does not make them again. Since, as previously discussed, many of the mistakes that the app would make would be systematic, the correction would allow for the correct decision to be made both in the current situation as well as all future decisions that relied on the issue that was corrected. For instance, if the app mistakenly was operating as though a certain charity was effective when it was really not, that would be a mistake that would be repeated for many users. If one of those users noticed the mistake and then questioned the app, the app would be able to change that belief within the app and then the mistake would not be made for any user going forward.

Both of Mill's arguments show that it is important for the user to be able to question the expert. The two arguments illustrate examples of why, even if the expert has good reasons, people should question experts, and so suggest that users of the app should be able to 'question' the app.

²⁰ Mills says that "If the opinion is right, they are deprived of the opportunity of exchanging error for truth: if wrong, they lose what is almost as great a benefit, the clearer perception and livelier impression of truth, produced by its collision with error" (Mill, John Stuart, "On Liberty". *Yale University Press*, 2003, 19). Here Mills is arguing that there is utility gained by people contradicting even well established beliefs both because there is the chance that the well established belief is wrong, and also because there is utility gained on both sides by the 'collision with error'.

III. The Cost of Questioning

Unfortunately, due to the functionality of the app, it would be extremely difficult to converse with it and so the utility of being able to ‘question’ the app may be diminished. In Mill’s arguments, the individual is always questioning another *person*; as a result, there is always the possibility for discussion. In the case of the app, the user could ask specific questions and could try to determine if they are right or if the app is right through that discussion. However, one might reasonably worry about whether that kind of “conversation” with the app would be helpful or even possible. We all have had the experience of asking AI questions on some companies website and failing to get the responses that we need and so calling in order to talk to a person. This worry is not simply about whether the technology for “understanding” the content of what users say could be developed; it is also a worry about the nature of the “conversational” engagement with the user. Although there currently is some level of AI technology that is able to support this kind of communication, they are not particularly effective at teaching or responding to questions in a way that would allow the user to internalize the moral advice they were being given. One element of this problem is that the app is not able to have empathy. Empathy is important for being able to teach the user in two ways. First, empathy is important for being able to understand the emotions that motivate the user. By understanding the user’s motivations, the app would better be able to explain things to the user in such a way that they would understand and then also act on that understanding. Another is also necessary for understanding the user’s emotions when they have messed up. It is inevitable that there will be times when the user will not do the right thing even with the app. In these cases, the user may feel guilt, anger, embarrassment, or any other emotion in response to their failure to follow the app. When failure and the emotions accompanying that failure happen, the user could get stuck in those feelings

rather than being able to move on to continue doing what the app says. In those circumstances it would be helpful for the user to have the app relate to them on a personal level over past mistakes in order for the user to be able to move on. Discussing past mistakes with someone who has also made similar mistakes would allow the user to feel more comfortable in that past mistake and therefore able to move forward to doing more utility maximizing actions.

Finally, even if that was a technical possibility, it would still not be time efficient. If the user is taking the time to question the app, then it seems to defeat the purpose of the app. The app is designed to generate utility and minimize the amount of wasted time, energy, money, etc that it takes to determine what action is best. If the user spends time questioning the app, then that time that was being saved is wasted again. Moreover, to the extent that the user believes that there is a possibility that they are more right than the app, then they are going to not do what the app says. Most of the time the app is going to be more right than the user; there would thus be very few cases where the user would catch a mistake that the app would make. As a result, more utility would be wasted by having the user take time to look for those mistakes than would be gained from catching those mistakes. Rather than creating an opportunity for productive discussion and learning, the choice to 'question' the app would create a very high likelihood that the user is going to do something that does not actually maximize utility.

The above considerations suggest that the ability to question the app may not be a worthwhile feature for the app to have. The technology does not exist and the utility gained from developing it would not be worth the cost of developing it. Therefore, we need to be able to show that the app is especially good at not making mistakes, even more so than a human expert. If we can do this, we can create a basis for trust of the app so that there would be no reason to want to question it.

IV. Reliability of Algorithms

In order to determine how reliable the app would be the app can be analyzed by how well it minimizes the same issues that humans have with making correct decisions: bias and lack of information.

In the first case, since the app is a computer program it cannot have any wants or desires that it would value more than maximizing utility. Therefore those are not a basis for bias for the app. As mentioned earlier in this chapter, the principle of equal consideration of interests is important for utilitarianism. An individual's pleasures are not more valuable simply because they are their own. However, humans often operate as if that is the case. In the case of the app as a moral expert, there is no issue of the app having a self-serving bias because the only motivation that app has is to maximize total utility. Therefore, since the app does not have that kind of self serving bias, it would be less affected by bias than the user would and would thus make more correct decisions.

One could argue that the programmers could program into the app's algorithms their own wants and desires and therefore the app will still be affected by this kind of desire based bias. However, since the app would be programmed by people who are focused on the app doing good, I am not going to focus on the possibility that the programmers' biases could affect the algorithm. One way to accomplish this goal of having unbiased programmers would be to create a system of checks and balances within the programming system.

However, there is still another kind of bias that the app could have that would affect its ability to know the truth. The app could rely upon unrepresentative samples to come to conclusion, which would result in a kind of bias. If the app relied on a set of data that was not representative of what it claimed to be representative of, then it would not be maximizing utility.

For example, imagine there is a government worker named Alex. Alex is in charge of organizing road construction in their city and so it is their job to decide what roads are going to get repaved. There is a set budget this year and so only some of the roads will get repaved and Alex has to decide which. Alex is a user of the ‘Do Good’ app and so they ask the app what to do. The app tells them that the ‘Forest Hills’ neighborhood is where they should divert the limited money and resources to fix the roads. So Alex decides to fix the roads in the ‘Forest Hills’ neighborhood. However, neither Alex nor the app realized that the data used to determine which area would generate the most utility by being fixed was not representative of what it was claiming to represent. The app was relying on data that recorded the amount of ‘bumps’ detected by a smartphone while in a moving car. The area with the most ‘bumps’ would presumably be the area that would benefit the most, if that was the area with the most traffic and the most potholes. However, this data is coming only from smart phones and therefore the data is overrepresented by wealthier people who could afford those phones. Therefore, while the data appeared to be representing which area had the most car traffic and the most potholes, it was really representing which area had the most number of smartphones. The app’s decision to tell Alex that fixing the roads in the ‘Forest Hills’ neighborhood would generate the most utility was wrong. It was wrong because it relied on a data set to represent something that it did not.

This problem of unrepresentative samples is a very real obstacle for the app’s ability to make decisions that maximize utility.²¹ However, there are ways to mitigate the effect that these unrepresentative samples could have on the app’s ability to maximize utility. The app could use human experts to check the use of data within the app. These experts would be experts in using

²¹ Other examples of these kinds of unrepresentative samples are crime reports and the use of the SAT. Crime reports are often used to show the amount of crime that is happening in the area. However, they more accurately predict the amount of policing in an area rather than the amount of crime that is happening (Lum, 2016). Additionally, the SAT, ACT, or any other standardized testing are often understood as a measure of intelligence when in fact they are more accurately a measure of the amount of time and resources one has available to put into studying for the test.

data samples and would be able to check to ensure that the app was only using data to represent what that data actually represented. One way that these experts could ensure that the data is truly representative is through ensuring that the experts themselves represent a wide range of experiences and expertises. Having experts with a variety of different perspectives, interpretations, heuristics, and predictive models allows for cognitive diversity. This kind of cognitive diversity has been shown to be as important, if not more important, than individual intelligence.²² This variety allows for there to be different kinds of people that would be able to catch different kinds of inaccuracies in the understanding of the data. For instance, in the case of Alex and the road construction decision, the problem was that the app was misrepresenting data. The data really only represented where there was the most traffic of people with smartphones rather than where there was the most damage to the road. The app misrepresented the data because it was not aware that not everyone had a smartphone. By having a cognitively diverse group of experts analyze the data, someone in that group would have the experience of not having an iPhone and therefore would be more likely to notice that kind of misrepresentation.

There is a cost to this method. It would involve paying or someone otherwise incentivizing experts to use their time to work on the app. It would also take time and potentially money to organize those experts in that work. However, those costs would be outweighed by the benefit of the app having more accurate information. That information would then be used for many users of the app and therefore would result in much more utility when applied correctly.

Additionally, even if these mitigating methods were not applied, the app would still be no worse at handling unrepresented data bias than the user would on their own. Humans often also take evidence to mean something that it does not. If shown the data of the streets with potholes,

²² Hong, Lu, and Scott E. Page. "Groups of diverse problem solvers can outperform groups of high-ability problem solvers." *Proceedings of the National Academy of Sciences* 101.46 (2004): 16385-16389.

many users would come to the same conclusion that the app did. However, the app is significantly better than the user at handling issues of desire based bias. Therefore, the app is better at overcoming the bias obstacle than humans are.

Lack of information is the second obstacle that humans have to knowing the truth. In the case of the app, it has far more information than the user. The app has access to all of the information on the internet and can process it all much faster than a human could. However, the app may not have more relevant information than the user. The app may not know specifics about the user's situation or about their motives. The user knows what is going on around them, whereas the app does not. What matters for making the correct decision is not just having more information but having all the information that is necessary to understand the situation. It may seem like the user has more of that kind of information than the app. However, the app could be altered to fix that problem. The app could be developed to have a camera, earpiece, etc that would allow it to gather more information both about the user and about the current situation and decision they have to make.²³ In this way there would be very little information that the user had about a situation that the app did not and far more information that the app knows that the user does not. Additionally, if there was any information that the app was missing that the user might have, the app could ask for that information. In this way there would be no information that the user had that the app could not also have. Therefore, the app has more and better information than the user does on their own and thus the app is better able to make correct decisions.

