



On the Origin, Content, and Relevance of the Market Failures Approach

Jeffrey Moriarty¹

Received: 1 November 2018 / Accepted: 8 January 2019 / Published online: 17 January 2019
© Springer Nature B.V. 2019

Abstract

The view of business ethics that Christopher McMahon calls the “implicit morality of the market” and Joseph Heath calls the “market failures approach” has received a significant amount of recent attention. The idea of this view is that we can derive an ethics for market participants by thinking about the “point” of market activity, and asking what the world would have to be like for this point to be realized. While this view has been much-discussed, it is still not well-understood. This paper seeks to remedy this problem. I begin by showing, against some recent commentators, that McMahon’s view and Heath’s view are fundamentally the same. Second, I clarify the sense of “efficiency” at work in the market failures approach. Finally, I argue that, in its current form, this view has little relevance to the real world of business. I conclude by sketching two ways of modifying it to fit our world.

Keywords Efficiency · Heath · Ideal theory · Market failures · McMahon

The view of business ethics that McMahon (1981, 2013) calls the “implicit morality of the market” and Heath (2006, 2014) calls the “market failures approach” has received a significant amount of recent attention (see, e.g., Jaworski 2014; Martin 2013; Néron 2016; Norman 2011; Silver 2016; Singer 2018; Smith 2018; von Kriegstein 2016). The idea of this view is that we can derive an ethics for market participants by thinking about the “point” of market activity, and asking what the world would have to be like for this point to be realized. While this view has been much-discussed, it is still not well understood. The goal of this paper is to remedy this problem. In doing so, I show that, in its current form, the theory has little relevance to the real world.

I begin by arguing that McMahon’s and Heath’s views are fundamentally the same: they share the same core features, and it is not clear how, if at all, they differ. Heath’s view has been described as “novel” and “new” (see, respectively, Norman 2011; von Kriegstein 2016), but it is better seen as a variation of McMahon’s view. Heath himself does not conceal the connection between his view and McMahon’s, but some of Heath’s interlocutors are less careful. It is worth

understanding the connection between Heath’s view and McMahon’s view because, as we will see, they reach different conclusions about its implications. Next, I clarify the concept of “efficiency” at work in the implicit morality of the market/market failures approach. McMahon and (especially) Heath say they are appealing to the Pareto criteria, but their use of this term is often looser, meaning something more like Kaldor–Hicks efficiency or aggregate welfare. This is not a mere terminological clarification but an important feature that goes to the nature and plausibility of their view. Finally I ask whether the implicit morality of the market/market failures approach has any implications for the real world. Most writers assume that it does, and that these implications are clear (cf. Steinberg 2017). Both McMahon and Heath are more circumspect. McMahon, in fact, calls his view “an ‘ideal’ theory of business morality” with “only indirect relevance in most actual business situations” (1981, p. 256). Heath acknowledges the force of McMahon’s reasoning, but thinks that, with further argument, the view can be made relevant for real economic agents. In this dispute, I side with McMahon. It follows that, if we want a theory of business ethics to tell us what real businesspeople ought to do—and I think we do—we must look elsewhere. This does not mean that we should consign the McMahon/Heath view to the dustbin of business ethics. I conclude by sketching two ways of modifying it to fit our world.

✉ Jeffrey Moriarty
jmoriarty@bentley.edu

¹ Department of Philosophy, Bentley University, 175 Forest St., AAC 109, Waltham, MA 02452, USA

The Implicit Morality of the Market View/ Market Failures Approach

In this section, I argue for my first point: that McMahon's implicit morality of the market and Heath's market failures approach to business ethics are fundamentally the same view. I begin with McMahon, then consider Heath.

Before we begin, a reminder about scope. We are considering the *content* and *implications* of McMahon's and Heath's view, not its justification. So we will consider McMahon's and Heath's explications of their view but not their arguments for it.¹

McMahon on the Implicit Morality of the Market

McMahon's summary of the implicit morality of the market goes by quickly. He says:

The implicit morality of the market consists primarily of the hypothetical imperatives which are generated by economic theory when the achievement of economic efficiency is taken as an end. Certain conditions must be satisfied if a free-enterprise system is to allocate resources to producers and distribute products to consumers in a Pareto optimal way. And from these conditions various requirements on the behavior of economic agents—they might be called “efficiency imperatives”—can be derived. (1981, p. 255)

This requires unpacking. First consider “economic efficiency.” McMahon defines “efficiency” in terms of Pareto optimality, so an efficient outcome is one in which no one can be made better off without at least one person being made worse off. The reference to “economic theory” is a reference to the first fundamental theorem of welfare economics. According to this theorem, perfectly competitive markets at equilibrium yield Pareto optimal outcomes. In a market, firms decide what to produce and how much to charge for it, and individuals decide what and how much of it to buy. In saying that market outcomes are Pareto optimal, what is meant is: given a certain initial resource distribution among firms and individuals, people will trade with each other in order to improve their positions until further improvements are impossible. But markets yield Pareto optimal outcomes only when certain conditions are met. These are the conditions of “perfect competition” and include conditions such as no market power, perfect information, homogenous products, frictionless movement of factors of

¹ Both give teleological arguments. Economic agents are justified in following the implicit morality of the market/market failures approach, McMahon and Heath say, because of the value of the outcome that obtains when they do.

production, and no barriers to entry or exit. These conditions are the source of the “hypothetical imperatives” McMahon refers to in the above passage. From these conditions, he says, “various requirements on the behavior of economic agents... can be derived” (1981, p. 255). Put another way, we can “read off” what economic agents should do from the conditions of perfect competition.

McMahon cites “conspiracy in restraint of trade” as an example. “The potential for economic efficiency associated with a free-enterprise system cannot be fully realized,” he says, “if economic agents create or maintain... monopolistic or oligopolistic practices” (1981, p. 255). So the implicit morality of the market prohibits agents from engaging in these practices. Another example McMahon gives concerns information. Markets achieve Pareto optimal outcomes only if agents have perfect information. As a result, the implicit morality of the market contains “various requirements” regarding information. First, market participants should not lie about “the nature or quality of what is sold.” Second, they should not attempt to deceive trading partners “by methods other than lying, such as deceptive packaging of products” (1981, p. 257). Third, they should not exploit “the ignorance of those who are misinformed” (1981, p. 257). Instead, they “must... correct any mistaken beliefs which those with whom [they deal] may have about the properties or market value of what is being exchanged” (1981, p. 257). Similar results can be derived for the other conditions of perfect competition.²

