A Market Failures Approach to Business Ethics

Professor Joseph Heath
Canada Research Chair in Ethics and Political Economy
Université de Montréal

"Business ethics" is widely regarded as an oxymoron. The only way to be a good soldier in an unjust war is to disobey orders, or maybe even to desert. Many people believe, along similar lines, that the only way to maintain one's ethical integrity in business is not to go into business. The reasons for this are not hard to find. Students are still routinely taught in their introductory economics classes that in a market economy, when engaged in market transactions, individuals act out of self-interest – whether it be by maximizing profits as producers, or by maximizing satisfaction as consumers. This sets up an almost indissoluble link in people's minds between "profit-maximization" and "self-interest." As a result, anyone who thinks that the goal of business is to maximize profits will also tend to think that business is all about self-interest. And since morality is widely regarded as a type of constraint on the pursuit of individual self-interest, it seems to follow quite naturally that business is fundamentally amoral, if not immoral.

The problem is that the association between profit-maximization and self-interest so often taken for granted is based upon a naïve and inadequate theory of the firm. Profit-maximization and self-interest are not the same thing, and the failure to distinguish adequately between the two can be a source of enormous confusion. Business ethics, as a subject, is essentially concerned with the moral responsibilities of managers. Managers often find themselves placed in circumstances in which the imperative to "maximize shareholder value" conflicts with their self-interest. Thus there are many cases in which profit-maximization should be viewed as a managerial obligation, not as an expression of self-interest.

Because of this somewhat elementary confusion, there has been a marked tendency in the business ethics literature to dismiss out of hand views that take the profit motive seriously. In particular, Milton Friedman's classic article "The Social Responsibility of Business is to Increase its Profits," is more often treated as a piece of apologetic than as a serious piece of moral reasoning.[1] This is unfortunate, since the moral laxity on display in Friedman's work is not so much a symptom of an inadequate normative framework as it is a consequence of specious economic reasoning. Or so I will attempt to show.

The more serious consequence of this confusion is the widespread perception that, in order for business ethics to be genuinely ethical, it must extend managerial responsibility to groups other than shareholders. This is, I believe, often the intuition underlying "stakeholder" theories of managerial responsibility. In this paper, I will argue that such efforts are misguided. Profit-maximization, understood as an obligation, rather than as an expression of self-interest, provides a perfectly legitimate platform for the development of a robust moral code. However, if profit-maximization is an obligation, the question naturally arises where this obligation stems from. It is in seeking to justify the profit motive that we discover that the appropriate form of managerial responsibility is not to maximize profits using any available strategy, but rather to take advantage of certain specific opportunities for profit. In many cases, the set of conditions under which profit-seeking is permissible is reflected in the legal environment in which firms operate. I will argue that business ethics is best understood as a set of additional constraints that preclude legally permissible, but not normatively justifiable, profit-maximization strategies.

[1] Milton Friedman, "The Social Responsibility of Business is to Increase its Profits," *New York Times Magazine*, (Sept. 13, 1970).

1. The profit motive

Andrew Stark's controversial 1993 *Harvard Business Review* article, "What's the Matter with Business Ethics?" argued that conventional business ethics was "largely irrelevant for most managers," because it failed to offer them any "practical" advice.[2] "Moral philosophy," he argued, "tends to value *altruism*, the idea that an individual should do good because it is right or will benefit others, not because the individual will benefit from it."[3] As a result, business ethicists have had too little to say about "the potential conflict between ethics and interests," and in particular, how managers should handle such conflicts when they arise.

This article had many people nodding their heads in agreement. But to see just how peculiar the claim is, suppose that the subject had been medical ethics instead of business ethics. Substitute "doctors" for "managers" throughout. Now imagine criticizing medical ethics on the grounds that it fails to offer doctors any "practical" advice on what to do in cases where the imperatives of patient care conflict with their self-interest. Suppose the patient doesn't really need an operation, but the doctor could make a lot of money by performing it anyway. What to do, what to do?

I would suggest, *pace* Stark, that we don't need professional ethicists to tell us where our obligations lie in such cases. Everyone knows that when there is a straightforward conflict between our self-interest and our moral obligations, the moral obligations win, at least from the moral point of view. This is not "ethical absolutism," as Stark maintains, it is simply the logic of moral justification. The question of when we may be forgiven for disregarding our moral

---

[2] Andrew Stark, "What's the Matter With Business Ethics," *Harvard Business Review* (May/June, 1993): 38-48.
[3] Ibid, p. 40.

obligations, i.e. acting immorally, is a separate one, and is in no way specific to the domain of business ethics.

