



Does Effectiveness Matter?

Measuring Consumer Value for Corporate Social Responsibility

by

Michelle Nee

Thesis Advisor: Justin Frake – Assistant Professor of Strategy

A thesis submitted in fulfillment of the requirements of the
Michigan Ross Senior Thesis Seminar (BA 480), 4/22/21

Abstract

Actions and discussions around for-profit corporations' duty of "Corporate Social Responsibility" (CSR) have grown exponentially in the last decade. Simultaneously, there has been a growing discourse in the philosophy community around the concept of "effective altruism": using evidence and data to determine the most effective ways of benefitting others. However, the effective altruism conversation is rather siloed to charitable organizations, while CSR typically relates to corporates and their core business interests. This paper serves to connect these two thought processes. In an experimental model, individuals' willingness to pay for consumer goods will be measured against two variables: the generosity and the effectiveness of that good/brand's CSR initiatives to determine if effectiveness, in terms of number of lives saved, has a positive effect on consumers' willingness to pay.

Contents

Introduction	pg # 1
Corporate Social Responsibility	pg # 2
Effective Altruism	pg # 2
Business Implications	pg # 4
Problem Statement and Justification	pg # 6
Literature Review	pg # 7
Corporate Social Responsibility Literature	pg # 7
Effective Altruism Literature	pg # 11
Theoretical Framework	pg # 14
Critical Assumptions	pg # 14
Theory of Planned Behavior	pg # 16
Proposed Methodology	pg # 18
Tradeoffs	pg # 18
Logistics	pg # 19
Pilot Study	pg # 20
Methods	pg # 20
Results and Discussion	pg # 21
Primary Study	pg # 25
Methods	pg # 25
Results	pg # 27
Discussion	pg # 37
Limitations	pg # 39

Considerations for Future Research	pg # 41
Conclusion	pg # 42
Appendices	pg # 43
References	pg # 57
Acknowledgements	pg # 59

Introduction

“Organizations can no longer be viewed merely as economic machines designed for technological progress and personal benefit for those who control them. Instead, they must be seen as sociotechnical systems responsive to human needs both in their external and internal environments. As human systems, organizations must develop a moral obligation to respond to the needs of consumers, minority groups, and others in their external environments. They must also respond to the social and altruistic needs of their members.” (Conger & Rabindra, 1993)

We are no longer standing at the crux of the Milton Friedman era. In his 1970 paper, “The Social Responsibility of Business is to Increase its Profits,” Friedman argued that socially responsible actions from corporations may hurt firm profits and curb the market. In his eyes, a free, capitalistic market is the most effective path towards optimal human welfare and should be prioritized over charities and socially responsible initiatives (Friedman, 2007). Over the past decade, however, consumers have become less enthralled with Friedman’s concept of “shareholder primacy” and have begun to investigate corporations’ overall impact on society. In fact, there is a positive relationship between consumers’ purchase intention scales and firms’ “Corporate Social Responsibility” (CSR) activities (Lee & Shin, 2010).

Separately, the rise of “effective altruism” in the world of non-profits and charities has created additional discourse around how to create more societal good with our given resources. This discussion is essential because the difference of effectiveness within charities is quite significant. In fact, experts estimate that the most effective charities are up to 100x more effective than other charities (Caviola, et. al, 2020).

Corporations that now face immense pressure from consumers who are looking for companies that recognize the aforementioned moral obligation for societal good can learn a lot from the effective altruism community. Similarly to charities, CSR initiatives are also variant in their effectiveness, and it is essential to understand this variance in order to maximize societal

welfare. However, one question still remains: will the introduction of an effective, altruistic mindset to CSR benefit the bottom line for these corporations? This paper hopes to tackle one component of this complex question by investigating consumers' willingness to pay for products that are produced from companies with these types of CSR.

Corporate Social Responsibility

Broadly, CSR can be considered, "a concept whereby companies integrate social and environmental concerns in their business operations and in their interaction[s] with stakeholders on a voluntary basis. It is about enterprises deciding to go beyond the minimum legal requirements and obligations stemming from collective agreements in order to address societal needs" (Smith & Rønnegard, 2016). Standing in stark contrast to the shareholder primacy model, CSR initiatives are supported by the social contract theory: the concept that a corporation's existence should be valued based on its positive contributions to society (Appendix 1). Recent discourse within the CSR community has centered around how *firms* and *workers* benefit from these types of initiatives, but there has been far less conversation around how *consumers* view these initiatives, aside from a generally positive trend in their purchase intention scales.

Effective Altruism

Separately within the world of charitable giving, the philosophy of "effective altruism" has become more prevalent with help from scholars like Peter Singer and William MacAskill. At its highest level, the movement aims to determine how to "do the most good" with each dollar and hour that we have (Alternative definitions seen in Appendix 2). Manifestations of the movement's principles are expansive, from pledging a percentage of personal income to effective charities, to adopting a vegan diet, or to crafting impactful careers at think tanks or high-paying finance roles that create the financial freedom to donate to effective charities. The

movement does not attempt to define aspects of morality; it rather intends to remain broad as to be useful to people with differing moral views, but determinate enough to make a positive, net impact on the world. Within this broad definition, effective altruism attempts to follow the subsequent pillars of thought: (MacAskill, 2019, pg. 14)

- *Non-normative.* Effective altruism consists of two projects, rather than a set of normative claims.
- *Maximizing.* The point of these projects is to do as much good as possible with the resources that are dedicated towards it.
- *Science-aligned.* The best means to figuring out how to do the most good is the scientific method, broadly construed to include reliance on careful rigorous argument and theoretical models as well as data.
- *Tentatively impartial and welfarist.* As a tentative hypothesis or a first approximation, doing good is about promoting wellbeing, with everyone's wellbeing counting equally. More precisely: for any two worlds *A* and *B* with all and only the same individuals, of finite number, if there is a one-to-one mapping of individuals from *A* to *B* such that every individual in *A* has the same wellbeing as their counterpart in *B*, then *A* and *B* are equally good.²⁶

Inspired by this philosophy, charities and philanthropic organizations are beginning to leverage data and logic to determine which how much “good” they are actually creating. Peter Singer presents a vivid example of this comparison when discussing the organization Make-A-Wish. Consider a five-year old child that has been through three years of chemotherapy and now wishes to be “Bat-kid” for a day and ride around San Francisco in a “Batmobile.” It sounds like a great dream come true, but it is essential to consider that the average “wish” costs Make-A-Wish around \$7,500. Even though effective altruists are emotive and would likely love to see this wish completed, they also understand the opportunity cost of that \$7,500 spend. Upon hearing this story, they would consider that the money could be instead donated to organizations like the

Against Malaria Foundation to supply families with malaria bed nets, potentially saving the lives of multiple children (GiveWell, 2020). By prioritizing the effectiveness of these funds, effective altruists would look past their emotional pull to the “Batkid” saga and instead donate to the Against Malaria Foundation.

Business Implications

It is clear that effective altruism has the potential to increase societal good if metrics and incentives are aligned. So, why shouldn't this same mentality be applied in the corporate sphere? Despite the growing interest from supporters of the social contract theory, many corporates are still engrossed with maximizing shareholder value above all else. Due to this, most CSR initiatives end up being closely tied to a corporation's core business strategy or activities because of that firm's efficiency in those activities. For example, companies like TOMS offer a “one-to-one” business model that promises for every pair of shoes sold, the company gives one pair of shoes to a child in a developing nation. In this case, TOMS has chosen an initiative that is closely tied to their most efficient core competency: the production of shoes. Due to its expertise in this space, it is likely that every dollar spent at a company like TOMS would be more efficiently used to produce shoes compared to a non-profit with far less experience in shoe production.

However, it is possible that the funds spent on the shoes for children in developing nations could do “more good” if they were allocated towards other initiatives the development of healthcare infrastructure in the region, funding vaccinations, or other medical expenses for example. This would be a choice made on the basis of effectiveness, rather than efficiency. Within the scope of the effective altruism community, effectiveness is typically measured by the number of lives saved, since it one of the most objective measures available. This measure of “lives saved” has some limitations because it may discredit entities like universities and research

labs that take years to show tangible lifesaving results. However, when looking for actionable research results that can be implemented quickly, the three main methods of economic analysis including the cost-effective analysis (CEA), the quality adjusted life year (QALY), and the cost benefit analysis (CBA) all utilize the number of lives saved as a main component of their calculations (World Health Organization, n.d.).

Using these metrics, it is clear that the “one-to-one” model employed by TOMS has prioritized potential brand benefits and efficiency of shoe production over the genuine, effective well-being of children in developing nations. This dilemma is seen in a similar sense with charities as well. Efficient charities tend to be seen as those that allocate high percentages of donated dollars directly to causes rather than overhead expenses. The missing piece in this evaluation is the true effectiveness of these charities (i.e. how many lives are actually being saved).

Even if individuals within an organization like TOMS support the allocation of CSR dollars in support of more effective initiatives, it is unlikely that TOMS would adopt this mindset unless it proves to bring in more profits for the firm. If it is possible to demonstrate that allocating CSR spend towards more effective projects can also increase profits for the firm, there is immense potential to increase the amount of “good” created in the world. One way to do this is to test if customers are willing to pay more for products from brands that prioritize *effective* CSR initiatives. If the hypothesis is true that customers value effective CSR, corporations can create value by simply re-allocating their resources towards effective initiatives. With a higher willingness to pay from customers, corporations can charge higher prices and receive better margins, while the amount of “good” created is also simultaneously maximized as CSR funds are re-allocated.

Problem Statement

The purpose of this research is to understand differences in consumers' willingness to pay for products based on the effectiveness of a corporation's CSR initiatives. To determine these differences, an experiment will be leveraged that forces consumers to make a spending choice on a hypothetical product with varying levels of CSR effectiveness on the firm's backend. To do so, random assignment will be leveraged to account for individual's varying preferences. The current hypothesis of this paper expects that there will be a positive effect on consumers' willingness to pay for products that come from corporations with effective CSR initiatives, in terms of the number of lives saved.