²³ There are, of course, potential privacy concerns with an app having access to this much data from the user. However, the only privacy concerns that would matter to the utilitarian are ones that would show how this lack of privacy would generate some kind of negative utility. Assuming that the app was able secure the data in such a way that the 'Do Good' app was the only corporation to have the data, there would be no concern of other companies using this data in ways that may create negative utility. Additionally, the app would only use the data in ways that would maximize utility, so it would be able to factor in ways that this lack of privacy may affect the user's ability to generate positive utility.

V. Objection 2: The Robot Objection

One potential concern with the ‘Do Good’ app is that if it is really going to be working to maximize utility, the user would just be turned into a robot. If the user is not required or even allowed to do any of the prompting or thinking behind any of their decisions, then they would essentially be acting as a robot. The app would have all of the information it needed from the internet, listening devices, and cameras, so it would require no inputs from the user. The app is also going to be better at using that information because it has better processing power than a human and no desire based biases. Therefore, the user’s only job is to actually do the decision that the app tells them to do. They do not have to think or discern at all. There is something intuitively wrong with the user being turned into a robot. This intuition could be motivated by virtue ethics or Kantian concerns about the value of autonomy or necessity of that autonomy in what it means to be human. I will address those concerns in later chapters, but first I will address a utilitarian response to this concern.

VI. Mill’s Higher Pleasures

At first one might think that the objection would be more of a problem for virtue theorists or Kantians because, - as I will discuss in later chapters- they emphasize the importance of an individual’s motivations in evaluating if an act has moral worth. However, this intuition turns out to be in line with one kind of utilitarianism. Mill would argue that the kind of life that you would live by using the app is a life that would not be generating “higher pleasures”. Higher pleasures are things such as intellectual debates, academic pursuits or activities that engage the higher faculties. In his book *Utilitarianism*, Mill famously says “It is better to be a human being dissatisfied than a pig satisfied; better to be Socrates dissatisfied than a fool satisfied.”²⁴ This is

²⁴ Mill, “On Liberty”, 13

because the human is able to achieve those higher pleasures whereas the pig is not and simply being content does not generate as much utility as being happy does.

Mill argues that we know there is this distinction between lower and higher pleasures because when faced with a choice between the two a “competent judge” would choose the higher pleasure. Mill explains this determination of preferences by saying,

“If I am asked what I mean by difference of quality in pleasures, or what makes one pleasure more valuable than another, merely as a pleasure, except its being greater in amount, there is but one possible answer. If one of the two is, by those who are competently acquainted with both, placed so far above the other that they prefer it, even though knowing it to be attended with a greater amount of discontent, and would not resign it for any quantity of the other pleasure which their nature is capable of, we are justified in ascribing to the preferred enjoyment a superiority in quality so far outweighing quantity as to render it, in comparison, of small account”. (II 5)

Here Mill is saying that because competent individuals choose higher pleasures over lower pleasures, higher pleasures have more value and therefore greater utility.

Mill would argue that by using the app in such a way that you become a ‘robot’ you essentially miss out on all those higher pleasures. By not reasoning through decisions yourself and engaging intellectually with the world, you are essentially turning yourself into the pig and no matter how satisfied that may make you or how much utility that generates for other people, it does not compare to the higher pleasure that comes from being a dissatisfied human. A version of the app that does all of the reasoning and prompting for you (thereby turning you into a robot)

would not maximize utility because it would miss out on all of the utility generated from higher pleasures. Since the app seems to not allow the user to do any of the thinking for themselves the user misses out on higher pleasures and therefore it is better to not use the app.

VII. Emergency Situation

If this response works, then it is only because what would be gained from the higher pleasures is greater than what would be gained by the app doing all of the decision making. Even if there might be some utility that could be generated from individuals exercising higher faculties and as a result enjoying higher pleasures, that utility would have to be greater than the utility that would be gained if the users followed blindly. This is not the case. As of 2018, about 1.9 billion people were malnourished.²⁵ As of 2017, about 2.1 billion people did not have safe drinking water on their premises and 3.6 billion people did not have access to basic health care.²⁶ As of 2018, four billion people have no high school education.²⁷ Not having access to adequate nutrition, water, health care, and education generates negative utility. Billions of people are currently suffering and they are suffering from things that are preventable through donating.²⁸

The amount of utility that would be gained from people no longer making uninformed and biased decisions would outweigh the degree of higher pleasures that they might feel they have gained by living a more fulfilling life. Our world is currently in an emergency situation where very little money and effort can have a very big impact on other people. The app leading

²⁵ Food and Agriculture Organization of the United Nations, "Preventing micronutrient malnutrition a guide to food-based approaches- Why policy makers should give priority to food-based strategies" *International Life Sciences Institute* (1997).

²⁶ "Progress on Drinking Water, Sanitation and Hygiene" *World Health Organization*, (2017)

²⁷ Global Education Monitoring Report: Target 4.1- Primary and secondary education (2015)

²⁸ In addition to the billions of people who are currently suffering and generating negative utility, most brands of utilitarianism also argue that animals can generate negative utility. About three billion animals are killed daily for food for humans (Zampa, 2018). Those deaths are counted in the utilitarian calculus for most brands of utilitarianism. Therefore, for those brands of utilitarianism, there is even more preventable negative utility occurring in the world and even more of a reason to resort to sacrificing higher pleasures.

people to that conclusion, even blindly, would create more utility than allowing those people to continue to suffer or to die.

One way to better understand the effects that donating can have is through the law of diminishing returns. This law states that the more you have of something, the less additional utility you get from one more of that thing. For instance, eating one cookie may generate three units of utility, the second maybe two, the third one, and by your fourth cookie you're not generating any utility. But, if you gave your cookies away, then more people would be eating their first cookie and more utility would be generated. By eating all four cookies you generate six units of utility, but if you were to give three away, then twelve units of utility would have been generated. The same is true with our money. Most Americans are spending past the point of generating more utility than they would if they gave their money away.

One possible complication to this idea that saving lives generates positive utility is that not all lives go to experience high quality pleasure or even to have lives with positive utility. Therefore, the app should not focus on diverting resources in this way and it is instead better for the individual to make their own choices. However, even if it is the case that not all lives generate positive utility, donating in order to minimize the number of people who are malnourished, suffering from lack of water, and/or dying from preventable diseases is still going to maximize utility. People suffering generates negative utility and donating to prevent that creates positive utility. Therefore, more utility would still be generated by the app having more control in order to divert those resources. This is especially true if many people are using the app. If a lot of people were donating, then the amount of money donated could very possibly get people past the point of living with suffering. The app could get people to the point where more people are able to more fully use their higher capacities and therefore generate higher pleasures.

Therefore, use of the app in this way could not only save lives but also make those lives much better. The app would be able to take higher pleasures into account so that money could be allocated to those who are more likely to be able to enjoy high quality pleasures.

IX. The Problem of Motivation

The two options then seem to be either the more autonomous option, which has a higher rate of error because the user is making more of the decisions or the more robotic option, which allows for the user to experience higher pleasures. Given the emergency situation that we are in the utilitarian choice seems to be the robotic option. However, a utilitarian may also argue that this robotic option would likely have a high rate of the user abandoning the app or listening arbitrarily. This abandonment could be likely if individuals do not like being told what to do. In some cases, being told what to do elevates stress, but often, people do not like being told what to do, especially if it conflicts with what they want to do. Therefore, if the app tells them something they do not want to hear, they may not list to the app or they may delete it.

If that is how the robotic version would work in reality, that version may not maximize utility. Therefore, it is important to consider how the app is going to motivate people to do what it says. One option to motivate people is to have them decide for themselves what to do, but that option would result in a significant amount of error.

There are possibly other ways that the app could encourage people to download it while still keeping the robotic version. One route would be that the app could advertise the ways in which this kind of use of the app would minimize personal stress. By not having to make any decisions, the user would be free of the burden of having to think and stress and worry because the app would handle all of that. The app could also show you photos of children dying to further

motivate you to opt for the robot setting. By showing you the benefits for other people of you becoming a robot, the app may be more able to convince more people to opt for that setting.

However, ultimately, different people would have different tolerances for how much of the robotic version they would be able to tolerate. Therefore, the app would be tailored to each person specifically. The app is going to try to maximize the amount of control it can have while also keeping you using the app. To determine what to do in a particular case, it also has to consider if and when it might be possible for some level of higher pleasures to be achievable without sacrificing accuracy. In the next section I will develop in more detail the ways that the app would balance accuracy, higher pleasures, and keeping people engaged with the app.

X. Balancing Accuracy, Engagement, Higher Pleasures, and Cost of Mistakes

The problem now is balancing the maximum amount of control that the user will tolerate the app having with maximizing utility. The more control the app has the higher chance the decision will be correct, but the less control the app has the higher the chance the user will continue to use the app and the higher the chance the user will achieve higher pleasures. We can imagine that there are three different kinds of control the app could have: the app makes all of these decisions without prompting, the app allows the user to prompt it with a question about what to do, or the app allows the user to make a decision fully on their own and then the user could check with the app. In order to determine which kind of control the app will have, the app will need to weigh the chances the user would be wrong with the cost of utility if they are wrong. One way to determine the chances of the user being wrong is through 'levels' within the app. As the user learns through experience with the app how utilitarian decisions are made, the app

would begin to allow the user to have more control in their decision making. Here is a graph outlining one way that this would work.