So why does Stark's argument sound even remotely plausible, whereas a comparable argument in medical ethics would be dismissed out of hand? The confusion has two distinct sources. The first arises from the way that introductory economics is usually taught. The standard microeconomics textbook starts out with the assumption that individuals maximize utility. When it comes to particular goods, these utility functions can be represented as a set of indifference curves. These indifference curves are then taken to provide the supply and demand curves. The thesis that individuals maximize utility is interpreted to mean that consumers will seek to maximize satisfaction, and suppliers will seek to maximize profits. Finally, in order to make the model more "realistic" consumers get aggregated together into "households," and suppliers into "firms" – each of which is thought to maximize some joint utility function.

While everyone understands that "the firm" is something of a black box in this analysis, the result is still an unhelpful blurring of the boundaries between the pursuit of self-interest and the maximization of profits. Stark, for instance, variously describes the conflict that managers face as one between "self-interest and altruism," "ethics and interests," "ethical demands and economic realities," "moral and financial costs," "profit motives and ethical imperatives," and even "consumer's interests" versus the "obligation to provide shareholders with the healthiest dividend possible."[4] Here we see a clear blurring of the distinction between self-interest, profit-maximization, and the obligation to shareholders.

We understand implicitly that the professional conduct of doctors is to be entirely governed by their obligations to their patients, and thus that they are not permitted to let considerations of self-interest intrude. Profit-maximization has precisely the same status for managers. To my knowledge, no one has ever tried to defend the managers of RJR-Nabisco, or Enron, on the grounds that they were simply acting in their own self-interest. Of course, if the incentive systems have been properly designed, managers will find it to be in their interest to maximize shareholder value (in the same way that doctors generally find it to be in their interest to cure their patients). But this is accidental and irrelevant from the moral point of view. In the case of a conflict, the obligations simply trump the relevant set of interests. Where things get interesting is when multiple obligations conflict, as in the case of a doctor who can improve a patient's chances of survival by lying to him about his condition, or of a manager who finds herself able to please investors by initiating an unnecessarily severe downsizing.

The second major source of confusion stems from the moral status of the objective sought by managers – profit maximization. The doctor's obligations to the patient flow quite naturally from the objective, which is to restore the patient to health. Health is widely regarded as a good thing, and thus the doctor's actions serve to promote a state of affairs that is morally desirable. This makes the doctor's actions directly justifiable, even intrinsically altruistic. Things are more complicated in the case of business. It is not clear that profits are intrinsically good. Furthermore, when a manager makes a decision that disadvantages workers in order to benefit owners, the profit maximization imperative generates a distributive transfer that is by no means morally sanctioned. In fact, under the typical set of circumstances, the transfer will be regressive, and thus problematic from the moral point of view.

The asymmetry arises from the fact that profit maximization is only indirectly justified. It is useful to note that this problem is one that business ethics shares with legal ethics. The adversarial trial system imposes upon lawyers an obligation to do whatever is in their power to defend or advance the interests of their client, even when these interests are highly refractory to

---

[4] Ibid, p. 44.

the concerns of justice. Thus the professional obligations of lawyers often conflict with the imperatives of everyday morality. What justifies their behaviour is the fact that they operate in the context of an institution with differentiated roles. The desirable outcome is a product of the *interaction* between individuals acting in these roles, none of whom are actually seeking that outcome. Justice is best served when there is both vigorous prosecution and vigorous defence.

Thus the effective trial lawyer "promotes an end which is no part of his intention." The adversarial system may, for example, maximize acquittal of the innocent, even though neither the prosecution nor the defence adopts that as their objective. As a result, neither lawyer's conduct can be justified by the intended outcome. It is justifiable only in through the consequences that the pursuit of this outcome leads to, when combined with the actions of the others.

The same can be applied to the case of managers. The manager should seek to maximize profits for the same reason that the defence lawyer should seek to have his client acquitted – not because the acquittal of his client would be a good thing, and even because his client *wants* to be acquitted, and is paying the bill, but rather because the adversarial trial system as a whole is taken to be the best form of institutional arrangement to serve its appointed function. This is why one cannot do legal ethics without a broader appreciation of how the legal system as a whole functions, and what valuable tasks the various roles are thought to discharge. Similarly, one cannot do business ethics without some appreciation of what justifies the system of private enterprise.

Thus the straightforwardly moralizing critique of the profit motive is jejune (comparable to attacking lawyers for "defending rapists and murderers"). We need to understand why criminals should be entitled to the best possible defence, in order to understand the responsibilities of lawyers. Similarly, we need to understand why corporations should be entitled to pursue profits, in order to understand the responsibilities of managers.