Heath on the Market Failures Approach

Like McMahon, Heath says that the purpose of markets is the achievement of an efficient allocation of resources. As Heath says, “the central rationale for having private profit-seeking firms is to establish competition among suppliers and consumers. This competition drives prices toward market-clearing levels, allowing society... to generate a more efficient allocation of its resources and labor time” (2013, p. 31). Also like McMahon, Heath understands efficiency in terms of Pareto optimality. The “Pareto principle,” he

² The implicit morality of the market has implications for market participants. As McMahon says, it is a set of “requirements on the behavior of economic agents” (1981, p. 255). But it also has implications for other members of society. Legislators and regulators must articulate and enforce the requirements for market participants, and otherwise ensure that the conditions of perfect competition obtain. To take but one example, governments must take steps to end pay secrecy, which is a significant source of information asymmetry in labor markets (Moriarty 2018). Moreover, when economic activity conforming to these requirements produces harm to individuals—e.g., when workers are laid off in an economic downturn—society may owe it to them to ameliorate the harm—e.g., by sponsoring welfare or job training programs.

says, “forms the normative core of... [the] ‘market failures’ approach to business ethics” (2013, p. 174). And finally like McMahon, Heath identifies the conditions of perfect competition as providing a guide to ethical market conduct. In particular, Heath says that market participants should not violate, or take advantage of deviations from, the conditions of perfect competition. For example, firms should not “[produce] pollution, or [sell] products with hidden quality defects” (2013, p. 88). They should not lie, nor should they engage in trades when they believe their trading partners have materially false beliefs. Heath calls these behaviors “nonpreferred,” by which he means *ethically impermissible*. These behaviors are incompatible with the conditions of perfect competition. Instead, they generate market failure, understood as an inefficient allocation of resources.³

Unlike McMahon, Heath identifies not just what market participants should *not* do, but what they—or at least productive organizations—*should* do. They should compete in the market by making new products, making their products better, or selling them for less. Heath calls these behaviors “preferred,” by which he means *ethically permissible* (2014, p. 88).⁴ When the conditions of perfect competition are met, Heath says, these “would be the only way[s] that firms could compete with one another” (2014, p. 88).⁵

It might be thought that Heath’s specification of what market participants should do (not just what they should not do) represents a deviation from McMahon’s view. But it is best seen as an elaboration of an idea that was latent in McMahon. McMahon believes that firms should not create or take advantage of market failures, including imperfect information or market power. It follows that they cannot compete with each other on this basis. Rather, they can only compete through behaviors that are consistent with the conditions of perfect competition, viz. they should try to make

their products cheaper or better, or they should try to make new products.

How are These Views Different?

To recap, according to McMahon’s implicit morality of the market, (a) the purpose of markets is to produce efficient, in the sense of Pareto optimal, outcomes. Markets produce these outcomes under certain conditions, viz. the conditions of perfect competition. (b) These conditions in turn are the source of ethical rules for market participants. (c) Market participants should not cause the conditions of perfect competition not to obtain, nor should they engage in transactions when these conditions do not obtain.⁶ Heath’s market failures approach accepts (a), (b), and (c). Given this, we can see the truth of Heath’s claim that McMahon’s implicit morality of the market forms the “normative core” of Heath’s own market failures approach.

We might think at this point that these are exactly the same views. What’s the difference? Heath supplies an answer. He describes McMahon’s view that efficiency imperatives can be derived from the conditions of perfect competition as “not overly controversial.” Heath goes on: “the claim that I want to make, the one that is controversial, is that these ‘efficiency imperatives’ are pretty much all there is to business ethics, at least with respect to market transactions” (2014, p. 174).

I do not think Heath is saying that the efficiency imperatives *themselves* are noncontroversial. These same imperatives are part of Heath’s market failures approach, and Heath thinks that some of them are extremely demanding and would be rejected by most business ethicists. This includes

³ For Heath, a market failure *just is* “a situation in which the competitive market fails to produce a Pareto efficient outcome” (2006, p. 549).

⁴ What’s the connection between ethics and law on the MFA? McMahon is not interested in this question, but Heath is. Heath’s answer seems to be that ethics provides the foundation for law, but not all parts of the MFA can be codified in law. Suppose that a prohibition on stealing is part of the MFA. We might turn this into a law against stealing, enforced by the state. But suppose that a prohibition on marketing to children is also part of the MFA. This might be impossible to translate into a formal law and/or have it enforced by the state. So it might remain (merely) an ethical rule. Since “the law is a somewhat blunt instrument... the deadweight losses imposed through use of the legal mechanism can easily outweigh whatever efficiency gains might have been achieved through the intervention. This often makes legal regulation unfeasible or unwise” (2014, p. 89).

⁵ This is not quite right. A standard assumption of perfect competition is “homogenous products.” Since there is no product differentiation in a perfectly competitive market, there is no “competing on quality.”

⁶ McMahon and Heath are not the only writers to articulate this view. In a little-cited article, Holley derives an ethics for salespeople from an efficiency-based justification of the market. Holley says that the “primary justification for a market system is that it provides an efficient procedure for meeting people’s needs and desires for goods and services” (1986, p. 3). The reason this is so, according to Holley, is that “people will efficiently serve each other’s needs if they are allowed to engage in voluntary exchanges” (1986, p. 4). However, for the exchanges people engage in to be truly voluntary, certain conditions must obtain. It must be the case that “[b]oth buyer and seller understand what they are giving up and what they are receiving in return. Neither buyer nor seller is compelled to enter into the exchange as a result of coercion, severely restricted alternatives, or other constraints on the ability to choose. Both buyer and seller are able at the time of exchange to make a rational judgment about its costs and benefits” (1986, p. 4). According to Holley, the behavior of salespeople must be adjusted in light of these conditions. Salespeople should not deceive their customers, either by act (lying) or omission (failure to disclose relevant information). “To behave in such ways is to undermine the conditions which are presupposed in teleological justifications of the market system” (1986, p. 5). For space reasons, and because Holley’s view differs subtly from McMahon and Heath’s, I do not include it in my analysis.

the requirement not to engage in transactions when there is asymmetric information. (To appreciate the demandingness of this requirement, note that it would require, e.g., a hedge fund manager to disclose everything she knows about the financial health of a firm to potential buyers before selling its stock.) Rather, Heath is saying that *the idea that efficiency imperatives can be derived from the conditions of perfect competition* is not overly controversial. What Heath thinks is controversial is his claim that these efficiency imperatives are “pretty much all there is to business ethics, at least with respect to market transactions.” The implication here is that McMahon thinks that there is something more to business ethics than Heath does. Heath limits business ethics to efficiency imperatives, and McMahon thinks that business ethics includes efficiency imperatives plus some other principles.