	The decision is easy to make correctly ²⁹	The decision is of medium difficulty to make correctly	The decision is difficult to make correctly
Little utility gained if the correct decision is made	Beginner- do it yourself, then check with the app Intermediate - do it yourself Expert-do it yourself	Beginner- ask the app Intermediate - do it yourself, then check with the app Expert-do it yourself	Beginner- ask the app Intermediate - do it yourself, then check with the app Expert-do it yourself
Some utility gained if the correct decision is made	Beginner- ask the app Intermediate- do it yourself, then check with the app Expert- do it yourself	Beginner- ask the app Intermediate- ask the app Expert- do it yourself, then check with the app	Beginner- ask the app Intermediate- ask the app Expert-do it yourself, then check with the app
Lots of utility gained if the correct decision is made	Beginner-ask the app Intermediate-ask the app Expert- do it yourself, then check with the app	Beginner- ask the app Intermediate- ask the app Expert- ask the app	Beginner- ask the app Intermediate- ask the app Expert- ask the app

Here the graph shows that the user who has reached expert level is given more ability to not rely on the app because there is a lower chance the expert would be wrong. Giving the user who has reached expert level more freedom allows that user to experience high pleasures while also not sacrificing utility by being wrong.

²⁹ where difficulty is a combination of how much information is needed/how many factors there are, how much time there is to make the decision, how much bias is inherently going to interfere with the decision making process

XI. Objection 3: The Overuse Objection

This sliding scale of different amounts of user control creates different kinds of objections. In any case where the app does not have full robotic control, one objection a utilitarian may have is that people would be encouraged to overuse the app, which **could** result in a significant loss of utility . Anyone who has reached a level where they have the ability to prompt the app needs to learn what kinds of things are worth prompting the app about. Prompting the app takes time and that time would be wasted if it was spent asking the app about things that generated less utility than was cost by asking. Prompting the app takes time because the user would have to physically input a prompt or question into the app, whereas if the app had all of the control, it would take no time to get the decision. Although it may initially appear to be a small amount of time that would be spent on the user asking the app, if the user had to do that for every decision that they made, that time would add up.

For instance imagine that I have to decide where I am going to buy my groceries. I can realistically choose between ordering them all from Amazon, buying from the health food store that is 30 minutes away, or going to the Walmart down the street. The first time I ever make this decision with the app I am completely unsure of what it is going to tell me. If the app is operating at its maximum control amount then it will cost about one ‘utility point’ to determine what to do.³⁰ The app does it automatically. If the app is at a control amount where I can prompt the app, then it will cost me five ‘utility points’ to take the time to prompt the app with my question about where to buy groceries. We can also imagine that the app could be at its minimum control amount where it would cost me fifteen ‘utility points’ to do the necessary research on my own to determine the correct answer to where to buy groceries. Then imagine that I have the

³⁰ The app takes some energy to run so we can imagine that some kind of utility is still lost but that it is marginal compared to what is gained.

right answer (that I have gained through whichever method) and I know that that choice will gain ten utility points. If the app was a maximum control then I gained those full ten points. If it was at the prompting amount of control then on the net I have gained five points. If it was at the minimum amount of control then five utility points were lost. In this case it's clear that the best option was to have the maximum control of the app.

However, we can assume that it was necessary for the app to be at the prompting level in order for the user to have the motivation to continue using the app. In that case five utility points were still gained and I, the user, continue to use the app. The next week I have to go to the store again. My week looks about the same as it did last week. Nothing significant has changed in my life. However, I still ask the app where I should go grocery shopping. I still sacrifice the five utility points to input that information and it gives me the same output and I still gain five points. However, this is what essentially every week of my life looks like. Every week I have about the same inputs that I am inputting into the app and every week I am sacrificing those five utility points. In this case, using the app is not maximizing utility. What would be better would be for me to stop prompting and get the full ten utility points from just going to the best grocery store. What would maximize utility is for me to form a habit of which grocery store I go to. That way no utility would ever be lost by me having to make a decision or prompt the app. The utilitarian could argue that this is what many decisions look like. Many decisions are repeat decisions or habits that we form throughout our lives. Therefore, utility is maximized when the app is used to form those habits, but then phased out.

Additionally, the app could also be subjected to overuse by being asked inconsequential questions. Not all good decisions will result in a positive net outcome if they have to be run through the app. For instance, imagine I have just finished washing my hands. I have to decide if

I want to use a cloth towel or let my hands air dry. In order to make that decision I input the necessary information into the app to ask it what to do and lose two utility points because of the time spent doing that. However, any decision that the app might give me is not going to generate any utility. It makes no meaningful difference if I let them air dry or use the towel. Any insignificant decision like this is going to result in negative utility if it is run through the app. By its nature, the app promotes this kind of overuse and users lose the ability to make even small decisions on their own.

Creating levels within the app solves both of these potential overuse problems. In this version of the app, the app would operate as a game with levels and points earned by reaching certain marks. Someone at the expert level would have the experience necessary to know when prompting the app maximizes utility and when it is a waste of time. The expert also knows how to make many correct decisions without prompting the app. Therefore, problems of overuse would be mitigated because the only people who would have the ability to prompt the app would have the experience to know when to. These experiences would come from the app eventually not giving you the answers. Once you reach certain levels, the app would stop giving you the answer right away. Based on your past experience the app would generate an answer such as “the probability of you getting this answer right is very high and the stakes are very low so figure it out on your own”. In determining when to give this answer the app would have to weigh the likelihood of the user getting the answer right based on their past experience versus the cost of the user getting the answer wrong. As the user moved further up in the ‘levels’ and gained more ‘utility points’ they would begin to get this answer more often. The app may also respond “this decision is not important” in cases such as how to dry one’s hands. Even if one option may be

slightly better than the other, the benefit is not enough for the user to get in the habit of identifying that kind of situation as one worth spending time to ask the app about.

The purpose of this adaptation would be to minimize unnecessary time spent on the app. The individual would learn and internalize many of the correct decisions to make so time is not wasted on the app. Additionally, throughout your use of the app it would generate habits for you to follow so that you do not even have to think about if you should ask the app what grocery store you should go to because it is always the same one. It would handle money in a way that looked at the user's life cohesively rather than decision by decision. For instance, it would set up things such as monthly donations at the maximum amount you could spare rather than deciding every day how much money you should donate.

There is also an additional benefit of this adaptation of the app: it would encourage people to use it who otherwise might not be interested. The idea of having levels and points and tasks to complete would appeal to a wide audience. It would maximize utility by giving people more personal incentive to do so. The app could also have features such as comparing points and levels with your friends to encourage a healthy amount of competition surrounding maximizing utility. It could also create tasks to do that may be analogous to 'quests' in some video games. For instance, 'help five elders cross the road: gain 10 extra utility points!' or 'say a kind word to a stranger: gain 5 extra utility points!' Again, these kinds of incentives would encourage people to use the app and thereby maximize utility.

Chapter Three: The Kantian Perspective of the ‘Do Good’

App

In this section I will first explain a Kantian objection to the utilitarian version of the app. Then I will argue how that objection against a utilitarian version of the app extends to be an objection of reliance on experts in general. Finally, I will show how the app could be modified to successfully respond to objections raised by Kantians.

I. Problems of Acting with Right Motivations

A Kantian might have many objections to the utilitarian version of the app. One objection would be that the utilitarian version of the app pushes the user to get pleasure out of advancing through levels. These levels were a necessary aspect of the app because they allow the app to distinguish between different amounts of control it allows the user to have, but the levels also help motivate the user to continue using the app. People tend to be motivated by competition and achieving short term rewards. The app provides both of these motivations by assigning points to different acts and moving the user through levels as they generate those points. This system could lead users to do good actions because the ‘points’ they get from doing so gives them pleasure, rather than because of a sense of moral duty.

For Kantians, the reason why someone performs a particular action is more important than the consequences of that action. Specifically, the actor has to have acted from a sense of moral duty: they have to have acted primarily because they knew that the thing they were doing was the right thing. Someone who acts from this kind of sense of duty displays what Kant calls the ‘good will’. According to Kant, the good will is incomparably higher in value than anything

else and also what makes anything else have moral worth. Any other virtue or reason for doing something only has value to the extent that it connects to the good will. Thus, happiness has no moral value in and of itself; any such value is contingent on its issuing in the right way out of an action that is done out of a good will. One common example that highlights this distinction of morally worthy actions versus non-morally worthy actions for Kantians is the case of a shopkeeper who is giving his clients their change.³¹ In one case, the shopkeeper may give change because they genuinely know that it is the right thing to do. Not giving the client change would be stealing and the shopkeeper knows that stealing is wrong so they give them the change. That act of giving change would have moral worth. In another case, the shopkeeper gives change because he wants the client to come back and buy from his business again. He knows that if he does not give the customer the change they will not come back and so he will lose money. Therefore, the shopkeeper chooses to give the client their change. This act of giving the change does not have moral worth because it was not done for the right reasons. The second shopkeeper's act of giving the change was therefore not in line with the good will. Therefore, a Kantian might object to the app because it could be inclined to encourage the user to act like the second shopkeeper, which would undermine the moral worth of the user's actions.

