



Publicly Accessible Penn Dissertations

2019

Three Essays On Strategy And Social Responsibility

Carson Everhart Young

University of Pennsylvania, ceverhartyoung@gmail.com

Follow this and additional works at: <https://repository.upenn.edu/edissertations>



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

Recommended Citation

Young, Carson Everhart, "Three Essays On Strategy And Social Responsibility" (2019). *Publicly Accessible Penn Dissertations*. 3227.
<https://repository.upenn.edu/edissertations/3227>

This paper is posted at ScholarlyCommons. <https://repository.upenn.edu/edissertations/3227>
For more information, please contact repository@pobox.upenn.edu.

Three Essays On Strategy And Social Responsibility

Abstract

The phenomenon of greenwashing receives significant public and academic attention. So do the related practices of bluewashing (regarding involvement with the United Nations Global Compact) and the use of corporate social responsibility (CSR) as mere window dressing. The first essay in this dissertation provides a conceptual and ethical analysis of these practices. It also criticizes some frameworks for using CSR as a tool for increasing a firm's economic performance. If CSR is to be taken seriously, rather than dismissed as a marketing gimmick at best or a veil for corporate malfeasance at worst, then we cannot endorse CSR decisions being reduced to considerations of economic performance.

On what basis should CSR decisions be made and evaluated? That is the research question that motivates this dissertation's second essay. I argue that any normative framework for CSR decision making must be based primarily on ethical considerations (as opposed to considerations of firm financial performance). As one possible such framework, I introduce effective altruism, a recent academic and social movement that advocates using strong evidence and rigorous analysis to increase the positive impact one makes through altruistic acts. Effective altruism provides a plausible framework for making and evaluating CSR whose main purpose is to benefit other people. However, there are other types of CSR that implicate alternative moral considerations—related to harms, rights violations, moral reconciliation, and reciprocity—that effective altruism does not capture.

How do we distinguish between ethical and unethical ways of pursuing profit? The Market Failures Approach to business ethics purports to provide an answer to this question. In the third essay of this dissertation, I argue that it fails to do so. This failure stems from its reliance on Pareto efficiency as a core ethical principle. Many ethically “preferred” tactics for seeking profit cannot be justified by appeal to Pareto efficiency. I argue that, rather than Pareto efficiency, we should look to the value of wealth creation to understand the ethical constraints on how market actors may pursue profit.

Degree Type

Dissertation

Degree Name

Doctor of Philosophy (PhD)

Graduate Group

Legal Studies & Business Ethics

First Advisor

Brian Berkey

Keywords

Business ethics, Corporate philanthropy, Corporate social responsibility, Philosophy of economics

Subject Categories

Ethics and Political Philosophy | Philosophy

THREE ESSAYS ON STRATEGY AND SOCIAL RESPONSIBILITY

Carson Everhart Young

A DISSERTATION

in

Legal Studies and Business Ethics

For the Graduate Group in Managerial Science and Applied Economics

Presented to the Faculties of the University of Pennsylvania

in

Partial Fulfillment of the Requirements for the

Degree of Doctor of Philosophy

2019

Supervisor of Dissertation

Brian Berkey, Assistant Professor of Legal Studies & Business Ethics

Graduate Group Chairperson

Catherine Schrand, Celia Z. Moh Professor, Professor of Accounting

Dissertation Committee

Diana Robertson, Samuel A. Blank Professor in Legal Studies, Professor of Legal Studies & Business Ethics

Thomas Donaldson, Mark O. Winkelman Professor, Professor of Legal Studies & Business Ethics

THREE ESSAYS ON STRATEGY AND SOCIAL RESPONSIBILITY

COPYRIGHT

2019

Carson Everhart Young

This work is licensed under the
Creative Commons Attribution-
NonCommercial-ShareAlike 3.0
License

To view a copy of this license, visit

<http://creativecommons.org/licenses/by-nc-sa/2.0/>

ACKNOWLEDGMENTS

I owe my biggest debts of gratitude to my wife, Cleo Young, and to my advisor, Brian Berkey. Without Cleo's love and support, I would not have completed my Ph.D. If Brian had not agreed to supervise this dissertation, I would not have been able to write it. It has benefitted enormously from his generous, thoughtful feedback.

Thanks to J.P Messina, Bas van der Vossen, James Stacey Taylor, Andrew Jason Cohen, and Lisa Tang, whose encouragement helped me persevere through difficult periods during my Ph.D journey. I am also grateful for the support I received from my doctoral student colleagues, including Suneal Bedi, Matthew Caulfield, Lauren Kaufmann, Vikram Bhargava, and Gaston de los Reyes.

The main essays that comprise this dissertation have been shaped by written comments from, or discussions with, Diana Robertson, Thomas Donaldson, Robert Hughes, Sarah Light, Julian Jonker, Vincent Buccola, Richard Shell, Scott Reynolds, Heather Dixon-Fowler, Keith Hankins, Nicolas Cornell, Jessica Flanigan, and Wayne Norman.

ABSTRACT

THREE ESSAYS ON STRATEGY AND SOCIAL RESPONSIBILITY

Carson Young

Brian Berkey

The phenomenon of greenwashing receives significant public and academic attention. So do the related practices of bluewashing (regarding involvement with the United Nations Global Compact) and the use of corporate social responsibility (CSR) as mere window dressing. The first essay in this dissertation provides a conceptual and ethical analysis of these practices. It also criticizes some frameworks for using CSR as a tool for increasing a firm's economic performance. If CSR is to be taken seriously, rather than dismissed as a marketing gimmick at best or a veil for corporate malfeasance at worst, then we cannot endorse CSR decisions being reduced to considerations of economic performance.

On what basis should CSR decisions be made and evaluated? That is the research question that motivates this dissertation's second essay. I argue that any normative framework for CSR decision making must be based primarily on ethical considerations (as opposed to considerations of firm financial performance). As one possible such framework, I introduce effective altruism, a recent academic and social movement that advocates using strong evidence and rigorous analysis to increase the positive impact one makes through altruistic acts. Effective altruism provides a plausible framework for making and evaluating CSR whose main purpose is to benefit other people. However, there are other types of CSR that implicate alternative moral considerations—related to harms, rights violations, moral reconciliation, and reciprocity—that effective altruism does not capture.

How do we distinguish between ethical and unethical ways of pursuing profit? The Market Failures Approach to business ethics purports to provide an answer to this question. In the third essay of this dissertation, I argue that it fails to do so. This failure stems from its reliance on Pareto efficiency as a core ethical principle. Many ethically "preferred" tactics for seeking profit cannot be justified by appeal to Pareto efficiency. I argue that, rather than Pareto efficiency, we should look to the value of wealth creation to understand the ethical constraints on how market actors may pursue profit.

TABLE OF CONTENTS

ACKNOWLEDGMENT..... III

ABSTRACT.....IV

LIST OF ILLUSTRATIONS.....VI

INTRODUCTION..... VII

**CHAPTER 1: WHAT’S WRONG WITH GREENWASHING?: THE ETHICS OF
CSR AS A MARKETING TOOL..... 1**

**CHAPTER 2: AN EFFECTIVE ALTRUIST APPROACH TO CORPORATE
SOCIAL RESPONSIBILITY36**

**CHAPTER 3: AGAINST PARETIANISM: A WEALTH CREATION APPROACH
TO BUSINESS ETHICS..... 73**

CHAPTER 4: CONCLUSION 119

BIBLIOGRAPHY 121

LIST OF ILLUSTRATIONS

Fig. 1: The Pareto Frontier.....	88
Fig. 2: Consumer and producer surplus.....	89
Fig. 3: Deadweight loss.....	90

INTRODUCTION

One disadvantage of the three-essay dissertation model is that it produces a fragmented final product. This introduction aims to remedy this, at least in part. It provides a brief overview of some of the main theoretical approaches and assumptions I have employed throughout these essays. I do not provide a full justification of my theoretical approaches and assumptions, but I do make them explicit, and in some cases I also indicate my main reasons for adopting them.

What is normative business ethics theory?

Philosophers of science (Giere, 1994; Suppe, 1972) distinguish between two different perspectives on what Harris, Johnson, & Souder (2013) call “model-theoretic knowledge accumulation” (442) in empirical scientific research. A *law statement perspective* understands theories to consist of falsifiable general statements about the real world. A *model-theoretic perspective*, by contrast, understands theories as models that represent certain aspects of the real world, applying to a greater or lesser degree depending on the circumstances. Schmitz (2006), a political philosopher, extends this model-theoretic perspective (without, to my knowledge, using that exact term) to normative theorizing about justice. He claims that we should think of theories not as arguments, as the law-statement perspective does, but as conceptual maps that represent a terrain. (If this is not identical to the model-theoretic perspective, it certainly bears strong resemblance to it).

Though I believe the law statement perspective can be valuable for normative business ethics, I myself adopt something closer to the model-theoretic perspective in the essays in this dissertation. I think about moral truths—including (though not limited to) particular facts about how a person should act in a given circumstance—as a terrain. Understanding ethics requires understanding the contours of this terrain. Moral theories are maps that represent the terrain. Theories, by necessity, involve simplification. This is how theories help us achieve a better understanding of their subject matter. But simplification also has a cost. Just as certain map projections represent some geographic areas well even as they distort others, certain ethical theories capture some parts of the ethical terrain well, even as they distort others. Of course, some theories (and projections) are just bad, and ought to be discarded. But the mere fact that a theory distorts a certain part of the ethical terrain does not mean that it has no value as a theory. An ethical theory can have value by capturing certain parts of the moral terrain especially vividly. Ethical truth, according to the model-theoretic perspective, resides not so much in ethical theories themselves as in the terrain that ethical theories represent.

Thus, I do not think of normative business ethics theory as consisting primarily of abstract principles providing necessary and sufficient conditions for moral rightness or goodness (though such principles can play an important role in normative theorizing). Business ethics theory should be *practical*. It should aid us in our attempts to understand how business organizations should be run, how managerial decisions should be made, and how people should act in situations that commonly arise in business. It should help practitioners understand what behaving ethically in business requires.

This discussion of a model-theoretic perspective for business ethics provides some context for claims I make about the wealth creation approach I defend in essay 3. The wealth creation approach identifies wealth as one important value for understanding a certain part of the terrain of business ethics: the restrictions on how market actors may permissibly pursue profit. The approach does not provide a law that tells us with certainty whether employing a certain set of strategic tactics is permissible. Instead, it offers a defeasible principle that highlights one important consideration for determining permissibility. It represents an important contour in the ethical terrain of business.

The model-theoretic perspective is also reflected in the pluralist approach I employ in essays 1 and 2. I do not choose to develop a normative theory of CSR by adopting a specific version of Kantianism, Aristotelianism, contractarianism, or utilitarianism and then applying that theory to CSR. Instead, I appeal to mid-level moral principles to justify the conclusions I defend: do not free-ride on others' investments in a common resource, do not manipulate others, do not seriously harm others or violate their rights, atone for (and compensate victims of) past wrongs, help the needy in effective ways. Not everyone will agree with each of these principles as I have formulated them, and people will certainly disagree about how they should be weighed against each other. But I would submit that each of these principles is significantly less controversial than any one comprehensive ethical theory. Arguments are generally stronger when they rely on less controversial premises: by definition, controversial premises fail to garner the endorsement of a significant proportion of one's audience. Thus, I believe that I can provide stronger arguments for the conclusions about CSR I defend by relying on a pluralist set of common moral principles than on a single comprehensive moral theory.

For precisely this reason, I sometimes go out of my way in the three essays to deny that endorsing the positions I defend requires endorsing a particular comprehensive moral theory. In essay 3, I deny that adopting my wealth creation approach requires endorsing welfare consequentialism: the value of wealth can be justified not only by appeal to welfare, but also by appeal to non-welfare considerations such as autonomy. Similarly, in essay 2, I argue that we should understand effective altruism not as a deeply consequentialist or utilitarian theoretical perspective—though it admittedly does require endorsing some thin consequentialist assumptions about moral aggregation—but rather as a reflection of the relatively commonsense idea that we should try to do more good through our altruistic acts.

An institutional approach to business ethics

Business does not occur in an institutional vacuum. It takes place in profit-oriented, bureaucratic organizations that compete against each other in a market. I strongly agree with Heath (2014) that an adequate theory of business ethics must take the institutional competitive market environment of business seriously. Trying to understand the ethics of business without understanding the normative structure of competitive market institutions is like trying to understand the ethics of being a lawyer without recognizing the adversarial nature of the judicial system, or trying to understand the ethics of a competitive team sport without recognizing that it involves a competition that each team tries to win.

I thus disagree in important respects with claims like the one in the title of John Maxwell's book, *There is No Such Thing as "Business" Ethics: There's Only One Rule*

for Making Decisions (Maxwell, 2007). We cannot develop an adequate understanding of business ethics through simple application of abstract theoretical principles to business contexts. However, I also do not go as far as Heath (2014), who appears to hold that business ethics is best understood by looking for norms that are implicit in various parts of the institutional environment of business, such as corporate law. I agree with the sort of view Berkey (2016) defends, according to which general principles of morality and justice apply to particular institutional contexts, including business and market contexts, and give actors in those contexts direct reasons for action. However, I also believe that institutional features can affect what general principles of morality and justice *imply* about how one should act in a particular circumstance. Thus, in order to understand business ethics, we must understand what (if anything) justifies competitive market institutions. Then we must understand how the structure of competitive market institutions—including the formal and informal norms governing the conduct of actors that inhabit them—serves the values that justify them.

Essay 3 reflects this institutional approach to business ethics by specifying and describing a certain value, wealth, that I believe serves as an important part of the overall justification for competitive market institutions. Understanding how profit-seeking actors in competitive market institutions create wealth can help us understand some important ethical limits on how profit in competitive market institutions may be pursued.

One can also see this institutional approach to business ethics reflected in essay 1. If businesses organizations, given the competitive environment in which they operate, have a general orientation toward pursuing profit, why should they not have a profit orientation toward corporate social responsibility (CSR)? Some prominent CSR scholars

believe they should. The business case for CSR is a popular idea. In my view, however, it is also a dangerous one, given how CSR is often understood. Adam Smith's (1776) invisible hand justification of the profit motive does not generally hold for CSR. If we view CSR primarily as way to gain or sustain competitive advantage or improve financial performance, rather than a practice whose proper purpose is fundamentally an ethical one, we risk countenancing what I call disingenuous CSR behavior—CSR that, like the notion of bullshit Frankfurt (2009) has analyzed, is undertaken with insufficient regard for the ends it purports to promote or protect. Viewing business firms as fundamentally profit-oriented entities that inhabit a competitive market environment, as an institutional approach to business ethics would have us do, helps to illuminate some potential dangers of CSR that otherwise might not be apparent.

CSR: a term of art

In a broad sense, all three of the essays in my dissertation are about CSR, since they all analyze and defend claims about how business firms should act toward certain stakeholders and other members of society. After all, it is natural to understand ‘CSR’ to refer to a set of *responsibilities* that businesses have toward certain stakeholders, society, or the environment. When Milton Friedman argued that “there is one and only one social responsibility of business—to use its resources and engage in activities designed to increase its profits so long as it stays within the rules of the game” (Friedman, 2002: 133), he was defending the position that business managers’ main responsibilities are owed to the firm’s shareholders, rather than to society in the form of “providing employment, eliminating discrimination, [and] avoiding pollution” (Friedman, 1970: 32).

Friedman advances a position that is explicitly normative. For Friedman's view to be a view about CSR, then, the term 'CSR' must be understood to refer to something like a firm's actual fulfillment of its social responsibilities, or a firm's successfully conducting itself in accordance with what its social responsibilities require. A given act will only qualify as CSR if it is an act that the firm has ethical reason to perform.

If we understand CSR in this way, a firm's corporate social responsibility refers to some subset of the overall set of ethical reasons that apply to that firm. Theories of CSR as responsibilities clarify what a firm's ethical responsibilities to society are, and which actions those responsibilities require. The subfield of scholarly literature that addresses CSR as responsibilities, then, is normative business ethics. Normative debates about corporate governance defending shareholder primacy (Friedman, 1970; Jensen, 2002; Marcoux, 2003) or normative stakeholder theory (Donaldson & Preston, 1995); about corporate ownership and its implications (Rodin, 2005; Strudler, 2017); about whether firms have a social responsibility to eliminate sweatshops from their supply chains (Arnold & Bowie, 2003, 2007; Zwolinski, 2007); about whether, how, and to what extent business firms ought to participate in political processes (Hussain & Moriarty, 2016; Scherer & Palazzo, 2007; Singer, 2018); these are the sorts of issues that CSR as responsibilities encompasses. Theories of CSR as responsibilities have as their central aim prescribing what businesses, from an ethical perspective, should do.

However, I was initially surprised to discover that many CSR researchers do not understand CSR in the way just described. CSR has become a term of art. Rather than referring primarily to responsibilities businesses have, CSR refers to a *category of activities* that businesses do. One clue that someone is using this category of activity

conception of CSR is that they speak or write of firms that ‘do CSR’—English syntax allows us to speak of ‘doing’ activities, but not of ‘doing’ responsibilities.

This conception of CSR is often expressed in terms of intention: CSR activities are those that are undertaken voluntarily (rather than in response to, e.g., a legal requirement) and are intended or designed to provide benefits or mitigate harms to stakeholders, society, or the natural environment (Barnett, 2007a; Mackey, Mackey, & Barney, 2007; Marquis, Glynn, & Davis, 2007; Rupp, Ganapathi, Aguilera, & Williams, 2006). But this notion of organizational intention is under-theorized (it is not clear what constitutes organizational intention) and difficult to measure (to the extent that organizational intention depends on organizational leaders’ internal mental states, it cannot be observed directly). It is perhaps for this reason that scholarly management literature often appeals to the notions of *appearance* and *reflection* of intention or social purpose, rather than making CSR ascription dependent on actual intention (e.g. Matten & Moon, 2008; McWilliams & Siegel, 2001). Thus, under the activities conception, we can understand CSR to refer to activities that (1) appear to promote some apparently social or ethical purpose beyond firm financial performance or (2) appear to reflect a commitment on behalf of the firm to that ethical or social purpose.

As someone who employs a normative, philosophical methodology, it is tempting to declare that this second conception of CSR represents a corruption of the term, and to insist that CSR be reserved for the responsibilities conception.

However, I think it is more interesting to take the activities conception of CSR seriously, instead of dismissing it. At any rate, this is what I attempt to do in essays 1 and 2. If CSR refers not to a firm’s ethical and social responsibilities, but instead to activities

a firm does that seem to promote or reflect a commitment on behalf of the firm to some ethical or social purpose, how should we develop normative theory about CSR?

As essay 1 demonstrates, adopting this activities conception of CSR has important implications for how firms may use CSR as a means for improving financial performance. Under the responsibilities conception, CSR is ethical by definition: a firm does CSR when it fulfills its normative responsibilities. The activities conception, however, severs the connection between CSR and substantive ethical considerations. This means that, perhaps surprisingly, certain kinds of CSR activity can be unethical. Indeed, viewing CSR primarily as a source of strategic advantage risks countenancing ethically problematic kinds of CSR such as greenwashing, bluewashing, and window dressing.

Decoupling CSR from substantive ethical considerations also opens up interesting questions about what positive view we should adopt regarding how CSR decisions should be made and evaluated. This is one of the main purposes of essay 2. Effective altruism provides a plausible source of prescriptive guidance about CSR whose primary purpose is beneficence. We must rely on other kinds of moral considerations—refraining from serious harms and rights violations, atoning for (and compensating victims of) past wrongs, reciprocity—for prescriptive guidance for CSR whose primary purpose is something other than beneficence.

In this introduction, I have attempted to make explicit some of the main methodological decisions I have made in the course of writing the three essays that follow. I have also highlighted some main themes that unify the essays. My hope is that these remarks will help readers understand how my dissertation essays relate to and

support each other, and how they contribute to normative business ethics theory more broadly.

CHAPTER 1: What's wrong with greenwashing?: The ethics of CSR as a marketing tool

I. Introduction

In the early 2000s, British Petroleum, colloquially known as BP, debuted a series of marketing campaigns rebranding itself as an environmentally conscious energy company. These campaigns celebrated its investments in green energy, encouraging the public to associate BP with the slogan “beyond petroleum” rather than the fossil fuel-evoking actual name of the company. The campaign was a success, winning the American Marketing Association’s golden Effie award for advertising effectiveness in 2007 (Solman, 2008). Critics, however, were suspicious of BP touting its green bona fides when the vast majority of its business remained in the extraction and processing of fossil fuel. BP’s 2010 Deepwater Horizon oil spill seemed to confirm their suspicions. In retrospect, the environmental activist group Greenpeace looks prescient for having awarded BP its (satirical) Emerald Paintbrush award for greenwashing in 2008 (Gray, 2008).

BP is one of the most notorious examples of corporate greenwashing, but many other firms have exhibited similar behavior, making public expressions that suggest a commitment to sustainability that is not backed up by substantive action. Indeed, greenwashing has achieved a sufficiently high profile to attract scholarly attention (Delmas & Burbano, 2011; Laufer, 2003; Lyon & Montgomery, 2015). Nor is sustainability the only domain in which companies have attempted to use the language of corporate responsibility for public relations purposes in a way that is misleading in light

of their other practices. First, there is the analogous phenomenon of bluewashing. This occurs when a firm uses its involvement with the United Nations Global Compact to deflect attention away from poor human rights practices. For example, consider an apparel firm that publicly touts its support for the UN Global Compact while sourcing clothing from sweatshops that routinely violate their workers' basic human rights. Then there are also more general (and older) terms, such as "window dressing," which is used to describe corporate social responsibility (CSR) activity that is primarily about creating the appearance of social responsibility (Friedman, 1970; Lin, 2010).

Intuitively, it seems obvious that there is something ethically problematic about greenwashing, bluewashing, and CSR as mere window dressing (hereafter referred to as "window dressing"). However, we lack a clear understanding of *why* behavior like greenwashing is especially ethically problematic. It makes sense that environmentalists would be critical of BP; environmentalists believe that we have strong reason to oppose all major fossil fuel companies, given the environmental damage associated with the extraction and use of fossil fuel. But why was BP, it seems, a special target of environmentalist ire? After all, BP at least made some attempts to adopt greener practices. Why should Greenpeace target BP rather than one of the other multinational companies in the fossil fuel industry, some of which (we might imagine) are likely to have engaged in practices at least as bad as (if not worse than) BP's, but without BP's token sustainability efforts?

Plus, even if it was hypocritical for BP to cloak itself in environmentalism while making as much money as it could from fossil fuels, it is not immediately clear why this

is ethically worse than causing similar environmental damage without an accompanying green public relations campaign. If the alternative is nihilism, perhaps hypocrisy is to be preferred. Indeed, hypocrisy may sometimes play a role in bringing about eventual improvements to firms' behavior if hypocritical organizations experience pressure to reduce discrepancies between talk and action (Christensen, Morsing, & Thyssen, 2013).

My aim in this paper is to clarify our conceptual and ethical understandings of phenomena like greenwashing, bluewashing, and window dressing. I also discuss some implications of my analysis for how we should think about CSR more generally, especially regarding the use of CSR to improve economic performance. As I show, because behavior like greenwashing is usually unethical, there are some important ethical limits on how a manager may use CSR to advance the firm's financial goals.

The paper proceeds as follows. Section II discusses the phenomena of greenwashing, bluewashing, and window dressing. I argue that these all qualify as specific instances of a more general phenomenon that I label *disingenuous CSR*: CSR activity that is undertaken with insufficient regard for the social or ethical end it purports to advance or protect. In section III, I offer four reasons for why disingenuous CSR—including greenwashing, bluewashing, and window dressing—is usually unethical: (1) it likely reduces the overall amount and quality of CSR in society, (2) it often constitutes morally impermissible free-riding, (3) it is often manipulative of stakeholders, and (4) it may increase people's propensity to engage in unethical behavior through the moral licensing effect. Some might take issue with my use of the term 'disingenuous CSR' for this sort of activity on the grounds that it seems odd to think of CSR as being unethical.

Is CSR activity not ethical by definition? Section IV argues that, given the conception of CSR that is most prominent in the CSR literature, CSR is in fact not ethical by definition. In other words, that an activity qualifies as CSR is neither a necessary nor sufficient condition for that activity being ethical. In section V, I turn my attention to strategic CSR: the idea that CSR can be a tool for gaining and sustaining competitive advantage. Some scholars, notably McWilliams & Siegel (2001) and Husted & Salazar (2006), have advocated an extreme version of strategic CSR according to which managers' decisions about CSR ought to be reduced to firm economic performance considerations. Managers, on this view, should pursue CSR investments that contribute to maximizing profit, and refrain from making CSR investments that are not expected to maximize profit. One reason that this extreme version of CSR is problematic, I argue, is that it countenances disingenuous CSR. Section VI addresses a specific argument from Husted & Salazar (2006), which I call the invisible hand defense of strategic CSR. I show why this argument fails to vindicate the view that managers ought to make decisions about CSR primarily or exclusively on the basis of firm economic performance considerations. Section VII concludes, offering some positive guidance about how managers can avoid ethically problematic greenwashing, bluewashing, and window dressing behavior. I also suggest some avenues for future research.