2. What justifies profit?

The right of corporations to earn profits is sometimes regarded as self-evident. This conviction usually stems from a set of broadly Lockean convictions, which suggest that individuals come naturally equipped with a set of property rights prior to the institution of government. Profit-maximization is then understood as the attempt to augment these holding through labour input or voluntary exchange – neither of which the state has any obvious authority to restrict.

The problem with this Lockean view – apart from the fact that the underlying conception of rights is deeply problematic – is that corporations are not individuals, they are highly artificial legal constricts. Furthermore, the corporate organizational form provides individuals with a number of very tangible advantages that they do not enjoy as private citizens. The most significant among these is limited liability – the ability to insulate their own private resources from those of the corporation, so that they cannot be pursued by creditors in the event of default. Because of this, creating a corporation is widely regarded as a privilege, not a right. This makes it legitimate for the state to impose certain obligations, in return for the privileges granted.

Many of the corporations chartered by the state are non-profit. They are specifically prohibited from showing more than a modest revenue surplus. So why permit an exception for other firms? To put it in Marxian terms, why should society tolerate the private appropriation of the social product?

The answer to this question is somewhat complex. Basically, it is that society wants to encourage competition between suppliers. This competition, when combined with competition between purchasers, will affect the prices at which goods trade. Under the correct circumstances, competition will push prices toward the level at which markets clear – i.e. suppliers will not be left with unsold merchandise, and consumers will not be left with any unsatisfied demands. When this occurs, it means that society has succeeded in minimizing the overall amount of waste in the economy. It means that fewer resources will have been spent producing goods that no one wants, at the expense of goods that people do want.

Thus the primary reason for introducing the profit motive into the economy is to secure the operation of the price mechanism. The price mechanism is in turn valued for its efficiency effects. It allows us to minimize waste. The formal proof of this is often referred to as "the first fundamental theory of welfare economics" (hereafter FFT), or else, in a nod to Adam Smith, the "invisible hand theorem." The central conclusion is that the outcome of a perfectly competitive market economy with be Pareto optimal – which means that it will not be possible to improve any one person's condition without worsening someone else's.

The importance of the price mechanism is often underestimated. Since the profit orientation of firms definitely has some adverse social consequences, this can sometimes make it difficult to see what the big gains are that justify our tolerance for the various abuses. In order to put things into perspective, it is helpful to consider the difficulties that we would face trying to make decisions in the absence of a set of prices. This is the situation that planners often confronted in the former Soviet Union. Imagine that one of your plants increases its production, so that you now have the capacity to produce an extra 500 tons of plastic. What to do with this material? You need to figure out where it is most needed. But how do you decide? Suppose, to simplify enormously, that there are two possible uses: to make toothbrushes or soup ladles. The question is, which do people need more of?

In a market economy, these needs will be expressed in the form of relative willingness to pay. If stores have too many ladles, and not enough toothbrushes, they will be willing to order more toothbrushes, and pay more for them. This in turn means that the toothbrush makers will be willing to pay more for the plastic. Thus if all firms sell to the highest bidder, the resources will be channelled toward the use for which there is the greatest need. But if there is not a competitive market for all these goods, not only will firms not have the incentive to engage in the necessary transactions, but the absence of prices will make it difficult for anyone even to determine which transaction should be occurring. Planners in the former Soviet Union used to get around this problem by sometimes looking at commodity prices in Western Europe and North America, and using these figures to do calculations for their own economy. In fact, they used to joke that in the event of a global communist revolution, it might be worthwhile to keep Hong Kong capitalist, so that everyone else would know what prices their goods should be trading at.

The joke has a very serious underlying point. Without prices, you simply cannot organize a complex economy, whether it be capitalist, socialist, or communist. And not just any prices will do. There are an enormous number of price points at which exchanges can occur. In cases where there is only one supplier or one consumer, this gives one side considerable power to dictate terms. Under such conditions, there is no reason to expect that the price level chosen will be the price that clears the market. Thus the price system will not induce efficiency. But when there is more than one supplier, or more than one customer, each one is in a position to undermine the negotiating power of the other. If one supplier insists on a price that is too high, the customer can go to the competition. The competitor is then able make a profit by undercutting the other one's price, making up for it through a larger volume of sales. The result is a race to the bottom among the suppliers, in which they competitively underbid one another until the market clears, and all profit disappears.