In addition to the objection that the app could incline users to act for the wrong reasons, a Kantian might also object that the app enables people to act without any reasons at all. One way for an act to not have moral worth is because the actor acts for bad reasons as in the case of the second shopkeeper. Another way for an act to not have moral worth is if the actor acts for no reasons. For Kantianism, it is vital that the actor be acting autonomously, and if someone does not have reasons for acting the way that they do, then that would mean that they are not acting

³¹ Kant, Immanuel. "Grounding for the Metaphysics of Morals". 3rd ed. Indianapolis: *Hackett Publishing Company*, (1993), 10.

autonomously; they are instead acting only on inclination or out of rote habit. For instance, imagine that instead of the shopkeeper giving change, there is a vending machine that gives change. The same act of the customer being given change is occurring, but this time the actor does not have any reasons behind it. Therefore, there is no moral worth to the vending machine giving the customer the right change. A Kantian could object that the app turns the user into the moral equivalent of that vending machine. The app enables the user to act without having their own reasons and therefore those acts lose moral worth in the same way the vending machine's acts do not have moral worth.

The Kantian objections to the Do Good App discussed above also apply more generally to expert advice: in both cases, the user is enabled to act without determining for themselves the relevant reasons for why they should act. Because the user has not reasoned to the conclusion for themselves, there is no guarantee that the user is acting out of an appreciation of moral reasons—or, really, for any reasons at all. Insofar as the expert only told someone what to do, the action would have no moral worth because the actor would not be acting for the right reasons.

II. A Modified Kantian Version of the App

The problems with human moral experts may be amplified by the use of the app. The first concern is that the user would be acting in order to get pleasure from winning points or moving through levels. Those motivations would be the wrong reasons for acting, but they could very easily become the user's primary motivation through repeated use of the app. The solution to this problem may be to just modify the app into a 'Kantian Version' that does not have levels or points so that those are not a temptation for the user.

However, even if that modification prevented the user from acting with the wrong reasons, there is still the concern that the user would be acting without reason. Although the user decides what moral principle they align with, the app does the rest of the reasoning and thinking. Since the user has not reasoned through what the right thing to do in the situation is, they are not doing the action for the right reasons and the action does not have moral worth. Therefore, the user is not performing a central aspect of what is required for the act to be a good action and the act does not have moral worth. An app like this would thus seem to help bring about a world that is not morally good, at least by the lights of a Kantian, because it helps bring about a world in which people are only robots who do not understand what they are doing or why they are doing it.

One might respond to this objection by arguing that the app still has the user do the most important kind of moral reasoning and therefore the app is permissible. The app requires you to formulate your first moral principle and then pushes you to clarify that during the initial stages of downloading the app. In this respect, one might argue that the individual using the app is still doing the critical ethical work; the app merely does the empirical work. For example, imagine the first case of the shopkeeper again. The shopkeeper intends to give the customer the correct change because they know it is the right thing to do. However, in order to determine what the correct change is, they need to use a calculator. Even though the shopkeeper used a tool in the process of his act, his act was still done for the right reasons and therefore still has moral worth. The same is true with the app. The app is a tool in the user's process of doing the right thing. Because they input their moral principle, they are acting for the right reasons. The app is only being used to assist the user in that process. Therefore, the individual's moral agency is still engaged in the way that matters.

A Kantian might not accept this answer for two reasons. First, even if the user is properly motivated by wanting to do the right thing (rather than by pleasure), their moral agency is still not being actively engaged in the reasoning that applies the ethical principles to the situations. This may still seem to at least diminish the moral worth of the action. Since the app gets rid of the need to be constantly questioning, developing, and applying your principles, a degree of moral agency is still being lost and the user is not actively choosing to do the act for the reasons that justify the action. Indeed, the user could also potentially have no understanding of the connection between their first moral principle and the act that the app is telling them to do. ‘Ethical work’ is still happening for every decision, but it is done by the app not by the individual.

Additionally, a Kantian may also object that when the user is not forming the reasons for themselves, it could be extremely tempting for the user to switch their mindset to acting just because the app told them to, and thus users’ actions would not have moral worth. For example, the app may determine that the correct thing for a user to do is to buy fair trade. That act, if done for the right reasons, would have moral worth. Someone who just started using the app and is in touch with their first moral principle would likely be able to do that act for the right reasons. They would understand that buying fair trade treats the producers of that item as ends rather than as merely means to the production of the product. However, because the user automatically relies on the app, rather than having to recall and reconsider her specific views when she acts, over time the user may forget the correct reasons for the action. This non-consideration of the moral reasons might also lend itself to the user beginning to see themselves as superior to other people because they are “doing the right thing” by buying fair trade and then only continue to buy fair trade in order to continue to feel morally superior. They could also potentially come to like fair

trade products more than products that are not fair trade and therefore continue to buy fair trade because they like those products more. This possibility of losing touch with the moral reasons that justify one's actions is possible in any case even without the app, but it may become easier in the case of the app because the app is not forcing the user to check in with the reasons motivating their actions regularly.

The problem with the above objections to expert advice, whether in the form of the app or not, is that the objections do not take into account that besides just telling you what to do, an expert can help you form and connect with the correct moral reasons behind the act. In response to this the app could be modified to require regular check ins with the first moral principles. In situations where the user has time to engage with the ethical work behind the decision, the app could ask questions and walk through the situation with the user rather than immediately giving them the answer. This adaptation would allow the user to still be able to engage with their ethical principles when they have time, but it would also have the accuracy benefits that come from using an expert. The goal of this kind of adaptation is for the user to gain moral understanding.

In cases where the user did not understand why the app was telling them to do something, the app could aid the user in gaining moral understanding by directing them in two potential directions depending on what would work best for the user. If the user signaled to the app that they do not understand the decision, then the app would direct them to one of two options: either the user could talk through the decision with someone else or they could take time to think through it themselves. Which option that app would give you would depend on what was most effective for you based on past experience with the app.

One option would be for the user to take time to think through the decision themselves. This time could look like reading philosophy or just taking time to personally reflect and think

logically through the decision that is to be made. When there are papers that are written on relevant issues to the decision that the user has to make, the app would tell the user to read those papers. The app could have an inventory of papers written at a level that the average user would be able to understand and that were pertinent to typical issues that users would face. After being given a paper, the user could take time to apply the reasoning and information in those papers to their own decision. Through reading how other philosophers have reasoned through moral decision making, the user could then apply those reasoning to their decision. This solution creates the possibility for mistakes, but it would still increase the user's ability to understand what and why they are doing something.

However, after doing this personal research the user may still have questions. The app would then direct the user to engage in a discussion with someone about the decision. Potentially that could be the app itself or with a human philosopher. In the former case, the user would be able to ask questions in the app itself about why the app made a particular decision. The app could then give the user the information it gathered to make the decision and a basic outline of the logic it used. The app discussion seems to be the most straightforward solution to the problem of the user not understanding, but as mentioned in chapter two there are problems with this approach. As in the case of the utilitarian app, there are practical barriers to the user being able to question the app. Although AI does exist that can answer questions, it is not clear that the technology could be developed that could actually teach the user effectively.

The app's inability to sufficiently respond to the questions and exhibit empathy makes the app not the best option for answering the user's questions. Another option that the app could do when the user does not understand what it is telling them to do is to direct the user to talk to an ethics philosopher. An ethics philosopher would be able to better and more clearly communicate

to the user than the app would given the current technology. The problem with the philosopher is that they may not be as accurate as the app would be. Although the philosopher could mitigate the problems of bias and lack of information, they are still a human so those problems would still exist for them.³²

All of these solutions have the potential problem of the user being wrong and they all also take a significant amount of time. However, for a Kantian, both of these potential drawbacks are worth the benefit of the user understanding why they are doing what they are doing. For a Kantian, the goal of the app is still to have the user make the correct decision but the right act is not really the right act unless the user makes that decision from a position of moral understanding. Therefore, mistakes and time invested into thinking are necessary conditions to making what are ultimately correct, *morally worthy* decisions.

The amount of time that this solution takes limits the amount of actions that the user would be able to take. However, failing to take actions that would have moral worth is not necessarily wrong for a Kantian. So long as the actions that are taken are done for the right reasons, the user can sit and think through their decision for as long as it would take.³³ This

³² Another problem with this option is how the app would recruit the philosophers. It is possible that philosophers may choose to use the app and then the app could tell them to communicate with other users. However, there would likely not be enough qualified philosophers who would use the app to make that work for every user who had questions. Another possibility is that users could pay for a subscription to use philosophers. The payment would serve as another incentive for more philosophers to download the app. However, as I discuss further in the conclusion, that practice would discriminate against those who could not afford it. There is no ethical reason why people who can afford to pay for a subscription are more worthy of having moral understanding than people who cannot afford it. Therefore, potentially users could donate to the app and then that money would be distributed to equally benefit any user who needs to use a philosopher to gain moral understanding. In the conclusion I will discuss other potential methods for funding these philosophers.

³³ This distinction illustrates the difference between doing versus allowing (Driver, 2007). Although this solution may mean that the user is allowing some bad things to happen, the things that the user is doing all have moral worth. For example, imagine that the user has three thousand dollars of excess money. If donated, that money could save someone's life. However, if the money was only donated because the user wanted the pleasure that came from doing something 'noble' or because it would make them feel proud, then that action would not have moral worth. Therefore, the Kantian would be justified in taking as much time as they needed determining what to do with that money in order to do it for the right reasons. They would be justified in doing that even if it meant that someone else suffered in the process because

solution may seem to undermine some of the initial motivations for the app discussed at the beginning of this thesis because the user still may potentially have to spend a significant portion of time coming to a decision. However, this solution would still save the user some time relative to the amount of time it would take the user to come to the decision on their own. Additionally, the app still improves the accuracy of the user's decision. Therefore, there are still noteworthy advantages to using this Kantian version of the app.