Heath does not say what McMahon thinks the “more” is to business ethics than efficiency imperatives. But a closer look at McMahon (1981) reveals that he thinks that some of the principles of “common morality”—like “respect for employee autonomy” and “do no harm”—can override the implicit morality of the market. Considerations of efficiency might require that managers make decisions in the least costly way possible, but respect for employee autonomy might require giving employees “a voice in policy formation” (1981, p. 276). Considerations of efficiency might require the totally free movement of factors of production, including labor. Employees should be permitted to leave employers on a moment’s notice, and employers should be permitted to “leave” (by firing) their employees, also on a moment’s notice. But common morality might require employers to soften the blow by giving workers warning when layoffs are coming (see also McMahon 2013, p. 133).

So Heath is correct that, for McMahon, doing the right thing in the context of business is about more than just complying with the market’s efficiency imperatives. It requires taking common morality into consideration. But Heath *himself* does not say that doing the right thing in the context of business is *only* about complying with efficiency imperatives. He says efficiency imperatives are “pretty much” all there is to business ethics, not all there is to business ethics *simpliciter*. Unfortunately, Heath does not explain the “pretty much” qualifier. So for all we know, Heath might endorse McMahon’s exact position. He could say that efficiency imperatives are *most* of business ethics, but the principles of what McMahon calls “common morality” can also be taken into consideration.

Let sum up. McMahon’s implicit morality of the market view and Heath’s market failures approach share the idea that the correct ethics for business will include “efficiency imperatives,” or ethical rules that derive from the conditions of perfect competition. Where they *might* disagree is about what else business ethics requires. McMahon is clear that

business ethics will take into account common morality in some cases. Heath suggests that business ethics may take into account some principles besides the MFA’s efficiency imperatives, but does not say which ones. Based on what these authors have said, we cannot say whether these views are mostly the same or exactly the same. But they are certainly very similar and could be exactly the same.⁷

Below I will inquire into the practical relevance of the efficiency imperatives that are common to the implicit morality of the market and the market failures view. But for convenience I will choose a single nomenclature, and I will choose Heath’s. So I will refer to these as the efficiency imperatives of the market failures approach, or the *MFA’s efficiency imperatives*. This choice makes sense because Heath’s terminology dominates the literature. But we should keep in mind that these same efficiency imperatives are found in McMahon.

Does the MFA have any Implications for the Real World?

Having sketched the MFA, my next question is: what are its implications for real people? This might seem like a silly question. Haven’t we just seen what the MFA’s implications are? Firms shouldn’t engage in price-fixing or other forms of collusion, they shouldn’t deceive customers, and so on. But in fact, McMahon and Heath both believe that the ethical rules generated by the MFA are not straightforwardly applicable—or not applicable at all—to the actual business world.

McMahon is explicit about this. He says that the MFA is an account of the duties of market participants “in an *optimally functioning* free-enterprise system—that is, under the ideal conditions assumed by the theory of the firm and general equilibrium theory” (1981, p. 256, emphasis in original). Since these ideal conditions do not obtain in the real world, McMahon says, “any conclusions which an investigation of this sort might reach would be of only indirect relevance in most actual business situations” (1981, p. 256). Heath agrees, to an extent (though as we will see he ultimately tries to rescue the MFA’s efficiency imperatives from mere ideality). He says: “we cannot use the [first fundamental

⁷ I have suggested that the strong similarity between Heath’s view and McMahon’s has not been appreciated by many recent commentators. I offer two reasons for this. First, when Heath introduced his theory in his 2006 article, “Business Ethics Without Stakeholders,” he did not reference McMahon’s work. His references to McMahon only come later, in (what I think is) the definitive statement of his view in his 2014 book *Morality, Competition, and the Firm*. Second, Heath does not use McMahon’s terminology; he gives his view a different name. It is possible, as I mentioned, that it is sufficiently different that it deserves a different name. But it is also possible that it is not sufficiently different, and should go by the same name.

theorem of welfare economics] to derive normative conclusions under real-world circumstances” (2014, p. 40). The first fundamental theorem tells us what economic agents ought to do when conditions are perfect, but not when conditions are imperfect, which in the real world, they always are.

Why does McMahon think that the MFA’s efficiency imperatives do not apply to economic agents in the real world? He explains as follows: “Results in the ‘theory of the second best’ seem to show that a set of efficiency imperatives for economic agents in the suboptimal conditions which actually obtain will not soon be found. The principal result in this area of welfare economics is a proof that partial satisfaction of the marginal conditions for a Paretian optimum will not necessarily result in increased welfare” (1981, p. 256). Heath gives the same explanation. He says that the problem with applying the MFA’s efficiency imperatives directly to the real world “arises from what is known as the ‘general theory of the second best’... [which] shows that in a situation in which one of the Pareto conditions is violated, respect for all of the other Pareto conditions will generate an outcome that is less efficient than some other outcome that could be obtained by violating one or more of the remaining conditions” (2013, p. 39).

These claims require explanation. We know that, under conditions of perfect competition, markets produce efficient outcomes. But now suppose that the conditions of perfect competition do not obtain. Indeed, suppose—as is often the case—that the market contains multiple imperfections. The theory of the second best says that we do not get closer and closer to an efficient outcome by eliminating more and more imperfections (Lipsey and Lancaster 1956).⁸ Indeed, we might get closer to an efficient outcome by adding a new imperfection than by eliminating an existing imperfection. An example will make this clearer. Suppose that the government subsidizes the production of natural gas, oil, and coal. These are market imperfections that encourage the overconsumption of these fuels, compared to other fuels such as solar power. The theory of the second best says that it might be better, from the point of view of efficiency, for the government to start subsidizing the production of solar power than for it to eliminate *one* of its existing subsidies for natural gas, oil, or coal.

McMahon and Heath take this result to imply that the MFA’s efficiency imperatives do not apply to economic agents in the real world. This means, I suggest, that they are thinking as follows. The MFA’s efficiency imperatives are justified *when and because* they lead to efficient outcomes. Compliance with these imperatives leads to efficient

outcomes in a perfectly competitive market. But compliance with them does not necessarily lead to efficient outcomes in an imperfect market. So they are *not justified* in an imperfect market. One of the MFA’s efficiency imperatives, according to Heath, is “do not seek tariffs or other protectionist measures” (2014, p. 37). This is justified in a perfectly competitive market because introducing a tariff in such a market creates an imperfection, making the market less efficient. But in an imperfect market—for example, in a market with multiple tariffs already in place—this is not necessarily the case. As Heath says: “if there is even one trade barrier or tariff in place, then minimizing the number of tariffs will not necessarily produce the best outcome—we may be better off imposing some additional tariffs” (2014, p. 40). The MFA’s efficiency imperatives are for ideal, not real, economic agents.