Thus the central rationale for having private profit-seeking firms is to establish competition among suppliers and consumers. This competition drives prices towards market-clearing levels, allowing society in turn to generate a more efficient allocation of its resources and labour time.

It should be noted that this concern with competitive markets, and market-clearing prices, is not simply an abstract philosophical theory about what might justify profit-maximization. The entire legal structure of the firm, along with the regulatory environment, has been organized in such a way as to promote not just competition, but the precise type of competition that is likely to generate market-clearing prices. This is true of everything from anti-trust to consumer protection law. In the past decade in Russia, corporations have been known to maximize profit by blowing up each other's factories and assassinating each other's chief executives. Much of the massive legal apparatus that governs corporate behaviour in more mature capitalist economies is designed to ensure that firms seek to maximize profits through a much more limited set of strategies – viz. those strategies that are likely to generate more efficient production, along with a more efficient allocation of goods and services in the economy.

Thus, if we ask what the obligations of managers are, the answer can be provided quite directly. The function of the market economy is to produce the most efficient use of our productive resources possible. This can be done, roughly speaking, by achieving the price level at which all markets clear. The role of the firm in that economy is to compete with other suppliers and purchasers for profits in order to drive prices to that level. Thus managers are obliged to do what is necessary in order for the firm to maximize profits in this way. Profits show that the balance of "needs satisfied" to "resources consumed" is positive, while losses show that the resources would have been put to better use elsewhere. Hence the old saying that if we penalize a man for making a profit, we should penalize him doubly for showing a loss.

3. Milton Friedman

The approach to business ethics that takes profit maximization as a central concern is often viewed with suspicion, since it has traditionally been used more as an apologetic for irresponsible behaviour than as a platform for a good-faith effort to develop a code of ethics. As we have seen, in order to be plausible, the profit-maximization approach to business ethics cannot identify profit-maximization with individual utility-maximization on the part of managers. The naïve version of the "invisible hand" view, according to which markets miraculously transform private vices into public virtues, has clearly become obsolete in the era of professional management.

Thus when Milton Friedman argued that the social responsibility of business is to increase its profits, his primary emphasis is on the fiduciary relationship between managers and shareholders.[5] The manager is in a similar position with respect to the shareholder that the lawyer is in with respect to a client – he is expected to advance the interests of the principal, not his own. This requires trust, and hence moral obligation, between the two parties. And of course, there are many ways in which the lawyer can exploit this relationship for private gain, as can the manager.

---

[5] Friedman, "The Social Responsibility of Business is to Increase its Profits," and Milton Friedman, *Capitalism and Freedom* (Chicago: University of Chicago Press, 1962).

This makes Friedman's view a genuine code of ethics, and not simply an apologia for self-interest. However, while Friedman is clear that managers are subject to genuine moral constraint, he is less than clear about the source of these obligations or constraints. At one point, he suggests that the manager is bound to assist the shareholder in the satisfaction of his or her desires, and that profits just happen to be what most shareholders want. This is clearly absurd – the manager is not the personal servant of the shareholder. The shareholder might like to have the manager do his laundry, and if he can supply appropriate incentives, he may even succeed in getting the manager to do it. But there is no sense in which the manager is morally obliged to do so, by the mere fact that the shareholder desires it. The manager's responsibility toward the shareholder is clearly restricted to the latter's investment returns. Or, as Friedman puts it when he is being careful, the responsibility of managers is "to make as much money for their stockholders as possible."[6]

However, even this more restricted concept of managerial responsibility is not enough to explain the source of the obligation. Simply making a promise is not enough to generate an obligation, in cases where the end in view is itself not morally justifiable. Promising to help a friend rob a bank does not generate an obligation to rob the bank. Thus the manager's obligation to help the shareholder maximize profits must be derivative of the latter's entitlement to do so. And since it is the FFT that justifies this entitlement, Friedman's argument derives managerial responsibilities from the efficiency argument for capitalism on the whole.[7]

This implicit dependence upon the FFT is discernible in a seemingly innocuous caveat that Friedman tacks onto the formulation of his central thesis. Here is what he says:

---

[6] Friedman, *Capitalism and Freedom*, p. 133.