IV. The Problem of Peer disagreement within Kantianism

In every moral theory there are going to be issues of peer disagreement over nuances of what the theory means. For instance, there are different kinds of utilitarians as well who differ on things such as the basis of what 'utility' is and exactly what the calculus should be focused on maximizing. In Kantianism one common issue of peer disagreement is over how far Kantian morality requires the actor to go. Kant argues that there is a duty to not treat individuals merely as means, but rather as ends in themselves. Some philosophers such as Robert Nozick argue that duty only requires one to avoid treating people merely as means.³⁴ However, other philosophers such as Barbra Herman take this duty to mean that people should support others fulfillment of their 'true needs' that they have in order to exercise rational autonomy.³⁵ This distinction may seem small but would have very different outcomes for what users of the app would be told to do. For instance, Nozick and Herman would argue that individuals are required to donate very different amounts of money. Not only does this peer disagreement cause differences in what the

of this distinction between doing and allowing. However, some Kantians have objected to this perspective and argued that there is a limit to the extent to which 'allowing' is morally permissible. Barbra Herman argues that individuals are obligated to prevent others from being treated merely as means rather than just being obligated to not treat others merely as means. This distinction means that a Kantian would not really be justified in choosing to sit and think through their decision excessively if that meant that someone else was being treated as a means.

³⁴ Nozick, Robert. *Anarchy, State, and Utopia*. Vol. 5038. New York: Basic Books, (1974), 30-31.

³⁵ Herman, Barbara. "Mutual aid and respect for persons." *Ethics* 94.4 (1984): 577-602.

user is told to do, but the correct reasons behind the act are going to be different depending on the theory. Therefore, if the user personally ascribed to one kind of Kantianism, but was listening to the app tell it what to do on the basis of another kind of Kantianism, then the user would not be doing the act for the right reasons and the act would lose moral worth.

There is some disagreement over whether or not the difference between Nozick and Herman is truly peer disagreement. The disagreement seems to be a peer disagreement for many reasons such as that both Nozick and Herman are qualified, educated, and informed experts in the same field. If this is peer disagreement, then the app would lead you through the process of determining for yourself which principle you think is right. In the same way the app learned whether you were Kantian, utilitarian, or a Virtue Ethicists, the app would also have to distinguish which kind of Kantian you were. The app gives you the choice in these situations because the choice gives you the opportunity to exercise more agency and ensures that you are acting for the right reasons when the app tells you what to do. Additionally, if the app did not give you the choice, it would have to choose for itself and that decision would be arbitrary if there is truly peer disagreement.

However, this disagreement may not be peer disagreement in which case the app would have to determine which of these theories was wrong. For instance, perhaps it could be shown that Nozick's perspective is not on the same level of Herman's because he is affected by a kind of first person exceptionalism which biases one to believe something that is going to be more beneficial for them to believe. If this is the case and there is no peer disagreement, then the app would be able to determine which was right. The app would then make decisions for the user based on the principle that was correct in the same way that it did in other cases.

V. Kantian Justification for the Utilitarian Version

One general problem for the app regardless of what version of the app is being used is how proponents of one version of the app would justify users using a different version of the app. For instance, if someone is utilitarian, how would they justify Kantians being allowed to use the Kantian version. Although this problem exists for each ethical theory's evaluation of the app, this problem is particularly difficult for Kantians because they emphasize the importance of one's understanding, and acting from, the right reasons. Kantians would argue that utilitarians not only use the wrong moral reasons in terms of their advice, but that version of the app also motivates people in the wrong way. Therefore, a Kantian would need to justify allowing people the option of using an app that motivates through pleasure rather than through doing things for the right reason.

One potential justification is that because it is important for the user to act on the basis of reasons they understand, it is best for the utilitarian to do what they see reason to do. The utilitarian presumably sees good reason to use the utilitarian version of the app and so the Kantian could support their use of the app because it is what is in line with their reason. If someone had moral understanding and therefore more reasons to use the utilitarian version of the app, then the Kantian could believe that the utilitarian actions have moral worth. The user would still have to be motivated by more than just the points or levels within the app. However, presumably there are some utilitarians who would choose to use that version of the app because they genuinely believe that utilitarianism is what is best and has moral worth. Those utilitarians may not understand the reasoning behind Kantianism and therefore not have moral understanding when it comes to the decisions that the Kantian version of the app would tell them to do. Therefore, the Kantian could support the utilitarian using the utilitarian version of the app.

Chapter Four: The Virtue Ethics Perspective of the 'Do Good' App

I. Introduction to Virtue Ethics

Virtue ethicists argue that Kantians and utilitarians place the focus of ethics in the wrong place. Rather than being focused on determining the right thing to do, virtue ethics views the question of ethics as being how to live a good life. Julia Annas argues that virtue ethics is about more than just doing the right thing. She argues that “issues of the best life to live, a good person to be and a good character to have are important along with doing the right thing”.³⁶ In order to answer this question about what kind of person we should be, virtue ethicists look towards people who are already doing the right thing: according to virtue ethics, the right thing to do is what the virtuous person would do. The virtuous person not only does the right thing, but also does it for the right reasons and - in particular - feels the right thing at the right time. For example, the virtuous person not only chooses to give to an effective charity but does so because they know it will help people, rather than for points on the app. They also empathize with who they are helping and enjoy the act of donating, rather than doing it from a place of guilt or a pure sense of duty.

In order to illustrate how this theory operates and how it differs from utilitarianism and Kantianism, imagine Sally again. In this situation imagine that she has to decide between going to the meeting where she will earn three thousand dollars or visiting her friend in the hospital. In order for Sally to make the right decision in this case according to virtue ethics, Sally has to do

³⁶ Annas, Julia. "Being virtuous and doing the right thing." *Proceedings and addresses of the American philosophical association*. Vol. 78. No. 2. American Philosophical Association, 2004, 66

the right thing for the right reasons and with the right feelings. It seems intuitive that there would be something missing ethically if Sally was only visiting her friend to get out of going to work or to make herself feel less guilty. This intuition is inconsistent with the utilitarian argument because utilitarians only care about the outcomes of an action rather than the reasons why someone chose to do that thing. Although Kantians do agree with virtue ethics that it is necessary to act for the right reasons in order for the act to have moral worth, there still seems to be some element of being good that is missing from the Kantian conception of good. In addition to doing the right thing and having the right reasons for acting, it also seems intuitive to think that Sally was not fully ethical if she was miserable the whole time that she was visiting her friend. If Sally went to visit her friend out of a strict sense of duty rather than because she wanted to help, then there seems to be something ethically missing from Sally. Some versions of Kantianism would actually see the actor being motivated at all by the pleasure they would get from visiting their friend as a detraction from that person's motivations to act on a pure sense of duty. According to them, the act would thus lose its moral worth. However, there seems to be something intuitively wrong about that. Virtue ethics captures this intuition to argue that in addition to someone needing the right reasons behind their action, the way they feel while doing that thing is also necessary to the moral worth of an act. Specifically, it is an important element of virtue ethics that the actor gets fulfillment or meaningful happiness out of acting the way the virtuous person does. This additional necessary element of doing the right thing fixes the alleged problem with Kantian ethics because it makes sense of the value of having the right feelings while acting. Aristotle called this result of living a good life 'eudaimonia'.³⁷ For virtue ethics, getting

³⁷ Although some virtue ethicists see eudaimonia as being central to determining what the right thing to do is, for the sake of this paper I will limit the understanding of eudaimonia to it simply being the result of acting virtuously.

fulfillment or happiness for an act is important to that act having full moral worth. Each of these elements is important for someone to be a good person according to virtue ethics.

In addition to arguing that focusing on being the right kind of person is necessary for ethics, virtue ethics also argues that only the virtuous person knows what the right thing to do is. The virtuous person is the only one who knows because, according to virtue ethics, there is no codifiable algorithm for determining the right thing to do. Ethics is not codifiable because the particulars of the situation matter a great deal for determining what the right thing to do is. Additionally, there are multiple and sometimes contradictory values that have to be balanced in each scenario. Virtue ethics can be seen as a type of ethical pluralism because it involves balancing multiple values and this pluralism makes it very difficult to codify the right thing to do. The values include things such as being just, honest, loyal, having wisdom, displaying temperance etc. Balancing values is important but difficult because there are situations where the values will conflict. For instance, displaying temperance might require not being honest. Although both values cannot be equally displayed in the circumstance, there is a correct way for them to be balanced. In Sally's case again, in order to determine what the right thing is, Sally has to be aware of the relevant values and how the particulars of each option affect how the values should be implemented and balanced. For instance, in order to understand how the value of friendship affects the two choices, she needs to know whether her friend would want her there. She needs to be able to perceive whether her visit would be appreciated as a gesture of support and friendship or whether it would be a violation of her friend's privacy. Additionally, in order to know whether it would be supportive or a violation of privacy, Sally needs to know what her friend is in the hospital for and how they would feel about her coming to visit. In addition to knowing about the option of going to visit her friend in the hospital, Sally also needs to be aware

of the values that factor into the decision to go to work. Only the truly virtuous person is able to be aware of all these factors and particulars and balance the values correctly.

