Optimality of Market-Based Regulations Depends Sensitively on the Profit Maximization Hypothesis

or

Command-and-Control Regulations Can Outperform Market-Based Mechanisms when Intra-Firm Behavior is Strategic

Robert L. Axtell[1]
The Brookings Institution
1775 Massachusetts Avenue, NW
Washington, DC 20036
(202) 797-6020
raxtell@brookings.edu

Abstract

Conventional analysis of regulatory mechanisms typically demonstrates that market-based approaches out-perform command-and-control policies. For example, in environmental policy, emissions taxes and tradable permits dominate command-and-control. Underlying these results is a model of firm behavior—profit maximization—that lacks significant empirical support. Here I show that the dominance of market-based regulatory instruments is sensitive to specifications of firm behavior. Specifically, richer, more realistic models of firm behavior that view an organization as composed of multiple agents engaged in team production (and therefore team pollution), lead to deviations from profit-maximizing behavior, and consequently to imperfect performance of market-based regulatory instruments. Certain kinds of command-and-control policies applied to such firms can generate higher social welfare than market-based policies. Given that most estimates of welfare gains associated with market-based instruments are indirect—they result from calculations based on profit-maximizing assumptions instead of direct cost savings data—they can be biased due to misspecification of firm behavior. The real-world is much too complex to permit simple generalizations such as 'market-based regulations always outperform command-and-control regulations.' Whether one regulatory approach outperforms another is inherently an empirical question, inextricably tied up with details of the production process and the social context. Only by moving inside the black box of the firm, focusing on the incentives facing agents and the interactions between agents, can the appropriateness of a regulatory regime be assessed.

---

[1] Research supported in part by the U.S. Environmental Protection Agency.

I Introduction: Choice of Regulatory Instrument

Market-based environmentalism has risen from theoretical obscurity (Montgomery 1972), to intriguing policy idea (Hahn and Noll 1982), to legislative dictate (e.g., Clean Air Act Amendments of 1990) to widespread implementation (EPA 2000). Today, market-based approaches to regulation are the darlings of the regulatory world, with applications across nearly all polluting media and political jurisdictions (Stavins 2003). Market-based policy instruments are claimed to dominate older command-and-control policies, and calculations of cost-savings ostensibly reveal that that significant savings have been realized (Newell and Stavins 2003).

However, it turns out that the burgeoning edifice of market-based environmentalism rests on certain specifications—as does much of neoclassical economic theory—that may or may not be realized in practice. Some of these assumptions concern the market environment, others are related to the details of the production process, while others are stipulations about behavior. For instance, imperfect competition, either for products (Maleug 1990) or for pollution permits (Hahn 1984) can diminish the performance of market-based instruments. Similarly, the advantages of market environmentalism can evaporate when monitoring and enforcement are uncertain (Hahn and Axtell 1995; Montero forthcoming) Some assumptions—particularly the behavioral ones—could, in principle, be tested in the field, but apparently are never subjected to empirical verification. Some of these assumptions are seemingly innocuous, and deviations from them in the real world are probably of minor import. For example, technical requirements on the shape of abatement cost functions—basically, that the cost of reducing a single unit of pollution rises smoothly with the amount of abatement—made to insure analytical tractability, are probably rarely satisfied in practice, where indivisibilities and other non-convexities surely obtain, at least locally. But departures from these technical requirements probably do little violence to claims concerning the superiority of market-based approaches to regulation.

One particular assumption, though, is a kind of keystone to the entire apparatus, the hypothesis that firm behavior is well-described by profit maximization.[2] This assumption is, on its face, crucial since the calculus by which the comparative performance of regulatory mechanisms is assessed begins with an expression for firm profit that is

[2] There are a few of papers that attempt to relax profit maximizing behavior—e.g., Tschirhiart (1984) in the context of firms maximizing sales or staff instead of profit—but in none of these is the unitary actor framing of firm behavior relaxed and therefore these models remain within the conventional mathematical programming paradigm of operations research applied to the firm.

then maximized (or equivalently, cost is minimized). If firm behavior is substantially different from this specification then all bets are off vis-à-vis the ultimate superiority of market-based instruments. For in such a case the alternative model of firm behavior may have little in common with the conventional model, and so there would be little *a priori* reason to believe that the salutatory properties of the market-based approach under profit maximization would stand up.