[7] Friedman also has a parallel argument concerning the role of markets in promoting freedom. But this line of thinking is, in my view, so riddled with fallacies that it does not merit serious consideration. Furthermore, it seems

> The view has been gaining widespread acceptance that corporate officials and labor leaders have a "social responsibility" that goes beyond serving the interests of their stockholders or their members. This view shows a fundamental misconception of the character and nature of a free economy. In such an economy, 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, which is to say, engages in open and free competition, without deception or fraud.[8]

Thus he argues that managers must maximize profits, not *tout court*, but rather subject to the "rules of the game," and in particular, subject to the constraint that they do so "without deception or fraud." The fraud constraint is unexceptional and redundant, since it is illegal. (It goes without saying, for instance, that one should not profit through theft or murder.) But why not deception? One is allowed to win a chess game through deception. In fact, deception is a common feature of strategic interactions. What's wrong with making money through deception?

The answer cannot be that the general moral imperative against lying is binding upon managers in all contexts. Everyday morality compels us to treat others as we ourselves would like to be treated, and yet the last thing we want a manager thinking about, before declaring a giant year-end clearance sale, is how she would feel if the competition did the same to her. More generally, price competition is an interfirm prisoner's dilemma – the outcome is suboptimal for all the competitors. Many moral norms have as their primary function the elimination of such collectively self-defeating interaction patterns. Yet in the case of businesses, we want them to

---

fairly obvious that Friedman's preference for market solutions to almost every social problem came from his conviction that governments were inefficient and markets were efficient.

remain stuck in the prisoner's dilemma. In fact, any agreements designed to eliminate these outcomes are specifically prohibited by law. So we cannot simply appeal to the fact that an action is prohibited by everyday morality as grounds for imposing this same prohibition upon managers, unless we want to adopt the very strict universalist view that morality does not permit any institutional differentiation.

Thus the problem with deception, in Friedman's view, cannot arise from any strict deontic prohibition. The problem with deception is that it violates one of the conditions needed for the economy to achieve an efficient outcome. It is these conditions that Friedman is adverting to as well when he talks about an obligation to engage in "free and open" competition.

The relationship between honesty and efficiency in market transactions requires very little demonstration. If suppliers lie to consumers about the character of the goods that they are acquiring, then the prices at which their exchanges are concluded are not going to reflect the actual need for the good in question. This will generate inefficiencies in the economy.

To take a very concrete case, consider the so-called "goulash capitalism" episode in Hungary. Shortly after the transition from communism to capitalism, Hungary was struck by a wave of lead poisoning. The source of the epidemic was eventually tracked down to paprika. After privatization, several paprika suppliers began adding ground-up paint – much of it lead-based – to the spice, in order to improve its colour. In other words, a competition developed to produce the best-looking paprika, not the best quality paprika. Needless to say, if consumers had been properly informed as to the quality of the goods they were purchasing, they would not have bought any. Thus the deception perpetrated by these firms resulted in a huge loss of welfare to consumers. Health authorities eventually had to step in and destroy the entire paprika supply in the country, in order to eliminate all the contaminated goods.

---

[8] Friedman, *Capitalism and Freedom*, p. 133.

This is a case of what economists call “market failure.” In order for the FFT to obtain, a set of very restrictive conditions must be satisfied. These are referred to as the Pareto-conditions. The state in which all the Pareto conditions are satisfied is often called, somewhat misleadingly, “perfect competition.” When one or more of the Pareto conditions are not satisfied, the competitive equilibrium of a market economy will be less than Pareto-optimal. When a Pareto-inferior outcome is realized, this is referred to as a market failure.

One of the Pareto conditions specifies that information must be symmetric. Each party to the transaction must have the same information (not only about the prices and goods that are directly relevant to the exchange, but about all other prices and goods in the economy as well). Thus what Friedman is suggesting, in effect, is that managers have no right to take advantage of market imperfections in order to increase corporate profits. The set of permissible profit-maximizing strategies is limited to those strategies that would be permissible under conditions of perfect competition.

In this view, there is a natural complementarity between law and morality. As mentioned, the primary function of the legal regulation of the market is to prevent market failures – both by ensuring that firms do not collude to escape the prisoner’s dilemma that competition imposes upon them, or by preventing them from displacing costs in a way that is not fully reflected in the price at which goods trade. In a perfect world, it would be possible to create perfect markets. However, in the actual world, the legal mechanism is a somewhat blunt instrument. In many cases, the state simply lacks the information needed to implement the necessary measures (sometimes because the information simply does not exist, but often because the state has no way of extracting it truthfully from the relevant parties). Even when the information can be obtained, there are significant administrative costs associated with record-keeping and compliance monitoring. Thus the deadweight losses imposed through the legal mechanism can easily outweigh whatever efficiency gains might have been achieved through the intervention. This makes legal regulation unfeasible.