II. Greenwashing, bluewashing, window dressing: CSR as a public relations tool

In his famous broadside against CSR, Friedman (1970) suggests that much of what people perceive to be socially responsible business conduct is actually ordinary, profit-oriented business activity that is presented to the public as socially responsible for public relations purposes. In the intervening years since Friedman's essay, some more specific terms have arisen to refer to similar phenomena that arise in certain domains of seemingly socially responsible business behaviors. *Greenwashing* occurs when a firm fosters a reputation for environmental sustainability as a façade to mask poor environmental performance (Laufer, 2003). *Bluwashing* occurs when a firm uses its association with the United Nations to help mask a poor human rights record (Berliner & Prakash, 2015). While the scholarly study of bluwashing remains in relatively early stages, an entire greenwashing literature has developed. This body of research attempts to clarify the factors that drive greenwashing. Candidates include industry conditions, firm performance, firm governance, the political-legal environment, and institutional demands (e.g. Delmas & Burbano, 2011; Kim & Lyon, 2014; Marquis, Toffel, & Zhou, 2016; Ramus & Montiel, 2005).

But to what, exactly, do the terms greenwashing, bluwashing, and window-dressing refer? A close reading of the literature on these concepts reveals several different ways of understanding them. Friedman (1970: 426) defines "window dressing" as the use of social responsibility as "a cloak for actions... that are entirely justified on [the firm's] own self-interest." In other words, a firm engages in window dressing when it represents ordinary business activity undertaken in order to improve bottom-line economic performance as being done out of a commitment to social responsibility.

Note the difference between this understanding of window dressing and several proposed definitions of greenwashing. Laufer (2003: 253) explains greenwashing as a form “of disinformation from organizations seeking to repair public reputations and further shape public images.” Other conceptualizations of greenwashing also identify it with public relations campaigns that use disinformation to paint in a misleadingly positive light the environmental performance of firms that fail to conduct themselves sustainably (Delmas & Burbano, 2011; Ramus & Montiel, 2005; Vos, 2009). Still, there are subtle differences between this disinformation-based conception of greenwashing and others that exist in the literature. For example, it seems that greenwashing does not *require* disinformation in the strict sense of a firm making false claims about environmental performance: Lyon & Montgomery (2015: 223) adopt a somewhat broader conception that encompasses “communication that misleads people into forming overly positive beliefs about an organization’s environmental practices or products.” This conception of greenwashing, unlike disinformation-based conceptions, encompasses the practice of making selective but true public claims about positive environmental actions while neglecting to disclose negative environmental actions (Marquis et al., 2016).

Additionally, some conceptualizations understand greenwashing to necessarily involve diverting public attention away from environmentally deficient underlying behavior (e.g. Delmas & Burbano, 2011), while others only require that greenwashing public communication paint an excessively rosy picture of the firm’s environmental performance (e.g. Lyon & Montgomery, 2015). The latter could involve behavior that is neutral from an ethical perspective (i.e. neither unethical nor especially ethically good),

and thus unproblematic when isolated from the accompanying public expression, but that is then embellished to suggest that the firm behaves more sustainably than it actually does. The former, on the other hand, require the underlying environmental behavior itself to be negative.

There is much less literature on bluewashing than greenwashing, but the various conceptions of bluewashing are roughly analogous to those we find in the domain of greenwashing. Some scholars define bluewashing as taking part in the UN Global Compact in order to distract public attention from conduct that is deficient from a human rights protection point of view—the underlying human rights conduct must be subject to negative assessment even when considered separately from the accompanying communication (e.g. Bigge, 2004). Others require only that there be a divergence between the firm’s underlying human rights conduct and its public presentation, which is consistent with the underlying conduct being ethically neutral (Berliner & Prakash, 2015).

a. Disingenuous CSR

One of the contributions of this paper is to unify the various types of conduct I am discussing—greenwashing, bluewashing, and window dressing—into a single broader category. This will give us a clearer picture of why these behaviors are ethically problematic.

My proposal is that we should understand window-dressing, greenwashing, and bluewashing to represent different instances of the more general phenomenon that I call disingenuous CSR. This term is inspired by philosopher Harry Frankfurt's (2009) pioneering work on a related idea. Frankfurt analyzes speech that is meant to persuade but lacks consideration for the truth. One who engages in speech in this way cares only about persuading their interlocutor, not about whether the content of the speech is true or false, or whether the speech induces the interlocutor to adopt true or false beliefs.

Frankfurt's focus is on a specifically discursive phenomenon. In other words, it applies to speech. However, I propose that we understand 'disingenuous CSR' to refer more broadly to CSR activity that is undertaken with insufficient regard for the social or ethical ends it purports to advance or protect. Thus, firms that exploit the marketing and public relations benefits of CSR without making a serious effort to deliver on the commitments their marketing and public relations efforts suggest do CSR disingenuously.

Understood in this way, disingenuous CSR provides a plausible conceptual account of greenwashing, bluewashing, and window dressing: each of these involves firms, in a certain domain, undertaking CSR activities with insufficient regard for the environmental, social, or ethical end that activity purports to advance or protect. Apply this definition to the "Beyond Petroleum" marketing campaign. BP presented itself as a firm that had embraced a genuine commitment to protecting the long-term health of the natural environment. However, BP's conduct strongly suggests that it did not, as an organization, actually take environmentally causes seriously in the way its marketing

both explicitly stated and implicitly implied. A United States federal judge found BP's conduct to have been, at various points in time leading up to the Deepwater Horizon explosion, reckless and grossly negligent, and that some of its most disastrous decisions were motivated by profit (Barbier, 2011). This is consistent with the perception that, aside from acquiring a relatively small solar energy company, BP was not giving environmental concerns any significant weight in its corporate decision making during the years when it was marketing itself as Beyond Petroleum. Its actions indicate that, as an organization, it did CSR disingenuously, failing to exhibit the organizational virtues or live up to the standards of conduct its CSR campaign portrayed. Thus, BP was guilty of greenwashing, because it was guilty of disingenuous CSR in the domain of environmental sustainability.

However, in addition to covering the conceptual space of greenwashing, bluewashing, and window dressing, disingenuous CSR also extends beyond it. In other words, there is an entire category of business conduct that is closely analogous to greenwashing, bluewashing, and window dressing, but that has not been clearly named in public and academic discourse in the way greenwashing and bluewashing have been. Prototypical examples of greenwashing and bluewashing are discursive (this is less clear for window dressing). A significant part of the Beyond Petroleum campaign was BP's corporate speech. But much disingenuous CSR is not primarily discursive in this way. I worry that our lack of a concept to refer to this non-greenwashing, non-bluewashing conduct may blind us to some of its ethically problematic aspects and impede our ability to identify it when it occurs.

To make my point concrete, consider the following example:

Bobs Boots

Bobs Boots specializes in making rugged outdoor footwear. Bobs' top management has recently attended a workshop on 'doing well by doing good.' They decide to put what they learned from the workshop into practice by investing in CSR that will contribute to Bobs' financial performance. Because Bobs has core capabilities related to designing and manufacturing boots, they decide it makes strategic sense for Bobs to adopt the popular "one-for-one" model of social impact according to which a firm gives away one of its products for free to a person in need for every product it sells (Marquis & Park, 2014). Bobs management hopes that this initiative will increase customer willingness to pay for Bobs Boots, provide material for effective marketing campaigns to raise brand awareness, increase motivation of Bobs employees, and improve Bobs' reputation among the general public.

After operating for several years, it becomes clear that Bobs' one-for-one CSR initiative contributes to Bobs' strategic goals in the way management had hoped. However, it also becomes clear that the initiative fails to make any meaningful social impact. Giving boots to people in poverty turns out not to be an effective way of helping them. Some recipients of boots from Bobs have pressing needs for food, medical supplies, or other basic necessities, and they sell the boots for a fraction of their cost in order to raise money to cover these expenses. For other recipients, though a new pair of boots makes them marginally more comfortable, the boots have no impact on their wellbeing or their ability to live dignified, minimally decent lives. There is little reason to judge that recipients of boots from Bobs are meaningfully better off than they would have been had they never received the boots at all. If Bobs were a charitable organization devoted to improving the lives of the poor, its boot donation program would have been judged a waste and a failure.

Suppose further that the lack of positive social impact of Bobs' one-for-one program does not come as a surprise for Bobs' management team. They were aware of arguments for why cash transfers are often much more effective means of alleviating poverty than donations of goods (Blattman & Niehaus, 2014a). They were also familiar with evidence indicating that a similar footwear donation program sponsored by TOMS Shoes had a negligible overall effect on the lives of recipients (Wydick, Katz, Calvo, Gutierrez, & Janet, 2016a). Indeed, at the time they granted final approval to go forward with the one-for-one boot donation initiative, Bobs' upper management team did not expect that the initiative would actually have any meaningful positive impact on the lives of its

beneficiaries. But they decided that the likelihood that the initiative would boost the firm's financial performance constituted sufficient reason for Bobs to pursue it.

There seems to be *something* unethical about Bobs' conduct. It may not be clear yet *why* it is wrong—I offer my account of that in the next section. However, the idea that it could be ethically acceptable for a firm to pour resources into a CSR initiative that has no or very little meaningful social impact merely for the purpose of increasing its financial performance rings false.

Notice how the Bobs Boots case does not quite fit as a poverty-alleviation analogue of greenwashing or bluewashing. First, though Bobs Boots' wrongful conduct may be partly discursive to the extent that its marketing efforts rely on misleading public statements about the social impact of its boot donation program, its wrongful conduct is not exclusively or primarily discursive. Bobs' overall conduct and presentation of its product—not just its speech, narrowly construed—creates a misleading connection in the minds of its stakeholders between its product and the cause of poverty alleviation. Second, unlike paradigmatic cases of greenwashing and bluewashing, Bobs is not using its CSR initiative to conceal other ethically dubious conduct from public scrutiny. It is not as though Bobs chooses to donate pairs of boots to people living in poverty to distract attention from some aspect of its business model that contributes to poverty. Aside from its disingenuous CSR, Bobs does nothing wrong. The ethical problems with Bobs' CSR stem from it being disingenuous in the sense that it is undertaken with insufficient regard for the end (in this case, poverty alleviation) that it purports to promote or protect.

One might object at this point that there seems to be a significant difference between some types of greenwashing, bluewashing, and window dressing that involve concealing or distracting attention from harmful activity, on the one hand, and cases like Bobs Boots, which does not involve underlying harmful activity, on the other. BP's conduct is in an important sense worse than the conduct of Bobs Boots because in addition to being disingenuous, BP also inflicted (unjustified) harm on the environment and various stakeholder groups.

I agree that BP's conduct involved wrongs that Bobs' conduct did not. However, I maintain that we should understand the wrong involved in BP's harmful conduct separately from the wrong involved in greenwashing and disingenuous CSR. If the harms BP inflicted on the environment and various stakeholder groups were indeed wrongful, they would have also been wrongful absent BP's Beyond Petroleum marketing campaign. But it is the Beyond Petroleum marketing campaign, not the harms themselves, that make BP's conduct qualify as greenwashing. Thus, when firms use greenwashing to cover up or distract attention from other harmful behaviors, we should understand the greenwashing and the harmful behaviors as separate kinds of wrongs. As a result, despite the additional wrong of inflicting harm present in the BP case that is absent from the Bobs Boots case, the analogy between BP's greenwashing and Bobs' disingenuous CSR holds up.

A final point of clarification: whether an act or tactic counts as disingenuous is intent-sensitive. Someone who makes a genuine effort to convey the truth but fails is not being disingenuous. Neither is a firm that makes a genuine effort to promote an ethical

end through a CSR initiative but fails due to some unforeseen circumstance. Thus, in the real world, where we often lack information about both individual and group intentions, it is usually difficult to determine for certain whether a given act qualifies as disingenuous or instead reflects the influence of bad luck or genuinely well-intentioned incompetence. (This is one reason why I have chosen to rely on the fictional Bobs Boots example rather than a real-world case for the purposes of my discussion.)

III. What's wrong with disingenuous CSR?

Now that it is clear what disingenuous CSR is and how it works, I will discuss the main reasons why it is unethical. One of my main claims in this paper is that we can understand greenwashing, bluewashing, and window dressing as specific instances of the more general category of disingenuous CSR. However, I do not believe that all disingenuous CSR is unethical for the same reason, or even that all disingenuous CSR is necessarily unethical—though, as the discussion that follows will suggest, most instances of disingenuous CSR probably *are* unethical. Thus, what follows is a description of some of the most important ways in which disingenuous CSR can be unethical.

a. Stakeholder trust and the market for CSR

A certain level of stakeholder trust (Pirson & Malhotra, 2008, 2011; Schnackenberg & Tomlinson, 2016) is required for an individual firm to derive benefits

from strategic CSR. If customers of the socially conscious footwear company TOMS shoes did not trust that TOMS would deliver on its stated commitment to give away one pair of shoes to a child in need for each pair sold, or if they learned the donated shoes were of such poor quality that they fell apart after a few normal days of wear, then they presumably would not be willing to pay more for TOMS Shoes than an otherwise identical product with no CSR attribute. Without the trust of its customers, TOMS would get no competitive advantage from its CSR program. An analogous effect would likely hold for certain other stakeholder groups: employees, suppliers, shareholders, and so on.

The first thing to notice about firms that sell products with CSR attributes is that their situation is analogous to the ‘lemons problem’ that Akerlof (1970) diagnosed for used cars. Akerlof showed how the information asymmetry between buyers and sellers of used cars creates an adverse selection problem in the used car market. As a result, high-quality used cars fail to reach the market. There is an analogous information asymmetry between firms and their CSR-valuing stakeholders. Generally, the firm will have much more access than its stakeholders to information relevant to assessing the quality of its CSR (Cho, Lee, & Pfeiffer Jr, 2013). After all, the firm will usually be responsible for administering its CSR initiatives, while stakeholders have to gain information provided by the media, third-party evaluators and standard-setting organizations, or, in many cases, the firm itself. To be sure, some stakeholder groups (e.g. employees) may have more access to information about CSR quality than others (e.g. customers), and information about some CSR initiatives may be relatively accessible. But in many cases, stakeholders have to rely chiefly on information from the firm’s own representations to assess quality

of CSR. As a result, it will often be difficult for stakeholders to gather enough information to determine whether a given CSR project is of high quality (a ‘peach,’ to use Akerlof’s folksy term for high-quality used cars) or low quality (a ‘lemon’).

As Akerlof’s (1970) analysis indicates, the likely effect of this situation is to inhibit high-quality CSR from being offered in the market. Assume that stakeholders make quality judgments about CSR: for an individual stakeholder, high-quality CSR effectively promotes a cause that stakeholder cares about, while low-quality CSR doesn’t promote a cause the stakeholder cares about or doesn’t promote its cause effectively. When stakeholders have the information necessary to make quality judgments about CSR, they will be willing to pay more for high-quality CSR than for low-quality CSR. However, in reality, stakeholders are likely to have difficulty distinguishing high-quality from low-quality CSR. As a result, they will not be willing to pay significantly more for high rather than low quality. Thus, firms interested in offering high-quality CSR will have difficulty obtaining a price premium for it. Assuming high-quality CSR requires greater investment on the part of the firm than low-quality CSR, firms will have little incentive to offer products with high-quality rather than low-quality CSR attributes. In sum, because of information asymmetries between CSR ‘producers’ (i.e. firms) and ‘consumers’ (i.e. stakeholders who value CSR), potential producers of high-quality CSR will be unable to enjoy strategic advantages of offering high-quality CSR on the market. This is a market failure that will cause high-quality CSR to be under-provided.

The second thing to notice about the market for CSR is that, because of the information asymmetry that gives rise to the lemons problem, stakeholder trust can be

seen as a collective good (Olson, 2009) from the point of view of firms seeking to derive strategic benefits from CSR. A given stakeholder's level of trust in a transaction will partially depend on features specific to an individual firm. Some firms have reputations for honesty and fairness, and some do not. In many cases, at least some reputational information is available. But in many cases, for many stakeholder groups, trust levels for a transaction will depend also on reputational factors of groups of firms in a given industry or region (Aldrich & Fiol, 1994; King & Lenox, 2000; Tucker, 2008). Thus, for many firms, the degree to which they can gain strategic benefits for CSR will depend not just on the reputation they cultivate individually, and the resulting level of stakeholder trust it elicits, but also on the overall level of trust various stakeholder groups have toward the industries and regions in which the firm does business.

Stakeholder trust functions as a collective good in the following sense (see also King, Lenox, & Barnett, 2002). As a group, firms seeking strategic advantage from CSR will benefit from higher levels of stakeholder trust toward firms in their regions and industries. When the level of stakeholder trust is high, stakeholders will give greater credence to firms' representations of the quality of their CSR projects. This will allow firms with high-quality CSR to overcome the information asymmetries that plague the market for CSR and successfully differentiate themselves from firms with low-quality CSR. As a result, firms will be able to generate greater strategic benefits through their CSR activities. The problem, however, is that it will be in the private interest of individual firms to under-invest in CSR quality, portray their CSR initiatives as higher quality than they actually are, and reap strategic benefits by free-riding off of the overall

level of stakeholder trust and exploiting stakeholders' lack of ability to distinguish between different quality levels of CSR.

This analysis of the market for CSR allows us to identify two ways in which CSR that is disingenuous but pursued out of strategic motivation risks countenancing unethical behavior. The first is that, by exacerbating failures in the market for CSR, disingenuous CSR both (a) contributes to the lemons problem dynamic by which high quality CSR is inhibited from being sold on the market and (b) undermines the collective good of stakeholder trust that facilitates firms having a strategic incentive to invest in CSR. As a result, the CSR market will provide less CSR, and of lower quality, than it would if it were not beset by these market failures. To the extent to which high-quality CSR is ethically valuable, this is a consequence that agents have ethical reason to avoid.

The second way that disingenuous CSR is likely to result in unethical behavior concerns fairness. Again, when a firm engages in disingenuous CSR, but does so strategically by publicly representing its CSR as high quality, it free-rides on the stakeholder trust to which other firms have contributed through honesty, fair dealing, and other trust-generating behaviors. By doing disingenuous CSR strategically, the firm gains the benefits from shareholder trust generated by firms that avoid disingenuous CSR, as well as the individual strategic benefits of stakeholders' perception that it is doing CSR (Tosi & Warmke, 2016).

Disingenuous CSR, then, leads both to bad consequences in the form of lower levels and quality of CSR, and to unfair free-riding on the collective good of stakeholder

trust generated by firms that do not engage in disingenuous CSR (on the latter, see Cullity, 1995).

b. Disrespect and manipulation

According to Strudler (2005: 459), it is “a sufficient even though not exhaustive characterization of manipulation that one person manipulates another when he intentionally causes that person to behave as he wishes through a chain of events that has the desired effect only because the manipulated person is unaware of that chain.” For example, in Shakespeare’s *Othello*, Iago manipulates Othello when he places Desdemona’s handkerchief in Cassio’s possession, causing Othello to adopt the false belief that Desdemona has been unfaithful to him (Noggle, 1996). Othello acts on this belief by murdering Desdemona in a fit of rage. Iago intentionally causes Othello to kill Desdemona in order to achieve the desired effect of destroying the lives of Othello and Desdemona. His scheme only works because Othello is unaware of the manipulation—had Othello known that Desdemona’s handkerchief came into Cassio’s possession because Iago planted it there, not because Desdemona had been unfaithful, Othello would not have slain his innocent wife.

Disingenuous CSR is likely to be manipulative in this sense to many stakeholders who are induced by the firm’s CSR to transact with the firm. Imagine Alex, a socially conscious boot collector who decides to buy her boots from Bobs rather than a rival boot maker in part because she has seen the publicity generated by Bobs’ one-for-one boot

donation initiative. Alex believes that combatting global poverty is an important cause, and she feels that she should support an organization that is stepping up and actually doing something about it. From Bobs' strategic point of view, their CSR investment has had the desired effect of raising the willingness to pay for boots from Bobs of customers like Alex.

But Bobs' CSR will only have the desired effect of causing Alex's willingness to pay for boots from Bobs to increase if Alex does not know pertinent information about Bobs' CSR initiative: For instance, if Alex knew that Bobs' CSR is disingenuous because Bobs is concerned only with the strategic benefits it derives from its CSR and not with its CSR's actual social impact, she would be even less willing to purchase Bobs products than she would have been had Bobs never established its CSR program at all. The reason why Bobs' CSR increases some consumers' willingness to pay for Bobs products is because those consumers care about the plight of the global poor. By harnessing that caring for its own private gain and failing to use its CSR activities to meaningfully advance the cause of poverty alleviation in the way a stakeholder like Alex would legitimately expect, Bobs manipulates those stakeholders. This kind of manipulation is unethical because, among other reasons, it is inconsistent with the requirement of Kantian moral respect for persons to always treat others as ends in themselves (Arnold & Bowie, 2003; Kant, 2002).

c. Moral licensing

Psychologists have identified a phenomenon called *moral self-licensing* whereby an individual's performance of a deed he believes is ethically good can increase his propensity to subsequently engage in conduct he believes to be unethical (Merritt, Effron, & Monin, 2010; Monin & Miller, 2001). Some studies on moral self-licensing suggest that CSR may interact with this psychological phenomenon in ethically troubling ways. In their experimental study of the effect of green products on ethical behavior, Mazar & Zhong (2010) find that participants who had purchased green products were less likely to exhibit altruistic behavior than participants who had purchased conventional products. In addition, List & Momeni (2017) conducted experiments that find a similar moral licensing effect for employees of firms that engage in CSR: employees assigned to work in CSR firms were about 24% more likely to cheat on subsequent tasks than employees assigned to non-CSR firms.

This research suggests a worrisome unintended effect of CSR: in some contexts, purchasing a product with CSR attributes or working at a firm that does CSR could make consumers or employees more likely to engage in unethical behavior, or less likely to engage in ethically good behavior, than they otherwise would be. If an individual's moral motivation or self-control are sensitive to how she perceives the ethical status of her behavior in the recent past, then exposure to CSR may displace ethical conduct in other areas of a person's life.

Moral licensing raises a host of important empirical and normative questions that I cannot resolve here. Empirically, we lack a detailed understanding of the explanatory mechanisms by which moral licensing occurs (Blanken, van de Ven, & Zeelenberg,

2015) and of how to render moral self-licensing effects compatible with evidence from social psychology about “foot-in-the-door” and related effects whereby individuals are driven to act in ways that are consistent with their past behavior (Freedman & Fraser, 1966; Mullen & Monin, 2016). Normatively, it is not clear to what extent individuals bear ethical responsibility for licensing effects (on themselves and others) of their behavior, or whether the permissibility of an act depends on its foreseeable licensing effects. One view that strikes me as initially plausible, but that I lack the space to defend here, would hold that it is permissible to displace ethical acts through moral licensing so long as the displacing acts are ethically better than the acts that are displaced.

If this is right, it would explain why non-disingenuous, meritorious CSR might be ethically permissible even if it has a licensing effect. It is one thing if a well-managed CSR initiative that contributes effectively to an important cause licenses future unethical behavior by stakeholders who interact with it. It is significantly more problematic if disingenuous CSR that fails to contribute effectively to an important cause also licenses those who interact with it to behave less ethically than they otherwise would. The more common and the greater the magnitude of the moral self-licensing effect, the stronger the ethical objection to strategic disingenuous CSR, which depletes stakeholders’ moral motivation without effectively advancing a good cause so that the firm can increase its competitive advantage. This is just the sort of socially destructive profit-seeking behavior that socially responsible businesses must avoid.

IV. How can CSR be unethical?

One of my main claims in this paper is that a certain category of CSR, disingenuous CSR, is usually unethical. It might strike some readers as odd that a category of CSR would be unethical. Is CSR not ethical by definition?

For the purposes of this paper, I have adopted the same definition of CSR that is most prominent in management and CSR literature. CSR, on this understanding, refers to activities that are undertaken voluntarily (rather in response to, e.g., a legal requirement) and are intended to provide benefits or mitigate harms to stakeholders, society, or the natural environment (Barnett, 2007a; Mackey et al., 2007; Marquis et al., 2007; Rupp et al., 2006). The notion of organizational intention that this definition of CSR invokes, however, is both under-theorized (it is not clear what constitutes organizational intention) and difficult to measure (to the extent that organizational intention depends on organizational leaders' internal mental states, it cannot be observed directly). It is perhaps for these reasons that definitions of CSR in the management literature often appeal to the notions of *appearance* or *reflection* of intention or social purpose, rather than making CSR ascription dependent on actual intention (e.g. Matten & Moon, 2008; McWilliams & Siegel, 2001). And when CSR activity is used as a variable in empirical studies, it nearly always treats the notion of intention as figurative rather than literal—CSR researchers do not typically attempt to ascertain the actual motives of the CSR decision makers in the firms they study in order to prove that they acted with the sort of intent that *real* CSR requires. If it looks like CSR, it is taken to be CSR. Thus, in effect, we can understand CSR to refer to activities that (1) appear to promote some apparently ethical or social

purpose beyond firm financial performance or (2) appear to reflect a commitment on behalf of the firm to that ethical or social purpose.