Many corporations are willing to put funds behind CSR projects, recognizing the trend that consumers are demanding more responsible practices from the companies they support. However, the deep-rooted shareholder primacy model limits the societal good that can be contributed to the world, thus often restricting projects to those closely aligned to efficient, core business activities. If the hypothesis of this report is supported, it has the potential to convince for-profit corporations that the re-allocation of their CSR funds to effective endeavors can both increase brand value and societal good.

Literature Review

To understand the logic behind this multifaceted research question, it is essential to first comprehend the existing studies that fall within the field. First and foremost, the core literature of corporate social responsibility will be analyzed to understand the moral grounding behind it. From there, existing studies around customer willingness to pay for goods that are produced more ethically will be brought to light. These studies differ from the scope of this paper in that they focus on the products themselves, rather than socially responsible initiatives that are conducted by the firms that support them. The scope of socially responsible initiatives that are separate from product creation are much more expansive, from charitable donations to changes to the supply chain. Finally, the existing literature around the use of effective altruism within charities and other organizations will be outlined, enabling this paper to properly apply these principles to the corporate sphere. The culmination of the literature from these various spheres will give a more complete image into the minds of both consumers and firms around both CSR and effective altruism to ground the hypothesis of this research.

Corporate Social Responsibility Literature:

Overall, the literature on the effectiveness of CSR has been largely indeterminate. There are strong calls to action from leaders in strategy like Michael Porter to align a firm's competitive advantage with its CSR initiatives, but these remain highly qualitative. For example, in the HBS piece "Strategy & Society," authors Michael Porter and Mark Kramer argue that, "the more closely tied a social issue is to a company's business, the greater the opportunity to leverage the firm's resources – and benefit society" (Porter & Kramer, 2006). Furthermore, this narrative creates a dichotomy between "strategic" and "responsive" CSR, recommending that firms opt for

the former (Appendix 3). What is missing in this argument is empirical data to support the impact that strategic CSR has for firms, if it has one at all.

CSR's Relationship with Firm Profits

Scholars Abigail McWilliams and Donald Siegel uphold the view that, “the analysis of CSR is still embryonic, and thus theoretical frameworks, measurement, and empirical methods [have] not yet been resolved” (McWilliams & Siegel, 2006). With regards to firm profits, the results are often inconclusive because of differences in research methods. For example, McWilliams and Siegel recognized one potential flaw in existing studies: the role of research and development (R&D). They noted that the misspecification of R&D in previous empirical studies may lead to overestimates of CSR’s financial impact, however once adjusted for, CSR often has a neutral effect on financial performance (McWilliams & Siegel, 2000).

To build upon this research, some scholars like Caroline Flammer have begun to run further regressions on the relationship between CSR and firm performance. Flammer’s best known research uses a regression discontinuity to examine CSR proposals that pass or fail by a small margin of votes – proposals known as “close calls.” In her regression, Flammer discovered that the adoption of these close call proposals led to superior financial returns of 1.77% on average, increasing value for the firm. However, Flammer did not determine if these strong financial returns were a result of effective CSR in her study. It is inconsistencies like these that create differing views within the field.

CSR's Relationship with Individual Behavior

Economists Roland Bénabou and Jean Tirole wanted to analyze the impact of CSR through another angle: the individual consumer. Drawing from both psychology and economics, they determined that prosocial behavior by consumers is driven by various motives including

intrinsic altruism, material incentives, and self-esteem concerns. Furthermore, they implored policy makers and activists to garner a strong understanding of these motives in order to leverage the public for certain types of corporate interventions (Bénabou & Tirole, 2010).

Outside of a firm centric context, Professor Orhun at the University of Michigan investigated this aforementioned relationship between altruistically and strategically motivated socially beneficial actions. Her experimental research demonstrated that when a first mover is likely motivated by strategic incentives rather than altruistic motives, the positive reciprocity triggered by the same beneficial action is lower. Within reciprocal interactions, the welfare gained depends both on the action and the degree of reciprocity that it triggers from the other party. Although this study focused on individuals in a two-stage reciprocity game, the findings can be applied to profit-maximizing firms who aim to make the most of the reciprocity they gain from individual customers and workers. Overall, this study highlighted that, “stronger incentives for beneficial behavior may not increase total welfare” (Orhun, 2018). In a similar vein, another study from Cassar & Meier acknowledged the ways in which CSR initiatives are able to motivate workers, specifically aiming to understand how the perceived intentions of the initiatives impacted workers effort levels. This case brought to light that charitable incentives that are instrumental to the firm backfire compared to non-instrumental incentives (Cassar & Meier, 2017). Coupled together, these two perspectives demonstrate that the prosocial behavior of individuals may depend also on the relatedness of the motivations behind socially beneficial actions. However, what they fail to address is whether or not the effectiveness of these actions also creates a material impact on individual behavior.

CSR's Relationship on Consumers' Willingness to Pay

Additionally, there has been some research in the field with regards to consumers' willingness to pay for products with strategically focused CSR. For example, mainstream examples of "greening" a company's supply chain or creating more ethically sourced products can have positive effects on consumers. Gregory Guagnano demonstrated this relationship in his study on willingness to pay for recycled paper products, using the Schwartz model. In this model, the relative influence of altruism and personal costs were weighed to better understand the contrast between altruistic behavior and self-interested behavior. In this case, Guagnano chose to measure pro-environmental behavior because it forced a trade-off between individual and collective benefits. Guagnano noted that many economists expect people to opt towards rational choice, especially in the market where self-interest typically dominates, however this study illustrated that market behavior is also often motivated by altruism. In fact, "over 86% of the respondents in the current analysis said they were willing to pay extra for a common household good made from recycled materials" (Guagnano, 2001).

Research by Lusk, Nilsson, and Foster (2007) took these insights one step further to understand the aforementioned trend of private goods that are affiliated with public good attributes. More specifically, they also critiqued the notion that consumers have purely selfish preferences and used psychometric scaling techniques to measure individuals' sense of altruism, as well as their propensity towards free riding. To do this, they employed an example-choice experiment focused on various brands of pork products (Appendix 4). Pork products were chosen because they typically do not carry many quality signals and they are often viewed as a homogenous commodity. Similarly to the findings of Guagnano, this study found that a consumer's choice to purchase a private good with public-good attributes like sustainably farmed

pork was, “not simply a result of individuals’ perceptions of the ability to mitigate private risks such as food safety, but that individuals are making private choices to affect public outcomes” (Lusk, Nilsson, & Foster, 2007). This conclusion has large implications in the policy sphere, since many legislators are forced to make tradeoffs between food labeling policies and bans related to the actual production processes for livestock products. This study would argue that private market incentives like labeling can actually help the public good without the need for government intervention. This thesis serves to dive deeper into these findings to better understand the relationship that the effectiveness of these private market incentives has with consumers’ willingness to pay.

Effective Altruism Literature:

Outside of the scope of the CSR literature above, there is another equally important body of work within the spheres of charities and non-profits with regards to effective giving. Overall, the literature examining the role of effective altruism within charities has demonstrated that individuals often fail to make optimally effective decisions because they are jaded by their own behavioral heuristics and understandings of charitable giving as a rightfully, subjective world. Furthermore, there is often pushback when corporates take an altruistic perspective because of the shareholder primacy model. Regardless of these moral qualms, altruism currently exists within the modern business world, and it is essential to understand it further.

Effective Altruism and Personal Agency

Normative models of altruism, like those described in the introduction from scholars like Peter Singer, demonstrate that individuals should opt to allocate their resources to the choices that create the most total welfare. However, in practice many people appear to be “distorted altruists” meaning that, “they care about welfare maximization, but without clear information to

make comparisons, they rely on their feelings to guide choice” (Loewenstein & Small, 2007). In a study from the Association for Psychological Science, impediments to effective altruism like these were explored and highlighted that individuals view charities as being largely subjective entirely. However, the findings in subsequent studies that are most relevant to this research were that 1) “people are more likely to override welfare maximization when choosing a charity than when choosing a financial investment” and 2) “individuals are less likely to license themselves and others to select an ineffective option when a decision maker assumes a position of responsibility” (Berzman, et al., 2018). The implications for these final two findings are critical, as they hint at the possibility for variant results of these findings if they were to be re-produced within the private sector, where consumers are the decision makers who make a financial investment as they consume, which is precisely what this paper hopes to explore.

Effective Altruism’s Role in Business: An Ethical Perspective

Although the concept of altruistic CSR may seem appealing upfront, there are some scholars who disagree with the idea from a moral standpoint. This argument is grounded in the perspective that public corporations owe a primary responsibility to shareholders. When evaluating CSR from a multifaceted perspective of utilitarianism rather than the firm’s financial position in isolation, these scholars argue that altruistic CSR unjustly seizes stockholder wealth in order to benefit the general welfare (Lantos, 2002). More specifically, the research cites “justice theory” and the notion that stockholders’ earnings, which are earned at their own risk, should not be taken away unless these groups are explicitly willing to sacrifice them for the cause. This viewpoint, however, does not claim that all CSR is immoral. In fact, it argues that *strategic* CSR is not only moral, but also commendable. Since strategic CSR has the potential to benefit both shareholders and the public, it does not violate the justice theory and it also creates

more public welfare, making it quite commendable. The research presented in this study aims to discover if effective CSR can still be categorized as strategic by increasing consumers' willingness to pay. If this potential increase in willingness to pay is enough to outweigh the costs of the newly implemented effective CSR, it may not seize shareholder value. This paper serves to explore if effective CSR may also be as commendable as strategic CSR from this perspective.

In stark contrast to this perspective, Luo and Kaul argue that it is critical to move past the siloed concept of, "does a firm benefit financially from being socially responsible?" and instead consider that there may be conditions (like when there is a need to develop novel solutions for social problems) that firms are comparatively efficient in improving total welfare. In their work, they aim to take a more holistic view of various change-making organizations to understand actions that enhance the most welfare (Appendix 5). The scope of this research will underscore this broader understanding of welfare efficiency, rather than one that sticks strictly to the shareholder primacy model.