Moral constraints, on the other hand, are subject to no such costs. Corporations, for instance, are often in a position where they can produce misleading advertising that stops short of outright falsity. In a perfect world, advertising would provide nothing more than truthful information about the qualities and prices of goods. However, the vagaries of interpretation make it impossible to prohibit anything but the most flagrant forms of misinformation. Thus misleading advertising stands to false advertising as deception does to fraud. It is something that would be illegal, were it not for practical (or perhaps even accidental) limitations on the scope of legal regulation. Profiting from such actions is therefore morally prohibited, because it runs contrary to the objectives that the market system was instituted to promote.

Friedman's view is often rejected on the grounds that it is morally lax. It basically lets business of the hook on the question of social responsibility. The above analysis shows, however, that Friedman's argument is not a Trojan horse for naked self-interest. Despite some confusion, it is clear that Friedman's managers have genuine ethical responsibility to shareholders, and that this responsibility is derived from the FFT. The problem is that Friedman arbitrarily limits the set of obligations to those that support only some of the many Pareto conditions.

For example, Friedman argues that pollution reduction is one of the illegitimate responsibilities pressed upon managers in the name of "social responsibility." But pollution is a negative externality – a cost associated with some economic activity that is transferred to a third party without compensation. These externalities exist because the set of markets is incomplete.

We cannot exercise property rights over the air that we breathe, for example. As a result, while we can charge people for dumping noxious substances on land that we own, we cannot do the same when they dump it in the air. For this reason, one of the Pareto conditions specifies that there must be no externalities. Any corporation that pollutes is essentially profiting from a market imperfection. This means that there is no difference, from the moral point of view, between deception and pollution – both represent impermissible profit-maximization strategies. Friedman's decision to prohibit deception, while giving the wink to environmental degradation, is arbitrary and unmotivated.

Figure 1 shows the basic structure of Friedman's normative framework. The overall set of profit maximizing strategies is partitioned into three categories, separating out the immoral and the illegal strategies from the normatively acceptable ones. The efficiency standard can be used to make both cuts. The "acceptable/unacceptable" distinction is imposed by the efficiency properties of the market system as a whole. The set of unacceptable strategies can then be subdivided into "immoral/illegal" using a transaction cost or regulatory cost analysis.

Profit maximization strategies

Acceptable	**Immoral**	**Illegal**
e.g. lowering price, improving quality	e.g. pollution, deceptive advertising	e.g. fraud, theft, embezzlement, false advertising

Figure 1. Friedman's normative framework

4. A Market Failures Based Code

The above reflections suggest that there is no reason to think that a business ethics focused on profit-maximization cannot deal with the obligations that have traditionally been described under the heading of "social responsibility."[9] What so often annoys people about corporations – and what gives profit-seeking a bad name – is the exploitation of one or another form of market imperfection. People generally have no problem with companies that make money by providing good service, quality goods, low prices, and so forth. In my opinion, if all companies fully internalized all costs, and charged consumers the full price that production of their goods imposed upon society, it would be impossible to make the case for any further "social responsibility" with respect to, for example, the environment.

In fact, one of the major advantages of the market failures approach to business ethics is that it is the only one that is able to pick out the "right" level of pollution. There can be no ethical imperative to eliminate pollution completely, since without some pollution there would be no economy. Society as a whole must be willing to accept some degradation of the environment in exchange for the goods produced. What is important is that the level of pollution be determined by people's actual preferences, not simply the subset of those preferences that happens to be legally enforceable. In other words, the cost of production should be the same as the social cost.

---

[9] In referring to "social responsibility" I am implicitly drawing a somewhat rough-and-ready distinction between internal and external obligations. A corporation is usually an organizational hierarchy. Thus the manager is involved with two very different types of relationships – those with other members of the organizational hierarchy, and those with individuals outside of it. Employees of a firm are "inside," for example, whereas customers are generally "outside" – precisely because the latter are not under the control of the organization in the way that the former are. Thus environmental issues, consumer relations, business partnerships, relations with rival corporations and government will all be classified as "external," and hence as falling under the rubric of "social responsibility." For the purposes of this paper, I will also classify shareholders as "external." This is an admittedly crude division of the conceptual terrain, but I think that it does separate out a set of quite distinct issues, which need to be addressed on their own terms. For example, I think that a number of fundamental norms of reciprocity that apply to internal relations do not apply to external ones.

This is precisely the state that would obtain if businesses derived no profit from displacement of costs that markets do not internalize.