Notice that, if we adopt this standard conceptualization of CSR, then the fact that an activity qualifies as CSR is no guarantee that that activity is ethical. After all, an activity can appear to promote an ethical or social purpose without *actually* promoting that ethical or social purpose. An activity can appear to reflect a commitment to an ethical or social purpose on behalf of the firm that engages in it without the firm *actually* being committed to that purpose. Appearances—and reflections—can be deceptive. Recall Bobs Boots. It is easy to imagine stakeholders interacting with a firm like Bobs and interpreting its boot donation initiative as a clear case of CSR. Yet, given the facts of the case, despite qualifying as CSR, what Bobs does is unethical.

Of course, we are not compelled to adopt this common understanding of CSR. We might define CSR in a way that conceptually guarantees that activity that counts as CSR is also ethical. We could, for example, define CSR in terms of a firm's actual fulfillment of its social or ethical responsibilities. Now, determining precisely what a firm's social or ethical responsibilities *are* is difficult. This is the function of normative ethical theory. Utilitarianism will imply that a given firm has a certain set of ethical responsibilities, Kantian deontology will imply a somewhat different set, contractualism yet another set, and so on. This is one drawback of tying CSR ascriptions to fulfillment of substantive ethical and social responsibilities: it requires resolution of deeply controversial questions of normative ethics in order to determine whether a given activity qualifies as CSR. (It is easy to see why management scholars studying the antecedents and consequences of CSR

activity have adopted a conception of CSR that does not require them to do this.) I do not take a position on whether this conceptual drawback outweighs the conceptual advantage of guaranteeing that activities that count as CSR are also ethical. My point is simply that, given how CSR has been conceptualized, unethical CSR is both a conceptual and empirical possibility.

V. Some reservations about strategic CSR

The idea that CSR can be strategically advantageous to a firm has become prominent over the last several decades. There is a large literature that studies whether greater levels of corporate social performance are associated with stronger or weaker firm financial performance (e.g. Hawn & Ioannou, 2016; Turban & Greening, 1997; Waddock & Graves, 1997). There is an even larger literature devoted to deepening our understanding of whether and how particular kinds of CSR provide firms with strategic advantages (e.g. Hart, 1995; Hart & Dowell, 2011; Henisz, Dorobantu, & Narthey, 2014; Jones, 1995; Porter & Kramer, 2002; Prahalad, 2006). In addition to its considerable influence among academics, strategic CSR has become a common topic in business school classrooms and has received considerable attention from practitioners (M. Porter & Kramer, 2011; Werther Jr & Chandler, 2010).

I do not want to overstate the worries about strategic CSR that I present in this section. I accept that it is generally a good thing for firms to take actions that constitute win-wins for both the firm and society. Managers ought to think rigorously and creatively

about how they can benefit society in ways that also contribute to their firms' gaining and sustaining competitive advantage. Academics ought to support this through both research and teaching.

However, it is also easy for strategic CSR to be taken too far. If ethics is equated with CSR—as I just argued, it should not be; nevertheless, it often is—then strategic CSR can mislead managers into thinking that considerations of ethics and social responsibility can be reduced to familiar financial performance metrics. Extreme versions of strategic CSR thus risk countenancing greenwashing, bluewashing, window dressing, and other forms of disingenuous CSR. Some of the most prominent strategic CSR scholarship encourages this way of thinking. It does not just urge managers to search for opportunities to do CSR that will also contribute to gaining and sustaining competitive advantage, but further argues that managers ought to base their decisions about CSR primarily or exclusively on strategic considerations.

The clearest articulation of the view that CSR ought to be reduced to strategic considerations comes from a seminal article on strategic CSR by McWilliams & Siegel. They suggest that their strategic CSR approach “enables managers to determine the appropriate level of CSR investment” (McWilliams & Siegel, 2001: 118). A firm’s “ideal level of CSR can be determined by cost-benefit analysis,” because the firm that invests in a profit-maximizing level of CSR “meets the demands of relevant stakeholders—both those that demand CSR (consumers, employees, community) and those that ‘own’ the firm (shareholders) (McWilliams & Siegel, 2001: 125). In the conclusion of the ‘managerial implications’ section of their paper, they declare, “In sum, managers should

treat decisions regarding CSR precisely as they treat all investment decisions” (McWilliams & Siegel, 2001: 125), i.e. as ordinary business decisions to be pursued or avoided based on their expected contribution to the firm’s bottom-line financial performance.

Endorsements of this extreme version of strategic CSR appear elsewhere as well. Husted & Salazar (2006), whose argument I discuss in greater depth below, support a similar position to that of McWilliams & Siegel. Siegel (2009: 5), in a practitioner-oriented article about sustainable management practices, asserts “that firms should adopt ‘green management’ practices only if such activities complement the organization’s business- and corporate-level strategies and ultimately enhance profitability or shareholder wealth.” One can also find strategic CSR advocates who do not explicitly avow versions as extreme as McWilliams & Siegel, but who nevertheless strongly suggest that managers ought to think about CSR from a primarily strategic perspective. Drucker (1984: 62) writes that “the proper ‘social responsibility’ of business is to... turn a social problem into economic opportunities and economic benefit.” More recently, the ‘creating shared value’ approach that Porter & Kramer (2011: 76) have famously advocated, which has become a fixture in business school classes, insists that managers should eschew CSR that does not yield strategic benefits in favor of “policies and operating practices that enhance the competitiveness of a company while simultaneously advancing the economic and social conditions in the communities in which it operates.” Again, although Drucker and Porter & Kramer do not clearly advocate a version of strategic CSR as extreme as that of McWilliams & Siegel—to their credit, Porter &

Kramer insist that CSR initiatives ought to solve social problems as well as contribute to attainment of competitive advantage, meaning that they do not countenance behaviors like greenwashing—they are at the very least prone to be interpreted as saying that CSR decisions ought to be reduced to strategic considerations.

The problem with the aforementioned extreme versions of strategic CSR, as well as the less-extreme-but-prone-to-misinterpretation versions, is that by reducing CSR to strategic considerations (or tempting careless readers to do so), they risk countenancing the greenwashing, bluewashing, and window dressing behaviors that I discussed earlier in the paper. Thus, although they are presented using the language of social responsibility, they encourage unethical behavior.

VI. The invisible hand defense of strategic CSR

One of the most famous justifications of market actors basing decisions upon their individual private interest, rather than the public interest, comes from Adam Smith's *Wealth of Nations*. Because of how markets align individual incentives with the public interest, even when market actors intend only to promote their own interest, they are “led by an invisible hand to promote” the public interest as well (Smith, 1904: 456). An advocate of an extreme version of strategic CSR might make a similar argument: although the decision procedure upon which managers ought to rely when making choices about CSR investments ultimately boils down to strategic considerations, the outcome for society of individual firms pursuing strategic CSR will be beneficial since

market pressures will induce profit-oriented firms to produce social benefits through their CSR initiatives. Indeed, Husted & Salazar (2006) make essentially this argument in an article advocating that firms pursue CSR out of strategic rather than altruistic motives. Firms tend to respond to appeals to their strategic interest better than they respond to appeals to corporate altruism. Thus, a “focus on corporate altruism would probably result in a lower overall social output by the business community than a focus on social strategy” (Husted & Salazar, 2006: 87).

This ‘invisible hand’ defense of strategic CSR fails for at least two reasons. First, invisible hand arguments for the market only vindicate the profit-seeking orientation of individual market actors in markets that at least roughly approximate conditions of perfect competition (Heath, 2014; Jones & Felps, 2013). But as the previous section on stakeholder trust and the CSR ‘lemons’ problem explains, there is at least one important condition of perfect competition that does not even approximately hold in many markets for CSR: that buyers have adequate information about products and prices to make sound decisions, given their preferences. Because ‘buyers’ of CSR lack information about the quality of CSR on the market, purely profit-seeking behavior on the part of firms that do CSR is unlikely to induce market actors to produce socially desirable levels of various types of CSR. Rather, individual firms’ purely strategic pursuit of CSR is likely to exacerbate the ‘lemons’ problem in the market for CSR.

Second, invisible hand defenses of the market require that buyer preferences as expressed in the market at least roughly approximate buyers’ actual preferences that we have reason to care about from an ethical point of view (Buchanan, 1985; Cohen, 1977).

If buyers make decisions in the market that do not advance their well-being—if they fail to behave rationally—then the market will fail.

It is important to take a moment to understand, at least in rough terms, some of the conditions that must obtain for an individual buyer's market behavior to reflect the preferences that might be thought to constitute her well-being. To start, the buyer must have knowledge about what her preferences are. If she is buying a couch, for her choice to reliably track her preferences, she must have engaged in at least minimal rational deliberation about what she values in a couch—comfort, durability, size, aesthetic qualities, and so on. Additionally, she must have invested in gathering a certain amount of information about which of the various couches available on the market will best satisfy these preferences.

For the invisible hand argument to apply to the market for couches, these two conditions about individual buyers' knowledge of their own preferences must at least approximately obtain. But even if we assume these features about consumer preferences generally hold for goods like couches, there is reason to expect that, in many cases, they likely will not hold for consumer (and other stakeholder) preferences about CSR. Judgments about the quality of a couch implicate fairly concrete practical questions: whether it is comfortable, its durability, its aesthetic qualities, and so on. It seems reasonable to suppose that individuals generally have a good idea of what makes a couch good or bad for them, given their circumstances. But judgments about CSR quality implicate more abstract issues about which people are likely to have less well-formed opinions than their taste in couches: e.g. relative assessments about which causes deserve

our attention and resources, what are the best means for advancing a given cause, and so on. Thus, the willingness of a typical consumer to pay for a product with a CSR attribute is likely to be much less responsive to the actual quality of the CSR than the typical consumer's willingness to pay for a couch is to the actual quality of the couch.

Additionally, the ethical relevance of individual consumer preferences for CSR is less evident than for goods like couches. Couches are largely a matter of individual taste. If Bob loves purple couches but Helga believes purple couches are ugly, the market can provide Bob with purple couches and Helga with other colors, and Helga has no legitimate grounds to object: assuming they do not share a living room, Bob's choices about couch color are none of Helga's business. It is difficult to imagine a case in which one would have a reasonable, serious objection to the market for couches providing whatever consumers happen to prefer. But CSR is different. CSR decisions often centrally implicate issues of ethics and justice in which we all have a stake.

Unlike for couches, we can imagine cases where one might reasonably and seriously object to the CSR market providing whatever consumers happen to prefer. Simler & Hanson (2017) argue that, in many areas of life, what people say they will do and what they are actually motivated to do diverge, and that people's actual motivations are often more selfish than their stated motivations. Suppose that this divergence between public presentation and private motivation frequently occurs when consumers purchase products with CSR attributes. For example, when a consumer purchases a green product, she mainly cares about signaling to others her concern for the environment rather than actually using her role as a consumer to promote and protect important environmental

causes. Her engagement in CSR is thus disingenuous in the way discussed above from the perspective of the firm. Hence, this consumer's preferences are fulfilled perfectly well when the firm whose product she buys does CSR disingenuously, so long as this does not undermine the positive signaling benefits of the CSR product purchase. Suppose that, when the market provides attractively-marketed, disingenuous CSR, it is providing precisely what many consumers really do want.

Even if this were the case, I maintain that we could still reasonably object to CSR that is disingenuous and that fails to effectively promote or protect an important cause. To be clear, my claim here is not that consumer preferences should be considered irrelevant to CSR decision making. Rather, my contention is that the fact that the market for CSR efficiently produces CSR that satisfies stakeholders' preferences is not sufficient for establishing that the resulting CSR initiatives are ones we would have reason to endorse from an ethical perspective. Because of this, as well as the possibility that stakeholder preferences for CSR may not be as fully developed as consumer preferences for ordinary market goods and services, those who put forward the invisible hand defense of strategic CSR need to show why stakeholders' revealed preferences for CSR represent preferences that have ethical significance. They also must demonstrate how the market for CSR can overcome the information asymmetries that lead to the CSR lemons problem. In the meantime, the claim that we can count on market mechanisms to induce profit-oriented firms to undertake ethically desirable levels and forms of CSR should be regarded with a healthy dose of skepticism.

VII. Conclusion and further research

In this paper, I have sought to clarify the concepts of greenwashing, bluewashing, and window dressing, while also directing our attention to a more general category of unethical CSR activity, disingenuous CSR, that contains many practices that are unethical for the same reasons greenwashing, bluewashing, and window dressing are unethical. As I have argued, disingenuous CSR is typically unethical. This is an important finding in light of the exhortations of some proponents of certain extreme versions of strategic CSR, which countenance disingenuous CSR. By reducing CSR decision making to financial performance considerations, and by understanding CSR in a way that is largely disconnected from substantive ethical considerations, the range of CSR activities that could be endorsed by a given firm on strategic grounds includes some conduct that is typically unethical.

Indeed, if it is the case that disingenuous CSR is, as an empirical matter, a fairly widespread phenomenon, then it is concerning that these extreme strategic CSR views have as much influence as they do. In the frameworks that management and CSR researchers promulgate in their more public works—such as textbooks, articles in popular venues, popular books—great care ought to be taken to make clear that CSR *cannot* be reduced to strategic considerations. As I have demonstrated, when managers undertake CSR primarily for strategic reasons, they must ensure that they do so within the bounds of certain constraints in order to avoid greenwashing, bluewashing, window dressing, and disingenuous CSR behaviors more generally.

My arguments suggest some moral constraints that strategic CSR initiatives must respect in order to be ethically acceptable. I cannot give a full explanation of these here. However, to conclude, I will briefly outline two important constraints that ethical strategic CSR must respect. In the course of this discussion, I also flag several questions for future research.

a. Worthiness and effectiveness

First, to be ethically acceptable, CSR must promote or protect a worthy cause. In an episode of the American television show *The Office*, Michael Scott, the branch manager, organizes a ‘fun run’ to raise money that will go toward raising awareness of the dangers of rabies. Clearly, this is not a worthy cause. The merits of awareness-raising campaigns are generally dubious—merely raising awareness often does little to advance a cause. Plus, though rabies is a serious disease, it is not clear what we as a society should do about it beyond what we are already doing. Effective rabies treatments exist, and people rarely contract it. Thus, even if a firm were able to gain strategic benefits from a CSR initiative devoted to an inconsequential cause such as raising awareness about the dangers of rabies, it would be ethically wrong for the firm to do so.

Second, in addition to advancing a worthy cause, CSR must promote or protect that cause at least minimally effectively. A good example of an ineffective attempt to achieve social impact is PlayPumps, a charity that installs water pump merry-go-rounds in some countries in Africa (MacAskill, 2015). The vision behind PlayPumps is that

children would use the merry-go-round for amusement while at the same time pumping much-needed clean water for their families. However, children find PlayPumps too tiring to be suited for play. As a result, adults end up having to spin them to pump water. Compared to the normal pumps they replaced, PlayPumps are more expensive, more difficult to maintain, and pump less water (MacAskill, 2015). Providing clean water to people who lack it may well be a worthy cause, but PlayPumps failed to promote it minimally effectively and ended up making worse the lives of the very people they were supposed to help. Even if a firm could gain strategic advantage from funding the construction and distribution of PlayPumps—picture glossy marketing materials with smiling children playing on a merry-go-round as crystalline water gushes forth—it would be wrong for it to do so. Strategic CSR must promote or protect the cause to which it is devoted at least minimally effectively.

I would suggest that a more developed account of these constraints would likely forbid most (if not all) forms of greenwashing, bluewashing, window dressing, and disingenuous CSR more generally. Clearly, though, there is much more to be said. I confess that a ‘fun run’ for rabies and PlayPumps are easy cases. They avoid some interesting, but challenging, questions that are raised when we attempt to determine where these ethical lines are drawn. How do we compare the worthiness of causes, and what is the minimal baseline of worthiness required for a cause to be the target of permissible strategic CSR? In a world in which the global poor die from preventable causes like malnutrition and curable disease, are common causes of corporate philanthropy like promoting opera, ballet, or sports justifiable as worthy causes? Are

managers obligated to pursue maximally effective means for promoting a worthy cause, or is the effectiveness requirement less demanding than this? Such questions require further theorizing about what makes a cause worthy and what level of effectiveness is ethically required.

If CSR is important as the attention it garners indicates, if we want people to take seriously the idea that CSR can change the world for the better and is not just a marketing gimmick at best and a veil for corporate malfeasance at worst, then clarifying our thinking about these questions is crucial.

CHAPTER 2: An effective altruist approach to corporate social responsibility

I. Introduction

In *Theory of Moral Sentiments*, Adam Smith observes that “fitness, [the] happy contrivance of any production of art, should often be more valued, than the very end for which it was intended” (Smith, 1982). His point is that it can be tempting to value attributes of the inner workings of machines over the ends whose furtherance is the reason why the machine is useful in the first place. When the inner workings of machines have aesthetic value, this may be perfectly appropriate. Smith gives the example of a man who values his watch for its mechanical perfection rather than its utility for telling time. Many people appreciate watches in this way. I see no reason to think their doing so is mistaken.

However, there are clearly cases where valuing machines in this way would be misguided. A watch that exhibits the highest virtues of fine craftsmanship, but is unreliable for telling time, would be a poor choice for use in, for instance, laboratory equipment that requires precise and accurate time measurements to function effectively.

It can be tempting for scholars of business organizations, as well as practitioners who rely on business scholarship, to make an analogous sort of mistake when thinking about corporate social responsibility (CSR). Many scholars take financial performance, or something similar, to be the goal according to which business managers’ decisions should be made and evaluated (Friedman, 1970; Jensen, 2002; Sundaram & Inkpen, 2004). Theories in the field of strategic management have proliferated to explain how firms can

achieve the sustained competitive advantage that is typically required to achieve superior financial performance (e.g. Barney, 1991; Grant, 1996; Porter, 2008). In particular, the resource-based view of the firm, or RBV, has emerged as a paradigm within the field of strategic management (e.g. Barney, 1991; Peteraf, 1993; Wernerfelt, 1984). According to the RBV, firms achieve sustained competitive advantage when they possess and control valuable resources and capabilities that their rivals cannot easily replicate or approximate. The RBV suggests a certain general logic for how managers should make decisions: in brief, they should identify their firm's valuable resources and capabilities, and make decisions based on what will best protect and extend the competitive advantage that those resources and capabilities generate (Grant, 1991). Thus, when deciding whether to diversify or divest, managers should generally favor related diversification that leverages resources and capabilities that the firm's current activities leave idle, and divest from lines of business that are unrelated to the firm's main resources and capabilities (Neffke & Henning, 2013).

When applied to CSR decisions, the RBV seems to imply that CSR should be *integrative*: the CSR activities a firm undertakes should be closely linked to the firm's valuable capabilities or resources (Battilana, Lee, Walker, & Dorsey, 2012; Dees & Anderson, 2003; Marquis & Park, 2014; M. E. Porter & Kramer, 2011). However, using the RBV as a source of prescriptive guidance for CSR decision making is problematic. As a theory of strategic management, the main purpose of the RBV is to explain how firms achieve superior financial performance. However, as I argue in the first essay in this dissertation, improved financial purpose cannot serve as a plausible basis for CSR

decision making. CSR decisions implicate ethical values that financial performance metrics do not capture. Thus, to use the RBV as a source of guidance for how managers should make decisions about CSR is to value the inner workings of business organizations and institutions at the expense of the ends that those organizations and institutions should promote. (Or, less charitably, it may reflect an excessive adherence to a shareholder primacy ideology.) Unlike Smith's watch aficionado, in the case of business organizations and CSR, this focus on internal workings at the expense of external function is highly problematic. If CSR is truly a way for business to change the world for the better, rather than a marketing gimmick or a veil for corporate malfeasance, then prescriptive guidance about CSR cannot come from theories of strategic management or views about organization design derived therefrom. Rather, prescriptive guidance about CSR needs to be based primarily on ethical considerations.

But which ethical considerations? After discussing the inadequacy of the RBV as a basis for prescriptive guidance about CSR in section III, I introduce effective altruism, a recent academic and social movement (Berkey, 2018; MacAskill, 2015; Singer, 2015), in section IV as a potential ethical framework for CSR. In brief, effective altruism urges people to use rigorous analysis and strong evidence to increase the amount of positive impact they have through their altruistic acts. To apply these theories to a concrete example of CSR, Section V discusses Toms Shoes' widely-praised Buy-One Give-One model for social impact. Toms being held up as a paradigm of CSR makes sense if we take a financial performance approach to evaluating CSR, like the one represented by the RBV, but not if we evaluate CSR using effective altruism.

Although effective altruism provides much better tools for guiding and evaluating CSR decisions than strategic management theories like the RBV, it is also, I argue, not entirely adequate as a framework for CSR decision making. This is because CSR is not just about beneficence, which is the focus of effective altruism. Some CSR reflects other moral considerations. Thus, in section VI, I propose that we distinguish between three categories of CSR: (1) beneficent CSR, which reflects moral reasons firms have to promote the good of others; (2) preventative/restorative CSR, which reflect moral reasons firms have to avoid committing harms and rights violations, and to make amends for and compensate victims of past harms and rights violations; and (3) reciprocal CSR, which reflects moral reasons firms have to foster and sustain reciprocal relationships. While effective altruism is not a good framework for decisions about preventative/restorative CSR and reciprocal CSR, I briefly show how the moral reasons underlying these categories of CSR suggest the general form that prescriptive guidance about them should take. Section VII concludes.

II. Financial performance-first versus ethics-first approaches to CSR

Moral philosophers commonly distinguish between moral reasons and prudential reasons for action. Agents have prudential reasons to act in ways that promote their own interests and goals. Agents have moral reasons to act in ways that morality commands or recommends. The categories of prudential and moral reasons are not mutually exclusive. Sometimes prudential and moral reasons overlap, recommending the same course of

action, though sometimes they of course diverge. The prudential-moral reasons distinction is useful for understanding the different kinds of normative considerations that apply to individual moral agents.

Similarly, it is useful to make an analogous distinction between the reasons that apply to business firms.¹ The firm analogue of prudential reasons are reasons based on financial performance considerations. When a manager chooses what action a firm should take, her deliberations will usually implicate both financial performance reasons and moral reasons. However, it is common to think about corporate decisions using what I will call a “financial performance-first” framework. This is what defenders of shareholder primacy advocate (Friedman, 1970; Jensen, 2002). Financial performance-first frameworks can still allow a place for ethical considerations. Advocates of tinged stockholder theory—which holds that managers should maximize profits within the constraints of not only obligations of law, but also ethics (Heath, 2014; Langtry, 1994; Norman, 2011)—are particularly explicit about this. But moral reasons are limited to a constraining role in financial performance-first decision frameworks. They delimit how the manager may pursue her firm’s financial performance-first goals. Moral reasons do

¹ In this paper, I write about business firms as if they are moral agents. In doing so, I do not intend to take a side in the debate over whether organizations qualify as moral agents in a deep metaphysical sense (e.g. Pettit, 2007; Velasquez, 1983). Either we can understand firms as literal moral agents, or we can understand talk of corporate moral agency as metaphorical. In the latter case, writing of firms being subject to certain reasons for action is shorthand for saying that the various individual human agents that act on behalf of a firm are subject to those reasons for action. The position one takes in this debate does not affect this paper’s main arguments, so I choose to remain neutral about it.

not themselves constitute the main goals that serve as the basis of managerial decision making.

Contrast financial performance-first frameworks with what I will call ethics-first frameworks. In ethics-first decision frameworks, the goals that the decision maker seeks to promote, and uses as a metric for comparing the desirability of alternative potential courses of action, reflect the ethical reasons that apply to the decision maker. For example, a decision maker who attempts to choose a course of action on the basis of a moral theory—e.g. Kantianism, utilitarianism, virtue ethics, contractualism—employs an ethics-first decision framework.

So long as the right ethical constraints are respected, financial performance-first decision frameworks are appropriate for many business decisions. This is because actors pursuing profit in well-functioning markets bring about important social benefits (such as wealth, which I discuss in this dissertation's third essay). There are some business decisions for which a financial performance-first decision framework cannot be justified. CSR, I claim, is one particular category of managerial decision that calls for an ethics-first, rather than financial performance-first, decision framework. Ethical reasons apply much more directly to decisions involving CSR than ordinary, non-CSR business decisions because CSR involves the appearance of promoting or being committed to some social or ethical goal beyond firm financial performance (see, e.g., Barnett, 2007; Mackey, Mackey, & Barney, 2007; Marquis, Glynn, & Davis, 2007). Additionally, for many (though not all) non-CSR business decisions, acting on the profit motive can be ethically justified by appealing to Adam Smith's invisible hand (Smith, 1776)

explanation of how markets align the incentives of individual market actors with the public interest. However, the invisible hand of market mechanisms cannot generally vindicate the profit motive in the domain of CSR (as I argue in this dissertation's first essay).

III. A financial performance-first framework: The resource-based view

As the field of strategic management developed in the 1980's and 90's, the RBV represented an important shift in emphasis from approaches such as Michael Porter's (2008), which focus on how a firm's positioning within its industry can create sustained competitive advantage by generating market power for the firm in product markets. The RBV, by contrast, focuses on how firms realize sustained competitive advantage from firm-specific resources and imperfections in factor markets (Lehrer, 2001). In other words, the RBV attributes the source of sustained competitive advantage to a firm's control of valuable, difficult-to imitate *resources*. These resources often take the form of *core competencies* or *capabilities* that are difficult for competitors to duplicate because of their complexity and dependency upon tacit knowledge (Nelson & Winter, 1982; Prahalad & Hamel, 1990).