As noted, McMahon, stops at this point. He calls his view “an ‘ideal’ theory of business morality” with “only indirect relevance in most actual business situations” (1981, p. 256). Heath, however, is not satisfied with this result. While agreeing that one cannot simply “read-off” ethical rules for real economic agents from the conditions of perfect competition, Heath thinks that, with additional argument, the MFA’s efficiency imperatives “could be justified in some form” (2014, p. 41). I will return to this claim below. But here it is necessary to say something about the notion of efficiency that is now being used in the MFA, especially in the appeal to the theory of the second best.

Efficiency: Pareto Optimality or Something Else?

In their discussion of the relevance of the second best theorem for the MFA, McMahon and Heath slide from using “efficiency” in the strict sense of *Pareto optimality* to using it in a looser sense. This is an important change which merits our attention.

The first fundamental theorem of welfare economics says that, under conditions of perfect competition, markets produce *efficient* outcomes. Here efficiency means *Pareto optimality*. This is the definition of “efficiency” assumed in the mathematical proof of this theorem (Mas-Colell et al. 1995). According to McMahon and Heath, the theory of the second best says that in a market with many imperfections, we do not get closer and closer to an *efficient* outcome by eliminating more and more imperfections. We might get a more efficient outcome by adding an imperfection (e.g., a tariff) into the market than by eliminating an existing imperfection. In the previous two sentences, “efficient” cannot mean Pareto optimality in the strict sense. It must mean something like *Kaldor–Hicks efficient* or *aggregate-welfare-enhancing*. This is because a more efficient outcome—a Pareto

⁸ Following McMahon and Heath, here I use the term ‘efficiency’ loosely, not in the strict sense of Pareto optimality. I explain this looseness in the next section.

improvement—is one in which at least one person is made better off and *no one* is made worse off. But almost any policy change, such as the introduction of a tariff, will create winners *and* losers.

Consider again the market with government subsidies for the production of natural gas, oil, and coal. Call this outcome O1. The theory of the second best, according to McMahon and Heath, says that we might get a more efficient result if we add an additional imperfection—e.g., a subsidy for solar power—than if we eliminate one of the existing subsidies for a fossil fuel. Suppose we do add a subsidy for solar power. Call this new outcome O2. O2 will almost certainly not be a Pareto improvement over O1. This is because, while a subsidy for solar power may benefit many people (e.g., producers of solar panels, people who live in places with poor air quality), it will hurt some others (e.g., oil companies). Yet O2 might be a Kaldor–Hicks improvement, or an improvement in aggregate welfare, compared to O1. A move from one distribution to another is a Kaldor–Hicks improvement, or an improvement in Kaldor–Hicks efficiency, if those who are better off in the new outcome could compensate those who are worse off in the new outcome and still be better off. (The winners don't actually have to compensate the losers for it to be a Kaldor–Hicks improvement. It just has to be the case that they *could* do so and still be better off.) If we understand “better off” and “worse off” in these claims in terms of welfare, then a Kaldor–Hicks improvement is an improvement in aggregate welfare.⁹

The point can also be made this way. In the real world, almost any distribution is Pareto optimal. For almost any move away from it is going to make someone worse off, even if it makes a lot of people better off, and makes those people much better off. But we can still compare the efficiency—in a different sense—of two states of affairs in the real world, both of which are Pareto optimal. We can compare them in terms of Kaldor–Hicks efficiency or aggregate welfare.

In fact, McMahon and Heath understand this, though it does not always come through clearly. Recall that McMahon says that the theory of the second best tells us that “partial satisfaction of the marginal conditions for a Paretian optimum will not necessarily result in increased *welfare*” (1981, p. 256, emphasis added). McMahon realizes that partial satisfaction of these conditions may result in a Pareto optimal state of affairs. For, as we said, all or almost all states of

affairs in the real world are Pareto optimal in the strict sense. But, as he says, this state of affairs may still be *worse* than a more complete satisfaction of these conditions in terms of *welfare*.¹⁰ Heath also understands that we have moved away from Pareto optimality in the strict sense when we begin to talk about “efficient outcomes” under conditions of imperfect competition. Heath says that the justification of the market is that “a well-structured, competitive market economy produces not just utilitarian gains (where some might benefit while others lose), but Pareto improvements (where everyone benefits).” He goes on: “Naturally this Paretianism must be hedged and qualified in numerous ways, in recognition of the fact that any action that affects a sufficiently large number of people is bound to produce both winners and losers” (2014, p. 197). Elsewhere Heath speaks of a commitment to “Pareto efficiency, modulated by a ‘realistic’ accommodation of the fact that literal Pareto improvements are few and far between” (2014, p. 198).

It is confusing, I submit, for McMahon and (especially) Heath to use the language of Pareto efficiency. If our “Paretianism” is, as Heath says, “hedged and qualified in numerous ways,” then it is no longer *really* Paretianism. If we concede, as we should, that “literal Pareto improvements are few and far between,” and we “modulate” our understanding of Pareto efficiency to accommodate this fact, then we are no longer *really* talking about Pareto efficiency. The continued use of the Pareto terminology would be understandable if alternatives were not available. But they are. We can compare outcomes in terms of Kaldor–Hicks efficiency or aggregate welfare. And in reality this is what McMahon and Heath are doing. Their terminology should reflect that.

The foregoing might seem like a minor and maybe even pedantic clarification of terminology but in fact it has important consequences for an evaluation of the MFA. As we noted, the MFA's efficiency imperatives appear to be justified by their contribution to certain outcomes. When these imperatives are followed, at least under certain conditions, then an *efficient* outcome is achieved. But then we must ask: what's so important about an efficient outcome? If it is said “that outcome is Pareto optimal,” then that is not a very compelling answer. All or almost all actual outcomes are Pareto optimal, since given these outcomes it will often be impossible to make anyone better off without making at

⁹ Suppose that in outcome O3 P has 10 units of welfare and Q has 15 units of welfare. In O4, P has 20 units and Q has 14. The move from O3 to O4 is not a Pareto improvement, because Q is worse off in O4. But it is a Kaldor–Hicks improvement, since (in principle) P could compensate Q—by transferring some of her resources to Q—and still be better off than she was in O3. Because P can do this, this means that O4 has more aggregate welfare than O3. In our example, O4 has 34 units total, whereas O3 has 25.