Regardless of the moral grounding behind altruism in business, the reality is that some businesses do act altruistically. In a study based on thematic interviews with small business owners, it was shown that, "some philanthropic decisions are based on mere willingness to contribute to the welfare of others" (Lähdesmäki & Takala, 2012). These findings suggest that there is room for the existence of altruism within small businesses. The scope of this paper will expand upon these findings to understand, from a consumer perspective, if there is space for altruism in larger, for-profit corporations as well.

Theoretical Framework

Critical Assumptions

One of the theoretical studies grounding this research stems from Luo and Kaul's research titled, "An economic case for CSR: The comparative efficiency of for-profit firms in meeting consumer demand for social goods," but there are a few key assumptions that will be altered. In their model, they look at the comparative efficiency of for-profits creating social goods relative to non-profits. The biggest pitfall of this research is that it fails to account for the *effectiveness* of these social goods, which is far different than efficiency. Again, a firm like TOMS may be "efficient" at producing shoes and donating them to the third world, but in reality, it may do "more good" for TOMS (and society at large) to instead donate to an effective healthcare charity that aims to save lives lost from diarrhea or other bacterial infections in children, for example.

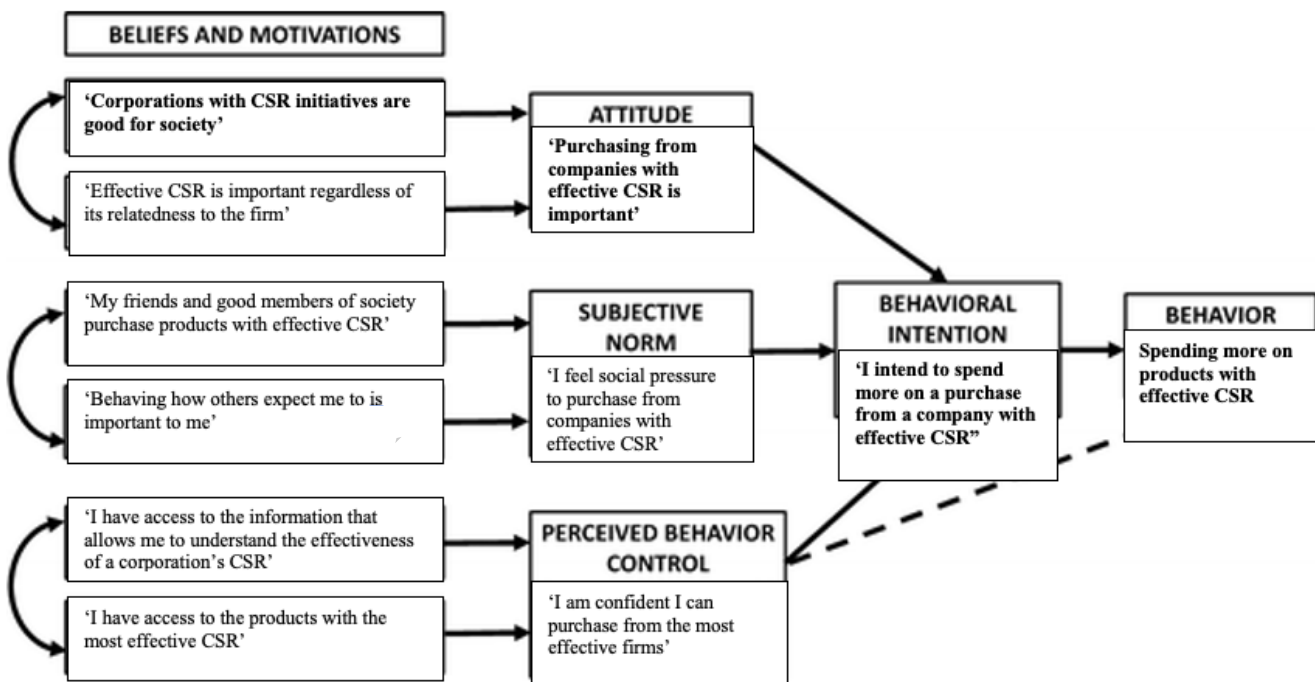
The main theoretical assumption that this paper aims to challenge from Luo, Kaul, and similar scholars is the distinction between efficiency and effectiveness. Just because an organization is efficient at completing a certain activity (achieving low manufacturing costs, for example) does not mean that it should make these activities a part of their CSR portfolio. Many of the world's largest multinational corporations serve relatively rich populations, which makes it incredibly unlikely that they have a competitive advantage in a value chain activity that directly benefits those who need it most (for example, it is very unlikely that a company producing electronics for US consumers has efficient distribution capabilities for highly effective vaccines in Ethiopia). There are a few other caveats to note within this assumption as well:

1. Updated Assumption: Efficiency vs. Effectiveness
 - a. Challenge: Organizations are similarly effective at "doing good"

- i. While the distinction between “symbolic” and “substantive” CSR brought to light by Luo and Kaul is an important one, the dichotomy should not end here. Delineating impact among these two factors alone discredits the notion that within the world of “substantive” initiatives, impacts are extremely variant. Various charities and for-profits firms are not equally effective in creating impact. In fact, most donors underestimate the differences in charities’ effectiveness (Caviola et al., 2020). Taking the effective altruistic view described earlier, this paper views impact as the ability to create the “most good” with a given set of resources.
- b. Challenge: Consumer Preferences
 - i. Luo and Kaul also assume that if a for-profit corporation is more efficient at providing a “social good” they should do the job rather than a non-profit. However, this logic assumes that consumers view the consumption of a “social good” equally whether it comes from a for-profit or non-profit corporation. This fails to account for the subjective preferences that individuals are prone to during decision-making processes. However, as mentioned in the Berzman, et al. piece, people tend to override welfare maximization more often when viewing charities than financial investments. This research aims to better understand this assumption by measuring the effect of consumers’ willingness to pay for these social goods.

Theory of Planned Behavior

To illustrate these challenged assumptions more clearly, the theory of planned behavior (TPB) model can be leveraged to demonstrate the hypothesis of this paper (Appendix 6 shows the generic framework). Below is the framework for the purpose of this research with the key beliefs bolded:



The graphic above shows several hypothesized beliefs and motivations that may lead to the desired behavior of increased willingness to pay for products with effective CSR via the TPB model. Each of these beliefs and motivations are grounded in the assumptions mentioned earlier. The beliefs and motivations above are not exhaustive, but they are illustrative of the types of motivations behind consumers' decisions to act effectively when they make a purchase. The goal

of this research is to understand if the bolded beliefs and motivations, with regards to the importance of effectiveness, compel consumers to spend more on products with effective CSR.

Proposed Methodology

The pilot study will take an experimental approach to gather empirical data that will be analyzed to determine the effects that CSR's effectiveness has on consumers' willingness to pay. For the focus of this study, effectiveness will be judged in terms of global poverty reduction by using the number of lives saved as the key metric of success, following the previously mentioned CEA, QALY, and CBA frameworks. To determine the impact of effectiveness, participants will receive the price of a product for a firm with no CSR initiatives and the effectiveness of one donation for the firm with CSR initiatives. The effectiveness of the CSR will be randomly assigned for each respondent. Furthermore, the context for the pilot study will be a pharmaceutical company with a lifesaving drug product because it is very simple for consumers to directly understand like the example below:

You are a customer picking up pain reliever at the store and the following claims regarding corporate social responsibility (CSR) are shown on the side of each product's label. Utilize the scale to describe how much are you willing to pay for the product in the scenario (in \$).

Keep in mind that the price for a generic skincare drug from a pharmaceutical firm with no current CSR initiatives, like Company X is \$10.

Company X: Price is \$10

Company Y: Uses a portion of its profits to donate and manufacture a drug to treat a deadly skin disease in the world's poorest populations.

Each dose of Company Y's donated medicine has an estimated X% chance of saving one life.
How much are you willing to spend on Company Y's pain reliever?

Although there are real world examples of specific drugs that could be used for this experiment, research has shown that within the realm of charity, "people frequently choose less effective charity options when those options represent more subjectively preferred options"

(Berman, et al. 2018). Assuming that this claim may be applied to the realm of CSR as well, a generic drug will be chosen for the survey description to eliminate potential bias. Participants will be randomly given a scenario like the one above and from there, they will be asked to select a dollar amount from a sliding scale with a set control price of \$10 as the median, which will be the anchoring point for a firm with no current CSR. To control for individuals' preferences and perspectives on certain goods, the treatment will be randomized. Although this research example focuses on a pharmaceutical company which offers a socially responsible good that saves lives, there are also implications for companies that offer less socially relevant products. If companies with less socially necessary products like TOMS reconsider the effectiveness of their CSR for products like cloth shoes, they may be able to create "more good" overall.

Logistics

The experiment will be run on Amazon's MTurk portal with a large, randomized sample to account for individual preferences. There will not be many external constraints aside from this randomization because a large enough data set should be able to account for confounding variables, while remaining representative of all types of consumers. Initially, a pilot survey will be run to test the proposed methodology above. Depending on the results of the pilot study, some alterations may be made before running the full survey. After the data is collected from the sample, it will be manipulated in Excel to understand the broad relationship between consumers' willingness to pay and the effectiveness of a firm's CSR with regression analyses. After the main analysis is complete, further post-hoc analysis can be completed to see if there are differing relationships based on participants socio-economic status, gender, morality, or other demographic factors.

Pilot Study

Methods

An experiment was constructed in which participants' willingness to pay for a consumer good was measured under different conditions of the effectiveness of the parent company's CSR initiatives. The manipulation question involved two different companies (X and Y), one that had CSR initiatives and one that did not. Participants were then asked to determine their willingness to pay for the good that came from the company (Y) with CSR initiatives, but respondents were presented with randomized levels of that CSR's effectiveness.