Whether a resource can generate sustained competitive advantage, according to the RBV, depends on the extent to which it is valuable, non-substitutable, rare, and difficult to replicate.

- A resource is *valuable* if it can potentially produce rents from a firm's assets—i.e. profits in excess of the assets' opportunity costs. A resource will often be valuable because it either increases what a firm's customers are willing to pay for its products or decreases the costs the firm bears to produce and sell its products (Porter, 2008).
- A resource is *non-substitutable* if it contributes to a firm capability that, in addition to being valuable, is more costly to acquire using an alternative resource (Barney, 1991; Dierickx & Cool, 1989).
- A resource is *rare* if it is costly or difficult for a firm's rivals to acquire, perhaps because it is firm-specific—i.e. only valuable in a context that is specific to the firm that uses it (Barney, 1991; Reed & DeFillippi, 1990).
- A resource is *difficult to replicate* when there are epistemic barriers to understanding its cause sufficiently for a firm to build the resource itself, perhaps because it relies on knowledge that is tacit or socially complex (Barney, 1991; Winter, 1997).

In other words, for a valuable resource to generate sustained competitive advantage, there must be some barrier preventing other firms acquiring the resource, or acquiring a reasonably close substitute resource, and competing away the economic rents it generates. As Grant (1991) shows, the basic explanatory tenets of the RBV have implications for how managers should make decisions aimed at achieving sustained competitive advantage. In brief, managers are advised to (1) identify their firm's main resources and capabilities; (2) estimate their potential to generate sustained competitive advantage (i.e. whether they are valuable, non-substitutable, rare, and difficult to replicate); (3) estimate whether the firm can capture the wealth they generate; (4) choose the strategy that best exploits the firm's resources and capabilities given the opportunities available to it; and (5) identify ways to protect and extend their positions of competitive advantage (Grant, 1991). Broadly speaking, the RBV perspective encourages managers to identify their firm's valuable resources and capabilities and base their corporate and

business strategies around what will best protect and exploit those resources and capabilities (e.g. Barney, 1995; Prahalad & Hamel, 1990; Williams, 1992). Thus, when deciding whether to diversify or divest, managers should generally favor related diversification that leverages resources and capabilities that the firm's current activities leave idle, and divest from lines of business that are unrelated to the firm's main resources and capabilities (Neffke & Henning, 2013).

The resource-based view's prominence in the field of strategic management has made a deep imprint in the field of CSR. Perhaps the best known example of this is Porter & Kramer's (2011) work on Creating Shared Value (CSV), which they propose as a new paradigm for understanding how business can contribute to alleviating social problems that should supplant CSR. One of the main features distinguishing CSV from CSR, as Porter & Kramer conceive of them, is that CSV initiatives bear a close connection to a firm's core capabilities: "CSR programs... have only a limited connection to the business.... In contrast, CSV... leverages the unique resources and expertise of the company to create economic value by creating social value" (Porter & Kramer, 2011: 16).

Porter & Kramer's work on CSV contains an especially clear link between the general ideas of the resource-based view and prescriptive guidance for firms pursuing CSR. But the suggestion that CSR activities ought to be closely linked to a firm's core business capabilities, activities, and strategies appears elsewhere as well. Scholars of hybrid organizations, which combine non-profit and for-profit missions, argue that managers should view "integration as the hybrid ideal," meaning that "everything [a

firm] does produces both social value and commercial revenue” (Battilana et al., 2012: 52). Advocates of integration claim that tightly integrating a social venture’s financial and social goals facilitates strategic decisions that further both types of goals simultaneously: reducing production costs while also reducing environmental impact (Hart, 1997), catering to under-served customers who occupy the bottom of the global economic pyramid while expanding the markets in which firms sell their products (Prahalad, 2009), increasing the amount customers who value the sort of impact to which the firm’s social mission is devoted are willing to pay (Dees & Anderson, 2003).

IV. An ethics-first framework: Effective Altruism

a. Moral theories

This section introduces an ethics-first framework for CSR decision making to contrast with the RBV-based, financial performance-first framework I just explained. But on what basis would one determine the prescriptive principles for such a framework? A moral theory—e.g. utilitarianism, Kantianism, virtue ethics—might seem like a promising place to start. Moral theories purport to be theories about what one should, in a moral sense, do. However, moral theories suffer from a problem of abstraction (Hasnas, 2013). The concepts that underlie most ethical theories—Kant’s categorical imperative, eudaimonia, the good (as employed by utilitarians and consequentialists)—are so abstract that their more precise content is itself a matter of significant debate. Thus, deriving prescriptive principles for CSR from the best-known moral theories would require

making controversial decisions about which version one adopts before the theory has been developed enough to generate actionable prescriptive principles. A valuable area for future research would be to articulate how plausible versions of the main ethical theories could provide actionable prescriptive principles for CSR decisions (of the sort the RBV generates encouraging that CSR be integrative). But that is not the approach I adopt here.

Alternatively, one might look to the main theories of CSR or business ethics as sources for an ethics-first framework for CSR. However, I do not think any of them are up to the task. Instrumental theories of strategic CSR (e.g. Husted & Salazar, 2006; McWilliams & Siegel, 2001), which focus on how CSR can improve a firm's financial performance, fail to provide a plausible guide to CSR decision making for reasons I explain in this dissertation's first essay. Normative stakeholder theory (Donaldson & Preston, 1995) may yield insights about one category of CSR (I discuss this below), but has not been theoretically developed enough to be a source of concrete practical guidance. Integrative social contract theory (Donaldson & Dunfee, 1994), in my view, is best understood as a theory about how and whether firms operating in international environments ought to adapt their practices to different cultures they inhabit. Thus, it is most valuable for understanding how CSR practices should be adapted to local cultures in which they take place, rather than for grounding a prescriptive framework for CSR more generally. The former issue is an important one, but not my focus here.

b. Effective altruism

I propose effective altruism as a better starting place than moral theories or business ethics and CSR theories for developing an ethics-first framework for CSR decisions. It is only a starting place: I alter it in some ways that, in my view, improve its suitability for the CSR context.

What is effective altruism? Philosopher Will MacAskill writes that 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. It... consists of the honest and impartial attempt to work out what’s best for the world, and a commitment to do what’s best, whatever that turns out to be. (MacAskill, 2015: 211)

Effective altruism insists that our altruistic decisions should be informed by the best available evidence about how to increase our social impact. It requires doing rigorous comparative analysis of one’s available altruistic courses of action. Berkey (2018) interprets effective altruism as containing four principles:

EA1: There are very strong moral reasons, grounded in fundamental values, for the well off to direct significant resources to efforts to address important moral issues (e.g. to alleviate the plight of the global poor).

EA2: These fundamental values include (but are not necessarily limited to) impartially promoting increases in welfare, or quality of life, for individuals, and the reasons provided by this value are at least fairly weighty.

EA3: There are strong reasons to prefer giving to efforts that will promote the relevant values most efficiently.

EA4: We should employ the best empirical research methods available in order to determine, as best we can, which efforts promote those values most efficiently.

I suggest that we view effective altruism somewhat more broadly than Berkey defines it here. If, as EA1 holds, the well-off have strong moral reasons to direct significant resources to important moral causes, then it seems that at least some successful corporations would have similarly strong reasons to devote resources to CSR. I find it plausible that firms that enjoy significant sustained economic rents (beyond their costs of capital), and thus (at least when they are not in rapid periods of expansion) have slack resources, do indeed have strong reasons to direct a portion of these rents toward effective measures directed at important moral causes. But I do not insist that one must accept this in order to qualify as an effective altruist for my purposes. Indeed, according to some presentations of effective altruist ideas, a commitment to giving effectively *if* one chooses to give in the first place is more fundamental than a commitment to giving in the first place (see Pummer, 2016). One could thus endorse an effective altruist approach to CSR decisions in the way I envision despite not fully embracing EA1.

Note also that effective altruism is not an exhaustive theory of moral decision making. As a prescriptive approach to CSR decision making, effective altruism requires taking principles EA2, EA3, and EA4 seriously. A moderate effective altruist might decide, in a given case, that some consideration not contained in EA2-4 has sufficient importance to weigh against the importance of reasons contained within EA2-4. (E.g. giving somewhat higher weight to the interests of members of the firm's own community when deciding upon which cause to select for its CSR program.) Of course, one does

have to give EA2-4 a reasonably significant level of weight in one's decision making for CSR to properly qualify as effective altruist. But they need not be absolute.

c. Effective altruist CSR

The tenets of effective altruism are open to some interpretation. Effective altruists try to increase the good they do through their altruistic acts, but what do they mean by “good”? Principle EA2 leaves the definition of “good” somewhat open, specifying that it includes promoting increases in welfare. But there are many different potential theories of welfare (e.g. Arneson, 2000; Harrison & Wicks, 2013; Jones et al., 2016). Different versions or interpretations of effective altruism will adopt different understandings of what “the good” is. Some effective altruists suggest following health economists’ lead and adopting Quality Adjusted Life Years (QALYs) as a metric for comparing the amounts of good done in alternative potential courses of altruistic action (MacAskill, 2015). These allow researchers to estimate the benefit of a health or charitable intervention in a way that weights years of extended longevity according to their quality—comfortable, enjoyable years get more weight than painful, miserable ones.

No plausible decision procedure can be reduced to a number. We must therefore guard against viewing any single metric as an all-encompassing gauge of an initiative’s social impact—i.e. the good that it accomplishes. But metrics are also indispensable for gathering information necessary for making informed comparative judgments between alternative potential courses of action. To take impact seriously in the way effective

altruism requires, we must base our decisions upon the best available evidence, and evidence often takes the form of metrics.

The effective altruist organization GiveWell attempts to put effective altruist principles into practice by curating a list of the most effective charities in the global poverty cause area. Their analysis is not value neutral, of course—it relies on decisions about how to weigh its various health, wellbeing, and economic indicators when doing comparative effectiveness evaluations of its recommended charities (see “Approaches to Moral Weights: How GiveWell Compares to Other Actors,” 2017). However, the analysis is useful, because it identifies some specific charities that GiveWell’s rigorous effective altruist analysis suggests are especially effective at saving and improving human lives. GiveWell’s 2019 list of recommended charities (“GiveWell List of Top Charities,” 2019) includes:

- Two organizations devoted to anti-malaria initiatives in sub-Saharan Africa, Malaria Consortium and Against Malaria Foundation;
- Four organizations that support deworming programs in low-income countries, Evidence Action’s Deworm the World Initiative, Schistosomiasis Control Initiative, Sightsavers, and END Fund;
- Helen Keller International, which provides vitamin supplements to reduce child mortality in sub-Saharan Africa; and
- Give Directly, which runs direct cash transfer programs targeting very poor people in Kenya and Uganda.

Examining GiveWell’s top charities list makes clear some implications of effective altruist principles. A marginal dollar is much more valuable to a poor person than a wealthier (even if poor by some absolute standard) person. Thus, resources are generally much more effectively spent targeting problems of the world’s poorest people. This is

why all of the recommended charities target extremely poor members of some of the world's poorest societies. Additionally, resources are more effectively spent on extremely harmful problems that can be alleviated relatively cheaply than on less harmful problems or problems that are more costly to alleviate. This explains why nearly all recommended charities (Give Directly is the exception) focus on basic health and medical problems. Some global health issues that reduce significant years (e.g. malaria) and quality (e.g. tropical parasites) from people's lives can be reduced and alleviated through strongly evidence-supported measures. Funding is the main bottleneck limiting how much of an impact these organizations are able to have. Thus, marginal dollars donated to these organizations contribute significantly to improving human welfare by averting deaths and increasing the wealth and income of some of the world's poorest people.

Effective altruism is a framework for thinking about CSR, not an algorithm. As I have indicated, effective altruists could weight principles of effective altruism and principles exterior to effective altruism differently and thus come to different decisions about how CSR ought to be pursued in a given case. What is required is a commitment to seriously considering which causes are the most important, and which means can address them most effectively, when choosing how to implement CSR initiatives.

d. The Easy Baseline Principle

Though I refrain from discussing too many details about implementing effective altruist thinking in CSR decision making, it is worth discussing one concrete prescriptive

effective altruist principle for evaluating and making decisions about CSR. This shows how effective altruist CSR would likely differ compared to other approaches to CSR decision making.

Direct cash transfers represent one of the simplest forms of charity. Donors simply give money to recipients. There is also strong evidence that cash transfers, like the ones managed by the GiveWell-recommended charity GiveDirectly, are among the most effective ways of significantly improving the lives of some of the world's least well off (Kapur, Mukhopadhyay, & Subramanian, 2008; MacAskill, 2015). Research shows that cash transfers have significant, long-term, positive effects on recipients' income: one study found that, five years after receiving cash grants, men in Sri Lanka enjoyed annual incomes between 64% and 96% greater than the original grant amount (De Mel, McKenzie, & Woodruff, 2012). Other research shows that cash grants can have large positive impacts on children's health (Gertler, 2004). Despite fears that poor recipients of cash grants will waste their money on gambling, alcohol, tobacco, and other drugs, studies have failed to find a positive correlation between receiving cash transfers and spending on these sorts of goods (Haushofer & Shapiro, 2013). Studies also have largely failed to confirm the related fear that cash transfers will reduce recipients' motivation to work: indeed, cash transfers often increase working hours by allowing recipients to move to places with better jobs (Blattman, Fiala, & Martinez, 2013).

Direct cash transfers, then, are a highly effective and simple way of making a large social impact. Thanks to organizations like Give Directly, anyone, including any firm, with resources to spare for charitable causes can do a huge amount of good by

giving to organizations that effectively facilitate direct cash transfers to the very poor. The simplicity, accessibility, and effectiveness with which organizations like Give Directly promote an indisputably important cause—contributing to alleviation of severe poverty by giving poor families the resources they need to raise themselves out of their penury—makes direct cash transfer charities a good fit for their role in what I will call the Easy Baseline Principle. As a heuristic, it will be useful to have an effective altruist baseline against which to weigh potential CSR initiatives. Any such baseline will be somewhat arbitrary by necessity. But fixing one is useful enough that we should do so anyway.

*Easy Baseline Principle*²: (1) The default form of CSR for an effective altruist firm should be donating to Give Directly or a similarly effective cash transfer charity; (2) a good reason is necessary to justify deviating from the Easy Baseline.

(1) is both easy to fulfil and extremely restrictive. It is easy to fill because any firm (or individual) with resources to donate is capable of fulfilling it. It is restrictive because it sets an extremely high bar for effectiveness.

It seems unlikely that many actually-existing activities that we would generally take to qualify as CSR³ could fulfil it, given how effectively direct cash transfers promote

² This principle is inspired by an article by Benjamin Todd, the founder of the effective altruist organization 80000 Hours: <https://80000hours.org/2015/12/whats-the-easiest-way-anyone-can-have-a-big-social-impact/>

³ For the purposes of this paper, I adopt a definition of CSR that is common in the management literature (Barnett, 2007a; Mackey, Mackey, & Barney, 2007; Marquis, Glynn, & Davis, 2007; Rupp, Ganapathi, Aguilera, & Williams, 2006). CSR refers to activities that are undertaken voluntarily (rather than in response to, e.g., a legal requirement) and appear to promote an apparently ethical or social purpose beyond firm financial performance or (2)

an undeniably worthy cause. Most real world donative CSR programs would not fulfil it; certainly not giving away shoes, as Toms famously does. Non-donative forms of CSR are even more hopeless when compared to this baseline. Consider CSR that takes the form of a firm treating its employees fairly and generously, as Costco has gained a reputation for doing. Consider CSR that takes the form of atoning for some past ethical transgression, such as Starbucks closing its stores for a day of racial bias training following its April 2018 incident (more on these below). It appears difficult for an effective altruist manager to justify doing any of these, or countless similar CSR examples, rather than the Easy Baseline Principle's default. This appears extreme. It might make effective altruism seem less appealing, even to those who agree that, as a general matter, we have reason to do things that contribute to bringing about more good and a greater social impact rather than less.

This objection contains an important kernel of truth, but also does not provide persuasive reason for rejecting effective altruism as a plausible, well-justified approach to making and evaluating CSR decisions. In brief, my response to the objection is that, to construct an ethics-first framework for CSR decision making, we must distinguish between three importantly different categories of CSR: (1) preventative and restorative CSR, (2) beneficent CSR, and (3) reciprocal CSR. I argue that effective altruism provides a plausible framework for beneficent CSR, but not the other categories of CSR. However,

appear to reflect a commitment on behalf of the firm to that ethical or social purpose (see the first essay of this dissertation).

before discussing that, it will help to tie together the above discussion of effective altruism and the RBV with a case.

V. TOMS Shoes as a CSR exemplar

Some examples of CSR become canonized in popular and academic writing, as well as business school classrooms. One business model that has gained significant buzz in this area is the Buy-One Give-One model most famously associated with the shoe company Toms and the eyewear company Warby Parker. Firms that employ the Buy-One Give-One model typically make a donation for each commercial sale. Toms shoes, for example, gives away one shoe to a child in need for each shoe it sells.

It should be clear, from the discussion in the previous two sections, that Toms Shoes makes a good deal of sense from an RBV-derived, performance-first framework for CSR decision making. Marquis & Park (2014) explain how Toms serves as a paradigmatic example of an integrative approach to CSR on the grounds that it links some of its core capabilities—relating to the design and manufacture of shoes—to the means by which Toms has its social impact. This creates a convergence between Toms' social goals and its economic goals. For example, improvements in its shoe design can result in both a better product that customers are willing to purchase at a higher price and a more functional item for recipients of donated shoes. The connection between Toms' social impact and a customer's decision to buy a pair of Toms shoes is both clear and simple, making it well-suited for an effective marketing message. To the extent to which

the capabilities underlying Toms' competitive advantage are difficult for competitors to build, acquire, or substitute for, the RBV considerations underlying the strategy-first framework might help companies like Toms achieve and sustain competitive advantage. In dispute, of course, whether this factor possesses much relevance to ethical decisions.

Toms' Buy-One Give-One program for shoes appears much more dubious when we examine it using an ethics-first framework. Effective altruists will generally applaud Toms' goal of improving the condition of children in need. However, they will question the means Toms employs to achieve that goal. Toms' shoe donation program has in fact been the subject of rigorous academic study, and the results were not promising.

Wydick, Katz, Calvo, Gutierrez, & Janet (2016) randomly assigned communities in El Salvador to either treatment groups, in which children received shoes from Toms, or control groups that did not receive shoe donations. Several months after the shoes were donated to the children in treatment group communities, the authors failed to find statistically significant differences between treatment and control groups *even for reductions in shoelessness*. The authors similarly found no evidence that the shoe donation program had any effect on overall health, foot health, or self-esteem, though they did find that boys who received donated shoes had slightly better school attendance records than boys who did not receive donated shoes (Wydick et al., 2016b). Of course, this is just one study, and it is possible that other studies in different contexts would find a greater impact. However, given the strong evidence supporting direct cash transfers as a relatively effective way of making a social impact (Blattman & Niehaus, 2014b), this troubling evidence about Toms suggests that, if we care about the potential of CSR to

make a positive social impact, rather than just contribute to improved firm financial performance, we should not embrace the Buy-One Give-One model as an exemplary model of CSR and social enterprise without seriously scrutinizing the impact of specific Buy-One Give-One programs.

My purpose here is not to vilify Toms. People, and organizations, make mistakes. Indeed, I believe that Toms deserves great credit for funding and supporting Wydick et al. (2016) study, despite its results seriously calling into question the program that put Toms on the map as a social enterprise in the first place. This suggests that Toms' management is not engaging in CSR disingenuously (see ch. 1 of this dissertation), but actually takes CSR seriously from the perspective of making a social impact. I believe we must all take the social impact of CSR similarly seriously, making and evaluating CSR decisions using an ethics-first framework, like effective altruism, rather than a financial performance-first framework, like the RBV. This could significantly increase the amount of good accomplished through firms' CSR activities, and lead to business doing a better job of contributing to making society better off and the world a better place, rather than just exploiting society and the world for the benefit of its shareholders (Jones et al., 2016; Jones & Felps, 2013).

VI. Effective altruism and CSR: some qualifications

The most committed effective altruist would likely hold that all CSR decisions should be based primarily or exclusively upon effective altruist principles, and that if this

rules out common but not maximally effective ways of doing CSR—e.g. Costco’s fairness and generosity toward its employees, Starbucks shutting down all of its stores for a day to implement racial bias training—then effective altruism requires that firms cease doing those forms of CSR.

As indicated above, I disagree with this more radical effective altruist approach to CSR. Although I strongly believe that effective altruist ideas and principles are in many cases vital for understanding how CSR decisions should be made, I also believe that they are generally only directly pertinent to what I will call beneficent CSR: CSR whose primary purpose is to do good or make the world overall a better place. But beneficent CSR is not the only category of CSR. There are at least two other important categories of CSR: preventative and restorative CSR (which, for reasons I explain below, I combine into a single category), and reciprocal CSR. The primary purpose of preventative and restorative CSR is preventing harms and rights violations, or compensating and atoning for past harms and rights violations. The primary purpose of reciprocal CSR is to establish and respect relations of reciprocity.

I do not use the term “primary purpose” to refer to the subjective goals of the decision makers who carry out CSR. Rather, I use the term to refer to the (objective) normative reasons for action that apply to a given situation, reflecting what an ethical actor should do in that situation (Alvarez, 2016). A proper understanding of the normative reasons for action underlying the various categories of CSR I identify will, as I show, provide more specific guidance about how CSR decisions within each category should be made and evaluated.

A proper understanding of the normative reasons for action underlying the categories of CSR will also be necessary for making sound decisions between the categories. For example, a firm may face a choice of whether to devote its CSR resources to preventative, reciprocal, or beneficent CSR. If devoting the resources to preventative CSR is necessary to keep the firm from committing a serious harm or rights violation, the firm will likely be morally obligated to use the resources on an effective preventative CSR initiative, given the moral gravity of committing serious harms or rights violations. Alternatively, there may be cases where beneficent CSR is morally required, e.g. when a firm has a unique capacity to rescue victims from a moral catastrophe (Dunfee, 2006). Ethical reasons must inform decisions about how to do CSR both within and between the three categories.

a. Preventative and restorative CSR

In April of 2018, Starbucks found itself at the center of a scandal after a viral video showed police, who had been called by a Starbucks store manager, arrested two black men in a Philadelphia Starbucks who video showed waiting peacefully at a table for a friend to arrive. The video exemplified the persistent racial inequities, especially those connected to the criminal justice system, that have plagued the U.S. throughout its history. Starbucks was perceived as complicit in perpetuating racial injustices due to its store manager's decision to call the police even though the two men who were arrested did not seem to be acting inappropriately, and certainly were not acting inappropriately

enough to justify involving the police. In the wake of the incident, the CEO of Starbucks issued a public apology to the two victims of the incident. Starbucks also closed 8000 of its U.S. locations for one day to implement racial bias training (Adams, 2018).

Any resources Starbucks spent on its public apology and racial bias training almost certainly fall short of the test suggested in the Easy Baseline Principle by bringing about significantly less good than simply donating to an effective direct cash transfer charity. Does this mean Starbucks made the wrong decision when it decided to undertake these CSR activities?

In my view, not necessarily. Starbucks, assuming it qualifies as a moral agent, has moral reasons to refrain from committing serious harms or rights violations (Mill, 1978; Nozick, 1974). These are weighty moral reasons—refraining from seriously harming or violating the rights of others represents one of the most basic kinds of moral obligations we, as moral agents, generally take ourselves to have. Firms thus can have strong moral reasons to invest in preventative CSR—CSR whose justifying purpose is the avoidance of serious harms and rights violations. An agent who has committed harms or rights violations in the past can face similarly strong moral reasons to compensate and atone for its past wrongdoing. Firms can thus face strong moral reasons to invest in restorative CSR—CSR whose justifying purpose is to atone for, and compensate victims of, a firm’s past wrongful acts.

The prevention of harms and rights violations involves deontic moral reasons (Nagel, 1986). Preventative and restorative CSR are thus sensitive to the identity of the

firm carrying out the CSR (as well as the identify of potential or actual victims of the firm's past wrongs) in a way that beneficent CSR is not. The Starbucks example provides an intuitive illustration of this difference. The purpose of Starbucks' racial bias training was to prevent further racist incidents in *Starbucks' own stores*, perpetrated by *Starbucks' own employees*, rather than to prevent further racist incidents in general. Starbucks' apology only made sense given a relational understanding of Starbucks' wrongful conduct: *Starbucks* making amends for its past wrongful act against *the two men who were arrested* (Radzik, 2004). Starbucks, because of its identity, history, and relationship with the two men who were arrested in its Philadelphia store, has a special reason to make amends to those men that does not apply to, say, Dunkin' Donuts or Tim Hortons.

Effective altruists evaluate potential actions mainly using effectiveness criteria. Given the amount of resources required, how much good would be produced by selecting course of action A versus the alternatives of B or C? To be sure, effectiveness considerations still bear some relevance to preventative and restorative CSR decisions. For racial bias training to be appropriately responsive to the moral reasons that applied to Starbucks, the training must be effective. Minimally, this requires using best practices based on the best available evidence about how to reduce racial bias. Reducing racial bias is a difficult task. In a recent meta-analysis of research on procedures that change racial bias, Forscher et al. (2016: 2) find that, although some procedures have a small effect on measures of implicit racial bias, these procedures "changed explicit measures less consistently and to a smaller degree than implicit measures and generally produced trivial changes in behavior." Even when anti-racial bias procedures have a measurable effect on

some kinds of individual racial bias, “those changes do not necessarily translate into changes in explicit measures or behavior” (Forscher et al., 2016). There is even a danger that racial bias training can inadvertently promote racial stereotyping and cause people to act in ways that reflect greater racial bias (Duguid & Thomas-Hunt, 2015).