What other sort of constraints does this approach impose? Imagine for a moment a deontically perfect world, in which everyone could be counted on to comply with all moral requirements. How should an ethical corporation behave in such a world? The answer is quite simple. The firm should behave as though market conditions were perfectly competitive, even though they may not in fact be. The following list of imperatives provides some examples of the restrictions that this would imply:

1. Minimize negative externalities.
2. Compete only through price and quality.
3. Reduce information asymmetries between firm and customers.
4. Do not exploit diffusion of ownership.
5. Avoid erecting barriers to entry.
6. Do not use cross-subsidization to eliminate competitors.
7. Do not oppose regulation aimed at correcting market imperfections.
8. Do not seek tariffs or other protectionist measures.
9. Treat price levels as exogenously determined.
10. Do not engage in opportunistic behaviour towards customers or other firms.

I think it is clear from this list that, rather than being morally lax, the market failures approach is actually quite restrictive. In fact, in the real world, any firm that began to unilaterally respect these constraints would be quickly eliminated from the marketplace. For instance, the requirement that firms compete only through price and quality excludes the use of non-informative advertising as a way of building market share. Advertising, as a form of non-productive competition, imposes significant deadweight losses on the economy. For example, Molson and Labatt spend $200 million per year on advertising. Studies have shown, however, that this competition is zero-sum. The amount of beer consumed has actually fallen over the years – thus the two companies are, at best, simply stealing customers back and forth from one another. This drives up the price of beer, a situation that is only sustainable because of market imperfections – viz. the significant economies of scale in the brewing industry, which constitute an effective barrier to entry.

Assuming that the nuisance value of beer ads exceeds their entertainment value, this means that society as a whole would be better off if the breweries stopped advertising. But it would be suicide for either company to do so unilaterally. The situation is identical to that of a country hoping to escape from an arms race through unilateral disarmament. Such a situation provides an ideal occasion for the old "they are doing it, so we have to do it too" defense of non-cooperation. (This is an argument used in favour of illegality as well, e.g. when foreign competitors are able to engage in business practices that would be considered corrupt in the home country.)

Of course, the fact that other people are not going to respect their moral obligations does not undo the obligation for everyone else. It may provide an excusing condition – a reason why one need not respect one's moral obligation in this case. At the same time, one is still obliged to do what is necessary in order to bring about the conditions under which the obligations could be fulfilled. And it cannot be argued that these demands are too onerous in principle, since the demands simply articulate the way that capitalist economies are supposed to function in the first place. Thus it is only the possibility of unethical behaviour by others that could justify non-compliance.

There are a variety of different ways in which businesses might try to bring about the conditions under which they could satisfy these ethical demands. The first is that they might engage in “experiments in trust,” – build up cooperation through reciprocity over time. We are already familiar with this process from the dynamic of arms negotiations. Thus, for example, firms might all agree to scale back their advertising expenditures by a fixed percentage every year, until they are eliminated completely. Compliance in the first round of cuts would help to build confidence going into the second.

Firms might also enter into agreements to restrict unethical conduct outside the framework of formal law. Anti-trust concerns create an environment in which legislators are very suspicious of such agreements – especially those that would limit competition. However, it is worth distinguishing between productive and non-productive forms of competition. Firms governed by self-interest, given the opportunity to collude, will eliminate the former, whereas firms governed by moral principles will eliminate the latter. One can imagine the development of an environment, through trust-building exercises, in which corporations could demonstrate their commitment to ethical conduct, and thus earn the trust of legislators. In such an environment, corporations could enter into binding agreements with one another to enforce ethical conduct.

Finally, there is the point sometimes made in the literature that firms which actively profit from market imperfections are, in effect, tempting legislators and regulators to intervene. And when the state does intervene, the costs associated with compliance usually leave all of the firms involved worse off than they had been prior to their exploitation of the imperfection. Thus companies sometimes pressure one another to respect moral principles using the “stop it or you’ll get us all caught" appeal. This sometimes provides an incentive structure that is able to secure the desired pattern of behaviour even in the absence of regulation (although fans of "industry self-regulation" have a tendency to overestimate the number of circumstances in which such incentives are present).

5. Further directions

The market failures approach to business ethics elaborated here shows that a very robust moral code can be developed out of the idea that the fundamental obligation of managers is to maximize shareholder value. It has always been accepted that managers must do so within the framework of the law. The suggestion here is simply that an ethical manager is one who does so while respecting not only the letter of the law, but also its spirit – which is to create the conditions necessary for private enterprise to generate an efficient allocation of goods and services in the economy.