One might argue that Sally's example only demonstrates that ethics is extremely difficult to codify, but not that it is impossible. Although there are lots of particulars and those particulars take a virtuous person to be able to notice, the virtuous person may still just be the person who is best at understanding how to apply the correct, complex ethical "code" amidst the numerous variables. However, in her article "Being Virtuous and Doing the Right Thing" Julia Annas raises an objection to the codifiability of ethics which she calls the 'objection from the lonesome advisor'.³⁸ For this objection Annas offers the example of someone who has a horrible moral character but offers extremely good moral advice. This hypothetical person may torture kittens but still be able to give advice that is correct. There seems to be a problem with this kind of situation where the advice that someone gives could be correct despite their character being morally repugnant. However, if ethics was truly codifiable, then there should not be anything strange about this lonesome advisor example. If it is true, as it intuitively seems, that a bad person can not give good moral advice, then that suggests that ethics is not codifiable. If you need a person with a good character to get the right advice, then that would suggest that ethics is not codifiable.

II. The Codifiability Objection

If ethics is not codifiable, then this poses a serious problem for the app. The app can only operate with codes and algorithms and so if there is more to determining the right thing to do than can be boiled down into those codes, then the app is not going to be able to correctly advise

³⁸ Annas, Julia. "Being virtuous and doing the right thing." *Proceedings and addresses of the American philosophical association*. Vol. 78. No. 2. American Philosophical Association, (2004), 64

the user. In Section III, I will consider objections to the ways that the app cannot handle the other aspects of ethics that are necessary for virtue ethics (concerning acting for the right reasons and with the right feelings), but this section will focus on whether the app can determine the right thing to do.

One possible response to the problem of uncodifiability is to have the app do a statistical analysis of all virtuous people to determine some kind of algorithm or equation from that data as to what the right thing to do is. The app could be fed an almost limitless amount of data of what virtuous people have done in past situations. From that data the app could identify patterns and consistencies between the virtuous people. When the user asks the app a question, the app would pull from its data set and find what someone in the past had done in a similar situation. If there was no data from a virtuous person's past similar situation, then the app could use the patterns that it found to come to a possible answer as to what the user should do.

One problem with this solution is that it requires identifying virtuous people.³⁹ This is a problem that is innate to virtue ethics and so I will leave that to virtue ethics to solve; for the sake of understanding the 'Do Good' app I will assume that we are able to correctly identify who the virtuous people are to the extent to which the app is able to have an accurate set of data to reach conclusions about what the virtuous person would do in a given circumstance. An additional, more pressing problem concerns using a collection of data in which virtuous people disagree, or at least appear to disagree. It might be that there is no real disagreement between virtuous people, either because there are very few virtuous people or because their acts are all so case specific that

³⁹ Identifying the virtuous people is a potential problem because different versions of virtue ethics place fundamental priority on different elements of the theory. Whether the virtuous person makes the act good or doing the good act makes someone virtuous is a problem that is innate to virtue ethics and so I will leave that to virtue ethics to solve. In particular the "priority problem" objects that defining the virtuous act as the act that the virtuous person does puts the priority in the wrong place. Rather than defining the virtuous person as someone who balances values correctly, most versions of virtue ethics argue that values are balanced correctly because the virtuous person did them that way.

anything that may look like disagreement is really just a variation based off of extreme specifics in the situation. One benefit of having the app do this data collection and sorting for the user is that the app can narrow down exactly what the 'act' is to be extremely specific. Therefore, it may be that no two people have ever really done the same 'act' given how case specific the act is and therefore there would be no disagreement.

However, even if there was true disagreement amongst virtuous people, Shafer-Landau in his book *The Fundamentals of Ethics* offers a solution to this problem.⁴⁰ According to Shafer-Landau, assuming that all of the people that have been deemed virtuous are really virtuous, if there is disagreement between them, then that affects how morally necessary or permissible we view the act that they are disagreeing about. If all virtuous people would do something, then that act is morally required. If some virtuous people would do something, then that act is morally permitted. If no morally virtuous person would perform an act, then that act is morally forbidden.⁴¹ The app could draw on Shafer-Landau's approach when determining what to tell the user in cases where virtuous people disagree. Therefore, the app would only tell users that doing something was morally necessary if all virtuous people agreed that it was the right thing to do.

III. The Necessity of Correct Emotions Objection

Even if the app is able to accurately tell the user what to do based on the data set of virtuous people, there is still something missing from the app's ability to actually assist the user in correctly aligning with virtue ethics. As discussed in the introduction to this chapter, according to virtue ethics, the focus of ethics needs to be who we are and what we feel at the time that we

⁴⁰ Shafer-Landau, Russ. "The fundamentals of ethics." (2009), 288

⁴¹ Ibid.

are acting, not just what we do. It is necessary that users of the app be good people, not just that they do the right thing. Initially, one might think that you could solve this problem by having people memorize the right reasons. If people have these reasons memorized and are always acting from those reasons, then it may seem like they are morally good. Although this might ostensibly meet the requirements from Kantianism, it is not enough for virtue ethics.

Philosophers refer to the fixation on doing things ‘for the right reasons’ as moral fetishism.⁴² Moral fetishism is the idea that simply doing the right thing because you know it is the right thing is not enough for the act to have moral worth. Rather, the act needs to come from the right kind of character that directly appreciates the moral features that are at stake in the situation.⁴³ This kind of appreciation involves emotions that are appropriately fitted to the situation rather than just cognitive abilities. That emotional response requires experience with the different values that are at stake. One specific way experience affects our ability to feel the correct emotions is that having experienced harm allows us to feel and understand the weight of harm in other cases. Having experienced harm, we then can feel the badness of that harm when it happens in other cases.

In Sally’s case, had Sally been in the hospital before then she would know what it would feel like to be in the situation that her friend was in and she would have experienced the kinds of harms that her friend would suffer as a result of her not visiting. That past personal experience of having been in the same situation as well as the past experience that she has with her friend so that she knows what would be best for them, are necessary for Sally to make the correct decision. Those experiences would make Sally want to go visit her friend. Those kinds of experience form

⁴² Smith, Michael R. "The moral problem." (1994), 76

⁴³ Chappell, R., 2019. "Why Care About Non-Natural Reasons?" *American Philosophical Quarterly*, 56(2): 125–134.

the kind of character that is able to feel the right things at the right time and also for the right reasons.

It does not seem possible that the app would be able to inculcate the kind of necessary dispositions of feeling into the user of the app. Simply being told the right way to feel is not going to be enough for someone to actually feel that way. One might argue that since the app is not going to be able to get you to feel the necessary way, the app is inevitably inadequate as a moral 'advisor', and so should not be used. Since having the correct feelings at the correct time is a necessary part of having a complete moral life, the app cannot make you a good person without being able to do that element.

However, the app does not have to cover every aspect of what is necessary for living a good moral life in order for it to be worthwhile for the user to have the app. The app would be a tool that the user could use to the extent to which was possible, and then the user would manage the rest of what was necessary to live a full moral life on their own, as they would before the possibility of the app.

One problem with this response to the objection is that the app may actually *distract* the user from being virtuous. Rather than helping, the app could easily *undermine* the user's ability to do the things that are necessary for being virtuous. The app accessibility, its ease of use, and its potential for being addictive would increase the chances of the user overusing the app and then becoming complacent with that kind of ostensibly 'ethical' life. Since the user would be able to do the right thing according to the app, they may be even more easily inclined to forget the other elements that are important for being virtuous than they would if they were not using the app--for instance, they may be even more easily inclined to forget the importance of inculcating the correct moral dispositions of feeling associated with virtue.

This problem could be solved by the app directing the user towards the experiences necessary for developing the correct emotions. These experiences would include both the user watching a virtuous person experience those emotions and the user practicing being in situations where they started feeling the correct emotions. For instance, imagine Sally again choosing where to donate to charity. At first, Sally may not want to donate. However, the app could direct Sally towards witnessing other people who are donating and getting fulfillment from it. Once Sally is able to witness that kind of fulfillment from other people, she may be more inclined to donate and to do so for the right reasons and with the correct emotional dispositions. Even if those kinds of generous people did not already exist in her life, the app could show Sally movies, documentaries, or other stories of people who have donated generously. The app's instructions or recommendations would lead the user towards experiences that would give the user an intellectual and emotional appreciation for what the right thing is, rather than just telling the user what to do.

Despite all the features that a virtue ethics version of this app could offer a user, a virtue ethicist might still object that the app still does not do enough to help the user feel the right kinds of feelings at the right time. That may be true. The app may not be able to fully make someone virtuous. However, the main goal is that the app needs to simply not distract the user from being able to gain those feelings. So long as the app helps in that it aids the user in making the right decision, and it does not harm the user in other ways, the app would still be worthwhile. To address this problem, the app could be programmed to have limitations in the amount it is able to be used or exactly how much it feels the user such that the user not reach the point where they forget about or get distracted from the importance of forming the right feelings.

IV. Conclusion: The Incompleteness Objection

Someone may object that even with the kinds of adaptations I have discussed in this chapter, the app is still not as good as having a real, personal moral exemplar in your life. However, not everyone has a moral exemplar. Moral exemplars are difficult to find and identify and it is very possible that most people do not have one in their life. In those cases, the app is able to be a stand-in moral exemplar. Exactly what kind recommendations that app would give the user would depend on what other tools they already have in their life. If someone already has a moral exemplar in their life that they look up to, then the app would not need to use its database on information about moral exemplars. Instead it would instruct the user to talk with the moral exemplar that they have in their life. Possibly someone has a moral exemplar in their life but they are not aware of it. The app could serve as a tool for the user to determine who if at all the moral exemplars are in their life. If the user does not have a moral exemplar in their life, then that database of information about moral exemplars could be used as a kind of moral exemplar for the user when they have specific questions about what to do. If the user is already involved in experiences that are assisting them in becoming a virtuous person, then the app will not need to recommend any kind of shift in that area. However, if the user would become better through experiences such as volunteering, working for a specific organization, etc. then the app could recommend those kinds of life changes to the user. As mentioned in the last paragraph, throughout all of these uses the app would also remind the user that the app is not all that there is to being a virtuous person. If the user were to begin to show signs of becoming dependent on the app (such as using it too frequently) then the app could tell the user to stop using it all together.