To ensure proper comprehension of the survey, respondents were asked a screening question later in the survey where they had to remember a specific detail from the manipulation question. Participants that answered incorrectly were removed from the rest of the survey. After this comprehension question, a series of demographic-focused questions were asked, as well as questions to test the moral values of participants using the Oxford Utilitarian Scale (OUS). The OUS scale is able to analyze participants' moral stances on both instances of impartial beneficence and instrumental harm to others by asking a series of moral statements and asking respondents to describe their alignment to those statements (Kahane, et al., 2018).

50 participants who passed the comprehension test completed the survey for the pilot study. 30 of these participants were male and 20 were female. Additionally, 42 of the respondents identified as Caucasian. The survey was restricted to participants who took the survey within the US and were above 18 years old.

Respondents accessed the survey through MTurk. At the beginning of the survey, they read an informed consent notice and were required to consent to the survey in order to advance to the next section (Appendix 7). Additionally, respondents had to click through a reCAPTCHA

box to ensure that they were real, human participants. To test the hypothesis that consumers are willing to pay more for goods from companies with effective CSR, all respondents were given the same initial prompt that described two pharmaceutical companies. It read that, “*Company X provides a generic pain reliever at a price of \$10, while Company Y provides a generic pain reliever for an undisclosed price and also uses a portion of its profits to manufacture and donate a drug to treat a deadly skin disease in the world’s poorest populations.*” In this scenario, Company X served as the control for a firm without any current CSR initiatives, setting the anchoring price at \$10. After reading about these two companies, respondents were randomly assigned various levels of Company Y’s CSR effectiveness through a prompt that read, “*Each dose of Company Y’s donated medicine has an estimated X% chance of saving one life. How much are you willing to spend on Company Y’s pain reliever?*” The range of variables was coded on the backend to reveal a random integer between 0-25% . Participants were then forced to decide how much they were willing to spend on Company Y’s pain reliever by using a sliding scale (Appendix 8).

Results and Discussion

Given the nature of the small sample size of the pilot, the goal was to test the mechanics of the survey and garner a general sense of consumers’ sentiments towards the effectiveness of CSR. Therefore, the analyses completed are mainly to refine the study and understand the overall correlation between effectiveness and willingness to pay.

First, the data was cleaned to remove participants that did not successfully pass the screening question. From there, a regression was run to analyze the relationship between effectiveness and willingness to pay. The p-value for this regression was quite high at 0.69, making it insignificant, as it is greater than 0.05. The regression also showed that for every 1%

increase in the effectiveness of lives saved, willingness to pay actually *decreased* by 2 cents (Appendix 9). Since there was such a small p-value for the key regression, it was determined that additional analyses into the relationship between this regression and the other demographic information, as well as the OUS scale were not needed at this point. With alterations to the pilot in the next iteration of the survey these analyses will be completed. The results of this pilot study prompted further examination into why the hypothesized effect was incorrect. After ample discussion, it was concluded that the two most likely reasons were that 1) participants misunderstood the manipulation/the survey was unclear or 2) consumers' willingness to pay genuinely is not affected by the effectiveness of CSR.

To address the first concern, adjustments were made to the wording of the manipulation and the screening question was altered to determine if participants understood the manipulation completely. First, the quantity of profits donated by the firm was added into the survey as a new variable which will be called "generosity." After analyzing the data, it was determined that some participants may have viewed the measure of effectiveness in isolation as synonymous with the overall altruism or generosity of the firm. However, this research aims to get at the core argument of effectiveness, in terms of number of lives saved. This research does not want to become confused with the warm glow of giving or generosity of a firm abstractly. To differentiate these two ideas, it was essential to add enough additional context to participants to be able to translate overall generosity and effectiveness into number of lives saved on the backend. To do so, a new initial prompt was created for participants that read:

Imagine you are a customer considering purchasing one of the two pain relievers below. The chemical formula, number of pills, and dosage of each bottle are identical. Both companies earn \$1 million in profits annually. The only difference is that Company Y contributes a portion of its profits to manufacture and donate a drug to treat a deadly skin disease in the world's poorest populations. This donated drug costs \$100 to manufacture.

After this message, participants would receive randomly assigned levels of both effectiveness and generosity which fell into one of the main categories below:

Generosity <i>Low → High</i>	Effectiveness <i>Low → High</i>	
	Percentage of Company Y's profits donated: 13-25% Chance of one dose of the donated medicine saving a life: 0-12%	Percentage of Company Y's profits donated: 13-25% Chance of one dose of the donated medicine saving a life: 13-25%
	Percentage of Company Y's profits donated: 0-12% Chance of one dose of the donated medicine saving a life: 0-12%	Percentage of Company Y's profits donated: 0-12% Chance of one dose of the donated medicine saving a life: 13-25%

Now participants have received 1) the total amount of profits 2) the generosity of those profits 3) the cost of each dose donated 4) the effectiveness of those doses to save one life. With these four key pieces of data, it is possible to determine the proposed number of lives saved for each manipulation precisely. Although respondents were not expected to complete this math themselves, it is helpful to understand the exact number of lives saved because it isolates effectiveness and takes out the chance that respondents assumed a certain level of firm generosity. Future studies could consider sharing the total number of lives saved upfront, however, this study chose to only reveal the percent of effectiveness explicitly to respondents since this is a commonly used marketing tactic.

To address the aforementioned second concern that consumers' willingness to pay genuinely is not affected by the effectiveness of CSR, the next version of the survey aimed to be run with a much larger sample size to hopefully become more statistically significant. It is very possible that consumers truly do not care about the effectiveness of CSR, which would also be a valuable insight for future studies.

Primary Study

Methods

Based on the results from the pilot survey, the survey was altered to include generosity as a variable, additional introductory information of total profit size, and the cost of a donated dose as described above (Appendix 10). Additionally, one pilot survey respondent mentioned difficulty setting the profit slider to exactly \$10, so alterations were made to make the user experience smoother. As mentioned above, participants were controlled to receive one of the four manipulations:

Generosity <i>Low → High</i>	Effectiveness <i>Low → High</i>	
	Less Effective and Highly Generous	Highly Effective and Highly Generous
	Less Effective and Less Generous	Highly Effective and Less Generous

Participants were equally placed into each of these manipulations at 100 respondents each. Since this new launch included a new variable, the two variables of generosity and effectiveness were coded to appear in a random order as well; roughly half of the participants saw the effectiveness of the dose appear before the firm generosity and the other half saw the reverse. This randomization helped account for any effect that the ordering of the variable presentation may have had on respondents.

400 respondents consented to this study and passed the comprehension screening question. The participants were 54.25% male (217 respondents), 44.5% female (178

respondents), and 1.25% (5 respondents) were undisclosed. The survey was again restricted to adults over the age of 18 who were currently located in the United States. The average age of all participants was 41.89 years old and the median age was 39 years old. The youngest participant was 18 years old and the oldest participant was 79 years old. 309 of the 400 participants (77.25%) identified as Caucasian, 37 (9.25%) identified as Asian, 22 (5.5%) identified as Hispanic or Latino, 21 (5.25%) identified as African American, and 11 (2.75%) chose not to disclose or identified with two or more identities. In terms of respondents' highest level of education, 215 respondents (53.75%) have earned a bachelor's degree, 116 respondents (29%) have earned a high school degree, 58 respondents (14.5%) have earned a master's degree, 7 respondents (1.75%) have earned a PhD or higher, and 4 respondents (1%) chose not to disclose or had completed some high school. 173 respondents (43.25%) were married and 225 (56.25%) were not. Additionally, 228 respondents (57%) did not have children and 168 respondents (42%) had at least one child. 55 respondents (13.75%) earned less than \$25,000 in terms of household income, 105 respondents (26.25%) earned between \$25,000-\$50,000, 138 respondents (34.5%) earned between \$50,000-\$99,999, 53 respondents (13.25%) earned \$100,000-\$149,999, 18 respondents (4.5%) earned \$150,000-\$199,999, and 20 respondents (5%) earned over \$200,000 annually.

Aside from these demographic based questions, respondents were also asked questions regarding their personal charitable giving, religion, and political views. Using the Oxford Utilitarian Scale (OUS), participants were asked nine questions about their agreement with questions regarding their morality (Appendix 11). The OUS scale is able to, “dissociate individual differences in the ‘negative’ (permissive attitude toward instrumental harm) and ‘positive’ (impartial concern for the greater good) dimensions of utilitarian thinking as

manifested in the general population,” which other common types of morality judgments, like the popular “trolley problem”, are unable to (Kahane, 2018). Similar to the analyses above, the OUS responses were assigned numerical values on a scale of 1-7, with scores of 7 representing the most utilitarian response. The nine numbers each participant received were then added together to create an overall utilitarian score.

Results

The primary analysis of this study was to understand if consumers generally altered their willingness to pay for products based on the effectiveness of a firm’s CSR. To determine this effect, a “number of lives saved” metric was calculated for each participant. To do this, the number of doses donated was calculated by multiplying together each participant’s given random percent of profits donated with the total amount of firm profits (\$1 million). This was then divided by the cost of each dose (\$100) to get the total number of doses. After this was calculated, the number of doses was multiplied again by each participant’s randomly assigned level of effectiveness to get a projected number of lives saved for each respondent. From there, a regression was run to test the independent variable of willingness to pay against the dependent variable of number of lives saved.

The results of this main regression were statistically insignificant (Appendix 12). With a p-value of 0.73 and coefficient of -0.00, effectiveness of CSR initiatives in terms of the number of lives saved did not appear to have a relationship with consumers’ WTP

$$\text{[Equation: WTP} = -11.23 + -0.00 \cdot \text{effectiveness]}$$

This study failed to reject the null hypothesis that consumers do not respond to increased effectiveness in firms’ CSR.

Next, a post-hoc analysis was completed to determine if there were any statistically significant relationships between willingness to pay, number of lives saved, and certain consumer demographics and characteristics. Prior to this study, there were no hypotheses developed for this analysis, so the results below should all be considered post-hoc.