Taking effectiveness seriously is necessary in order to take appropriate action in response to the moral reasons underlying preventative and restorative CSR. But the relevance of effectiveness is limited. Suppose a manager must choose between three alternative courses of CSR action, A, B, and C. A involves taking measures necessary to prevent the manager’s firm from causing severe harms or rights violations (or, alternatively, A involves taking measures necessary to atone for past wrongs the firm has committed and compensate the victims of those wrongs). Assume B and C are, by relevant measures, more effective than A: they bring about a greater amount of good relative to the resource investments they require. But the effectiveness of B and C comes at the cost of failing to prevent harms or rights violations (or failing to atone for and compensate victims of past wrongs). Absent unusual circumstances, a firm engaged in preventative or restorative CSR cannot justify doing B or C rather than A, even if B and C are more effective.

This discussion goes to show that preventative and restorative CSR reflect moral reasons the fulfillment of which depends on factors that differ from the main moral considerations of effective altruism. Prescriptive guidance for preventative and restorative CSR requires looking elsewhere. For preventative CSR, research on human rights (Beitz & Goodin, 2009; Caney, 2010), as well as the associated duties of

businesses (Arnold, 2010; Wettstein, 2012), will be relevant, though preventative CSR of course applies in local and national contexts, as well as the international contexts that are usually the focus of human rights discussions. Restorative CSR requires understanding what must be done to atone for past wrongs and what duties are owed to victims of past wrongs. Academic research relevant for developing a prescriptive framework for restorative CSR includes literature on reconciliation (Griswold, 2007; Radzik & Murphy, 2015) and corrective justice theory (Coleman, Hershovitz, & Mendlow, 2015; Weinrib, 2012).

b. Reciprocal CSR

Costco has developed a reputation for CSR thanks in large part to its famously generous labor practices (Cascio, 2006a) (which are often compared favorably to the stingy labor practices associated with Walmart's Sam's Club, Costco's main rival). Costco pays its employees a high hourly wage and provides generous health care and retirement benefits (Cascio, 2006b). Clearly, Costco employees benefit from these policies, at least compared to less-generous alternatives. But Costco benefits as well, because these employment practices are an integral part of Costco's overall human resources strategy, which depends on securing a loyal, productive workforce with low turnover rates (Cascio, 2006a).

As with preventative and restitutive CSR, effective altruism seems ill-suited as a framework for making good reciprocal CSR decisions. To be sure, effectiveness matters

for reciprocal CSR. In Costco's case, doing reciprocal CSR well requires understanding what Costco employees value and successfully providing them with employment policies and benefits that match what they value. But effectiveness is not among the main considerations that should be used to choose between different potential courses of reciprocal CSR action. Costco's generous employment policies likely fall short of the standard in the Easy Baseline Principle. Any resources Costco spends on its employees would almost certainly do more good if they were donated to an effective direct cash transfer charity. But although Costco may have reasons of beneficence to do what would be best from an impartial point of view, these reasons do not necessarily outweigh its reasons to return benefits to the partners with whom it shares reciprocal relationships, or its reasons to foster reciprocal relationships in the first place.

To understand the moral reasons underlying reciprocal CSR requires a basic understanding of reciprocity. Reciprocity refers to the idea that when one party A benefits from interacting or relating with another party B, it is at least a good thing for A to return some of that benefit to B (Schmidtz, 2006). This understanding of reciprocity illustrates how the relationship between Costco and its employees is a reciprocal one: Costco benefits its employees significantly more than it would if it treated its workforce more opportunistically, and in turn Costco's employees contribute to Costco's bottom line significantly more than if they treated Costco more opportunistically (Jones, 1995; Williamson, 1979). Thus, reciprocal CSR principally involves (1) properly returning some of the benefits derived from ongoing reciprocal relationships and (2) seeking out opportunities to develop new reciprocal relationships.

Of the three categories of CSR I discuss, reciprocal CSR is by far the most theorized in business and management literature. This may be a reflection of the fact that reciprocal CSR can easily be assimilated into the dominant logic that business decisions ought to be justifiable in terms of bottom-line financial performance.

For example, instrumental stakeholder theory essentially elaborates how reciprocal relationships between a firm and its stakeholders can function as a source of competitive advantage (Donaldson & Preston, 1995; Freeman, 1984; Jones, 1995). Creating Shared Value, “defined as policies and operating practices that enhance the competitiveness of a company while simultaneously advancing the economic and social conditions in the communities in which it operates” (Porter & Kramer, 2011: 66), is nearly definitionally equivalent to reciprocal CSR (de los Reyes Jr, Scholz, & Smith, 2017). Bottom of the Pyramid strategies are rooted in the idea of reciprocity, urging managers and entrepreneurs to search for competitive advantage in yet-to-be developed markets by designing products aimed specifically at meeting the needs of consumers at the bottom of the global wealth pyramid (Prahalad, 2006).

The above perspectives are all valuable (at least to varying degrees). For business to realize its full potential to promote the interests of stakeholders and society, managers must think seriously, rigorously, and creatively about how to find and foster reciprocal relationships for their firms. Scholarship that contributes to our understanding of these reciprocal relationships—improving our understanding of how they work, which factors contribute to their failure or success, how we should locate them in our overall theoretical conceptualization of business and CSR—is valuable. However, as my categorization of

CSR demonstrates, we must be careful about equating reciprocal CSR with CSR more generally. We must be particularly careful not to allow the ideology of shareholder primacy to blind us to the fact that there are other domains of CSR that cannot be understood through the lens of mutual benefit. A related observation: we must not assume that a firm that only does reciprocal CSR is fulfilling all of its ethical and social obligations. (More on this below.) These truths can be obscured when advocates of reciprocal CSR endorse positions like the following: “Creating shared value (CSV) should supersede corporate social responsibility (CSR) in guiding the investments of companies in their communities” (Porter & Kramer, 2011: 76). Reciprocal CSR is important and worthy of attention. But there are other important categories of CSR as well, and they cannot plausibly be reduced to considerations of mutual benefit.

c. Beneficent CSR

I have argued that effective altruism does not serve as an adequate basis for an overall framework for CSR decision making because some categories of CSR—what I have called preventative/restorative CSR and reciprocal CSR—are backed by moral reasons against engaging in, and failing to make proper amends for, serious harms and rights violations, and by moral reasons in favor of returning benefit to those with whom one shares reciprocal relationships. Understanding the moral reasons that underlie the different categories of CSR gives us a basic picture of how CSR decisions within (and between) the categories should be made and evaluated.

Unlike preventative/restorative CSR and reciprocal CSR, effective altruism does provide a plausible basis for making and evaluating beneficent CSR decisions. Effective altruism, as I have explained it, urges use of the best available evidence and analysis to increase the beneficial impact one has through one's altruistic actions. Beneficence refers to actions, norms, and dispositions whose goal is to promote the good of other people (or, arguably, sentient beings more generally) (Beauchamp, 2019). Thus, effective altruism highlights the same underlying moral reasons as those favoring beneficent action.

Admittedly, the claim that effective altruism provides plausible guidance for beneficent actions is somewhat controversial. But I want to argue that it is not as controversial as some might think. Effective altruism as an approach to beneficent action is controversial because it involves at least a minimal amount of moral aggregation. When choosing between alternative courses of beneficent action, effective altruism tells us to estimate how much good is accomplished by each alternative, and to choose the alternative that (holding costs constant) produces the most good. We must assume that we can intelligibly compare the amount of good that each potential course of action accomplishes. But some moral philosophers deny this (Taurek, 1977; Wiggins, 1998). Beneficence, according to these philosophers, is not about producing as much good as possible. Instead, the "beneficent person is one who helps X, or rescues Y, or promotes this or that cause and does so *because* each of those thing[s] is in its own way an important and benevolent end" (Wiggins, 1998: 274).

I want to highlight some unattractive features of such radical anti-aggregation views of beneficence. First, they have some highly counter-intuitive implications. Our

intuitions about the extent to which numbers should count in our moral judgments may be somewhat hazy when we are faced with various versions of, for instance, the famous trolley problem dilemma, where one must choose between saving/killing an individual versus saving/killing a group of several people (Foot, 1967). But they become clearer with larger numbers. Suppose a trolley is driving toward a trip wire that, if activated, will trigger a series of nuclear explosions that will kill nearly all of the 7 billion human beings who inhabit the earth. You stand at a switch. If you pull the switch, it will divert the trolley onto a track with one unlucky innocent person strapped across it. This person will die if the switch is pulled. (Suppose that this person somehow would be spared from the ensuing nuclear apocalypse if the trolley hit the trip wire.) It seems clear to me that the right thing to do in such a case would be to pull the switch. The explanation for why this is the right thing to do is also clear: morally speaking, the numbers count. (They may not be determinative, but they do count.) Radical anti-aggregationists will bite the bullet and insist that, because the numbers do not count, the person at the switch should not divert the trolley, or should decide whether to divert on the basis of some fair procedure (e.g. rolling dice) that gives the unlucky innocent person a chance at being spared. This, in my view, is a significant bullet to bite.

Second, radical anti-aggregationist views fail to provide plausible guidance for beneficent actions. There are many potential ways for a person or firm to be beneficent. One might be beneficent by volunteering at one's local food bank, helping Habitat for Humanity construct affordable houses, donating money to an anti-poverty charity, and so on. Radical anti-aggregationists would have us believe that we generally lack a rational

basis for choosing one of these beneficent acts over another. But that does not seem right. Medical professionals working in disaster areas or war zones often must engage in triage to ensure that their limited resources go to the most pressing causes. Doctors and nurses try to avoid treating mild coughs at the expense of treating heart attacks, and avoid using all of their time and resources treating a single person while many untreated others suffer or die around them (MacAskill, 2015). Do we really want to say that medical professionals' in such cases lack a rational basis for their decisions? Do we really want to say that there is no way of making normative distinctions between available courses of beneficent action? Again, some will happily bite this bullet and respond in the affirmative. I, on the other hand, take these implications to cast significant doubt upon radical non-aggregationist moral theories.

Acceptance of moral aggregation is generally associated with consequentialist moral theories. It may thus be tempting to associate an effective altruist approach with moral consequentialism. However, I believe it is a mistake to equate effective altruism with moral consequentialism, especially given the overall view that I offer in this paper. First, consequentialism is typically a theory of morality in general: roughly, an act is right if and only if it brings about the best possible consequences. But the view I advocate here limits the proper domain of effective altruism to that of beneficent action. One reason to reject consequentialism is that it famously cannot account for deontic constraints such as those generated by moral rights. Consequentialists struggle to explain why it would be wrong for a doctor to abduct the postal worker delivering packages to the hospital and then kill him so that his organs can be distributed to five patients who otherwise would

die. But this sort of case does not pose a problem for the view I am advocating. Firms have obligations of preventative CSR not to seriously harm others or violate their rights. Serious harms and rights violations are prohibited even in cases where violating rights would allow one to do more good. Effective altruist principles are straightforwardly applicable only for choosing between beneficent acts, not for making moral decisions that hinge upon moral considerations other than beneficence.

Effective altruism encourages us to make a serious effort to do more good through our beneficent acts. The most influential versions of consequentialism, to once again put it in somewhat rough terms, demand that we devote all of our acts to doing the most good possible. There is an important difference between these stances. The former could, as I have tried to argue, be a plausible view even to those who do not harbor a deep commitment to consequentialism. Skeptics of consequentialism primarily criticize the latter sort of view, not the former. It would be a mistake to equate the two views and understand effective altruism to involve much more radical moral commitments than it actually does.

Return now to the case of Toms Shoes and its Buy-One Give-One model for beneficent CSR. The Easy Baseline Principle, the tenets of effective altruism more generally, and the empirical evidence about the social impact of giving away footwear to children in need should cause us to reject Toms footwear initiative as an exemplar of beneficent CSR. The evidence is overwhelmingly strong that Toms could have done significantly more to help the global poor if it had simply donated a portion of its profits to an effective direct cash transfer charity instead of giving away footwear. Business

scholars should not continue to encourage relatively ineffective forms of beneficent CSR by continuing to play up the dubious success of Toms.

VII. Conclusion

CSR refers to a broad category of corporate activity. It is thus not surprising that a plausible framework for making and evaluating CSR decisions cannot be reduced to a single ethical principle, consideration, or theoretical perspective. I have tried to show that it is nevertheless possible to specify several ethical considerations that seem especially pertinent to CSR decisions: refraining from serious harms and rights violations, reconciliation for past wrongful acts, reciprocity, and beneficence. Normative theory about CSR can make further progress by clarifying our understanding of how the moral reasons arising from each of the above considerations interact: in which cases do reasons from one of these moral considerations outweigh or exclude reasons from the others? Future research can also contribute to normative theory about CSR by shedding further light on how decisions within the CSR categories I introduce should be made and evaluated. Finally, normative theory about CSR would potentially benefit from consideration of ethical perspectives and factors that my account does not consider. CSR is difficult to understand from a normative point of view, and this paper is obviously not the last word on the subject. It is clear, however, that normative research on CSR requires engagement with normative ethical theoretical perspectives; financial performance-focused management theories such as the RBV cannot serve as the primary basis for

prescriptive guidance about CSR (see Donaldson, 2012). Understanding the normative aspects of the social responsibilities of business firms requires grappling with normative ethics.

I. Introduction

Debate in the field of business ethics has historically centered around the question of whose interests ought to receive priority in managerial decision making. Advocates of shareholder primacy, such as Friedman (1970) and Jensen (2002), argue that the interest of shareholders ought to be prioritized over those of other corporate constituency groups. Advocates of stakeholder theory, most notably R. Edward Freeman, argue that managers ought to weight the interests of non-shareholder constituency groups, such as customers, suppliers, and employees, at least as much as the interests of shareholders (Freeman, Harrison, Wicks, Parmar, & De Colle, 2010).

Although this debate about corporate purpose was (and remains) an important one, it diverted the field's attention away from adjacent questions of equal theoretical and practical importance. One of these questions is: setting aside the issue of who the beneficiary of a firm's profits ought to be, how can we distinguish between ethical and unethical ways of pursuing profit in a market? After all, nobody—not even the most dyed-in-the-wool shareholder primacist—believes that firms should profit by murdering and pillaging. There must be some ethical constraints on how firms may seek profit. But what are they? Surprisingly little work in twentieth-century business ethics addressed this question in a systematic way (McMahon, 1981 is an important exception; see also Moriarty, 2019).

This started to change shortly after the turn of the millennium, when Joseph Heath began to develop his Market Failures Approach to business ethics (which I will refer to as the MFA) (Heath, 2004, 2006). In my view, one of Heath's most valuable contributions to the field has been his success at getting business ethics scholars focused on the previously-neglected question of how to differentiate between ethical and unethical tactics for pursuing profit. His work has produced a burgeoning literature that addresses this under-theorized topic (Heath, 2014; Jaworski, 2014; Norman, 2011; Singer, 2018; von Kriegstein, 2016).

Heath deserves great credit for getting business ethics to focus on an important, neglected question. However, as I will argue in this paper, Heath's own answer to this question, which is contained in the MFA, is unsatisfactory. I am not the first to criticize the MFA: scholars have advanced external critiques of Heath's work, objecting to the MFA's narrow focus on economic efficiency and its neglect of other seemingly important ethical values (Hsieh, 2017; Martin, 2013; Silver, 2016; Singer, 2018).

My critique of the MFA, however, differs from many existing ones because it is internal, rather than external. I argue that the MFA fails on its own terms to provide a plausible, clear answer to the important question Heath raises of how we should distinguish between permissible and impermissible (or, to use Heath's terms, preferred and non-preferred) tactics for pursuing profit in a market. In recent work, Heath has helpfully clarified that the MFA should be understood as Paretian: it identifies the 'point' of economic market institutions to be the achievement of Pareto efficiency, and uses this to justify restrictions on how market actors may seek profit (Heath, 2013, 2014). I argue,

however, that Heath's emphasis on Pareto efficiency is misguided. Pareto efficiency cannot explain our basic intuitions about which sorts of profit-seeking tactics are permissible or impermissible, and it does not provide plausible guidance about the permissibility of tactics about which intuitions fail to provide clear judgment.

The paper does not just advance a negative criticism of the MFA. It also offers a solution (or at least the beginnings of one) to the problems it highlights. I propose that it is not Pareto efficiency, but rather the creation of wealth, that we should look to in order to understand some basic constraints on how market actors may pursue profit. My claim is not that wealth is the single unifying value that can explain everything about the ethics of business and markets—I am skeptical that any such value exists. Still, wealth is a significant value for business and market ethics, and it plays a large role in our understanding of the ethical constraints on pursuit of profit in market contexts. I defend a practical Wealth Creation Principle, which holds that firms may not employ combinations of strategic tactics that appropriate more wealth than they create on net. The upshot of the principle is that market actors should seek profit by engaging in productive, rather than unproductive or destructive, entrepreneurial activities (Baumol, 1996).

The paper proceeds as follows. Section II differentiates between Easy Cases (where intuition provides clear, plausible answers) and Hard Cases (where intuition fails to yield clear, plausible answers) concerning tactics for pursuing profit in the market. A good theory should cohere with our intuitions about Easy Cases (or, if it diverges with our intuition, provide compelling reason for why our intuition should be revised) and help us think through Hard Cases where our intuitions give out. Section III presents the MFA,

and identifies a principle that the MFA provides for distinguishing between ethical and unethical (or preferred/non-preferred) tactics. I call this the Perfect Competition Principle. Section IV argues that the Perfect Competition Principle is inadequate because it has implausible implications about Easy Cases and fails to illuminate Hard Cases. It also explains more generally why the MFA goes astray by placing so much emphasis on Pareto efficiency: this notion cannot explain the desirability of competitive tactics that disrupt the equilibrating processes of the market. I invoke rival conceptions of entrepreneurship from Joseph Schumpeter and Israel Kirzner to help illustrate this point. Section V explains my positive proposal, which is to replace Pareto efficiency with wealth creation as a guiding value for market competition. I tentatively propose a conception of wealth based on the Kaldor-Hicks standard from economics, as well as a Wealth Creation principle that places a strong ethical presumption against the use of profit-seeking tactics that appropriate more wealth than they create on net. Section VI argues why, despite reservations of skeptics such as Dworkin (1980) and Jones & Felps (2013), we should view wealth creation as ethically important, especially in the domain of business ethics. I argue that we have ethical reason as a society to pursue greater wealth because increased wealth constitutes increased positive freedom. Section VII concludes.

II. Ethical constraints on pursuing profit in a market

Consider the following tactics that market actors might employ in pursuit of profit. I will refer to these as Easy Cases because we generally have clear, plausible, intuitive judgments about their permissibility. These judgments may not be completely ironclad, but we would need to be provided with strong reasons to seriously consider revising them, given their intuitive plausibility. To borrow a metaphor from Quine & Ullian (1978), they are closer to the center than the periphery of our web of ethical beliefs.

- a) A firm invests in new manufacturing technology that lowers costs by improving production efficiency.
- b) A firm redesigns one of its products in order to improve quality and increase customer willingness to pay.
- c) A firm attempts to tacitly collude with its competitors by sending signals that it will follow its rival's lead setting prices in various local markets.
- d) A firm attempts to deter entry into the market through a form of contrived deterrence: it seeks profit by investing in excess manufacturing capacity—beyond the manufacturing capacity it would build if it were simply trying to invest up to the point at which marginal cost equals marginal revenue—with the goal of credibly threatening to flood the market with products and drive down price in the event of entry.

Intuitively, both (a) and (b) are desirable: from the perspective of society overall, businesses competing with each other by improving production efficiency or product quality represent instances of the market working as it should, and of market actors behaving as they should. Equally intuitively, (c) and (d) are not desirable: from the perspective of society overall, businesses competing with each other through tacit collusion or contrived deterrence represent instances of the market failing to work as it

should, and of market actors behaving poorly. We can arrive at these judgments without the aid of an elaborate theory. Intuition gets us most of the way there.

Some cases, however, are not so easy. Call them Hard Cases. Consider:

Small Town Stores

Suppose a chain of stores specializes in a category of products that is generally not available in physical stores in small towns. This chain's management perceives that there is an untapped market opportunity in relatively small towns for the kind of stores in which it specializes—perhaps rival firms in the industry have wrongly judged that towns below a certain size couldn't support this type of store, given the minimal scale required to run the store efficiently. Management thus decides to quickly expand into small towns, rapidly opening locations in all towns it judges to be promising locations. This tactic pays off handsomely: there had been significant unmet demand for physical stores selling this category of products in these towns. In addition to being viable businesses in small town locations, these stores turn out to be sources of sustained competitive advantage for the chain precisely because of geographic features of their small-town locations: most of the towns only have enough population to support one store (which must be relatively large due to scale efficiencies). The chain's stores thus enjoy an advantageous competitive position due to a barrier to entry: their first-mover and incumbency advantages, combined with the towns only being able to support one store, makes building competitor stores an unattractive investment for the chain's rivals. The reduced competitive pressure resulting from the barrier to entry allows the chain to charge somewhat higher prices for its products than it charges in more competitive environments. It appears that the chain will reap enormous profit margins from these locations for the indefinite future.

How should we understand the ethical responsibilities of firms that occupy favorable competitive market positions such as this? On the one hand, it might seem that above-normal returns are the just rewards for a firm that has successfully out-competed rivals to occupy an attractive market position. It is not the fault of the chain in Small Town Stores that some of the factors contributing to the barrier to entry shielding it from competition—scale efficiencies, the small towns' relative geographic isolation—are

present. Indeed, the chain deserves some credit for having been the first among its rivals to identify the business opportunity that these small-town locations offered. Plus, the general social and economic effects of its move are (let us assume) positive, and its customers shop there because they correctly perceive that the store offers products at a good value that make their lives better. These factors indicate that the profit motive in this case is doing precisely what we would want it to do.

On the other hand, it also seems that there must be some limits on how firms in competitively advantageous positions are permitted to protect those positions and exploit them for profit. One justification for the practice of setting prices at a profit-maximizing level in competitive markets is that if other producers are capable of supplying an equivalent product for a better value, then customers are free to switch. But this justification is not available (or is at least deprived of some of its force) in cases like Small Town Stores. The barriers to entry impede the competitive market forces that bring about conditions in which consumers can easily defect to rival sellers in response to excessive price increases.

For these Hard Cases in which our intuitions fail to provide clear guidance, it would be helpful to have theoretical principles to aid us in identifying ethically salient considerations. Thus, to summarize, two important desiderata for an adequate theory of constraints on pursuing profit are (1) coherence with (or persuasive justification of deviations from) basic intuitions regarding Easy Cases, and (2) plausible guidance for our ethical judgments about Hard Cases.

III. The Market Failures Approach to business ethics

Joseph Heath presents his Market Failures Approach to business ethics (Heath, 2014) as providing just the sorts of principles that it seems we would need for understanding constraints on pursuing profit in a market. Heath introduces the terms *preferred* and *non-preferred strategies* to refer to the specific constraints on ethical pursuit of profit that the Market Failures Approach provides. “Under conditions of ‘perfect competition,’” he writes, “lower price, improved quality and product innovation would be the *only* way that firms could compete with one another. We can refer to these as the set of *preferred* competitive strategies” (Heath, 2006: 549). However, because markets in the real world at best approximate and frequently depart significantly from the conditions of perfect competition (e.g. many buyers and sellers, homogeneous goods, perfect access to information), “firms are often able to make a profit using *non-preferred* competitive strategies, such as producing pollution, or selling products with hidden quality defects” (Heath, 2006: 550).

The question, then, becomes: what principle or set of principles does the Market Failures Approach (MFA) provide that can help us understand more specifically and precisely which sorts of tactics for pursuing profit are among the preferred, and which are among the non-preferred?

One interpretation of the above quoted passage from Heath is that lower price, improved quality, and product innovation are the *only* preferred tactics (I use the term

‘tactic’ where Heath uses ‘strategy’). But that seems unduly restrictive. Does Heath believe that firms are prohibited from competing with each other by hiring away each other’s employees or developing more efficient production technologies? Presumably not. Hence, it is best to understand lower price, improved quality, and product innovation as *examples* of preferred tactics, but not as *exhaustive* of the complete set of preferred tactics. Therefore, to understand how the MFA distinguishes between preferred and non-preferred tactics, we will need to look elsewhere.

a. Pareto improvement

In recent refinements of his work on business ethics, Heath has insisted that his approach should be understood as “Paretian” (Heath, 2013: 50, 2014: 5). Heath sometimes even appears to suggest that what is Paretian about his approach is that it is built upon the principle of Pareto improvement: “the principle that, whenever it is possible to improve at least one person’s condition without worsening anyone else’s, it is better to do so than not” (Heath, 2014: 9–10). To avoid confusion, however, it is important to understand that Heath himself explicitly rejects Pareto improvement as the standard that market actors’ actions must meet in order to be judged desirable or permissible according to the MFA. He is clearly right to do so—requiring that market actors’ actions meet a Pareto improvement standard would render impermissible conduct that is clearly desirable, like a firm lowering its prices in order to compete away customers from rival firms. (When successful, price competition of this sort will make a

firm's rivals worse off and thus fail to qualify as Pareto improving.) Whatever is the principle the MFA uses to distinguish preferred from non-preferred tactics for pursuing profits, then, it cannot be Pareto improvement.

b. Pareto efficiency

A state is Pareto efficient (or Pareto optimal) if nobody can be made better off without making someone else worse off. Pareto efficiency is thus related to Pareto improvement in the following way: a state is Pareto efficient if and only if it contains no potential Pareto improvements.