¹⁰ McMahon takes a similar position in his (2013). When introducing the market failures approach, he defines ‘efficiency’ as Pareto optimality, which is how efficiency is understood in the first fundamental theorem of welfare economics (p. 114). But he subsequently connects efficiency with “social prosperity” (p. 117), which he defines as “the enjoyment, by the members of a polity, of goods and services that have been produced... with the polity's own resources” (p. 133). This is the looser sense of efficiency which I have said is Kaldor–Hicks efficiency or aggregate welfare.

least one person worse off. But if it is said “that outcome is Kaldor–Hicks optimal” or “aggregate welfare is maximized in that outcome,” then that is a much more compelling answer. It implies that that outcome is, along one dimension, the best of a large set of possible outcomes.

To be clear, I am not pressing a criticism of the MFA’s *substance*. It is a criticism, and more importantly, a clarification, of the terminology its proponents employ. Indeed, this clarification should make the MFA *more plausible*. The Pareto criteria are weak. Kaldor–Hicks and aggregate welfare are much stronger. If the MFA’s imperatives could be shown to promote efficiency in these senses, then that would be a powerful reason to comply with them.

Heath’s Attempt to Save the MFA’s Efficiency Imperatives

We paused to examine the meaning of “efficiency” in the MFA. Now let us return to the question of whether the MFA’s efficiency imperatives apply in the ideal world only, as McMahon believes, or in the real world as well, as Heath believes. In this section, I examine Heath’s explanation of why the MFA’s efficiency imperatives apply in the real world, and argue that it doesn’t work.¹¹ I quote at length from Heath’s text, as his reasoning is subtle.

Heath’s Attempt

Heath says that the theory of the second best “does not mean the efficiency standard is deprived of all normative force. It simply means that we cannot make... big sweeping generalizations... Moral reasoning in a business context must be a more contextual affair” (2014, p. 40). I take this to mean: we cannot say that it is always correct—i.e., always conducive to efficiency—for businesspeople not to (e.g.) lobby for special tax breaks or not to collude with each other. Instead, whether it is permissible for businesspeople to do these things is a “contextual affair.” That is, it depends on the nature of the specific circumstances in which they find themselves.

So far this is mostly just a statement of the problem that we sketched in the previous section. The question is what the MFA’s implications are, given this problem. But this statement also provides a clue as to how Heath wants to proceed. He directs our attention away from universal principles for market actors—the sorts of rules that the MFA yields for

conditions of perfect competition—to the individual transactions that comprise “the market.” Heath observes:

Every voluntary exchange generates a Pareto improvement. It is through these tangible, incremental efficiency gains that the private market system has established its merit. Thus, instead of offering a “top-down” justification of profit-seeking – through appeal to the general equilibrium of the economy as a whole, one could adopt a more bottom-up strategy, which would appeal to the particular efficiency gains that the firm is able to realize among its shareholders, its employees, and its customers. (2014, p. 40)

Heath characterizes this as a “resource custodianship” perspective of business ethics. He says “the ultimate goal of the economy as a whole is to satisfy human needs,” so we should think of “all productive resources as being ‘earmarked’ for the satisfaction of needs.” Thus, “whenever the firm uses these resources in a way that does not contribute to welfare, but rather imposes deadweight losses on the economy as a whole, it is acting as a poor custodian of these resources” (2014, pp. 40–41).

Heath believes that, “using this sort of “‘bottom-up’ reasoning... all of [the MFA’s efficiency imperatives] could be justified in some form” (2014, p. 41). This includes the rules against deception and collusion, as well as the rules against lobbying for tariffs and other sorts of governmental protections. In this framework, Heath continues, “the Pareto conditions would function as a set of heuristics, allowing us to determine what type of conduct, in general, is likely to constitute an illegitimate source of gain. However, actually making the case requires a more detailed analysis, one that examines the specific conditions of the market in question” (2014, p. 41).¹²

¹² Interestingly, while Heath thinks that the MFA’s efficiency imperatives do apply to real economic agents, he thinks that people are sometimes *excused* from following them. He says that “[u]nder the assumption of full compliance, any deviation from the deontology prescribed by the market failures view is unethical” (2014, p. 202). But “because of the competitiveness of the market economy, noncompliance by one firm can put very serious pressure on all of its competitors” not to comply (2014, p. 202). So if, for example, “all of one’s competitors are exploiting a particular regulatory loophole, this does not make it right to do so, but it may provide one with a reasonable excuse for acting wrongly” (2014, p. 202; see also p. 37). According to these passages, a person who exploits this type of loophole, or who otherwise undermines the conditions of perfect competition, does something wrong, but she is not to be blamed for doing it. This seems to be Singer’s reading of Heath. Singer says that “much of what is required by managers under the MFA still winds up being overly demanding in the context of the actual conditions of the market economy” (2018, p. 50). Singer’s point is: given that it is *too* demanding, managers are excused from doing what the MFA requires, though the MFA is still what morality requires.

¹¹ Steinberg (2017) is a rare author who recognizes that the MFA’s efficiency imperatives do not apply in the real world, for the reasons outlined above. But Steinberg does not seem to see that Heath recognizes the problem and tries to rescue these imperatives, using arguments I discuss below.

Problems with Heath's Attempt

I do not think Heath's "bottom-up" reasoning can show that the MFA's efficiency imperatives are applicable in the real world. This approach tells us to focus on individual transactions in a market. As Heath says, we should "appeal to the particular efficiency gains that the firm is able to realize among its shareholders, its employees, and its customers" (2014, p. 40). Each of these transactions, he says, "generates a Pareto improvement." This seems right. Under normal circumstances, you would not engage in a voluntary transaction—a transaction that was free and informed—unless it made you better off. But what follows about the MFA from this?

You might think that Heath is saying is that, in the real world, the MFA requires economic agents to make their trading partners better off. As a seller of a good or service, for example, you do the right thing insofar as the person who buys it is made better off.

That seems like a reasonable requirement. But I do not see how it yields in the real world anything like the MFA's efficiency imperatives. In the real world, we can make our trading partners better off through all kinds of activities that run afoul of these imperatives. Suppose I engage in collusion with another producer to fix the price of a good that we both sell at a certain high level. Our collusive activity makes us, the producers, better off. And the people who voluntarily buy the good that we sell are also made better off. If they were not, then they would not have bought the good. But the MFA's efficiency imperatives expressly prohibit collusion. Or suppose I lobby successfully for a tariff on the goods of a foreign competitor. Presumably the way I do this is by persuading a politician that this is the right thing to do. I might have to make a contribution to the politician's reelection fund to gain access to her. When this "transaction" is complete, the politician is no worse off, and may be better off. And if people are still buying my goods at the higher price I can now charge, then they are better off too. But the MFA's efficiency imperatives expressly prohibit this activity. The requirement to benefit your trading partner cannot even justify the efficiency imperative prohibiting fraud. For it may be possible to make your trading partner better off through fraud—say if he is confused about the value of what he is buying or if he is not thinking clearly about his own interests.