However, there is a significant complication with this view, one that merits further discussion. The problem arises from what is known as the "general theory of the second best," or the "second-best theorem" for short.[10] This theorem 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. In other words, while perfect competition generates a perfectly efficient outcome, a situation that is as close as possible to perfect competition will not generate an outcome that is as close as possible to perfect efficiency.

---

[10] Richard Lipsey and Kelvin Lancaster, "The General Theory of the Second Best," *Review of Economic Studies*, 24 (1956): 11-32.

The second-best theorem blocks a line of analogical reasoning that has long appealed to economists. Everyone understands that Newtonian physics, for instance, employs a number of idealizations. We also understand that the more closely the real world resembles these idealizations, the more closely the objects at our disposal will respect these laws. So while we do not have access to a frictionless plane, we can often substitute a very smooth tabletop in order to illustrate a variety of principles. Furthermore, the smoother the tabletop, the more closely the objects on it will conform to the predictions of ideal theory.

People sometimes like to extend this sort of analogy to economics. Perfect competition, according to such a view, is like a frictionless plane. It is an idealization. But the more closely the real world resembles this idealization, the more closely the various predictions will obtain. Friedman is, like many others, tempted by this form of reasoning. He writes, for instance, that:

> Of course, competition is an ideal type, like a Euclidian line or point. No one has ever seen a Euclidian line – which has zero width and depth – yet we all find it useful to regard many a Euclidian volume – such as a surveyor's string – as a Euclidian line. Similarly, there is no such a thing as "pure" competition. Every producer has some effect, however tiny, on the price of the product he purchases. The important issue for understanding and for policy is whether this effect is significant or can properly be neglected, as the surveyor can neglect the thickness of what he calls a "line."[11]

On the basis of this analogy, we may be tempted to conclude that if perfect competition generates perfect efficiency, then near-perfect competition should generate something as close as possible to perfect efficiency. The second-best theorem shows that this line of reasoning is unsound. If one of the Pareto conditions is violated, then the closest approximation to perfect competition will produce an outcome that is less efficient – and thus worse for society – than some more distant alternative. This has massive consequences. It means, for example, 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. Similarly, if one sector of the economy is subject to monopolistic pricing, then having prices in all the other sectors determined by competition will produce an outcome that is inferior to some other outcome that would result if these prices were not competitively determined.

Of course, the kind of information that would be required in order to figure out how to achieve the second-best outcome is almost always unobtainable. The second best theorem is primarily a limitative result. It shows us that we cannot use the FFT to derive normative conclusions under real-world circumstances. Thus the second best theorem basically blocks the line of reasoning the Friedman develops. It also presents a very fundamental challenge to the market-failures based approach to business ethics being mooted here. It suggests that ethical behaviour, in the absence of complete reciprocity, may be bad not only for the firm that sticks its neck out, but for the rest of society as well.

Of course, this does not mean that the efficiency standard is deprived of all normative force. It simply means that we cannot make the big sweeping generalizations that were the stock-in-trade of economists of Friedman's generation. In particular, it means that the properties of general equilibrium models are not going to be relevant to the normative evaluation of actual economies. Moral reasoning in a business context must be a more contextual affair. We cannot simply adopt the best competitive strategy, then hope that the invisible hand will take care of the

[11] Friedman, *Capitalism and Freedom*, p. 120.

rest. Even if we are in perfect conformity with both the spirit and the letter of the law, profit-maximization may still generate an inferior outcome.

There are several responses that suggest themselves at this point. The first is that the FFT specifies the conditions under which a Pareto optimum is attainable. But in day-to-day life, this optimum is irrelevant. 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.

We can think of this approach as a "resource custodianship" perspective. The ultimate goal of the economy as a whole is to satisfy human needs. The demand for various goods is an expression, however imperfect, of the intensity of these needs. The function of the price system is to channel resources toward the satisfaction of the most important of these needs (not according to an objective measure, of course, but rather according to each individual's own assessment of his or her needs). Thus the firm purchases a bundle of productive inputs in order to satisfy these needs, and profit – when earned under the correct conditions – is the reward that is enjoyed for having done a *better* job at satisfying these needs than any of its rivals.

Thus we can think of all productive resources as being "earmarked" for the satisfaction of needs. The managers and shareholders are the custodians of these resources. Their job is to convert these resources into consumer welfare – and when they do, they are rewarded with a profit. As a result, 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.

Using this sort of “bottom-up” reasoning, I believe that all of the constraints outlined in section 4 could be justified in some form. In this framework, the Pareto conditions would function as a set of heuristics, allowing us to determine what sort 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. These remarks are clearly unsatisfactory. The more general research program, however, is one that I believe has considerable promise.