Dependent Variable: Willingness to Pay

Independent Variable	Coefficient (Standard Error)
Generosity N = 400	
Generosity	-5.68 (3.321)
Generosity X ¹	0.043* (0.018)
Likelihood of Saving a Life N = 400	
Likelihood of Saving a Life	2.453 (3.234)
Age N = 400	
Age	-0.020 (0.017)
Age X	4.70E-05 (8.057E-05)
Gender N = 395	
Female	0.320 (0.465)
Female X	0 (0)
Male	0 (0)
Male X	0.000 (0.002)
Oxford Utilitarian Scale N = 400	
OUS	0.043 (0.024)
OUS X	1.64E-05 (0.000)
Race N = 398	
Caucasian	-2.142 (1.134)
Caucasian X	0.007 0.006

Asian	-1.686 (1.322)
Asian X	0.006 (0.007)
Latino/Hispanic	-1.655 (1.468)
Latino/Hispanic X	0.008 (0.008)
African-American	0 (0)
African-American X	0 (0)
Two or More Races	-1.344 (2.385)
Two or More Races Interaction	0.001 (0.013)
Employment Status N = 384	
Full Time Employment	1.154 (0.951)
Full Time Employment X	-0.003 (0.007)
Seeking Employment	1.401 (1.278)
Seeking Employment X	8.567E-05 (0.007)
Student	0.748 (2.143)
Student X	0 (0)
Retired	0 (0)
Retired X	-6.559E-05 (0.008)
Part Time	1.241 (1.080)
Part Time X	-0.003 (0.007)
Religion N = 331	
Buddhism	0.0802 (0.981)
Buddhism X	-0.001 (0.0156)
Judaism	2.252 (3.318)
Judaism X	-0.002 (0.014)
Hinduism	6.518 (5.289)
Hinduism X	-0.056 (0.0341)
Islam and Islam X	0 (0)

Catholicism or Christianity	0.986 (3.035)
Catholicism or Christianity X	0.002 (0.013)
Other	0.523 (3.052)
Other X	0.002 (0.013)
Birthplace² N = 397	
North America	0.688 (0.819)
North America X	-0.007 (0.014)
Europe	0.227 (4.230)
Europe X	-0.009 (0.018)
Asia	1.091 (3.351)
Asia X	-0.007 (0.017)
Highest Education Status² N = 398	
High School	0 (0) [p-value unable to be estimated]
High School X	0 (0) [p-value unable to be estimated]
Bachelor's	-1.142 (0.515) [p-value unable to be estimated]
Bachelor's X	0.004 (0.002) [p-value unable to be estimated]
Master's	-0.428 (0.713)
Master's X	0.002 (0.003)
Political Views N = 384	
Very Liberal	0.402 (0.736)
Very Liberal X	0.001 (0.002)
Slightly Liberal	-0.314 (0.672)
Slightly Liberal X	0.003 (0.002)
Slightly Conservative	0.250 (0.730)
Slightly Conservative X	0 (0)

Very Conservative	0 (0)
Very Conservative X	-0.001 (0.003)
Number of Children N = 396	
Number of Children	-0.014 (0.001)
Number of Children X	0.000 (0.001)
Percent of Income Donated to Charities Annually N = 383	
0%	-3.141 (1.869)
0% X	0.007 (0.008)
1-2.9%	-1.799 (1.823)
1-2.9% X	0.006 (0.007)
3-4.9%	-1.990 (1.950)
3-4.9% X	0.007 (0.008)
5-6.9%	-1.394 (2.029)
5-6.9% X	0.003 (0.008)
7-9.9%	-3.001 (2.105)
7-9.9% X	0.013 (0.009)
Greater than 10%	0 (0)
Greater than 10% X	0 (0)
Annual Household Income N = 389	
Less than \$25,000	-1.095 (1.355)
Less than \$25,000 X	0.014 (0.008)
\$25,000-\$49,999	0.715 (1.295)
\$25,000-\$49,999 X	0.014 (0.008)
\$50,000-\$99,999	0.313 (1.265)
\$50,000-\$99,999 X	0.124 (0.007)
\$100,000-\$149,999	-0.711 (1.344)
\$100,000-\$149,999 X	0.018* (0.008)

\$150,000-\$199,999	0 (0)
\$150,000- \$199,999	0.143 (0.008)
\$200,000+	0.842 (1.517)
\$200,000+ X	0 (0)
Marriage Status N = 398	
Married	0 (0)
Married X	0 (0)
Unmarried	0.0434 (0.466)
Unmarried X	-0.000 (0.002)
Likelihood of Telling Friends About the Purchase N = 400	
Very Likely	4.817*** (0.866)
Very Likely X	0 (0)
Somewhat Likely	2.576*** (0.558)
Somewhat Likely X	0.004 (0.003)
Somewhat Unlikely	1.367* (0.585)
Somewhat Unlikely X	0.004 (0.003)
Very Unlikely	0 (0)
Very Unlikely X	0.005 (0.004)
Moral Alignment with Company Y N = 400	
Very Aligned	-1.274 (0.674)
Very Aligned X	0.004 (0.003)
Somewhat Aligned	-1.418* (0.644)
Somewhat Aligned X	0.002 (0.003)
Not Very Aligned	-1.076 (1.428)
Not Very Aligned X	0.004 (0.008)
Not Aligned at All	-2.683 (1.802)
Not Aligned at All X	0.006 (0.009)

Excitement About the Purchase	
N = 400	
Excitement About Purchase	0.495*** (0.078)
Excitement About Purchase X	-0.000 (0.000)

[Notes: Standard errors are reported in parentheses

, **, * indicates significance at the <.05, <.01 and <.001 level respectively*

X after the variable title indicates the variable was run in an interaction with the number of lives saved.

1 – Number of lives saved is a function of generosity. The generosity X interaction between generosity and number of lives saved should be interpreted with caution due to this.

2 – There was not a large enough sample size of responses for the birthplaces of Australia, Caribbean Islands, South America, and Central America, as well as the education levels of some high school and PhD so the regression gave responses of 0

The sample sizes vary to account for responses that were removed when participants selected “prefer not to respond.” Additionally, number of children and excitement about purchase variables do not have any interactions listed, since the variables were assumed to be continuous and linear.]

As seen above, a plethora of post-hoc analyses were completed on the data. To complete these regressions, the data was analyzed in Excel using SWITCH functions to change the qualitative survey responses into numerical responses. Additionally, participants that responded “prefer not to say” were withdrawn from that specific analysis.

Statistically Insignificant Results

In isolation, the likelihood of saving a life was not significantly correlated with consumers’ willingness to pay. When filtering the data for strictly male responses and separately strictly female responses against the number of lives saved and willingness to pay, the results were again statistically insignificant. Similarly, age, marriage status, and number of children did not have a statistically significant effect on willingness to pay.

To determine the effects of political affiliation, percent of income donated, race, OUS, religion, and birthplace each categorical variable was transformed to dummy variables. For example, “very liberal” was defined as 1 if the participant self-identified as “very liberal” and zero otherwise. This method was followed for each possible response. Once each of these categories were transformed to categorical variables, they were interacted with the corresponding number of lives saved for each respondent to create interactive variables as well. From there, all of those various x-variables were run against the willingness to pay each participant recorded for themselves. No statistically significant main effects of interactive effects were found for political affiliation, percent of income donated, race, OUS, religion, education level, or birthplace.

Statistically Significant Results [Appendices 13 and 14]

A. Generosity

When the interactive variable of generosity (i.e. percentage of corporate funds donated) was tested against willingness to pay, the results showed a positive and statistically significant coefficient (beta = 0.043, p-value = 0.018). This means that as a company increases its CSR by one percent, people increased their willingness to pay by \$0.043, which is about 4 percent over the given price of \$10 for a product with no CSR. However, consumers did not similarly respond to increases in the effectiveness of those dollars. In fact, the impact of effectiveness when tested was small, negative, and statistically insignificant. The results suggest that consumers increase their willingness to pay as firms become more generous, but do not respond to increases in the effectiveness of the firm’s generosity. It should be noted that the number of lives saved is a function of generosity, so the statistically significant outcome of the generosity interaction variable discussed here should be viewed with caution and isolated thoroughly in following studies.

B. Household Income

When annual household income was tested against the willingness to pay, the only income level that significantly related to WTP was the interactive variable at the \$100,000-\$149,999 level. This means that those who earn \$100,000-\$149,999 annually are more responsive to the number of lives saved at a statistically significant level, relative to wealthier respondents.

However, this significant finding may be spurious since there is no theoretical reason why this income group should respond more positively than people in lower and higher income groups.

C. Excitement, Moral Alignment, and Likelihood of Sharing with Friends

The three other questions that were statistically significant revolved around respondents' excitement about the proposed product, their moral alignment with the firm, and the likelihood that they would share their purchase with friends (Appendix 15). Those who responded to the survey that they were, "very likely," "somewhat likely," or "somewhat unlikely" to tell their friends about the purchase responded more positively to the number of lives saved. In terms of moral alignment with the firm, only those who responded "somewhat aligned" showed statistical significance at the .05 level. This indicates that those who are somewhat aligned with the firm's morals respond more positively to the number of lives saved or the effectiveness of CSR.

Finally, on a scale of 0-10 respondents were asked to rank their excitement about their hypothetical purchase. Assuming a continuous and linear model, those who self-reported higher excitedness about the proposed purchase demonstrated a positive response to the number of lives saved at the 0.001 level.

Although these results were statistically significant, upon further discussion it was noted that the questions regarding excitement and sharing the purchase with friends are simultaneously determined with WTP; they may be viewed as intermediate outcomes and should be interpreted

with caution. To account for this, the responses to these two questions were also run as dependent variables with the number of lives saved run as the independent variable, as seen below. Results from this regression would determine if respondents may be more excited or likely to tell their friends about the purchase in response to the number of lives saved. The additional tests found below were statistically insignificant.