Perhaps Heath would say at this point that I'm *stuck* at the bottom. He is using bottom-up reasoning, so we have to look at what the impact of a certain behavior (e.g., collusion, lobbying, deception) would be not just on one's trading partners, but on society as a whole. In support of this, Heath says that "whenever the firm uses... resources in a way that does not contribute to welfare, but rather imposes deadweight losses on the *economy as a whole*, it is acting as a poor

custodian of these resources" (2014, p. 41, emphasis added), and is doing the wrong thing. Here the focus is on social welfare, not individuals making trades. And it is plausible to suppose that collusion among producers is generally worse for society than competition among them. It is worse for society as a whole if firms lobby for protective tariffs than if they do not. And so on for the MFA's other imperatives.

This response doesn't work. Once we move "up" from the bottom—i.e., we move from a consideration of whether our trading partner is better off if we act in a certain way to a consideration of whether society as a whole is better off if everyone acts in that way—we confront the original problem, viz. the problem of the second best. In a world in which no one is lobbying for tariffs, then it is worse for the economy as a whole if one firm does. But in the real world, in which this behavior is common, then it is not necessarily worse for the economy as a whole if one firm lobbies for a tariff. It could well be better. So it is false that following the MFA's efficiency imperatives always leads to efficient outcomes.

In fact, as we have seen, Heath *acknowledges* that this is false. While the MFA prohibits seeking tariffs, he notes that "if there is even one trade barrier or tariff in place, then minimizing the number of tariffs will not necessarily produce the best outcome—we may be better off imposing some additional tariffs" (2014, p. 40). This means that complying with the MFA's efficiency imperatives (in this case, refraining from seeking a tariff) may not, in such a market, lead to an efficient outcome. It is curious, then, that Heath should claim that bottom-up reasoning, or indeed any kind of reasoning, could show that compliance with the MFA's efficiency imperatives always leads to efficient outcomes.

Go Rule?

Perhaps this curiosity can be explained, however, as a misinterpretation. It might be argued that we should see Heath as saying not that complying with the MFA's efficiency imperatives *always* leads to more efficient outcomes, but that it *usually* leads to more efficient outcomes. Perhaps he is saying that complying with the MFA's efficiency imperatives leads to more efficient outcomes than complying with any other set of ethical rules. And, Heath might further be saying, market participants should comply with the MFA's efficiency imperatives even if, in a certain case, complying with them does not lead to a more efficient outcome. This is how rule consequentialism is typically understood. On this view, it is wrong (e.g.) to kill innocent people, because a rule which prohibits the killing of all innocent people has better consequences than a rule which allows killing in some cases (Hooker 2000). And this is so even if in an unusual case—the sort of case that philosophers often highlight—killing

an innocent person would have better consequences than not killing him.

This line of reasoning is suggested by Heath's claim that the MFA's efficiency imperatives should be understood as *heuristics* for the real world, i.e., as determining "what type of conduct, *in general*, is likely to constitute an illegitimate source of gain" (2014, p. 41, emphasis added). Heath might be suggesting that the MFA's efficiency imperatives should be understood as rules of thumb—as justified by the good consequences they bring, compared to all other ethical rules for business. And this is so even if, in unusual cases, complying with them does not produce efficient outcomes, i.e., even if acting in a different way would produce a more efficient outcome.

Is this claim true? Heath's "bottom-up" reasoning does not provide evidence that it is true. It simply provides a way to think about the value of individual exchanges. And he provides no other evidence that it is true. So we must conclude that it is unjustified.

This does not mean, of course, that it *could not be* true, only that it has not now been proven to be true. How might we go about determining whether the MFA's efficiency imperatives lead to more efficient results than any other set of ethical rules for business? Not, I suggest, by philosophical or economic theorizing on the "ground" level about individual transactions. We need empirical evidence from the "top" level about policies, regulations, and practices. We need evidence that, when people follow certain rules, certain results occur. As we search for evidence, we should assess the MFA's efficiency rules individually, not as a set. We might find evidence that, when people follow the rule requiring truth in advertising or the rule prohibiting collusion, people are in general better off. But we might not find that rules requiring the elimination of all information asymmetries make people better off. (A case in point: some people—investment analysts, pharmaceutical companies, market researchers—make a living *discovering* valuable information. If they were required to disclose this information to their trading partners before engaging in a trade, then they would have little incentive to discover the information in the first place, and society would be worse off.) Our findings might be relative to all sorts of factors: industry type, level of economic development, type of economic system, and so on. Perhaps some evidence can be provided that compliance with the MFA's efficiency imperatives—one of them, some of them, or all of them (some or all of the time)—leads to efficient results, i.e., results that are more efficient than any other imperatives. Or perhaps we will discover that compliance with a different set of rules produces more efficient results.

Both McMahon and Heath believe that the MFA's efficiency imperatives apply in an ideal world, i.e., to agents under conditions of perfect competition. Our question in this

section has been whether they also apply in the real world. McMahon says "no"; Heath says "yes." I have unpacked Heath's argument for "yes," and argued that it fails. I also identified what would need to be done to show that the MFA's efficiency imperatives apply to real economic agents.

The MFA as Merely Ideal

Our discussion to this point has left us with a theory of business ethics that applies in the ideal world only. It might now be asked: is this a problem? Heath's efforts to show that the MFA's efficiency imperatives apply in the real world suggests that he thinks it is. But perhaps he is wrong, and there is no problem. This section explores this possibility.

The natural worry about an ideal theory of business ethics is that such a theory would be pointless. McMahon, who is satisfied with the MFA as an ideal theory of business ethics, resists this conclusion, saying "despite its descriptive and predictive inadequacies, general equilibrium theory"—the theory from which the MFA's efficiency imperatives derive—"provides a basic framework for comprehending economic activity... Consequently if we wish to understand the place of morality in business, it is appropriate to begin within the framework. The result can be considered an 'ideal' theory of business morality, construction of which is a necessary preliminary to work on a (possibly more relevant) 'nonideal' theory" (1981, p. 256). Here McMahon borrows from Rawls (1971, 2001). Rawls offers his justice as fairness as an ideal theory, i.e., one that makes certain idealizing or simplifying assumptions. "The reason for beginning with ideal theory," Rawls says, is that it provides "the only basis for the systematic grasp of the more pressing problems" of nonideal theory (1971, p. 9). Similarly, McMahon suggests that the construction of an ideal theory of business ethics is a "necessary preliminary" to work on a nonideal theory of business ethics for our imperfect circumstances.