Although the real moral exemplar may be better at helping the user through some parts of the moral experience, there are other regards in which the app may be more even helpful than

real exemplars. One particular way that the app would significantly help the user is in forming habits. Habits are an important element of becoming a virtuous person because they help the virtuous person establish the kind of character that is consistent with being a virtuous person. In many cases, doing the morally correct thing is not going to feel right until the person has done it multiple times. However, after doing it over and over, the user begins to experience the feelings that are appropriately connected with doing the right thing. Nancy Sherman argues that, "practice yields pleasure to the extent to which it exhibits increasingly fine powers of discernment".⁴⁴ One example of how this works is with exercise. The first time you go on a run it hurts, it's not enjoyable, and you are really only doing it because you know that it will make you healthy. However, after running regularly for a couple of months, you begin to be running because you enjoy the actual process of running. The same is true with acting virtuously. Think back to the case of Sally having to decide how to spend her money. At first, Sally may really not want to donate it. A new phone or a fancy vacation may be extremely appealing to her initially. However, once Sally starts regularly donating her money, she can begin to enjoy the process of donating. She no longer feels like she needs the new phone or the fancy vacation and she gets pleasure from knowing that her money is helping other people. Even if this new habit did not result in Sally having more happiness than before, there is reason to believe that her overall happiness would at least not decrease significantly. Humans have the tendency to revert back to a stable level of happiness after these kinds of big changes. This tendency is often referred to as the hedonic treadmill.⁴⁵ Although the fancy vacation may make Sally happy for a moment, she would return to her previous state of happiness after having made that change. The same is true for Sally deciding to give to charity more often. Although it may require what initially feels like

⁴⁴ Sherman, Nancy. *The Habituation of Character*, in: Sherman, Nancy. *The Fabric of Character*. 1989.

⁴⁵ Rosenbloom, Stephanie (August 7, 2010). "But Will It Make You Happy?". *The New York Times*.

sacrifices, Sally would go back to generally feeling the same amount of satisfaction from her life as she was before giving away a lot of her money. The app can help Sally form that habit by giving her reminders in a more consistent and attentive manner than any human moral expert could.

The app can help many people in many ways like it helps Sally donate money. Whether that is determining where to donate, what grocery store to shop at, or what kind of products to buy, the app can help people form habits so that they develop a sense of enjoyment around doing the right thing.

Conclusion

Throughout this thesis I have argued that there is a need for us to be able to make better ethical decisions and that this need can be met through the use of the ‘Do Good’ app. In addition to helping users make better decisions for themselves, the app also has potential consequences for leadership. Leadership is in many ways a method of one person helping to influence another’s decision making process in some ways. James MacGregor Burns defines leadership “as leaders inducing followers to act for certain goals that represent the values and the motivations – the wants and needs, the aspirations and expectations – of both leaders and followers”.⁴⁶ Followers will often have broad or undefined moral goals, but leaders use their influence to ‘induce’ those followers in specific ways in order to pursue those goals. However, leaders, like everyone else, make mistakes, and therefore the ‘Do Good’ app would be able to better help leaders make decisions that are best for their followers. The app also gives us insight into what extent leaders are able to help their own followers make ethical decisions. Throughout

⁴⁶ Burns, James MacGregor. “Leadership” *Harper and Row Publishers*, (1978), 19

this thesis I have outlined the important considerations to be made when someone is being helped with their decision making process. It is vital that the leader know how much and in what ways they can or should help their followers make ethical decisions for themselves. Therefore, this same framework for assessing the app's ability to help the user could be used to assess a leader's ability to help their followers.

Although I have at certain points addressed some elements of the feasibility of an app of this sort, I in many ways assumed that it would be something that could be created and as well as funded. I am still going to assume that the technology does or could exist to be able to support an app of this sort. However, in this conclusion, I will offer some possible means of funding an app of this sort, none of which are perfect, but all of which offer some framework for thinking about how to actualize the 'Do Good' app.

The concern about funding is a genuine problem that would have to be overcome for the app to be actualized. First, the app consists of a considerable amount of coding that would require people to be hired to do that work. Second, there would need to be philosophers employed during that coding process to ensure that it was done correctly. Third, there are issues of potentially needing additional technology to implement the app to its fullest extent for each of the theories and that additional technology would cost money to create. Finally, some of the theories also require having on call philosophers available to confer with the user when they have questions. All of these factors create a cost for creating the app which means that someone is going to need to pay that cost. For each possible option for funding it is important to consider three different things: what is lost by that funding being directed towards the app rather than towards something else, what opportunities for corruption, biases, or unfairness are created by

the app being funded in that way, and what is the actual feasibility of the source being willing to fund the app.

One option is to have users pay to use the app through the use of a subscription. However, the creation of a necessary paid subscription for the app would make it inaccessible for many people who want to use the app. This inaccessibility is unfair because it allows people of a certain wealth standard to be able to act more ethically than people who cannot afford to buy the subscription. Each moral theory that I have outlined in this paper would agree that acting morally should not be something that you have to pay in order to be able to do. Additionally, this method of funding would not be able to cover the start up costs of creating the app. The initial money that it would cost to code the app would not be able to be covered by money generated by people downloading the app. Therefore, since the necessary subscription would unfairly leave some people out of being able to use the app and the subscription could not cover start up costs, there should be no subscription needed to download and use the app.

One possible modification of this idea is to have the app factor in its own upkeep into its considerations of where the users should donate. The app is already aiding the user in deciding where to donate. If paying for the upkeep of the app is necessary for the app to continue being able to help the user, then the app could take a small portion of the money it is already having the user donate. However, having the app tell users to donate to itself creates an issue of the app having a self serving bias. Should the creators of the app recognize that they personally could make money from programming the app in such a way as to get users to donate more money to the app, then the app would no longer be serving its purpose of helping users do good. This method of funding also directly takes away from the money that would be used to help other people.

Therefore, we need another solution to funding the app that does not rely on the users donating money to the app.

One option is that large corporations could fund the app through a certification program similar to Fair Trade or other ethical accreditations. These certification programs operate by large companies paying the certifiers to ethically audit their company. The certifiers check to determine the efficacy of the company's decisions such as how they hire their employees, their supply chain, how they advertise, etc. The audit is then completed by a private, independent consulting firm which determines if the company meets the ethical requirements of the accreditation. The 'Do Good' could implement this method by having corporations pay the app to audit their company. As mentioned throughout this paper, one major role of the app in people's lives is going to be to help them decide how to spend their money. Therefore, the app is already going to be doing the kind of audit of companies necessary to determine if they are ethical enough to have users buy from them. Since this work is already happening, the app could simply add the feature of some products or companies being accredited as 'Do Good' certified. One concern with having the app take money from corporations is that the app could become biased by the money. Brands could offer to pay a lot of money for the certification even if they did not meet the requirements. Brands would be motivated to do this because having this kind of accreditation would make their product more appealing to consumers.

The creators of the app may be vulnerable to taking these bribes because of how badly they may need the money in order to keep the app running. When the app takes bribes from businesses that it is accrediting, it loses the legitimacy of the accreditation and therefore the app will have lost legitimacy as well. This potential for bias is a real concern, but it assumes that the creators of the app would be willing to take such bribes. Corruption is always possible, but

assuming that the creators of the app had the genuine intention of trying to help people become better, it seems unlikely that the company would take bribes that would undermine that goal. However, this is still a real concern because it would require the 'Do Good' app to become a kind of corporation in and of itself which makes it vulnerable to corruption and capitalism. For instance, if the app became a corporation there would have to be leaders within that corporation. Once the corporation is making money, those leaders could become corrupted by the possibility of themselves making money off of the app. If the app becomes focused on making money for itself rather than helping users do good, then the app becomes entirely pointless. Therefore it seems we need another solution.

Another option is for the government to fund the app. The government could give a grant to the 'Do Good' app that would provide for the initial costs of creating the app including the costs of forming any new technology, hiring people to code, and philosophers to ensure that it is all happening correctly. The government could be motivated to do this through an argument similar to the 'personal is political' argument. Susan Okin and other second wave feminist have argued that decisions that are made in private affect the public political achievement of justice. Therefore, 'the personal is political' and there is an extent to which the government should be involved in people's personal life choices. Okin and others who first started using this phrase were referring more specifically to issues concerning the effects of the patriarchy within the home. However, the same idea can be extended to other personal issues that begin as personal but extend into the political realm. The 'Do Good' app helps users make better personal decisions and, under the same framework that Okin and others have used, these decisions become political decisions because they have the potential to affect people's relation to politics. For instance, our personal decisions are often made with a racial bias. This racial bias is able to

seep into many of our decisions especially those that are made at a political level. One example of racial biases in the political sphere is the disproportionate incarceration rates among black men compared to any other demographic.⁴⁷ This disproportionate incarceration rate is due in part to racial biases present in the decisions of juries, judges, and police officers. The app would be able to filter out this bias in those individuals' decisions to an extent and therefore prevent this personal problem from becoming a political one.