Dependent Variable: Likelihood of Telling Friends

Independent Variable: Number of Lives Saved
N = 400
-0.000 (0.000)

Dependent Variable: Excitement About Purchase

Independent Variable: Number of Lives Saved
N = 400
-0.000 (0.000)

Discussion

The purpose of this study was to understand if consumers broadly are willing to pay more for goods with effective CSR, based on the number of lives saved from that CSR. This study failed to reject the null hypothesis that consumers do not respond to increased effectiveness in firms' CSR. This was demonstrated by running a regression of each respondents' willingness to pay against the hypothetical number of lives saved that was presented during their simulation.

Despite the lack of support for the relevance of effectiveness in consumers' willingness to pay overall, there were some more niche insights that were gleaned from the post-hoc analysis. First, the reported generosity of the firm was significantly correlated with respondents' willingness to pay. This insight coupled with the insignificant results of effectiveness on consumer's willingness to pay implies that firms may not be detailing the effectiveness of their CSR initiatives with good reason. Consumers in this sample were not responsive to the effectiveness of CSR initiatives, but were responsive to changes in overall firm generosity in terms of the percent of dollars donated, which is not necessarily optimal for human welfare. However, the number of lives saved variable was a function of generosity in this study, so the findings from this interaction should be interpreted with caution. If this finding were to be found significant in subsequent studies, there would be vast implications; consumers' lack of responsiveness to the effectiveness of CSR initiatives empowers corporations to continue to focus only on their generosity (or percent of funds donated) when they market their CSR. Actions like these have the capacity to perpetuate "greenwashing" and the advancement of CSR initiatives that heavily prioritize firm profits over general welfare. There are a variety of social trends that may describe why consumers feel this way, and further research should expand upon these sentiments. If information problems are hypothesized to be a main driver of these

consumer sentiments for example, it may be advantageous to pursue interventions that teach consumers the value of effective charities in order to improve human welfare.

Respondents' income within the \$100,000-\$149,999 range was also correlated with their willingness to pay, relative to wealthier respondents. This could imply that those with more wealth begin to care less about the effectiveness of CSR dollars. However, given the large number of specifications and the post-hoc nature of the analysis, these results may be spurious and require further studies to either confirm or deny this significance. Additionally, consumers self-reported excitement about the product, their moral alignment with the product, and their likelihood of sharing the product with friends were also statistically significant, but may be intermediate outcomes. When tested as dependent variables, they were also insignificant.

Limitations

This study investigated the interactions between consumers' willingness to pay and the effectiveness of corporations CSR initiatives. One of the limitations with this research is that purchase intentions for consumers in real life “not only depend on the composite product and a budget constraint, but also on alternative product offerings, so-called reference products” (Breibert, 2006). This study only compared two different companies (X & Y), which do not fully represent the array of products or choices a typical consumer would have in an actual store. Furthermore, participants are not spending real money or taking any real possession of the “purchased” goods. This artificial set up could lead to differing results in the real world.

Another limitation to this research is the overall generalizability of the results. In this study, a pharmaceutical company was used as the example company. However, pharmaceutical companies have very simple value propositions that relate directly to the metric of number of lives saved. These results should not be blindly generalized towards all for-profit corporations because consumers will likely view other industries with a different perspective. Additionally, the results should not be generalized towards all firms even within the pharmaceutical industry. This is because consumer sentiment is multi-faceted, and it is likely that consumers have different sentiments regarding certain, real-life pharmaceutical companies. For example, a consumer may view a pharmaceutical company like Purdue Pharma that had ties to the opioid epidemic negatively and would be willing to pay less for their products, regardless of the effectiveness of the firm's CSR.

Additionally, a potential limitation to this research is the population of the respondents on MTurk. It is hard to say if the sample population is representative of all consumers. More specifically, it may not capture those who donate the most to charities or high net worth

individuals who do not participate in the MTurk portal. Although questions were crafted to understand the demographics of those participating in the survey, it is difficult to know if this population is representative of consumers at large. Additionally, certain firms likely have a different consumer base than others. For example, a luxury good company likely has a vastly different customer than an everyday low-price grocer, leading to differing results for each hypothesized firm.

The sample size of this study was 400 participants which may cause another limitation. In analyses like these it is important to garner as many responses as possible from a diverse array of respondents. With a larger sample size, perhaps the results could highlight additional insights.

Considerations for Future Research

This study's research and resulting discussion can contribute to the fields of CSR and effective altruism, respectively. However, the primary results of this particular study were not statistically significant. Future research should craft hypotheses regarding the statistically significant components of the post-hoc analysis and run various tasks to better understand their relationship with consumers' willingness to pay. Additionally, future studies with larger budgets should consider employing a larger sample size, as well as manipulations that span various industries. Future studies could also consider leveraging real, in-person simulations to understand consumers' reactions to real, tangible packaged goods. This design may pose challenges considering the uncontrollable biases consumers may have with pre-existing brands and products, however the insights could still prove to be useful.

The results of this research could lead to two possible conclusions 1) that consumers truly do not care about the effectiveness of CSR initiatives with regards to willingness to pay, or 2) this specific research design was underpowered, flawed, or misconstrued by participants. Considering the importance of the research in this field, it would be worthwhile for future researchers to continue to explore the field and build upon this research. As alluded to earlier, future studies should consider interventions that inform and indoctrinate consumers with the literature of the effective altruism community to see if there is an effect on their willingness to pay. Effectiveness is a crucial measure when discussing human welfare and saving human lives. If the first conclusion is true that consumers do not care about the effectiveness of CSR, then perhaps the discussion should be broadened to consider *if* corporate social responsibility should be pursued rather than *how* it should be pursued.

Conclusion

The aim of this study was to test the hypothesis that consumers are willing to pay more for goods from firms with effective CSR, filling a gap in the existing literature. This was tested through an experimental model that forced tradeoffs between two variables: the generosity and the effectiveness of that good/brand's CSR initiatives through a simulation regarding a hypothetical pharmaceutical company. These two variables were then translated into a measure of effectiveness in terms of the number of lives saved and compared to consumers' willingness to pay. There was no significant evidence to demonstrate that the effectiveness of a firm's CSR initiatives has an effect on consumers' willingness to pay. However, post-hoc analysis revealed statistically significant relationships between a variety of niche consumer segments. This study serves as a helpful springboard for future research on the relationship between morality, willingness to pay, and effectiveness of corporate social responsibility initiatives. If corporations and consumers alike understand the benefit of allocating CSR funds to the most effective initiatives, there is the potential for a greater impact on the "amount of good" created in the ecosystem. Perspectives like these will further the discussion on welfare efficiency rather than exclusively viewing CSR through the narrow lens of shareholder primacy.

Appendices:

Appendix 1: Lee, K. H., & Shin, D. (2010). Consumers' responses to CSR activities: The linkage between increased awareness and purchase intention. *Public Relations Review*, 36(2), 193-195.

Table 1 Theories of corporate purpose

Corporate responsibility theory	Main idea
Shareholder theory (SPN)	Shareholder centric: the purpose of the corporation is to act as a vehicle for furthering the interests of <i>shareholders</i>
Stakeholder theory	Stakeholder centric: the purpose of the corporation is to act as a vehicle for furthering the interests of <i>stakeholders</i>
Social contract theory	Society centric: the legitimacy of the corporation's existence is based on its positive contributions to <i>society</i>

Appendix 2: MacAskill, W. (2019). The Definition of Effective Altruism. *Effective Altruism: Philosophical Issues*, 2016(7), 10.

(1) To us, "effective altruism" means trying to do as much good as possible with each dollar and each hour that we have.¹⁸

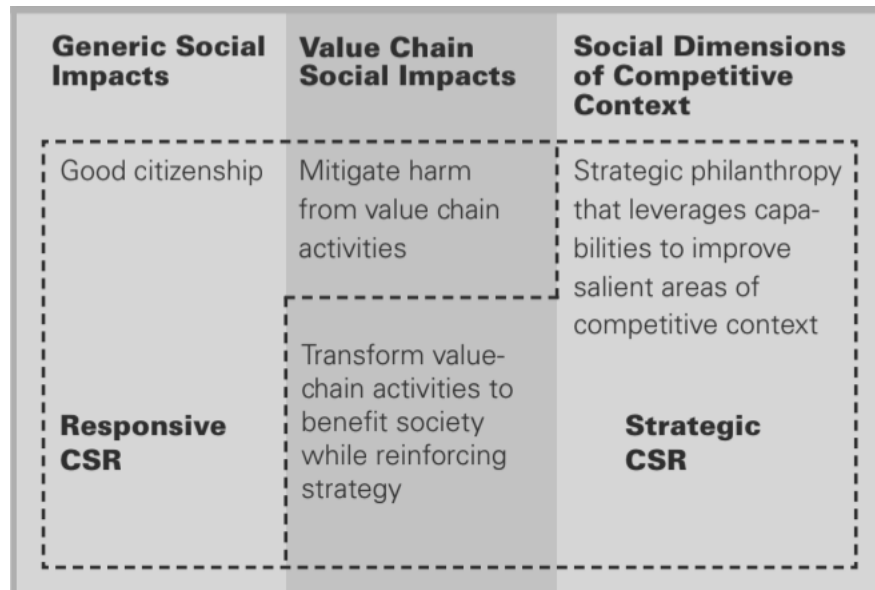
(2) Effective altruism is about asking, "How can I make the biggest difference I can?" and using evidence and careful reasoning to try to find an answer.¹⁹

(3) Effective altruism is based on a very simple idea: we should do the most good we can ... Living a minimally acceptable ethical life involves using a substantial part of our spare resources to make the world a better place. Living a fully ethical life involves doing the most good we can.²⁰

(p.13) (4) Effective altruism is a research field which uses high-quality evidence and careful reasoning to work out how to help others as much as possible. It is also a community of people taking these answers seriously, focusing their efforts on the most promising solutions to the world's most pressing problems.²¹

(5) Effective altruism is a philosophy and social movement that uses evidence and reason to determine the most effective ways to benefit others.²²

Appendix 3: Porter, M. E., & Kramer, M. R. (2006). *Strategy & Society*. Harvard Business Review, 84



Appendix 4: Lusk, J. L., Nilsson, T., & Foster, K. (2007). Public preferences and private choices: effect of altruism and free riding on demand for environmentally certified pork. *Environmental and Resource Economics*, 36(4), 499-521.