Heath (2014: 5) claims that “the market is essentially a staged competition, designed to promote Pareto efficiency.” Heath relies on the first fundamental theorem of welfare economics to illustrate how we can distinguish between preferred and non-preferred tactics for pursuing profit (Heath, 2014: 29). The first fundamental theorem proves that, under conditions of perfect competition (i.e. market actors are price takers, there is perfect information, markets are complete, property rights are perfectly defined and enforced, and so on), any competitive market equilibrium is Pareto efficient (Arrow & Debreu, 1954; Blaug, 2007). This allows us to see that, under “conditions of perfect competition, lower price, improved quality, and product innovation would be the only way that firms could compete with one another. We can refer to these as the set of *preferred* competitive strategies” (Heath, 2006: 549). Heath appears to endorse what I will call the *perfect competition principle*:

Perfect competition principle: A strategic tactic is *preferred* if and only if it is possible for a market actor to employ under conditions of perfect competition, and *non-preferred* if and only if it is not possible for a market actor to employ under conditions of perfect competition.

Two examples of non-preferred tactics that Heath offers, producing pollution and selling products with hidden quality defects (Heath, 2006: 550), support this way of understanding the MFA. The conditions of perfect competition make it impossible for market actors to pursue profit by selling products with hidden quality defects because the condition that all market participants have perfect information (including about product features relevant to quality evaluations) preclude sellers from duping buyers in this way. Similarly, the conditions of perfect competition make it impossible for market actors to pursue profit by polluting—at least polluting beyond the level where the marginal social cost of pollution is greater than the marginal social benefit. Perfect competition relies on an idealized conception of property rights that are perfectly defined and enforced, and that grant a property owner the ability to prevent other actors from behaving in ways that negatively affect the use-value of their property. If burning a fire on my land causes smoke to reduce the quality of the air above your land, then under perfect competition, you can prevent me from burning fires without your assent. I, the polluter, would thus have to bargain with you, and others who suffer negative effects from my pollution, until a mutually agreed upon price on pollution is reached. This makes it impossible to profit by polluting more than the socially optimal amount.

Because of complications arising from the “theory of the second best” (Lipsey & Lancaster, 1956), it is not clear whether the tactics the perfect competition principle judges preferred/non-preferred also ought to be judged permissible/impermissible (Heath, 2004; Steinberg, 2017). However, my main point is distinct from second-best complications. Hence, I will set them aside and assume that preferred is equivalent to permissible, and non-preferred equivalent to impermissible. Ignoring second-best complications, then, the perfect competition principle appears to be just the sort of theoretical device we need to account for our intuitions about preferred/permissible and non-preferred/impermissible tactics in Easy Cases and to help us think through Hard Cases. As I will show, however, this appearance is deceiving: the perfect competition principle fails in an important way to fulfil either of these desiderata.

IV. Why the perfect competition principle must be rejected

Product quality improvement is an Easy Case. It clearly must be judged preferred/permissible. When firms compete with each other by making products that better meet consumer needs, they are obviously pursuing profit in a way we as a society want them to. Heath agrees, and he even insists that product quality improvement, along with lowering prices and product innovation, are the *only* preferred strategic tactics (Heath, 2006: 549). However, the judgment that product quality improvement is a preferred tactic contradicts the perfect competition principle.

a. Implications of the perfect competition principle

Perfect competition requires that there be (infinitely) many buyers and sellers exchanging homogeneous goods. These conditions preclude individual market actors from having any degree of market power, i.e. power to set the prices of the goods they buy and sell. Suppose that these and the other conditions of perfect competition all obtain. A firm, call it F, producing a homogeneous good discovers a way to produce a higher quality version of the good. F begins selling this improved version of the good at a price that will cause some customers to defect from F's rivals to buy F's new, higher quality product.

The conditions of perfect competition are so restrictive that it is difficult to see how this product quality improvement strategy could be possible to pursue when they obtain. Recall that information under perfect competition is costless. This will enable F's rivals to copy the improvement instantaneously and underbid F if F prices the improved good any higher than the price that will prevail when the perfectly competitive market reaches equilibrium (where price is equal to marginal cost of production, and economic profit is zero). It is thus not clear how competing on product quality is possible under the conditions of perfect competition in a way that non-preferred tactics are not. This makes it difficult to see how the perfect competition principle, which requires us to conduct thought experiments about which tactics would be possible for a firm to pursue under conditions of perfect competition, can illustrate the difference between preferred and non-preferred tactics in the way Heath seems to envision.

One might attempt to object to this point by arguing that there is an understanding of ‘possibility’ in the perfect competition principle that avoids this problem. For example, there is a sense in which market actors under conditions of perfect competition may *try* to profit from product quality improvement. The relentless tendency of perfectly competitive markets to move toward Pareto-efficient equilibrium renders this effort futile, but market actors can still attempt it.

Contrast this with attempting to shift costs from pollution onto third parties. This tactic appears to more inherently involve a violation of one of the conditions of perfect competition than competing on property quality, given that it entails a violation of property rights. One condition of perfect competition is that property rights are perfectly defined and enforced. This means that, in the world of competition, one person cannot impose costs of pollution on another in the way they often can in the real world. If I dispose of my trash by dumping it in my neighbor’s pristinely manicured yard, I have avoided the cost of trash disposal by imposing it on my neighbor. Clearly, I have violated my neighbor’s property rights by doing this—my neighbor can no longer benefit from the use of her yard in the way she was able to when it was in its former pristine state. Moreover, there is a clear analogy between this case and other standard pollution cases. Imagine if my neighbor depended on a stream flowing downhill from my land to hers for drinking and watering her farm animals, and I dumped pollution into the stream that made the water unsuitable for my neighbor’s purposes. Imagine if a factory on my land emitted smoke that blew to my neighbor’s land, befouling the air that she and her farm animals breathed. These polluting activities all involve violations of property rights.

Thus, under conditions of perfect competition, they are not possible. Perhaps this points to a way of distinguishing employing tactics that involve pollution from tactics involving product quality improvements.

However, the conditions of perfect competition do not actually preclude the possibility of attempting to profit from shifting costs of pollution onto third parties in the way proposed above. What the conditions require is that polluters may not impose costs upon third parties *without compensating them*. The costs and benefits of pollution are thus allocated in the same way the costs and benefits of any product or good are allocated: through a competitive bidding process between buyers and sellers. Properly understood, then, the perfect competition principle forbids polluting beyond the level at which one would pollute if the total social costs of pollution (rather than just the private costs borne by the polluter) were factored into the polluter's decision about how much to pollute. There is a sense in which a polluter could attempt to pollute at a level beyond the amount a perfectly competitive market would bear, just as the seller of any other product could (like the seller pursuing profit by improving product quality) attempt to sell the product at a price above the competitive market price. But the result in either case would be the same: the polluter/seller would be left in possession of an unsold unit of pollution/goods, since no buyer will agree to accept it when other units of pollution/goods are available at a more attractive price. The firm that seeks profit by imposing costs of pollution on third parties, like the firm that seeks profit by improving quality, can attempt to profit through this tactic, but will be unsuccessful due to the nature of the conditions of perfect competition. Hence, it does not appear that this way of distinguishing different

notions of possibility under perfect competition will allow us to separate the product quality case from the pollution case in a way that will salvage the perfect competition principle or a version of the MFA that relies on it.

b. Schumpeterian versus Kirznerian views of entrepreneurship

If the perfect competition principle cannot account for the ethical difference between competing on product quality and competing by polluting (beyond the socially optimal amount), then it clearly fails as a principle for identifying preferred and non-preferred strategic tactics. But perhaps the perfect competition principle is the wrong way to interpret the MFA. My own close reading of Heath's work has convinced me that it has the strongest textual support of any alternative interpretation. But maybe I am missing something. Maybe there is another way of interpreting them that yields a more promising principle, or maybe there are other passages in Heath's work that point toward the kind of principle we are looking for.

However, I am skeptical that the MFA will yield a better principle (unless we interpret it to already contain a principle resembling the wealth creation ones that I suggest below). The MFA, at least insofar as it emphasizes the conditions of perfect competition that guarantee Pareto efficiency when the market reaches equilibrium, only focuses on a certain sense in which market competition is beneficial. It will be helpful to invoke a distinction from the entrepreneurship literature to illustrate this point.

Kirzner (1997) argues that we should understand entrepreneurship as a process that pushes markets toward a hypothetical competitive equilibrium state. Entrepreneurs discover market inefficiencies—products that consumers would want and are possible to make but are not being provided, producers who value a scarce resource more than those who are currently using it. These inefficiencies represent entrepreneurial opportunities, and the entrepreneurs who seize them are rewarded with profits as they iron out markets in which they participate toward competitive equilibrium.

Schumpeter (1947, 2010), by contrast, identifies entrepreneurship as a process that *disrupts* market equilibrium. Entrepreneurs carry out innovations—new goods, new methods of production, new forms of organization—that shock the markets in which they are introduced out of their equilibrium states. Entrepreneurship is an example of the creative destruction (to use Schumpeter’s term) process by which market systems facilitate the displacement of outmoded, inefficient ways of doing business by new, improved ways of doing business.


The MFA adopts a perspective about markets and how they produce benefits that is similar to the Kirznerian understanding of entrepreneurship. Having actors in competitive markets pursue profit is desirable because it eliminates wasteful inefficiencies from the market, bringing the system as a whole closer to Pareto-efficient equilibrium. Eliminating inefficiencies is indeed a valuable feature of competitive markets; this I do not dispute. However, as Schumpeter teaches us, eliminating inefficiencies is not the *only*, or even the *main*, valuable feature of markets. Competitive markets are also valuable because they foster innovation. But we cannot account for the

value of innovation by appeal to the desirability of getting closer to the Pareto frontier. Schumpeter emphasizes how the market disruptions innovation causes often result in the market being drawn, at least temporarily, *away* from its Pareto-efficient competitive equilibrium. Equilibrium-seeking, Kirznerian market processes will tend to draw the market back toward Pareto-efficient equilibrium in the wake of disruption. But the innovation itself does not facilitate this—indeed, Pareto-efficient equilibrium is the very thing it disrupts. If we want to be able to explain the value of innovative activities, such as a firm’s discovery of a new product design that allows it to produce a better quality product than its rivals, then we need to reference something other than the desirability of getting closer to Pareto efficient market equilibrium.

c. The Pareto frontier

This point can be put slightly differently by appealing to the notion of a Pareto frontier. A Pareto frontier is a graphical representation of all possible Pareto efficient allocations (see Fig. 1). By identifying the goal or ‘point’ of the market economy as the achievement of Pareto efficiency, the MFA justifies the restrictions it imposes on how market actors may pursue profit on the grounds that market actors’ adherence to them “gets us as close to the Pareto frontier as possible” (Singer, 2018: 111). Desirable sets of restrictions will be those that move the economy from Pareto sub-optimal states like *n* in Fig. 1 closer to the Pareto frontier.

Fig. 1: The Pareto frontier



Suppose the current state of the economy is at a Pareto sub-optimal state like n in Fig. 1. Application of the First Fundamental Theorem would have us understand this as a failure of markets to reach their perfectly competitive equilibrium: if the perfectly competitive equilibrium is achieved, the First Fundamental Theorem implies that the state of the economy should lie somewhere on the Pareto frontier, rather than below it. Of course, there are many reasons, under the First Fundamental Theorem, why a given Pareto sub-optimal economy fails to reach the Pareto frontier. One might be that competitive market processes have not yet had time to iron out existing inefficiencies.

Perhaps firms in a market with reasonably competitive features have not yet had time to determine their optimal pricing strategy, and there are appropriable economic rents available for firms that do a better job of competing on price.

Basic economics allows us to demonstrate how firms competing on price in a market like this will move the state of the economy closer to the Pareto frontier. (Note that this analysis is subject to second-best caveats, but again, for the sake of simplicity, I am ignoring those for now.)

Fig. 2: Consumer and producer surplus





Fig. 3: Deadweight loss



Suppose that our imaginary Pareto-sub optimal market can be represented by Fig. 3.

Prices in the market are at level p_1 , which is higher than p , the price level in the perfectly competitive market depicted in Fig. 2. When price is p_1 , quantity exchanged is q_1 , which

is lower than the quantity q exchanged when price is p . Fig. 2 and Fig. 3 illustrate the consumer and producer surplus realized when goods are exchanged in this market. This surplus represents the difference between the highest price that consumers occupying various points on the demand curve are willing to pay for goods and the lowest price at which producers occupying various points on the supply curve are willing to sell the good.

Note that Fig. 3 has deadweight loss in the triangle to the right of the vertical line representing q_1 , whereas in Fig. 2 this area is consumer surplus. This represents the fact that there are consumers in Fig. 3 (occupying points on the demand curve that lie between q_1 and q) who are willing to buy goods at a price at which sellers occupying the corresponding points on the supply curve would be willing to sell them. However, because in Fig. 3 competitive market pressures have not yet pushed the price from p_1 to p and quantity from q_1 to q , there are unrealized mutually-beneficial transactions in Fig. 3. The deadweight loss triangle in Fig. 3 represents these unrealized transactions. This analysis shows us why competing on price is a preferred strategy: It pushes markets occupying inefficient states like the one represented by the Fig. 3 graph or by point n in Fig. 1 toward the efficient states represented in Fig. 2 or by the Pareto frontier in Fig. 1. The Paretian approach thus seems to be able to justify competing on price.

This sort of justification, however, is not available for product quality improvement. Improving product quality generally does not eliminate the deadweight loss in markets like the one depicted in Fig. 3 or move the economy from sub-optimal points like n in Fig. 1 to the Pareto frontier. The economic effect of attempting to profit

by introducing an improved-quality product is captured in Fig. 1 by moving the Pareto frontier outward to pf_2 . We cannot account for the value of this by appealing to First Fundamental Theorem's explanation of how markets at perfectly competitive equilibrium are Pareto efficient. Indeed, if sub-optimal point n in Fig. 1 represents the state of economy, and successful employment of a product quality improvement tactic pushes the Pareto frontier outward to pf_2 without simultaneously moving point n outward by the same amount, then the state of the economy post-product quality improvement will be *further* from the Pareto frontier than it was pre-product quality improvement. A firm in a perfectly competitive market like the one depicted in Fig. 2 that successfully pursues a product quality improvement will create a situation with deadweight loss, like in Fig. 3, where none existed before (assuming rivals cannot instantly copy the improvement and compete away the entrepreneurial profits of the firm that innovated the improved-quality product).

Let me be clear: I do not wish to claim that Heath holds the (absurd on its face) position that product quality improvement is not a preferred tactic. His explicit claim to the contrary is strong evidence against interpreting him in this way. My claim, rather, is that the Paretian principles Heath relies on to distinguish preferred from non-preferred tactics are inadequate. Heath claims that competitive market institutions are valuable because they eliminate deadweight loss, move the market closer to the Pareto frontier, and "ensure the smooth operation of the price system" (Heath, 2006: 541). These are all different ways of appealing to the desirability of reaching Pareto efficiency. But market institutions, I claim, are valuable not just because they tend toward Pareto efficiency, but

also because they *push the Pareto frontier outward*. The Paretian principles Heath emphasizes cannot account for this, and they thus cannot account for the desirability of preferred tactics like product quality improvement that have this effect. Addressing this problem requires a different approach.

V. Wealth creation

My positive proposal in this paper is that we should look to the value of wealth creation, rather Pareto improvement, to explain much of what the MFA purports to illuminate: the point of the market economy, what justifies the profit motive in the context of a market economic system, and how we should differentiate between preferred and non-preferred tactics for pursuing profit.

Readers who are sympathetic to the wealth creation view I propose, but who want to be as charitable as possible to Heath, might wonder whether Heath himself actually accepts something like the wealth standard I recommend. Heath, however, comes close to explicitly rejecting a wealth standard in the introduction of his book on the MFA. Underscoring the Paretian foundation of his approach, he writes that economic “efficiency is often conflated with... the wealth-maximization standard proposed by ‘law and economics’ scholars” such as Posner (1973), and insists that “I use the term ‘efficiency,’ by contrast, in the strict Pareto sense, to refer to the principle that, whenever it is possible to improve at least one person’s condition without worsening anyone else’s it is better to do so than not” (Heath, 2014: 9–10). The wealth-based approach I propose

in this paper differs from Posner's; among other things, it does not countenance wealth maximization. Nevertheless, the fact that Heath only invokes a wealth standard to caution against what he views to be a misreading of his view would seem to provide ample reason against interpreting him as advocating a wealth standard.

It is possible that Heath rejects wealth maximization as a general standard in the way Posner uses it, but would be amenable to the more pluralistic, less all-encompassing wealth-based approach I advocate below. As I show, we retain many of the attractive features of Heath's Paretian MFA by adopting wealth creation as a guiding value for business ethics. Thus, I want to avoid overstating the tensions between my approach and Heath's—perhaps he would agree with much of what I advocate. I do believe, however, that we can infer from Heath's written work that the MFA he envisions differs in at least some important ways from the wealth creation approach I explain and defend here.

a. What is wealth?

Wealth refers to the all-purpose economic means that people may use to pursue their ends, usually by engaging in exchange with others. Thus, something counts as wealth if it possesses economic value. Tangible examples of wealth include consumer goods, productive assets, and land. But wealth also includes intangible things like ideas, productive processes, social capital, and human capital.

Because my discussion here focuses on the idea of creating wealth, it will be useful to identify two different ways we might try to measure wealth for the purpose of

defining what it means to have more or less of it. We can measure wealth by focusing on either *exchange value* or *use value* (Bowman & Ambrosini, 2000; Lepak, Smith, & Taylor, 2007). The exchange value of a good is determined by the price it commands on a given market. The use value of a good is determined by the maximum amount that someone is willing to pay for it. Suppose you pay \$1 for an apple at the supermarket, and that you value the apple at \$2 (i.e. you are indifferent between having \$2 and having the apple). The exchange value of the apple in this example is \$1, and the use value is \$2.

Wealth is a stock rather than a flow. An individual's wealth is determined by what they own, not what their annual income is. A country's wealth is determined by what its members own, not by its gross domestic product. Of course, how much an individual or group owns will usually be a function of how much value flows to them in the form of income. But strictly speaking, this income stream does not itself constitute a person's wealth.

The general position I am defending is that wealth creation is an important guiding value for understanding ethical constraints on how profit may be pursued in markets. The specific content of this general position will depend on the precise version of 'wealth' one adopts. In the discussion that follows, I adopt a conception of *wealth* as potential preference satisfaction as specified by the Kaldor-Hicks standard, and a conception of *wealth creation* as Kaldor-Hicks improvements. This view involves a simple, though significant, modification of Heath's Paretian MFA. What I propose can thus be understood as a friendly amendment to the MFA. There are, however, some persuasive objections to the conventional economic understanding of welfare as

subjective preference satisfaction upon which the conventional Kaldor-Hicks standard relies. As I show, a plausible view emerges if we accommodate some of these objections.

b. The Kaldor-Hicks standard

Kaldor-Hicks improvements are defined in terms of hypothetical Pareto improvements:

Kaldor-Hicks improvement: A state of economic distribution B is a Kaldor-Hicks improvement over another state A if agents whose utilities are higher in B than in A could hypothetically compensate agents whose utilities are lower in B than A so that B would qualify as a Pareto improvement compared to A.

In somewhat more colloquial terms, B is a Kaldor-Hicks improvement over A if there are enough economically valuable things in B relative to A that, through hypothetical redistribution, everyone in B could be made at least as well off as they would be in A.

The Kaldor-Hicks standard gives us a criterion for comparing the amount of wealth between different states. If state B represents a Kaldor-Hicks improvement over state A, then there is more wealth in B than A. Wealth is thus a function of agents' relative preference satisfactions, weighed against each other in the way the Kaldor-Hicks standard specifies. If an agent prefers a set of possessions Y over an alternative set of possessions Z, then that agent has more wealth if she has Y than Z. If a state of economic

distribution B is a Kaldor-Hicks improvement over another state A, then there is more wealth in B than in A.

One inconvenient feature of the Pareto standard is that many states are Pareto non-comparable (Buchanan, 1985; Hausman, McPherson, & Satz, 2016; Sen, 1991). For example, if one agent has a higher utility in A than B, and another agent has a higher utility in B than A, then neither is a Pareto improvement over the other. The Kaldor-Hicks standard, on the other hand, can make pairwise comparisons between all states of economic distribution that are inhabited by the same agents.

c. Reservations about Kaldor-Hicks

However, the conventional Kaldor-Hicks standard also has some drawbacks due to its reliance on the problematic theory of welfare that underlies much of modern economics. Thus, given that I am looking for a standard that will inform ethical judgments about permissible and impermissible tactics for pursuing profit in a market, Kaldor-Hicks does not provide an entirely adequate understanding of wealth creation.

Conventional welfare economics measures welfare in terms of subjective preference satisfaction. However, if we understand wealth in this way, we cannot distinguish between the different ways people acquire their preferences. Such distinctions are crucial for any conception of wealth creation that bears ethical significance. For example, there is clearly an ethically relevant difference between a preference that cannot secure the rational endorsement of the preference-holder (e.g. because it is the result of

manipulative advertising) versus a preference that is rationally accepted (or at least not rejected) by the preference-holder (Crisp, 1987). If Kaldor-Hicks cannot distinguish the former from the latter, it will overweight the ethical importance of fulfilling preferences that are irrational, non-autonomous, or whose fulfillment would be against the interests of the preference-holder.

Thus, a fully developed wealth creation approach will need a theory of welfare that can make these sorts of distinctions. I cannot provide such an account here, but I am confident that one exists, at least in principle. One option, outlined by Arneson (2000), would be to adopt an objective list of things that contribute to welfare—Arneson offers “engagement in relations of love and friendship and intellectual and cultural achievement” (514) as examples. One might add to the list some items that intersect more with economic life than Arneson’s examples. For instance: living a normal-length human life, maintaining good health, finding fulfillment, achieving valuable things, acting autonomously and rationally. (This list roughly tracks some of the central human capabilities from [Nussbaum, 2001], though her approach has some deeper theoretical differences from theories of welfare as traditionally construed.) Alternatively, one might choose to hew closer to a conventional Kaldor-Hicks account of wealth creation by limiting which subjective preferences are relevant, rather than discarding them altogether. For example, we could exclude subjective preferences that result from manipulative advertising, or that the individual preference-holder would herself reject.

d. The Wealth Creation Principle

I have spoken of product quality improvement as a tactic a firm may use to pursue profit. Strictly speaking, however, pursuing profit through product quality improvement involves at least two distinct tactics: (a) improving some aspect of a product in the eyes of some customers so that the willingness of those customers to pay for the product increases, and (b) choosing the price at which the product is offered for sale. The former is a *wealth creation* tactic, while the latter is a *wealth appropriation* tactic (Lepak et al., 2007). Wealth creation tactics create wealth by increasing customer willingness to pay for a firm's products or by increasing the efficiency (thereby decreasing costs) associated with production and distribution. Wealth appropriation tactics appropriate wealth on the part of the firm that would otherwise accrue to some other party.

In many cases, for it to be in a firm's financial interest to implement a wealth creation tactic, that tactic will need to be combined with a wealth appropriation tactic. In the long run, a firm will only gain financially from selling a superior-quality product if it is at least able to set its prices high enough to cover its costs. In the likely event that the firm's wealth creation tactic involves undertaking risky investments, the firm will also need to appropriate an additional risk premium for the investments to be *ex ante* financially justified. To pursue profit successfully, firms must adopt tactics for both wealth creation and wealth appropriation.

If our goal is to identify a wealth-based principle that distinguishes permissible tactics for pursuing profit from impermissible ones, we cannot consider wealth creation

tactics and wealth appropriation tactics in isolation of each other. Wealth appropriation tactics often reduce overall wealth relative to if they were not employed at all. Consider the wealth appropriation tactic of increasing prices. In perfectly competitive equilibrium, all firms have total revenues that are just high enough to cover their costs (including costs of capital), but no higher. (Under perfect competition, the economic rents appropriated by a firm that brings in revenue in excess of its costs will be competed away by rivals or new entrants.) In other words, firms under perfect competition achieve zero economic profit (i.e. profit in excess of costs of capital). Imagine a zero-profit firm that prices its products just high enough to cover its costs. Suppose market conditions change so that the conditions of perfect competition no longer hold, allowing this firm to raise its prices and achieve positive economic profit. Absent specific market conditions, this price increase will cause deadweight loss: potential mutually-beneficial transactions will not occur, because buyers whose maximum willingness to pay was greater than the price level when the firm was just barely covering its costs, but lower than the increased price, will no longer purchase the product. There is less wealth after the firm employs its wealth appropriation tactic than before. Thus, if we apply a wealth standard to wealth appropriation tactics in isolation, we risk judging all or most wealth appropriation tactics impermissible.