This solution would result in the app being paid by tax dollars which would require people to vote in politicians who supported the app. Using tax dollars would mean that the money would be redirected from some other source. Assuming that the government is already using tax money to better its citizens in some way, funding the 'Do Good' app would require taking away support in another area. Although this may initially be a concern, the 'Do Good' app would eventually result in less need for government social support. For instance, the 'Do Good' app would probably tell the majority of people to quit smoking because it is an unnecessary cost and bad for their health. If people quit smoking, then there is the potential that government funding for health care could go down because smoking related hospital stays would be far less common. The same could be said for the app telling people to eat less meat or drive slower or donate to homeless shelters. Although the app and the government have different goals, in many ways they could be seen to be working towards similar outcomes. Therefore, although the app may cost the government money in the short term, in the long run the cost would pay off.

One drawback to the app being government funded is that it is probably very unlikely that the government would be willing to make this kind of short term investment for the long term reward. Although it is possible, there is no real guarantee that the app would result in

⁴⁷ Alexander, Michelle. *The new Jim Crow: Mass incarceration in the age of colorblindness*. The New Press, 2020. Chapter Three

outcomes that would be cheaper for the government, so it would be a real risk for people to take. Additionally, many people do not trust the government and would not respond well to being told exactly what to do by an app that was funded by the government.

A final potential option for funding the app is through Universities. Universities already have time and money set aside for their professors to be involved in personal research. If Universities could be properly motivated to encourage some of their professors and research students to focus their attention on developing the app, then the app could be funded through Universities. Undergraduate and graduate students in particular are an excellent resource for universities to utilize because the programs already exist to get unpaid research and writing labor out of those students in exchange for class credit. Students quite literally pay universities to give the university this labor which creates an opportunity for Universities to use that labor in ways that better the University. The University would only need to organize those individual research projects happening between students and professors across the university so that they could work together to create the app. Universities could be motivated to divert research attention in the direction of the app because of the positive press that would come from developing something this monumental. If the app were to exist to the extent that I have laid out in this paper, it would be an avenue for massive change across the world. This kind of opportunity for change would garner significant positive publicity for the university that was able to get it off the ground. Therefore, universities could be motivated to do the organizing and funding necessary to pull students and professors together to create the app. After the initial creation of the app, there would still be the costs of providing philosophers for the users who use a version of the app that allows them to communicate with a philosopher. The universities would be able to cover these costs as well through the use of their ethics professors. By adding that work into the contracts

that professors sign with the university, the university would be able to supply the necessary philosophers to answer those questions. It may be that one university would not be able to provide enough resources to cover the initial or continuing costs. If that were the case, multiple universities could pool together to supply the professors and students necessary for creating the app.

None of these potential funding sources would be perfect, but they are all a starting point in beginning to see the 'Do Good' app as a real potential tool for people becoming better moral decision makers. It is incredibly important that even if we do not become perfect moral decision makers we at least attempt to become better ones. As mentioned in Chapter Two, we are currently facing a kind of emergency situation where even becoming slightly better moral decision makers will make an incredible difference. There are currently 2.2 million cases of malaria each year which lead to about half a million deaths a year, and these deaths are easily preventable through the use of mosquito nets and low-cost treatment programs.⁴⁸ Additionally, about 25% of the world's population is malnourished, 29% lack safe drinking water on their premises, 50% lack basic health care, and climate change is continuing progress at an exponential rate in part due to the production of three billion animals that are killed daily for humans to eat.⁴⁹ By making better decisions with the aid of the app, we would be able to mitigate the effects of these kinds of global, large scale atrocities. By helping guide us towards

⁴⁸ World Health Organization, "Malaria" *Global Health Observatory*, (2017), <https://www.who.int/data/gho/data/themes/malaria>

⁴⁹ World Health Organization, "Progress on Drinking Water, Sanitation and Hygiene" (2017) <https://www.who.int/mediacentre/news/releases/2017/launch-version-report-jmp-water-sanitation-hygiene.pdf>

World Health Organization, "Tracking Universal Health Coverage: 2017 Global Monitoring Report" (2017) <https://pubdocs.worldbank.org/en/193371513169798347/2017-global-monitoring-report.pdf>

Zampa, Mathew. "How many animals are killed for food everyday." *Sentient Media* (2018).

donating, consuming less, and purchasing more ethically, the app could make the world a significantly better place.

References:

Alexander, Michelle. *The new Jim Crow: Mass incarceration in the age of colorblindness*. The New Press, 2020. Chapter Three

Annas, Julia. "Being virtuous and doing the right thing." *Proceedings and addresses of the American philosophical association*. Vol. 78. No. 2. American Philosophical Association, 2004.

Barrett, William. "In Pictures: America's 10 Least Efficient Large Charities" *Forbes* (2010) https://www.forbes.com/2010/11/17/charity-10-personal-finance-philanthropy-10-least-efficient-charities_slide.html?sh=37ab74c4556e

Burns, James MacGregor. "Leadership" *Harper and Row Publishers*, (1978).

Brown, Jessica. "Knowledge and practical reason." *Philosophy Compass* 3.6 (2008): 1135-1152.

Chappell, R., 2019. "Why Care About Non-Natural Reasons?" *American Philosophical Quarterly*, 56(2): 125–134.

Christensen, David. "Disagreement as evidence: The epistemology of controversy." *Philosophy Compass* 4.5 (2009): 756-767.

Clifford, William K. "The Ethics of Belief." *Contemporary Review*, (1877).

Driver, Julia. "Ethics: The Fundamentals." (2007).

Esgate, Anthony; Groome, David . “An Introduction to Applied Cognitive Psychology”. *Psychology Press*. (2005) p. 201.

Food and Agriculture Organization of the United Nations “Preventing micronutrient malnutrition a guide to food-based approaches- Why policy makers should give priority to food-based strategies” *International Life Sciences Institute* (1997).
http://www.fao.org/docrep/x0245e/x0245e01.htm#P38_2721

Global Education Monitoring Report: Target 4.1- Primary and secondary education (2015)
<https://gem-report-2017.unesco.org/en/chapter/primary-and-secondary-education/>

Gutting, G. “Religious belief and religious skepticism.” *University of Notre Dame Press*, Notre Dame, IN (1982)

Hawthorne, John, and John P. Hawthorne. *Knowledge and lotteries*. Oxford University Press, 2004.

Hawthorne, John, and Jason Stanley. "Knowledge and action." *The Journal of Philosophy* 105.10 (2008): 571-590.

Herman, Barbara. "Mutual aid and respect for persons." *Ethics* 94.4 (1984): 577-602.

Hills, Alison. "Moral testimony." *Philosophy Compass* 8.6 (2013): 552-559.

Hong, Lu, and Scott E. Page. "Groups of diverse problem solvers can outperform groups of high-ability problem solvers." *Proceedings of the National Academy of Sciences* 101.46 (2004): 16385-16389.

Jones, Karen. "Second-hand moral knowledge." *The Journal of Philosophy* 96.2 (1999): 55-78.

Lackey, Jennifer. "12. What Should We Do When We Disagree?." *Oxford Studies in Epistemology Volume 3* (2010): 274.

Landemore, H el ene. *Democratic reason: Politics, collective intelligence, and the rule of the many*. Princeton University Press, 2017.

Le Guin, Ursula K. "The ones who walk away from Omelas." (1975).

Lum, Kristian, and William Isaac. "To predict and serve?." *Significance* 13.5 (2016): 14-19.

Mill, John Stuart, “On Liberty”. *Yale University Press*, (2003)

- Nozick, Robert. *Anarchy, state, and utopia*. Vol. 5038. New York: Basic Books, 1974.
- Rosenbloom, Stephanie (August 7, 2010). "But Will It Make You Happy?". *The New York Times*.
- Rowland, Richard. "The epistemology of moral disagreement." *Philosophy Compass*. 2017;12: e12398. <https://doi.org/10.1111/phc3.12398>
- Singer, Peter. "The Drowning Child and the Expanding Circle." (1997).
- Singer, Peter. "Practical ethics". *Cambridge university press*, 2011.
- Shafer-Landau, Russ. "The fundamentals of ethics." (2009).
- Sherman, Nancy. *The Habituation of Character*, in: Sherman, Nancy. *The Fabric of Character*. 1989.
- Smith, Michael R. "The moral problem." (1994).
- Stanley, Jason. *Knowledge and practical interests*. Clarendon Press, 2005.
- Sylla, Ndongo Samba. "Fairtrade is an unjust movement that serves the rich" *The Guardian*, (2014)
- Walzer, Michael. "Political action: The problem of dirty hands." *Philosophy & public affairs* (1973): 160-180.
- Wieland, Jan Willem, and Philip Robichaud, eds. *Responsibility: The epistemic condition*. Oxford University Press, 2017.
- World Health Organization, "Malaria" *Global Health Observatory*, (2017), <https://www.who.int/data/gho/data/themes/malaria>
- World Health Organization, "Progress on Drinking Water, Sanitation and Hygiene" (2017) <https://www.who.int/mediacentre/news/releases/2017/launch-version-report-jmp-water-sanitation-hygiene.pdf>
- World Health Organization, "Tracking Universal Health Coverage: 2017 Global Monitoring Report" (2017) <https://pubdocs.worldbank.org/en/193371513169798347/2017-global-monitoring-report.pdf>
- Zampa, Mathew. "How many animals are killed for food everyday." *Sentient Media* (2018).