Table II. Example choice experiment question


Please choose A, B, C, D, or E					
Option/Attribute	A	B	C	D	E
Brand name	Hormel	Tyson	Store brand	No brand	I would not purchase any of these products
Price (\$/lb)	\$4.00	\$3.60	\$3.00	\$3.30	
Environmentally certified	✓	✓	✓		
Certified for animal well being	✓		✓		
Certified free of antibiotics	✓	✓		✓	
I would choose ...	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Appendix 5: Luo, J., & Kaul, A. (2019). Private action in public interest: The comparative governance of social issues. *Strategic Management Journal*, 40(4), 476-502.

TABLE 2 Mapping the comparatively efficient governance arrangement

Externality	Ex post information asymmetry	Low creative or absolute uncertainty		High creative or absolute uncertainty	
		Low commercial co-specialization	High commercial co-specialization		
		Column (1)	Column (2)	Column (3)	
Low	Low	Row (1) ^a	Market <i>e.g., Financial spot markets, farmers' markets</i>	For-profit governance <i>e.g., GE, Walmart, Google, Apple</i>	
	High	Row (2)	Service Nonprofit <i>e.g., Doctors without Borders, animal shelters, non-profit nursing homes</i>	Non-profit certification / partnerships <i>e.g., BBB, Underwriters Laboratories, Tom's shoes, Microsoft's partners in learning</i>	Social entrepreneurship <i>e.g., Method products, Drinkwell, Thorn technologies, benefit corporations</i>
Bounded	Low	Row (3)	Self-governing collective <i>e.g., Irrigation collectives, labor unions, PTAs, credit unions, retail co-ops</i>	BoP initiatives <i>e.g., Commercial micro-finance initiatives, Project Shakti</i>	Sharing economy <i>e.g., AirBnB, Kickstarter, Goodreads</i>
	High	Row (4)	Member Nonprofits <i>e.g., Churches, American Bar Association, Academy of Management</i>	Social activism <i>e.g., Fairtrade, HRC, before they book</i>	Social platforms <i>e.g., Edustar, Ushahidi</i>
General	Low	Row (5)	Govt. provision / maintenance <i>e.g., NYPD, National Park Service, IRS</i>	Govt. regulation / subsidies <i>e.g., OSHA, EPA, Medicare</i>	Govt. contracting <i>e.g., Road maintenance, NIH, NSF</i>
	High	Row (6)	Govt. sponsorship / political activism <i>e.g., Charter schools, NRDC, ACLU</i>	Public good certification / PPP / others <i>e.g., Rainforest alliance, World Food Programme, Project XL</i>	Public entrepreneurship / PPP / others <i>e.g., Hybrid prisons, EMS in India</i>

Appendix 6: Theories of Behavior and Behavioral Change. (n.d.). Retrieved November 12, 2020, from https://ebrary.net/12670/environment/theories_behavior_behavioral_change [reference for the graphic in-text] Ajzen, I. (1991). The theory of planned behavior. *Organizational behavior and human decision processes*, 50(2), 179-211. [reference for graphic below]



Appendix 7: Informed Consent

Thank you for participating in this research study in conjunction with the University of Michigan.

Before starting, please read the information below and, if you choose, provide your consent to participate in this study.

The goal of this survey is to ask basic questions about you. Your participation is voluntary. You are free not to answer any question or to withdraw from the study at any time.

Payment will be conditional on successful completion of the survey. If you do fail the attention checks, you will not be able to complete the survey.

The survey should take about 5 minutes.

Since you are enrolling in this research study through the Amazon Mechanical Turk (MTurk) site, we need to let you know that information gathered through Amazon MTurk is not completely anonymous. Any work performed on Amazon MTurk can potentially be linked to information about you on your Amazon public profile page, depending on the settings you have for your Amazon profile. Any linking of data by MTurk to your ID is outside of the control of the researcher for this study. We will not be accessing any identifiable information about you that you may have put on your Amazon public profile page. We will store your MTurk worker ID separately from the other information you provide to us. Amazon Mechanical Turk has privacy policies of its own outlined for you in Amazon's privacy agreement. If you have concerns about how your information will be used by Amazon, you should consult them directly.

Appendix 9: Pilot Regression

SUMMARY OUTPUT								
<i>Regression Statistics</i>								
Multiple R	0.057365526							
R Square	0.003290804							
Adjusted R Square	-0.017473971							
Standard Error	2.669222854							
Observations	50							
<i>ANOVA</i>								
	<i>df</i>	<i>SS</i>	<i>MS</i>	<i>F</i>	<i>Significance F</i>			
Regression	1	1.129131164	1.129131164	0.1584801	0.6923249			
Residual	48	341.9880308	7.124750642					
Total	49	343.117162						
	<i>Coefficients</i>	<i>Standard Error</i>	<i>t Stat</i>	<i>P-value</i>	<i>Lower 95%</i>	<i>Upper 95%</i>	<i>Lower 95.0%</i>	<i>Upper 95.0%</i>
Intercept	10.86095022	0.750766919	14.46647414	3.9996E-19	9.35143216	12.3704683	9.35143216	12.3704683
Effectiveness	-0.019542377	0.049089659	-0.398095586	0.6923249	-0.1182438	0.079159	-0.1182438	0.079159

Appendix 10: Updated Survey Question



Please be careful in reading the information below. You will receive a comprehension question. If you do not answer correctly, you will not be able to complete the survey.

Imagine you are a customer considering purchasing one of the two pain relievers below. The chemical formula, number of pills, and dosage of each bottle are identical. Both companies earn \$1 million in profits annually. The only difference is that Company Y contributes a portion of its profits to manufacture and donate a drug to treat a deadly skin disease in the world's poorest populations. This donated drug costs \$100 to manufacture.

Chance of one dose of the donated medicine saving a life: 9%

Percentage of Company Y's profits donated: 7%



How much are you willing to spend on Company Y's pain reliever?

0 \$ 20
10

Willingness to pay



Appendix 11: Oxford Utilitarian Scale Questions



How much do you agree or disagree with the following statements?

If the only way to save another person's life during an emergency is to sacrifice one's own leg, then one is morally required to make this sacrifice.

Strongly disagree	Disagree	Somewhat disagree	Neither agree nor disagree	Somewhat agree	Agree	Strongly agree
-------------------	----------	-------------------	----------------------------	----------------	-------	----------------

From a moral point of view, we should feel obliged to give one of our kidneys to a person with kidney failure since we do not need two kidneys to survive, but really only one to be healthy.

Strongly disagree	Disagree	Somewhat disagree	Neither agree nor disagree	Somewhat agree	Agree	Strongly agree
-------------------	----------	-------------------	----------------------------	----------------	-------	----------------

From a moral perspective, people should care about the well-being of all human beings on the planet equally; they should not favor the well-being of people who are especially close to them either physically or emotionally.

Strongly disagree	Disagree	Somewhat disagree	Neither agree nor disagree	Somewhat agree	Agree	Strongly agree
-------------------	----------	-------------------	----------------------------	----------------	-------	----------------

It is just as wrong to fail to help someone as it is to actively harm them yourself.

Strongly disagree	Disagree	Somewhat disagree	Neither agree nor disagree	Somewhat agree	Agree	Strongly agree
-------------------	----------	-------------------	----------------------------	----------------	-------	----------------

Appendix 11 (cont.)

It is morally wrong to keep money that one doesn't really need if one can donate it to causes that provide effective help to those who will benefit a great deal.

Strongly disagree	Disagree	Somewhat disagree	Neither agree nor disagree	Somewhat agree	Agree	Strongly agree
-------------------	----------	-------------------	----------------------------	----------------	-------	----------------

It is morally right to harm an innocent person if harming them is a necessary means to helping several other innocent people.

Strongly disagree	Disagree	Somewhat disagree	Neither agree nor disagree	Somewhat agree	Agree	Strongly agree
-------------------	----------	-------------------	----------------------------	----------------	-------	----------------

If the only way to ensure the overall well-being and happiness of the people is through the use of political oppression for a short, limited period, then political oppression should be used.

Strongly disagree	Disagree	Somewhat disagree	Neither agree nor disagree	Somewhat agree	Agree	Strongly agree
-------------------	----------	-------------------	----------------------------	----------------	-------	----------------

It is permissible to torture an innocent person if this would be necessary to provide information to prevent a bomb going off that would kill hundreds of people.

Strongly disagree	Disagree	Somewhat disagree	Neither agree nor disagree	Somewhat agree	Agree	Strongly agree
-------------------	----------	-------------------	----------------------------	----------------	-------	----------------

Sometimes it is morally necessary for innocent people to die as collateral damage—if more people are saved overall.