Rawls's claim that ideal theorizing is a *necessary* condition of nonideal theorizing is problematic. In his well-known critique of Rawls, Sen argues that "investigation of different ways of advancing justice in a society..., or of reducing manifest injustices that may exist, demands comparative judgments about justice, for which the identification of fully just social arrangements is neither necessary nor sufficient" (2006, p. 217). In order to assess the heights of two nearby mountains, Sen explains, we don't need to know what the tallest mountain in the world is, and knowing this information does not even help us to determine which of the two nearby mountains is tallest. Similarly, Sen says, in order to assess the relative justice of two arrangements, we don't need to know what ideal justice requires. And a "characterization of spotless justice does not entail any delineation

whatever of how diverse departures from spotlessness can be compared and ranked" (2006, p. 220).

This criticism can be easily adapted to the MFA understood as an ideal theory. We might conclude that an ideal theory of business ethics is neither necessary nor sufficient for figuring out how to make the business world more ethical. We can determine which of two alternatives is ethically superior without knowing how agents in the ideal world would behave, and knowing how ideal agents would behave would not tell us which of these alternatives is ethically superior.

Sen believes that we should do away with ideal theorizing altogether. But others, such as Simmons (2010), believe that this goes too far. True, if we want to climb the taller of two nearby mountains, it is neither necessary nor sufficient to know what the tallest mountain in the world is. But what if we want to keep climbing taller and taller mountains? What if we want ultimately to climb the tallest mountain in the world? Then surely we need to know what the tallest mountain is. Simmons suggests that the search for justice is like this. We should want not only to improve our present condition with respect to justice, but to keep improving it until it is the best that it can be (see also Stemplowska and Swift 2012). To do this we need a target, which is provided by an ideal theory. Something similar might be said in defense of the MFA understood as an ideal theory. That is, it might be said that we want not only to make business practice less unethical than it now is, we want to keep reducing its unethicality. We want to make it the best that it can be. For this we need a target, and the MFA provides one.

This response to Sen strikes me as correct, up to a point. It seems correct to say, with Sen, that ideal theorizing is not necessary for determining how to improve things in the here and now (with respect to justice or business ethics). And yet it seems correct to say, with Simmons, that there is a point to ideal theorizing. It provides a target at which to aim, and in doing so, helps us to determine which improvements are better than others. This conclusion may give some hope to the defenders of the MFA understood as ideal theory. But I do not think it should give them too much hope. This is because, as I will now argue, the MFA is a *highly* ideal theory.

In his theory of justice, Rawls makes two main idealizing assumptions. One is *strict compliance*. That is, Rawls assumes that "(nearly) everyone strictly complies with, and so abides by, the principles of justice" (2001, p. 13). The second is *favorable conditions*. Rawls assumes that people have the "economic means," "education," and various "skills needed to run a democratic regime" (2001, p. 47). Among other things, people are not faced with extreme resource scarcity and are not deeply suspicious of each other. Rawls insists that, in constructing a theory of justice on the basis of these assumptions, he is not constructing a theory that is fit

only for supernatural beings, the "saint or the hero" (1971, p. 479). He says that his theory is "realistically utopian: it probes the limits of the realistically practicable, that is, how far in our world (given its laws and tendencies) a democratic regime can attain complete realization of its appropriate political values" (2001, p. 13).

By contrast, the MFA is not, I suggest, "realistically utopian." It *exceeds* the limits of what is realistically practicable. Like justice as fairness, the MFA also makes assumptions about human behavior and social circumstances. With respect to behavior, the MFA assumes, as does justice as fairness, strict compliance. In particular, it assumes strict compliance with the efficiency imperatives derived from the conditions of perfect competition. This is implied by the assumption of perfect competition. If some economic agents do not comply with these imperatives (e.g., by deceiving trading partners or engaging in collusion), then these conditions do not obtain. The MFA makes an additional behavioral assumption, viz. people seek to maximize their welfare in the market (with producers trying maximize profits). There are of course many examples of actual economic agents not seeking to maximize their welfare in the market, but it seems *possible* for them to do so. Where the MFA departs significantly from what is "realistically practicable" is in the social circumstances it assumes. The MFA assumes that the conditions of perfect competition, including perfect information, frictionless movement of factors of production, no barriers to entry or exit, and no market power. These circumstances are not simply hard to bring about; they are impossible to bring about. There is nothing realistic about them.

In making these claims, I am not saying something that economists—whose theories are also based on a model of perfect competition—don't know. They know that the real world is not a world of perfect competition (Mas-Colell et al. 1995). Their explanations and predictions must allow for the ways that the world is imperfect. Moreover, McMahan and Heath are aware of the ways that the real world diverges from the world of perfect competition. What they—or at least McMahan, who is satisfied with the MFA understood as an ideal theory—fail to see is the way that this challenges the value of the ideal theory that they have constructed. Ideal theory is useful when the idealizing assumptions do not take us too far away from reality. But the idealizing assumptions of the MFA take us very far away. Rawls's justice as fairness is sometimes criticized for being too ideal. What we have shown is that the MFA is a lot more ideal than Rawls's theory.

It might be replied that this is not a significant problem. Yes, it might be admitted, the MFA is more ideal than Rawls's theory, but this just means that *more work* must be done to translate it into something usable for the real world. We must be sensitive to *even more* ways that the real world differs from the world assumed in the theory. In the next

section, we will consider whether anything can be extracted from the MFA for the real world. But I would just say now that we are heading in the wrong direction. There is something odd, in my view, about constructing an *ideal* theory of right behavior in an area of *applied ethics*. It is always risky to say what a particular field of inquiry is “about.” But I will take this risk and say that the field of applied ethics is about describing the moral features of certain domains, with a view toward telling actual people what they should do in those domains. This is what most theories of business ethics try to do. The MFA understood as an ideal theory does not do that. It tells people living in a world that does not exist what to do, and this world is far removed from our own. The MFA may serve some purpose as ideal theory. It may not be necessary to help us improve the current state of business practice, but it may give us a target at which to aim. (I have not said anything here about the value of hitting that target.) But if this is all the MFA does, it is reasonable to question the volume of attention that has been paid to it. Justifying this attention would demand a re-conception of what the field of business ethics is about.