This seems overly restrictive. Perhaps there is something to be said, ethically speaking, for a firm that voluntarily forgoes profit and keeps its prices at a level just high enough to cover its costs. But making this an ethical requirement seems both unintuitively demanding and prevents us from making the sorts of ethical distinctions

between permissible and impermissible tactics for pursuing profit that motivated this paper in the first place. Fortunately, we can avoid this result by evaluating tactics for appropriating wealth in the context of the wealth creation tactics that often accompany them, rather than in isolation.

How wealth appropriation and wealth creation tactics should be combined is a tricky question. Precise answers are illusive. However, to gain some clarity on this issue, I propose that we group tactics together into what I will call *strategic sets*. A strategic set of tactics contains one or more wealth creation tactics, plus wealth appropriation tactics that are necessary for the wealth creation tactics in the set to be in the strategic interest of the firm to employ. Thus, when a firm pursues profit by improving the quality of one of its products, the resulting strategic set of tactics will include (a) the wealth-creating tactic of improving the quality of the product and (b) the wealth appropriating tactic of setting the price of that product so that the firm captures a portion of the value it creates. These two tactics should be evaluated together. They should also be evaluated separately from other strategic sets of tactics the firm employs. For example, if the firm undertakes another quality improvement for an unrelated product, the two strategic sets of tactics that constitute the two different product quality improvements should be evaluated separately.

Having made these preliminary remarks, I propose the following principle for distinguishing permissible and impermissible tactics for pursuing profit:

Wealth Creation Principle: It is impermissible for a firm to employ strategic sets of tactics that are expected to appropriate more wealth than they create.

The Wealth Creation Principle can explain our intuitions regarding Easy Cases. As the above remarks suggest, product quality improvements are judged permissible because, absent serious market failure (e.g. a quality improvement that introduces or exacerbates an especially costly negative externality), a firm that implements a product quality improvement will create more wealth than it appropriates. The wealth created by increasing some customers' willingness to pay will be greater than the wealth the firm appropriates through its pricing (see Brandenburger & Stuart, 1996).

Conversely, strategic sets of tactics involving collusion and contrived deterrence will typically be impermissible under the Wealth Creation Principle. When executed successfully, these tactics destroy wealth overall by creating deadweight loss and failing to create additional wealth. Because firms that employ collusion and contrived deterrence successfully appropriate some wealth while creating negative wealth, they appropriate more wealth than they create and run afoul of the Wealth Creation Principle.

I am also optimistic that the Wealth Creation Principle points us in a plausible direction for analyzing Hard Cases. Evaluating whether the firm in Small Town Stores acted permissibly would require identifying the firm's wealth creation and appropriation tactics and placing them into strategic sets, and then gathering the information necessary to compare appropriated wealth to created wealth for each set. This will not be an easy task, and the final ethical judgments will depend on the specifics of the case. Nevertheless, the Wealth Creation Principle seems to contain a plausible and tractable way to analyze even Hard Cases.

The Wealth Creation Principle does have some important limits. It provides neither a necessary nor a sufficient condition for a strategic set of tactics being ethical overall. It is rather a supporting condition: the fact that a strategic set of tactics satisfies it counts in favor of its ethical permissibility, and the fact that a set violates it counts against its permissibility. This is as it should be. The principle is based on wealth, but wealth clearly cannot explain everything in business ethics. Lying is usually wrong, and a strategic set of tactics that included a serious lie would usually be wrong to employ even if it created wealth overall. Strategic business decisions will often implicate considerations apart from just wealth. Making such decisions in an ethical way will require more than simply abiding by the Wealth Creation Principle. Still, a wealth creation approach makes sense in the domain of business ethics because, as I discuss in the next section, wealth is especially important for explaining and justifying ethical norms in business and market contexts. Part of the explanatory value of a wealth creation approach is that much of what it explains is counterintuitive. People often have difficulty understanding how competitive, profit-oriented market activity could possibly be ethical—this may explain the common suggestion that business ethics involves an inherent contradiction or oxymoron. An appreciation of the ethical value of wealth creation, and the way in which profit-seeking actors in reasonably competitive market institutions facilitate wealth creation, helps resolve this apparent tension. It is true that wealth creation cannot explain everything in business ethics. On the other hand, what it does explain is important.

Moral philosophers distinguish between ethical principles as criteria of the right versus ethical principles as decision procedures. Criteria of the right tell us what the conditions are for an act to count as morally right (e.g. that it maximizes overall utility, or that it follows from Kant's Categorical Imperative). Decision procedures provide practical guidance to people that is supposed to induce them to engage in better moral behavior (Brink, 1986; Railton, 1984; Rawls, 1951). As I understand the Wealth Creation Principle, it falls somewhere between the poles of this spectrum. Unlike a pure decision procedure, the Principle does not give detailed prescriptive guidance that could plausibly replace an individual's moral judgment regarding what to do and care about. Like a decision procedure, the Principle does serve as a source of practical guidance. It highlights certain contours of the moral terrain that are especially important for decisions in business and market contexts. Unlike a pure criterion of the right, the Principle does not provide necessary or sufficient conditions for an act to qualify as ethically right. However, as I show below, the Principle does track moral considerations that are intrinsically important. It is not a noble lie whose aim is purely to induce better behavior in people who adopt it (Plato, 2004). The Wealth Creation Principle identifies an important moral consideration that is relevant to what agents ought to do, particularly in business and market domains. Competing moral considerations can outweigh it, but absent competing moral considerations, compliance with the Principle becomes morally obligatory.

Another objection to the Wealth Creation Principle is that it requires a baseline, and that I have failed to provide one. The Wealth Creation Principle says that doing

business ethically means refraining from opportunities to profit that destroy wealth. But destroy wealth relative to what? Consider:

State-Sanctioned Monopolist: a firm in a crony-capitalist economy enjoys regulatory privileges that give it an effective monopoly in its industry.

Reformers take over the government and eliminate these regulatory privileges. However, the firm has enough advantages, even in the absence of regulatory protection, to effectively deter rivals from entering its industry.

If this firm merely refrains from employing further strategic sets of tactics that appropriate more wealth than they create, it will still benefit from the privilege it enjoys as a vestige of its former state-sanctioned monopolist status. Shouldn't the firm be required to do more than this to rectify this unjust heritage of its crony-capitalist past? And won't determining how much more the firm must do require the sort of baseline that the objection envisions?

My response is that, properly understood, the Wealth Creation Principle already contains such a baseline. The principle insists that firms refrain from employing strategic sets of tactics that appropriate more wealth than they create. This does not just implicate sets of tactics the firm is considering implementing in the future; it also implicates the sets of tactics the firm currently has in place. To determine what the Wealth Creation Principle would require of the firm in State-Sanctioned Monopolist, one would need to define the firm's strategic set of tactics and estimate whether, for each individual set, it appropriates more wealth than it creates. If it does, then the Wealth Creation Principle demands that the firm give up that set of tactics. Notice how strict the requirements are

that this standard imposes on incumbent firms that enjoy a significant degree of monopoly power. The Wealth Creation Principle requires that wealth-appropriating tactics (such as price increases) be matched with wealth-creating tactics (such as quality improvements). But firms with significant monopoly power often engage in wealth appropriation without accompanying wealth creation. For example, a monopolist will often raise its prices far above the level that would prevail under conditions of perfect competition, thereby causing and exacerbating deadweight loss. For the Wealth Creation Principle to judge such price increases permissible, the monopolist would be required to combine them with wealth-creating tactics in strategic sets of overall tactics where, for each individual set, that set is expected to create more wealth than it appropriates. The wealth creation principle thus forbids simple cases of monopolists raising their prices simply because they can. Simple price increases appropriate wealth without creating any wealth at all. If the firm in State-Sanctioned Monopolist engaged in such price increases in the past, and has sustained them without any change in its other operations that might satisfy the Wealth Creation Principle, then the Wealth Creation Principle will prohibit the firm from maintaining those monopoly prices.

These refinements aside, I hope it is clear that the Wealth Creation Principle at the very least succeeds where the MFA fails: it provides a way of distinguishing between preferred and non-preferred tactics for pursuing profit in a market. Preferred strategic sets of tactics create more wealth than they appropriate, while non-preferred strategic sets of tactics appropriate more wealth than they create.

VI. The ethical importance of wealth

I have argued that wealth creation, rather than any Paretian principle, provides the best explanation for what distinguishes preferred and non-preferred tactics for pursuing profit. If wealth considerations play such an important role in justifying ethical norms governing pursuit of profit in a market, then wealth must have significant ethical importance. But does it? Some are skeptical (Dworkin, 1980; Jones et al., 2016; Jones & Felps, 2013). However, I maintain that, as a society, we have stronger ethical reasons than critics recognize for promoting wealth creation through our economic institutions and the business activity that occurs within them.

a. The overlapping consensus in support of wealth creation

Libertarian political philosophers often appeal to wealth creation as a reason for endorsing institutions like private property and free markets (e.g. Schmidtz & Brennan, 2011; Van der Vossen & Brennan, 2018). For this reason, it may be tempting to conceive of a wealth creation approach to business ethics as a specifically libertarian way of understanding business ethics. However, while a strong emphasis on wealth creation is likely to appeal to many libertarians, as well as the specific sort of economic libertarian traditionally associated with the fields of economics and finance, it does not require one to adopt any particularly controversial or strong libertarian premises.

Indeed, John Rawls, one of the most important left-liberal political philosophers of the twentieth century, identified wealth creation as a sort of fundamental political

value. According to Rawls, it is in the interest of individual members of society to have greater access to primary social goods. For the purpose of shaping society's political institutions, we evaluate how well-off individual members of society are under various institutional regimes by evaluating how much access to primary goods they enjoy (Wenar, 2017). Rawls' list of primary social goods includes "[i]ncome and wealth, understood as all-purpose means (having an exchange value) generally needed to achieve a wide range of ends whatever they may be" (Rawls, 2001: 58–59). Of course, Rawls places strict limits, contained in his two principles of justice, on how primary goods may be distributed. But he still endorses wealth as something citizens have a fundamental interest in having more of, and as something political institutions should promote so long as the two principles of justice are satisfied.

That both the left-liberal Rawls and libertarians like Brennan, Schmidtz, and Van der Vossen affirm the importance of wealth creation as a goal for political and economic institutions suggests that wealth creation can be endorsed from a range of moral and political perspectives. Those who disagree on deep questions about value in political philosophy might nevertheless arrive at substantial agreement about the value of fostering wealth creation. A similar overlapping endorsement of wealth creation is possible starting from the premises of various prominent moral theories. There is a strong consequentialist case for promoting wealth creation: wealthier people generally report greater levels of subjective well-being (Stevenson & Wolfers, 2013), and societies that generate large amounts wealth have a greater ability to protect and promote the well-being and functioning of their members (Cahill, 2005; McGillivray, 1991). But one who follows

Rawls (2009) in rejecting consequentialism can still find ample reason for endorsing wealth creation as an important moral value to the extent to which societies with greater levels of wealth are better able to promote and protect values like autonomy or human flourishing. (Pluralist consequentialists can endorse these values as among those worth promoting, but they tend to take a more central role in approaches to political philosophy whose adherents take themselves to be opposed to consequentialism, e.g. autonomy in Rawls' Kant-inspired theory of justice.) Rigorously demonstrating the connection between greater wealth and the increased realization of the values of autonomy and human flourishing is a task that extends beyond the scope of this paper. My point here is simply that the ethical importance of wealth can be supported by reference to a range of deeper ethical values. Even those who disagree about *why* wealth is valuable can agree *that* it is valuable.

b. Positive freedom

There is one value whose justificatory relationship to wealth creation is worth drawing out a bit more fully: positive freedom. Positive freedom is understood by contrast to negative freedom, which refers to the absence of coercive interference. As Elizabeth Anderson explains it, "If you have positive freedom, you have a rich menu of options effectively accessible to you" (Anderson, 2015: 100). A person has positive freedom to the extent that they enjoy the effective power to act or pursue their ends (Gaus, Courtland, & Schmidtz, 2015).

If positive freedom is intrinsically valuable, then wealth must at least be of instrumental value. Greater wealth correlates with greater positive freedom: holding all else equal, the more wealth a person has, the more options she has available to her, and the greater her ability to realize her conception of the good life.

Is wealth's value *merely* instrumental? I think not. Indeed, there are strong reasons for thinking that wealth is not merely of instrumental importance for positive freedom, but that it actually *constitutes* positive freedom. Following G. A. Cohen, we might compare having wealth to having tickets that allow us to do things we otherwise would lack the freedom to do. Wealth is “nothing but a highly generalized form of... a ticket” that gives its owner “a license to perform a disjunction of conjunctions of actions—actions, like, for example, visiting one’s sister in Bristol, or taking home, and wearing, the sweater on the counter at Selfridge’s” (Cohen, 1995: 58). A person who does not have enough wealth to trade in exchange for a department store sweater lacks the positive freedom to take home and wear that sweater. If he tries to take it home and wear it, he can expect to face physical coercion that prevents him from succeeding. But if he has a sufficient level of wealth to trade for the sweater, ‘taking home and wearing the sweater’ is added to his menu of potential options. People with greater wealth enjoy more real options about what sort of life to lead, what sort of person to become, and what to do at a given moment (Brennan, 2012). Wealth, then, is not just an instrument to achieving greater positive freedom; in an important sense it *is* a form positive freedom. Thus, if we have reason to shape our political and economic institutions in a way that promotes

positive freedom, we have reason to shape those institutions in a way that promotes the creation of greater levels of wealth.

c. The importance of wealth in business ethics

Even those who accept that our political and economic institutions ought to promote positive freedom, and accept that wealth constitutes positive freedom, may balk at the emphasis I place upon wealth in this paper. After all, wealth is clearly inadequate as a general standard of justice. Consider the following example adapted from Émile Zola's *Germinal*, a novel about a harsh coal miner's strike in 19th century northern France:

Rich Get Richer: The fabulously wealthy Grégoire family owns the Voreux coal mine, which employs hundreds of workers who spend their lives in poverty and toil in hazardous conditions in the mine. The Grégoires change the mine's compensation policy in a way that increases overall production, as well as profits for the Grégoire family, while decreasing each worker's take home pay. The change in compensation increases the wealth of the Grégoire family by more than it reduces the overall wealth of the workers. Thus, it creates wealth.

One struggles to imagine what kind of ethics- or justice-related reason could support the change in compensation in Rich Get Richer. The deprivation of the coal miners is worsened for the benefit of the Grégoires' already vast fortune. Rich Get Richer appears to function as a *reductio ad absurdum* of wealth as a general standard of justice.

I agree wholeheartedly that wealth must be rejected as a general standard of justice. We cannot infer from the fact that society E has more wealth than society F that E

is a more just society than F. However, wealth can still be a *component* of justice without being a *general standard* of justice. Wealth is a component of justice so long as greater wealth tends to contribute to greater justice.

Greater wealth indeed tends to contribute to greater justice, all else equal. In large-scale, modern societies, a certain minimal level of wealth is required to enable members of society to meet their basic needs—e.g. for water, nutrition, basic education, health care, and shelter. Beyond basic needs, as the previous section argues, wealth contributes to justice by promoting the positive freedom of the individual members of society. People who have greater wealth have greater effective ability to pursue their passions and projects, enjoy their leisure time, and spend time with their loved ones. Of course, all else is often not equal. Wealth can be problematic from the point of view of justice when it allows some people to buy political power and corrupt a society's government in order to further promote ends. More generally, large inequalities in wealth may allow citizens with great wealth to dominate and oppress citizens who lack wealth. But the existence of such cases does not undermine the place of wealth creation in the set of overall goals that modern societies ought to pursue. They only specify certain instances of wealth creation in certain contexts where countervailing considerations override reasons in favor of promoting the creation of wealth. They do not show that society should not pursue wealth as a general matter.

I have argued that firms have a presumptive duty to conform their conduct to the Wealth Creation Principle. That this duty is presumptive means that it can be overridden by competing considerations. The nature and strength of countervailing considerations

will depend upon, among other things, how just the society is within which a firm operates. In an ideal society, one in which we could trust that political, social, and economic institutions function well enough to ensure justice, for-profit firms would perhaps be justified pursuing profit maximally within the bounds of the minimal limits contained within the Wealth Creation Principle and other basic ethical principles and values: e.g. honoring commitments, avoiding fraud and deception, and refraining from wrongful harms and rights violations.

In the non-ideal real world contexts, of course, political, social, and economic institutions do not function perfectly, and firms cannot simply depend on other institutional actors to protect justice and equality. There will be times when, for reasons Singer (2018) explains, justice or equality-based reasons will be the countervailing considerations that override the pro tanto desirability of certain instances of wealth creation. Rich Get Richer would likely qualify as one such case: the moral goodness of wealth being created is outweighed by the potential moral badness of an already powerful economic actor who exploits dominating relationships with vulnerable others for its own benefit gaining even more power, allowing it to exercise even greater domination (Anderson, 2017). But again, the existence of such cases does not undermine the Wealth Creation Principle as a mid-level principle for business ethics that provides actors with presumptive reasons to act in certain ways. Nor does it undermine wealth's status as an important guiding value for business ethics.

VII: Conclusion

In this paper, I have argued that business ethics cannot be Paretian. Paretian principles fail to provide a plausible way of distinguishing between permissible and impermissible tactics for pursuing profit in a market. Focusing on wealth creation provides a better way of understanding the ethics of market competition.

Of course, one could accept all of this and still maintain a broadly Paretian view about the overall political and economic institutions of society. Over time, these institutions should generate broad prosperity, finding ways to make most members of society (especially the less well-off) better off, and not doing so at the expense of any particular individual or group. Having a society that repeatedly makes Pareto improvements, however, requires a capacity to create wealth. Creating wealth requires market institutions. And market institutions cannot function effectively if market actors are only permitted to employ tactics that cause Pareto improvements.

If this is right, then having a broadly Paretian society requires something closer to the Wealth Creation Principle than any principle that is recognizably Paretian for determining the ethical constraints on market competition. The Wealth Creation Principle may seem lax in comparison to Paretian principles, but it places important restrictions on the tactics managers may use to pursue profit for their firms. Though it is by no means a sufficient principle for business ethics, it does forbid some especially problematic kinds of business behavior that scholars of institutions have shown can inhibit societies' efforts to become peaceful, prosperous, and just. The Wealth Creation Principle requires that

business firms conduct themselves as inclusive institutions that that produce wealth for society, rather than extractive institutions that use their power to appropriate wealth that others produce (Acemoglu & Robinson, 2013). It requires that business firms seek economic success through productive, rather than unproductive or destructive, entrepreneurship (Baumol, 1996). To comply with the principle, managers must ensure that their firms' business activities contribute to the overall prosperity of society, rather than immiserating it for private gain. That may not be all that business ethics is (or should be) about, but it is surely an important part of it.

CHAPTER 4: Conclusion

By way of conclusion, I will briefly note what I consider to be three of the most important lessons that can be taken from the above three essays.

First, as essay 3 argues, an adequate ethical understanding of business requires understanding the means by which profit-oriented business activity contributes to the creation of wealth. Individual market transactions often create winners and losers in a way that seems intuitively objectionable. Understanding how these transactions facilitate the creation of wealth overall contributes to our understanding of why these intuitively objectionable market transactions are often justified, as well as why they sometimes are not.

Second, as essays 1 and 2 argue, when we adopt a definition of CSR that is decoupled from substantive ethical considerations, thinking about CSR primarily as a source of strategic advantage can be dangerous. Even if we adopt a non-moralized conception of CSR, normative theorizing about CSR must be based primarily on substantive moral considerations, rather than on instrumental firm performance considerations.

Third, as essay 2 argues, CSR programs often fail to live up to their potential because they are rarely selected on the basis of comparative estimates (based on strong evidence and rigorous analysis) of social impact. Normative CSR theory must recognize and account for the vast amount of additional good that CSR initiatives could accomplish if they took insights from effective altruism more seriously. Approaching CSR in this

way need not be burdensome; any firm with resources to devote to beneficent CSR causes can accomplish an enormous amount of social good simply by allocating its CSR resources to effective charitable organizations.

BIBLIOGRAPHY

- Acemoglu, D., & Robinson, J. A. 2013. *Why nations fail: The origins of power, prosperity, and poverty*. Crown Books.
- Adams, R. 2018, April 17. Starbucks to close 8,000 US stores for racial-bias training after arrests. *The New York Times*.
- Akerlof, G. A. 1970. The Market for “Lemons”: Quality Uncertainty and the Market Mechanism. *The Quarterly Journal of Economics*, 84(3): 488–500.
- Aldrich, H. E., & Fiol, C. M. 1994. Fools rush in? The institutional context of industry creation. *Academy of Management Review*, 19(4): 645–670.
- Alvarez, M. 2016. Reasons for action: Justification, motivation, explanation. (E. N. Zalta, Ed.) *Stanford Encyclopedia of Philosophy*, Winter 2017 Edition.
- Anderson, E. 2015. Liberty, Equality, and Private Government. *Tanner Lectures in Human Values*, 61–122.
- Anderson, E. 2017. *Private Government: How Employers Rule Our Lives (and Why We Don’t Talk about It)*. Princeton University Press.
- Approaches to Moral Weights: How GiveWell Compares to Other Actors. 2017, November. *GiveWell Blog*. <https://www.givewell.org/how-we-work/our-criteria/cost-effectiveness/comparing-moral-weights>.
- Arneson, R. J. 2000. Welfare should be the currency of justice. *Canadian Journal of Philosophy*, 30(4): 497–524.
- Arnold, D. G. 2010. Transnational corporations and the duty to respect basic human rights. *Business Ethics Quarterly*, 20(3): 371–399.
- Arnold, D. G., & Bowie, N. E. 2003. Sweatshops and respect for persons. *Business Ethics Quarterly*, 13(2): 221–242.

- Arnold, D. G., & Bowie, N. E. 2007. Respect for workers in global supply chains: Advancing the debate over sweatshops. *Business Ethics Quarterly*, 17(1): 135–145.
- Arrow, K. J., & Debreu, G. 1954. Existence of an equilibrium for a competitive economy. *Econometrica: Journal of the Econometric Society*, 265–290.
- Barbier, C. 2011. *In re Oil Spill by the Oil Rig “Deepwater Horizon” in the Gulf of Mexico, On April 20, 2010*.
- Barnett, M. L. 2007a. Stakeholder influence capacity and the variability of financial returns to corporate social responsibility. *Academy of Management Review*, 32(3): 794–816.
- Barnett, M. L. 2007b. Stakeholder influence capacity and the variability of financial returns to corporate social responsibility. *Academy of Management Review*, 32(3): 794–816.
- Barney, J. 1991. Firm resources and sustained competitive advantage. *Journal of Management*, 17(1): 99–120.
- Barney, J. B. 1995. Looking inside for competitive advantage. *Academy of Management Perspectives*, 9(4): 49–61.
- Battilana, J., Lee, M., Walker, J., & Dorsey, C. 2012. In search of the hybrid ideal. *Stanford Social Innovation Review*, 10(3): 50–55.
- Baumol, W. J. 1996. Entrepreneurship: Productive, unproductive, and destructive. *Journal of Business Venturing*, 11(1): 3–22.
- Beauchamp, T. 2019. The principle of beneficence in applied ethics. (E. N. Zalta, Ed.) *Stanford Encyclopedia of Philosophy*, Spring 2019 Edition.
- Beitz, C. R., & Goodin, R. E. 2009. *Global basic rights*. OUP Oxford.
- Berkey, B. 2016. Against Rawlsian institutionalism about justice. *Social Theory and Practice*, 42(4): 706–732.
- Berkey, B. 2018. The institutional critique of effective altruism. *Utilitas*, 30(2): 143–171.

- Berliner, D., & Prakash, A. 2015. “Bluewashing” the Firm? Voluntary Regulations, Program Design, and Member Compliance with the United Nations Global Compact. *Policy Studies Journal*, 43(1): 115–138.
- Bigge, D. M. 2004. Bring on the bluewash: A social constructivist argument against using Nike v. Kasky to attack the UN Global Compact. *Int’l Legal Persp.*, 14: 6.
- Blanken, I., van de Ven, N., & Zeelenberg, M. 2015. A meta-analytic review of moral licensing. *Personality and Social Psychology Bulletin*, 41(4): 540–558.
- Blattman, C., Fiala, N., & Martinez, S. 2013. Generating skilled self-employment in developing countries: Experimental evidence from Uganda. *The Quarterly Journal of Economics*, 129(2): 697–752.
- Blattman, C., & Niehaus, P. 2014a. Show them the money: why giving cash helps alleviate poverty. *Foreign Affairs*, 93: 117.
- Blattman, C., & Niehaus, P. 2014b. Show them the money: Why giving cash helps alleviate poverty. *Foreign Affairs*, 93(3): 117–126.
- Blaug, M. 2007. The fundamental theorems of modern welfare economics, historically contemplated. *History of Political Economy*, 39(2): 185–207.
- Bowman, C., & Ambrosini, V. 2000. Value creation versus value capture: towards a coherent definition of value in strategy. *British Journal of Management*, 11(1): 1–15.
- Brandenburger, A. M., & Stuart Jr, H. W. 1996. Value-based business strategy. *Journal of Economics & Management Strategy*, 5(1): 5–24.
- Brennan, J. 2012. For-profit business as civic virtue. *Journal of Business Ethics*, 106(3): 313–324.
- Brink, D. O. 1986. Utilitarian morality and the personal point of view. *The Journal of Philosophy*, 83(8): 417–438.
- Buchanan, A. E. 1985. *Ethics, efficiency, and the market*. OUP Oxford.