Strongly disagree	Disagree	Somewhat disagree	Neither agree nor disagree	Somewhat agree	Agree	Strongly agree
-------------------	----------	-------------------	----------------------------	----------------	-------	----------------

Appendix 12: Survey Main Regression

SUMMARY OUTPUT								
<i>Regression Statistics</i>								
Multiple R	0.016677348							
R Square	0.000278134							
Adjusted R Square	-0.00223373							
Standard Error	2.909294237							
Observations	400							
<i>ANOVA</i>								
	<i>df</i>	<i>SS</i>	<i>MS</i>	<i>F</i>	<i>Significance F</i>			
Regression	1	0.937201903	0.9372019	0.11072811	0.739491919			
Residual	398	3368.669198	8.46399296					
Total	399	3369.6064						
	<i>Coefficients</i>	<i>Standard Error</i>	<i>t Stat</i>	<i>P-value</i>	<i>Lower 95%</i>	<i>Upper 95%</i>	<i>Lower 95.0%</i>	<i>Upper 95.0%</i>
Intercept	11.23048469	0.228146005	49.2249894	2.388E-171	10.78196281	11.6790066	10.7819628	11.6790066
Number of Lives Saved	-0.000348455	0.001047172	-0.33275833	0.73949192	-0.002407134	0.00171022	-0.00240713	0.00171022

Appendix 13: Statistically Significant Results

Dependent Variable: Willingness to Pay			
Independent Variable	Intercept	Coefficient	Standard Error
Primary Regression: Effectiveness (# of Lives Saved)	-11.230	-0.000	(0.001)
Generosity X	12.058	0.043*	(0.018)
\$100,000-\$149,999 X	11.208	0.018*	(0.008)
Somewhat Aligned with Firm Morals	12.362	-1.418*	(0.644)
Excitement About Purchase	8.461	0.495***	(0.078)
Likelihood of Telling Friends About Purchase	9.445	Very Likely 4.817***	Very Likely (0.866)
		Somewhat Likely 2.576***	Somewhat Likely (0.558)
		Somewhat Unlikely 1.367*	Somewhat Unlikely (0.585)

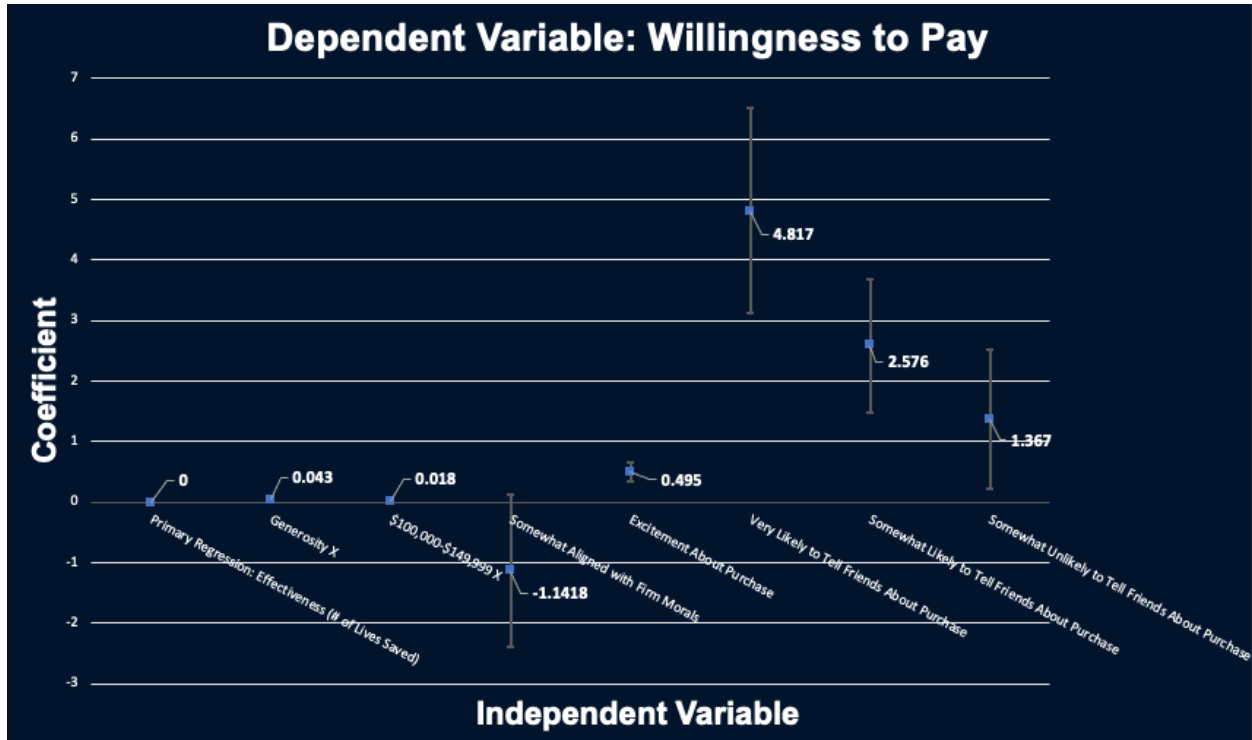
*,**,*** indicates significance at the <.05, <.01 and <.001 level respectively

Sample size for \$100,000-\$149,999 X was 389 rather than 400

Independent variable were each run in separate regressions

X after the variable title indicates the variable was run in an interaction with the number of lives saved

Appendix 14: Statistically Significant Results [visual format]



References

- Bénabou, R., & Tirole, J. (2010). Individual and corporate social responsibility. *Economica*, 77(305), 1-19.
- Berman, J. Z., Barasch, A., Levine, E. E., & Small, D. A. (2018). Impediments to effective altruism: The role of subjective preferences in charitable giving. *Psychological science*, 29(5), 834-844.
- Breidert, C., Hahsler, M., & Reutterer, T. (2006). A review of methods for measuring willingness-to-pay. *Innovative Marketing*, 2(4), 8-32.
- Cassar, L., & Meier, S. (2017). Intentions for doing good matter for doing well: The (negative) signaling value of prosocial incentives (No. w24109). National Bureau of Economic Research.
- Caviola, L., Schubert, S., Teperman, E., Moss, D., Greenberg, S., & Faber, N. (2020). Donors vastly underestimate differences in charities' effectiveness.
- Flammer, C. (2015). Does corporate social responsibility lead to superior financial performance? A regression discontinuity approach. *Management Science*, 61(11), 2549-2568.
- Friedman, M. (2007). The social responsibility of business is to increase its profits. In *Corporate ethics and corporate governance* (pp. 173-178). Springer, Berlin, Heidelberg.
- Guagnano, G. A. (2001). Altruism and market-like behavior: An analysis of willingness to pay for recycled paper products. *Population and Environment*, 22(4), 425-438.
- GiveWell: Our Top Charities. (2020). Retrieved November 19, 2020, from <https://www.givewell.org/charities/top-charities>
- Kahane, G., Everett, J. A., Earp, B. D., Caviola, L., Faber, N. S., Crockett, M. J., & Savulescu, J. (2018). Beyond sacrificial harm: A two-dimensional model of utilitarian psychology. *Psychological Review*, 125(2), 131.
- Kanungo, Rabindra N. & Conger, Jay A. (1/8/1993). Promoting Altruism as a Corporate Goal. *The Academy of Management executive* (1993), 7(3), 37-48. Briarcliff Manor: Academy of Management.
- Kaul, A., & Luo, J. (2018). An economic case for CSR: The comparative efficiency of for-profit firms in meeting consumer demand for social goods. *Strategic Management Journal*, 39(6), 1650-1677.
- Lantos, G. P. (2002). The ethicality of altruistic corporate social responsibility. *Journal of consumer marketing*.

- Lähdesmäki, M., & Takala, T. (2012). Altruism in business—an empirical study of philanthropy in the small business context. *Social Responsibility Journal*.
- Lee, K. H., & Shin, D. (2010). Consumers' responses to CSR activities: The linkage between increased awareness and purchase intention. *Public Relations Review*, 36(2), 193-195.
- Loewenstein, G., Small, D. A. (2007). The Scarecrow and the Tin Man: The vicissitudes of human sympathy and caring. *Review of General Psychology*, 11, 112–126
- Luo, J., & Kaul, A. (2019). Private action in public interest: The comparative governance of social issues. *Strategic Management Journal*, 40(4), 476-502.
- Lusk, J. L., Nilsson, T., & Foster, K. (2007). Public preferences and private choices: effect of altruism and free riding on demand for environmentally certified pork. *Environmental and Resource Economics*, 36(4), 499-521.
- Michaeli, M. (2020, November 06). TDWI: Your Money or Your Life - The High Cost of Gleevac. Retrieved November 19, 2020, from <https://thedoctorweighsin.com/high-cost-life-saving-drugs/>
- McWilliams, A., & Siegel, D. (2000). Corporate social responsibility and financial performance: correlation or misspecification?. *Strategic management journal*, 21(5), 603-609.
- McWilliams, A., Siegel, D. S., & Wright, P. M. (2006). Corporate social responsibility: Strategic implications. *Journal of management studies*, 43(1), 1-18.
- Orhun, A. Y. (2018). Perceived motives and reciprocity. *Games and economic behavior*, 109, 436-451.
- Porter, M. E., & Kramer, M. R. (2006). *Strategy & Society*. Harvard Business Review, 84
- Smith, N. C., & Rønnegard, D. (2016). Shareholder primacy, corporate social responsibility, and the role of business schools. *Journal of Business Ethics*, 134(3), 463-478.
- World Health Organization (n.d.). Impact Assessment. Retrieved from <https://www.who.int/roadsafety/projects/manuals/alcohol/4-How%20to.pdf>

Acknowledgements

This senior thesis was completed for the course requirements of the BA 480 - Senior Thesis Capstone course at the University of Michigan Ross School of Business. This research would not have been possible without the help and mentorship from many important members of the Ross community. First and foremost, Assistant Professor of Strategy Justin Frake served as the advisor to this thesis. From the ideation phase all the way to the fine tuning of this thesis, Professor Frake has been absolutely instrumental in the creation of this research. Thank you for your constant support, advice, and patience throughout this past year. Additionally, the BA 480 course leadership of Professors Burcu Tasoluk and Francine LaFontaine, as well as PhD student advisor Paul Organ provided immeasurable guidance throughout the formation of this thesis, and for that I am immensely grateful. To Dr. Lillian Chen, thank you for your hands-on guidance throughout the survey phase of this study, your expertise of Qualtrics and IRB guidelines was essential to the crafting of the manipulation of our survey. Finally, thank you to the peers in my BA 480 cohort who never cease to amaze me with their intellect, critical feedback, and zest for research. Without the help of these contributors, this work would not have been possible.