We saw in the last section that Heath tries to show that the MFA’s efficiency imperatives apply in the real world. While it was wise to make this effort, it failed. This section explored McMahon’s idea that the MFA has value as an ideal theory. In response, we said that ideal theories have some value, but the MFA is arguably *too* ideal—much more so than the paradigmatic ideal theory, Rawls’s justice as fairness. And theories of applied ethics should tend toward the real. As a result, we should be unsatisfied with the MFA as a theory of business ethics until it can be shown to apply to the real world.

Conclusion: Two Ways of Modifying the MFA

This paper is an effort to clarify a view that has been the subject of much discussion, but which has not been properly understood. I first sketched McMahon’s implicit morality of the market and Heath’s market failures approach. I argued that these are fundamentally the same view. Next I examined the notion of “efficiency” at work in the MFA. While McMahon and (especially) Heath claim that they understand this concept in terms of Pareto optimality, in practice they often mean something closer to Kaldor–Hicks efficiency or aggregate welfare. Finally, and most importantly, I considered whether the market failures approach has implications for the real world. Against Heath but with McMahon, I argued that it does not. All of these are significant clarifications of the MFA and must be considered in any comprehensive evaluation of this view. I went on to argue that a theory of business ethics that is merely ideal is unsatisfactory.

What’s next? As I said at the outset of this paper, it would be a mistake to consign the MFA to the dustbin of business ethics. To conclude, I sketch two ways of modifying it to fit our world, highlighting the advantages and disadvantages of each modification.

The MFA offers us a purpose of markets—the efficient allocation of resources—and a set of imperatives—the efficiency imperatives—that are designed to achieve this purpose. In the real world, the purpose and the imperatives come apart. Conforming our behavior to the imperatives may not lead to the desired outcome; the desired outcome may be achieved by behaving in different ways. Given this, it seems that, if we want to fit the MFA to our world, we have two options: privilege the purpose of markets or privilege the efficiency imperatives.

Suppose, first, that we privilege the purpose of markets. Then the MFA might be understood to say: the right thing for market participants to do is *whatever* creates efficient outcomes. On this modification, we would not need to discard the MFA’s efficiency imperatives right away. But they would have no independent weight. At best they would serve as placeholders until empirical research reveals the proper rules for promoting efficient outcomes. An attractive feature of this modification is that it respects the justification that McMahon and Heath give for the MFA: that it leads to efficient outcomes. A drawback is that it eliminates what was distinctive about the MFA: the efficiency imperatives it derived from the conditions of perfect competition. In doing so, it turns the MFA into a rather mundane form of consequentialism about business ethics.

The second option is to privilege the MFA’s efficiency imperatives. Then MFA might be understood to say: the right thing for market participants to do is to comply with these imperatives. And this is so even if doing so will not lead to efficient outcomes. Put another way, market participants should act *as if* markets were perfectly competitive. An advantage of this approach is that it preserves the MFA’s distinctive efficiency imperatives. The main drawback, of course, is that the justification of these imperatives—that they lead to efficient outcomes—goes away. The MFA would be a set of rules without a justification for following them.

This is not the place to try to develop either possible modification of the MFA. Both seem worthy of development. Until this happens, however, we must conclude that the MFA does not tell us much about ethics in the actual business world. This, I believe, is a surprising conclusion, and calls into question the value of the market failures approach.

Acknowledgements Versions of this paper were presented at Northeastern University, Queen’s University, TU Dortmund University, and the University of Pennsylvania. I thank those audiences for instructive feedback. Thanks also to Julian Jonker, Santiago Mejia, Christopher McMahon, Alan Strudler, and an anonymous reviewer for this journal for perceptive comments on a draft of this paper.

Compliance with Ethical Standards

Conflict of interest The author declares that he has no conflict of interest.

Research Involving Human and Animal Rights This article does not contain any studies with human participants or animals performed by any of the authors.

References

- Heath, J. (2006). Business ethics without stakeholders. *Business Ethics Quarterly*, 16(4), 533–557.
- Heath, J. (2014). *Morality, competition, and the firm: The market failures approach to business ethics*. New York: Oxford University Press.
- Holley, D. M. (1986). A moral evaluation of sales practices. *Business & Professional Ethics Journal*, 5(1), 3–21.
- Hooker, B. (2000). *Ideal code, real world: A rule-consequentialist theory of morality*. New York: Oxford University Press.
- Jaworski, P. (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.
- Lipsey, R. G., & Lancaster, K. (1956). The general theory of second best. *The Review of Economic Studies*, 24(1), 11–32.
- Martin, D. (2013). The contained-rivalry requirement and a ‘triple feature’ program for business ethics. *Journal of Business Ethics*, 115(1), 167–182.
- Mas-Colell, A., Whinston, M. D., & Green, J. R. (1995). *Microeconomic theory*. New York: Oxford University Press.
- McMahon, C. (1981). Morality and the invisible hand. *Philosophy and Public Affairs*, 10(3), 247–277.
- McMahon, C. (2013). *Public capitalism: The political authority of corporate executives*. Philadelphia: University of Pennsylvania Press.
- Moriarty, J. (2018). Against pay secrecy. *Journal of Applied Philosophy*, 35(4), 689–704.
- Néron, P.-Y. (2016). Rethinking the ethics of corporate political activities in a post-Citizens United era: Political equality, corporate citizenship, and market failures. *Journal of Business Ethics*, 136(4), 715–728.
- 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.
- Rawls, J. (1971). *A theory of justice*. Cambridge: Harvard University Press.
- Rawls, J. (2001). *Justice as fairness: A restatement*. Cambridge: Harvard University Press.
- Sen, A. (2006). What do we want from a theory of justice? *Journal of Philosophy*, 103(5), 215–238.
- Silver, D. (2016). Competition, value creation, and the self-understanding of business. *Business Ethics Journal Review*, 4(10), 59–65.
- Simmons, A. J. (2010). Ideal and nonideal theory. *Philosophy & Public Affairs*, 38(1), 5–36.
- Singer, A. (2018). Justice failure: Efficiency and equality in business ethics. *Journal of Business Ethics*, 149(1), 97–115.
- Smith, J. (2018). Efficiency and ethically responsible management. *Journal of Business Ethics*, 150(3), 603–618.
- Steinberg, E. (2017). The inapplicability of the market-failures approach in a non-ideal world. *Business Ethics Journal Review*, 5(5), 28–34.
- Stemplowska, Z., & Swift, A. (2012). Ideal and nonideal theory. In D. Estlund (Ed.), *Oxford handbook of political philosophy* (pp. 373–389). New York: Oxford University Press.
- von Kriegstein, H. (2016). Professionalism, agency, and market failures. *Business Ethics Quarterly*, 26(4), 445–464.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.