- Cahill, M. B. 2005. Is the human development index redundant? *Eastern Economic Journal*, 31(1): 1–5.
- Caney, S. 2010. Climate change, human rights, and moral thresholds. *Climate Ethics: Essential Readings*, 163–177.
- Cascio, W. F. 2006a. Decency means more than “always low prices”: A comparison of Costco to Wal-Mart’s Sam’s club. *The Academy of Management Perspectives*, 20(3): 26–37.
- Cascio, W. F. 2006b. The high cost of low wages. *Harvard Business Review*, 84(12): 23.
- Cho, S. Y., Lee, C., & Pfeiffer Jr, R. J. 2013. Corporate social responsibility performance and information asymmetry. *Journal of Accounting and Public Policy*, 32(1): 71–83.
- Christensen, L. T., Morsing, M., & Thyssen, O. 2013. CSR as aspirational talk. *Organization*, 20(3): 372–393.
- Cohen, G. A. 1977. Labor, leisure, and a distinctive contradiction of advanced capitalism. *Markets and Morals*, 107–36.
- Cohen, G. A. 1995. *Self-ownership, freedom, and equality*. Cambridge University Press.
- Coleman, J., Hershovitz, S., & Mendlow, G. 2015. Theories of the Common Law of Torts. (E. N. Zalta, Ed.) *The Stanford Encyclopedia of Philosophy*.
- Crisp, R. 1987. Persuasive advertising, autonomy, and the creation of desire. *Journal of Business Ethics*, 6(5): 413–418.
- Cullity, G. 1995. Moral free riding. *Philosophy & Public Affairs*, 24(1): 3–34.
- de los Reyes Jr, G., Scholz, M., & Smith, N. C. 2017. Beyond the “Win-Win” creating shared value requires ethical frameworks. *California Management Review*, 59(2): 142–167.
- De Mel, S., McKenzie, D., & Woodruff, C. 2012. One-time transfers of cash or capital have long-lasting effects on microenterprises in Sri Lanka. *Science*, 335(6071): 962–966.
- Dees, J. G., & Anderson, B. B. 2003. For-profit social ventures. *International Journal of Entrepreneurship Education*, 2(1): 1–26.

- Delmas, M. A., & Burbano, V. C. 2011. The drivers of greenwashing. *California Management Review*, 54(1): 64–87.
- Dierickx, I., & Cool, K. 1989. Asset stock accumulation and sustainability of competitive advantage. *Management Science*, 35(12): 1504–1511.
- Donaldson, T. 2012. The epistemic fault line in corporate governance. *Academy of Management Review*, 37(2): 256–271.
- Donaldson, T., & Dunfee, T. W. 1994. Toward a unified conception of business ethics: Integrative social contracts theory. *Academy of Management Review*, 19(2): 252–284.
- Donaldson, T., & Preston, L. E. 1995. The stakeholder theory of the corporation: Concepts, evidence, and implications. *Academy of Management Review*, 20(1): 65–91.
- Drucker, P. F. 1984. Converting social problems into business opportunities: The new meaning of corporate social responsibility. *California Management Review*, 26(2): 53–63.
- Duguid, M. M., & Thomas-Hunt, M. C. 2015. Condoning stereotyping? How awareness of stereotyping prevalence impacts expression of stereotypes. *Journal of Applied Psychology*, 100(2): 343.
- Dunfee, T. W. 2006. Do firms with unique competencies for rescuing victims of human catastrophes have special obligations? Corporate responsibility and the AIDS catastrophe in Sub-Saharan Africa. *Business Ethics Quarterly*, 16(2): 185–210.
- Dworkin, R. M. 1980. Is wealth a value? *The Journal of Legal Studies*, 9(2): 191–226.
- Foot, P. 1967. The problem of abortion and the doctrine of double effect. *The Oxford Review*, 5: 5–15.
- Forscher, P. S., Lai, C., Axt, J., Ebersole, C. R., Herman, M., et al. 2016. *A meta-analysis of procedures to change implicit measures*.
- Frankfurt, H. G. 2009. *On bullshit*. Princeton University Press.

- Freedman, J. L., & Fraser, S. C. 1966. Compliance without pressure: the foot-in-the-door technique. *Journal of Personality and Social Psychology*, 4(2): 195.
- Freeman, R. E. 1984. *Strategic management: A stakeholder approach*. Boston: Pittman.
- Freeman, R. E., Harrison, J. S., Wicks, A. C., Parmar, B. L., & De Colle, S. 2010. *Stakeholder theory: The state of the art*. Cambridge University Press.
- Friedman, M. 1970, September 13. The social responsibility of business is to increase profits. *New York Times Magazine*, 32–33, 122, 124, 126.
- Gaus, G., Courtland, S., & Schmidtz, D. 2015. Liberalism. *Stanford Encyclopedia of Philosophy*.
- Gertler, P. 2004. Do conditional cash transfers improve child health? Evidence from PROGRESA's control randomized experiment. *American Economic Review*, 94(2): 336–341.
- Giere, R. N. 1994. The cognitive structure of scientific theories. *Philosophy of Science*, 61(2): 276–296.
- GiveWell List of Top Charities. 2019. *GiveWell*. <https://www.givewell.org/charities/top-charities>.
- Grant, R. M. 1991. The Resource-Based Theory of Competitive Advantage: Implications for Strategy Formulation. *California Management Review*, 33(3): 114–135.
- Grant, R. M. 1996. Toward a knowledge-based theory of the firm. *Strategic Management Journal*, 17(S2): 109–122.
- Gray, L. 2008, December 22. Activists ejected after presenting greenwash award to BP. *The Telegraph*. <https://www.telegraph.co.uk/news/earth/earthnews/3903090/Activists-ejected-after-presenting-greenwash-award-to-BP.html>.
- Griswold, C. 2007. *Forgiveness: A philosophical exploration*. Cambridge University Press.

- Harris, J. D., Johnson, S. G., & Souder, D. 2013. Model-theoretic knowledge accumulation: The case of agency theory and incentive alignment. *Academy of Management Review*, 38(3): 442–454.
- Harrison, J. S., & Wicks, A. C. 2013. Stakeholder theory, value, and firm performance. *Business Ethics Quarterly*, 23(1): 97–124.
- Hart, S. L. 1995. A natural-resource-based view of the firm. *Academy of Management Review*, 20(4): 986–1014.
- Hart, S. L. 1997. Beyond greening: strategies for a sustainable world. *Harvard Business Review*, 75(1): 66–77.
- Hart, S. L., & Dowell, G. 2011. Invited editorial: A natural-resource-based view of the firm: Fifteen years after. *Journal of Management*, 37(5): 1464–1479.
- Hasnas, J. 2013. Teaching business ethics: The principles approach. *Journal of Business Ethics Education*, 10: 275–304.
- Haushofer, J., & Shapiro, J. 2013. Household response to income changes: Evidence from an unconditional cash transfer program in Kenya. *Massachusetts Institute of Technology*, 24(5): 1–57.
- Hausman, D., McPherson, M., & Satz, D. 2016. *Economic analysis, moral philosophy, and public policy*. Cambridge University Press.
- Hawn, O., & Ioannou, I. 2016. Mind the gap: The interplay between external and internal actions in the case of corporate social responsibility. *Strategic Management Journal*, 37(13): 2569–2588.
- Heath, J. 2004. A market failures approach to business ethics. *The invisible hand and the common good*: 69–89. Springer.
- Heath, J. 2006. Business ethics without stakeholders. *Business Ethics Quarterly*, 16(4): 533–557.

- Heath, J. 2013. Market Failure or Government Failure? A Response to Jaworski. *Business Ethics Journal Review*, 1(8): 50–56.
- Heath, J. 2014. *Morality, competition, and the firm: The market failures approach to business ethics*. Oxford University Press.
- Henisz, W. J., Dorobantu, S., & Narthey, L. J. 2014. Spinning gold: The financial returns to stakeholder engagement. *Strategic Management Journal*, 35(12): 1727–1748.
- Hsieh, N. 2017. The Responsibilities and Role of Business in Relation to Society: Back to Basics? *Business Ethics Quarterly*, 27(2): 293–314.
- Hussain, W., & Moriarty, J. 2016. Accountable to whom? Rethinking the role of corporations in political CSR. *Journal of Business Ethics*, 1–16.
- Husted, B. W., & Salazar, J. de J. 2006. Taking Friedman seriously: Maximizing profits and social performance. *Journal of Management Studies*, 43(1): 75–91.
- Jaworski, P. M. 2014. An absurd tax on our fellow citizens: The ethics of rent seeking in the market failures (or self-regulation) approach. *Journal of Business Ethics*, 121(3): 467–476.
- Jensen, M. C. 2002. Value maximization, stakeholder theory, and the corporate objective function. *Business Ethics Quarterly*, 235–256.
- Jones, T. M. 1995. Instrumental stakeholder theory: A synthesis of ethics and economics. *Academy of Management Review*, 20(2): 404–437.
- Jones, T. M., Donaldson, T., Freeman, R. E., Harrison, J. S., Leana, C. R., et al. 2016. Management theory and social welfare: Contributions and challenges. *Academy of Management Review*, 41(2): 216–228.
- Jones, T. M., & Felps, W. 2013. Shareholder wealth maximization and social welfare: A utilitarian critique. *Business Ethics Quarterly*, 23(2): 207–238.

- Kant, I. 2002. *Groundwork for the Metaphysics of Morals*. (J. B. Schneewind, Ed.). Yale University Press.
- Kapur, D., Mukhopadhyay, P., & Subramanian, A. 2008. The case for direct cash transfers to the poor. *Economic and Political Weekly*, 37–43.
- Kim, E.-H., & Lyon, T. P. 2014. Greenwash vs. brownwash: Exaggeration and undue modesty in corporate sustainability disclosure. *Organization Science*, 26(3): 705–723.
- King, A. A., & Lenox, M. J. 2000. Industry self-regulation without sanctions: The chemical industry's responsible care program. *Academy of Management Journal*, 43(4): 698–716.
- King, A., Lenox, M. J., & Barnett, M. L. 2002. Strategic responses to the reputation commons problem. In A. Hoffman & M. Ventresca (Eds.), *Organizations, Policy and the Natural Environment: Institutional and Strategic Perspectives*.: 393–406. Stanford, CA: Stanford University Press.
- Kirzner, I. M. 1997. Entrepreneurial discovery and the competitive market process: An Austrian approach. *Journal of Economic Literature*, 35(1): 60–85.
- Langtry, B. 1994. Stakeholders and the moral responsibilities of business. *Business Ethics Quarterly*, 431–443.
- Laufer, W. S. 2003. Social accountability and corporate greenwashing. *Journal of Business Ethics*, 43(3): 253–261.
- Lehrer, M. 2001. Macro-varieties of capitalism and micro-varieties of strategic management in European airlines. *Varieties of Capitalism-The Institutional Foundation of Comparative Advantage*. New York: Oxford.
- Lepak, D. P., Smith, K. G., & Taylor, M. S. 2007. Value creation and value capture: a multilevel perspective. *Academy of Management Review*, 32(1): 180–194.
- Lin, L.-W. 2010. Corporate social responsibility in China: Window dressing or structural change. *Berkeley J. Int'l L.*, 28: 64.

- Lipsey, R. G., & Lancaster, K. 1956. The general theory of second best. *The Review of Economic Studies*, 24(1): 11–32.
- List, J. A., & Momeni, F. 2017. *When Corporate Social Responsibility Backfires: Theory and Evidence from a Natural Field Experiment*. Working Paper no. 24169, National Bureau of Economic Research. <https://doi.org/10.3386/w24169>.
- Lyon, T. P., & Montgomery, A. W. 2015. The means and end of greenwash. *Organization & Environment*, 28(2): 223–249.
- MacAskill, W. 2015. *Doing good better: How effective altruism can help you help others, do work that matters, and make smarter choices about giving back*. Penguin.
- Mackey, A., Mackey, T. B., & Barney, J. B. 2007. Corporate Social Responsibility and Firm Performance: Investor Preferences and Corporate Strategies. *The Academy of Management Review*, 32(3): 817–835.
- Marcoux, A. M. 2003. A fiduciary argument against stakeholder theory. *Business Ethics Quarterly*, 13(1): 1–24.
- Marquis, C., Glynn, M. A., & Davis, G. F. 2007. Community isomorphism and corporate social action. *Academy of Management Review*, 32(3): 925–945.
- Marquis, C., & Park, A. 2014. Inside the buy-one give-one model. *Stanford Social Innovation Review*, (Winter).
- Marquis, C., Toffel, M. W., & Zhou, Y. 2016. Scrutiny, norms, and selective disclosure: A global study of greenwashing. *Organization Science*, 27(2): 483–504.
- Martin, D. 2013. The Unification Challenge. *Business Ethics Journal Review*, 1(5): 28–35.
- Matten, D., & Moon, J. 2008. “Implicit” and “explicit” CSR: A conceptual framework for a comparative understanding of corporate social responsibility. *Academy of Management Review*, 33(2): 404–424.

- Maxwell, J. C. 2007. *There's no such thing as "business" ethics: There's only one rule for making decisions*. Faithwords.
- Mazar, N., & Zhong, C.-B. 2010. Do green products make us better people? *Psychological Science*, 21(4): 494–498.
- McGillivray, M. 1991. The human development index: yet another redundant composite development indicator? *World Development*, 19(10): 1461–1468.
- McMahon, C. 1981. Morality and the invisible hand. *Philosophy & Public Affairs*, 247–277.
- McWilliams, A., & Siegel, D. 2001. Corporate social responsibility: A theory of the firm perspective. *Academy of Management Review*, 26(1): 117–127.
- Merritt, A. C., Effron, D. A., & Monin, B. 2010. Moral self-licensing: When being good frees us to be bad. *Social and Personality Psychology Compass*, 4(5): 344–357.
- Mill, J. S. 1978. *On Liberty*. (E. Rapaport, Ed.). Hackett Publishing.
- Monin, B., & Miller, D. T. 2001. Moral credentials and the expression of prejudice. *Journal of Personality and Social Psychology*, 81(1): 33.
- Moriarty, J. 2019. On the Origin, Content, and Relevance of the Market Failures Approach. *Journal of Business Ethics*, 1–12.
- Mullen, E., & Monin, B. 2016. Consistency versus licensing effects of past moral behavior. *Annual Review of Psychology*, 67.
- Nagel, T. 1986. *The View From Nowhere*. New York, NY: Oxford University Press.
- Neffke, F., & Henning, M. 2013. Skill relatedness and firm diversification. *Strategic Management Journal*, 34(3): 297–316.
- Nelson, R., & Winter, S. 1982. *An Evolutionary Theory of Economic Change*. Cambridge, MA: Harvard University Press.
- Noggle, R. 1996. Manipulative actions: a conceptual and moral analysis. *American Philosophical Quarterly*, 33(1): 43–55.

- Norman, W. 2011. Business ethics as self-regulation: Why principles that ground regulations should be used to ground beyond-compliance norms as well. *Journal of Business Ethics*, 102(1): 43–57.
- Nozick, R. 1974. *Anarchy, state, and utopia*, vol. 5038. New York: Basic Books.
- Nussbaum, M. C. 2001. Symposium on Amartya Sen’s philosophy: 5 Adaptive preferences and women’s options. *Economics & Philosophy*, 17(1): 67–88.
- Olson, M. 2009. *The Logic of Collective Action: Public Goods and the Theory of Groups, Second printing with new preface and appendix*, vol. 124. Harvard University Press.
- Peteraf, M. A. 1993. The cornerstones of competitive advantage: A resource-based view. *Strategic Management Journal*, 14(3): 179–191.
- Pettit, P. 2007. Responsibility incorporated. *Ethics*, 117(2): 171–201.
- Pirson, M., & Malhotra, D. 2008. Unconventional insights for managing stakeholder trust. *Sloan Management Review*, Summer.
- Pirson, M., & Malhotra, D. 2011. Foundations of organizational trust: What matters to different stakeholders? *Organization Science*, 22(4): 1087–1104.
- Plato. 2004. *Republic*. (C. D. C. Reeve, Tran.). Hackett Publishing Company.
- Porter, M. E. 2008. *Competitive Advantage: Creating and Sustaining Superior Performance*. Simon and Schuster.
- Porter, M. E., & Kramer, M. R. 2002. The competitive advantage of corporate philanthropy. *Harvard Business Review*, 80(12): 56–68.
- Porter, M. E., & Kramer, M. R. 2011. The Big Idea: Creating Shared Value. How to reinvent capitalism—and unleash a wave of innovation and growth. *Harvard Business Review*, 89(1–2).
- Porter, M., & Kramer, M. R. 2011. Creating shared value. *Harvard Business Review*, 89(1/2): 62–77.

- Posner, R. A. 1973. An economic approach to legal procedure and judicial administration. *The Journal of Legal Studies*, 2(2): 399–458.
- Prahalad, C. K. 2006. *The Fortune at the Bottom of the Pyramid*. Pearson Prentice Hall.
- Prahalad, C. K. 2009. *The fortune at the bottom of the pyramid, revised and updated 5th anniversary edition: Eradicating poverty through profits*. FT Press.
- Prahalad, C. K., & Hamel, G. 1990. The core competence of the corporation. *Harvard Business Review*, (May-June): 79–91.
- Pummer, T. 2016. Whether and where to give. *Philosophy & Public Affairs*, 44(1): 77–95.
- Quine, W. V. O., & Ullian, J. S. 1978. *The web of belief*, vol. 2. Random House New York.
- Radzik, L. 2004. Making amends. *American Philosophical Quarterly*, 41(2): 141–154.
- Radzik, L., & Murphy, C. 2015. Reconciliation. (E. N. Zalta, Ed.) *Stanford Encyclopedia of Philosophy*, Summer 2015 Edition.
- Railton, P. 1984. Alienation, consequentialism, and the demands of morality. *Philosophy & Public Affairs*, 134–171.
- Ramus, C. A., & Montiel, I. 2005. When are corporate environmental policies a form of greenwashing? *Business & Society*, 44(4): 377–414.
- Rawls, J. 1951. Outline of a decision procedure for ethics. *The Philosophical Review*, 60(2): 177–197.
- Rawls, J. 2001. *Justice as fairness: A restatement*. Harvard University Press.
- Rawls, J. 2009. *A theory of justice*. Harvard university press.
- Reed, R., & DeFillippi, R. J. 1990. Causal ambiguity, barriers to imitation, and sustainable competitive advantage. *Academy of Management Review*, 15(1): 88–102.
- Rodin, D. 2005. The ownership model of business ethics. *Metaphilosophy*, 36(1–2): 163–181.

- Rupp, D. E., Ganapathi, J., Aguilera, R. V., & Williams, C. A. 2006. Employee Reactions to Corporate Social Responsibility: An Organizational Justice Framework. *Journal of Organizational Behavior*, 27(4): 537–543.
- Scherer, A. G., & Palazzo, G. 2007. Toward a political conception of corporate responsibility: Business and society seen from a Habermasian perspective. *Academy of Management Review*, 32(4): 1096–1120.
- Schmidtz, D. 2006. *The elements of justice*. Cambridge University Press.
- Schmidtz, D., & Brennan, J. 2011. *A brief history of liberty*. John Wiley & Sons.
- Schnackenberg, A. K., & Tomlinson, E. C. 2016. Organizational transparency: A new perspective on managing trust in organization-stakeholder relationships. *Journal of Management*, 42(7): 1784–1810.
- Schumpeter, J. A. 1947. The creative response in economic history. *The Journal of Economic History*, 7(2): 149–159.
- Schumpeter, J. A. 2010. *Capitalism, socialism and democracy*. Routledge.
- Sen, A. 1991. *On Ethics and Economics*. Wiley.
- Siegel, D. S. 2009. Green management matters only if it yields more green: An economic/strategic perspective. *The Academy of Management Perspectives*, 5–16.
- Silver, D. 2016. Competition, Value Creation and the Self-Understanding of Business. *Business Ethics Journal Review*, 4(10): 59–65.
- Simler, K., & Hanson, R. 2017. *The Elephant in the Brain: Hidden Motives in Everyday Life*. Oxford University Press.
- Singer, A. 2018. Justice Failure: Efficiency and Equality in Business Ethics. *Journal of Business Ethics*, 97–115.
- Singer, P. 2015. *The Most Good You Can Do: How Effective Altruism Is Changing Ideas About Living Ethically*. Yale University Press.

- Smith, A. 1776. *An inquiry into the nature and causes of the wealth of nations*. London: W. Strahan and T. Cadell.
- Smith, A. 1904. *An Inquiry into the Nature and Causes of the Wealth of Nations by Adam Smith, edited with an Introduction, Notes, Marginal Summary and an Enlarged Index by Edwin Cannan*. London: Methuen.
- Smith, A. 1982. *The theory of moral sentiments*. (D. Raphael & A. Macfie, Eds.). Liberty Fund.
- Solman, G. 2008, January 14. BP: Coloring Public Opinion. *Adweek*.
<https://www.adweek.com/brand-marketing/bp-coloring-public-opinion-91662/>.
- Steinberg, E. 2017. The Inapplicability of the Market-Failures Approach in a Non-Ideal World. *Business Ethics Journal Review*, 5(5): 28–34.
- Stevenson, B., & Wolfers, J. 2013. Subjective well-being and income: Is there any evidence of satiation? *American Economic Review*, 103(3): 598–604.
- Strudler, A. 2005. Deception unraveled. *The Journal of Philosophy*, 102(9): 458–473.
- Strudler, A. 2017. What to Do with Corporate Wealth. *Journal of Political Philosophy*, 25(1): 108–126.
- Sundaram, A. K., & Inkpen, A. C. 2004. The corporate objective revisited. *Organization Science*, 15(3): 350–363.
- Suppe, F. 1972. What’s wrong with the received view on the structure of scientific theories? *Philosophy of Science*, 39(1): 1–19.
- Taurek, J. M. 1977. Should the numbers count? *Philosophy & Public Affairs*, 293–316.
- Tosi, J., & Warmke, B. 2016. Moral Grandstanding. *Philosophy and Public Affairs*, 44(3): 197–217.
- Tucker, A. 2008. Trade associations as industry reputation agents: A model of reputational trust. *Business and Politics*, 10(1): 1–26.

- Turban, D. B., & Greening, D. W. 1997. Corporate social performance and organizational attractiveness to prospective employees. *Academy of Management Journal*, 40(3): 658–672.
- Van der Vossen, B., & Brennan, J. 2018. *In defense of openness*. New York: Oxford University Press. Google Scholar.
- Velasquez, M. G. 1983. Why corporations are not morally responsible for anything they do. *Business & Professional Ethics Journal*, 2(3): 1–18.
- von Kriegstein, H. 2016. Professionalism, agency, and market failures. *Business Ethics Quarterly*, 26(4): 445–464.
- Vos, J. 2009. Actions speak louder than words: Greenwashing in corporate America. *Notre Dame JL Ethics & Pub. Pol’y*, 23: 673.
- Waddock, S. A., & Graves, S. B. 1997. The corporate social performance-financial performance link. *Strategic Management Journal*, 303–319.
- Weinrib, E. J. 2012. *Corrective justice*. Oxford University Press.
- Wenar, L. 2017. John Rawls. *Stanford Encyclopedia of Philosophy*.
- Wernerfelt, B. 1984. A resource-based view of the firm. *Strategic Management Journal*, 5(2): 171–180.
- Werther Jr, W. B., & Chandler, D. 2010. *Strategic corporate social responsibility: Stakeholders in a global environment*. Sage.
- Wettstein, F. 2012. CSR and the debate on business and human rights: Bridging the great divide. *Business Ethics Quarterly*, 22(4): 739–770.
- Wiggins, D. 1998. The Right and the Good and WD Ross’s Criticism of Consequentialism. *Utilitas*, 10(3): 261–280.
- Williams, J. R. 1992. How sustainable is your competitive advantage? *California Management Review*, 34(3): 29–51.

Williamson, O. E. 1979. Transaction-cost economics: the governance of contractual relations.

The Journal of Law and Economics, 22(2): 233–261.

Winter, S. G. 1997. Knowledge and competence as strategic assets. *The strategic management of*

intellectual capital: 165–187. Elsevier.

Wydick, B., Katz, E., Calvo, F., Gutierrez, F., & Janet, B. 2016a. Shoeing the Children: the

impact of the TOMS Shoe donation program in rural El Salvador. *The World Bank*

Economic Review.

Wydick, B., Katz, E., Calvo, F., Gutierrez, F., & Janet, B. 2016b. Shoeing the Children: the

impact of the TOMS Shoe donation program in rural El Salvador. *The World Bank*

Economic Review, 32(3): 727–751.

Zwolinski, M. 2007. Sweatshops, choice, and exploitation. *Business Ethics Quarterly*, 17(4):

689–727.