In this paper we investigate the performance of regulatory instruments using just such an alternative model of firm behavior. It has some features in common with the conventional profit maximizing view, but is more closely related to modern game theoretic views of the firm, in which the conflicting incentives facing various intra-firm agents are explicitly considered. Specifically, we treat the firm as a multi-agent system in which each agent contributes to production non-cooperatively. This is so-called 'team production' (Holmstrom 1982) and is a richer specification of firm behavior than classical profit maximization insofar as it recognizes the inherently multi-agent character or real firms, in which the preferences and ambitions of all employees play out with each agent engaged in purposeful behavior. Team production models are a way to get inside the 'black box' of the firm, to relax the Panglossian perspective of the firm as a perfectly functioning profit engine. Such models have had important empirical success, in explaining firm size and growth rate distributions, for example (Axtell 1999). Profit maximizing models, based on U-shaped cost curve considerations, have vague empirical relevance, a point first made by Simon and Bonini (1958) long ago.

In what follows we first review the profit maximization hypothesis and recite the conventional analysis of regulatory instruments as applied to profit maximizing firms. In doing so important aspects of this conventional view, useful for developing an alternative perspective, are emphasized. Then, the scientific status of this hypothesis is assessed, both its logical standing, based largely on informal evolutionary arguments, as well as its empirical underpinnings. In each case there is precious little evidence that can be brought in favor of the hypothesis, while alternative hypotheses are legion. Specifically, any account of intra-firm strategic behavior seems to lead away from profit maximizing behavior toward a more behavioral theory of the firm (Cyert and March 1963).

I.A Regulation Under the Profit Maximization Hypothesis

In the conventional view of the perfectly competitive firm, profit maximization is equivalent to cost minimization (Varian 1992).

Therefore, when it comes to pollution abatement the profit maximization hypothesis reduces to abatement cost minimization (see Stavins (forthcoming) for a typical presentation). Typically, the abatement cost function is considered to be a continuous, twice differentiable, convex function of abatement, *c(a)*. A firm facing emissions tax rate τ will abate until the marginal cost of abatement is equal to the tax rate, i.e.,

$$\frac{dc(a)}{da} = \tau$$

In a tradeable emissions permit regulatory environment the same kind of result obtains, with the tax rate replaced by the permit price, i.e.,

$$\frac{dc(a)}{da} = p$$

The superior performance of market-based instruments is technically clear when firms are profit-maximizers. One way to think about these results is a follows. Compared to any command-and-control regime, market-based regulations add degrees of freedom to firm behavior. Therefore, it can only result in reduced costs—firms could always choose to abate at the command-and-control levels.

But in this presentation of the advantages of market-based approaches is also contained the germ of the main idea of this paper—that deviations from profit-maximization can lead to loss of efficiency of market-based instruments. When one admits that firms can operate away from global optimality then the additional degrees of freedom associated with the market-based approaches may be of no consequence since it is the intrinsic incentive problems facing the firm that command-and-control regulations might circumvent, and giving more behavioral freedom to firms does not solve the basic incentive problem. But before providing an example we first review the problems with conventional justifications of the standard view of the firm.

I.B Profit Maximization as an Hypothesis

Winter (1993) has argued that the conventional view of firms as unitary actors and profit maximizers is not only empirically flawed, but also ironic methodologically. For economists have pioneered, among all the social sciences, the view that the behavior of individual agents must be explicitly treated if the performance of a social systems overall is to be explained. This is how economists treat consumer behavior, markets, institutions, and so on. But curiously, when it comes to the firm this characteristic methodological individualism is jettisoned in favor of a unitary actor framework, not unlike what one finds in international relations theory in political science or macro-sociological theory.

In this section we take a critical view of the conventional view of firms, ultimately concluding the standard view has little to recommend it.

I.B.1 Logical Character of the Hypothesis

Early proponents of the profit maximization hypothesis argued for it on essentially logical grounds, based (informally) on a crude formulation of the rationality axiom more is better. Essentially they argued that among the many ways a firm *might* behave, maximizing profits would be a reasonable goal insofar as all other behaviors would result in less profit, and because firm participants preferred more to less.

There are many arguments against this flavor of the hypothesis, most of which will be brought up below. One particular one, ostensibly due to Simon (1947), is that for many organizations the determinants of profit are understood so poorly that it is not really possible to execute such a calculus.

I.B.2 Evolutionary Justification of the Hypothesis

A well-known defense of the hypothesis is due to Friedman (Friedman 1953), who argued, in essence, that firms which fail to profit maximize are doomed to extinction when in competition with ones that do. The utter simplicity of this argument led later to arguments against it, by eminent philosophers such as E. Nagel and fellow economists such as H. Simon .

More recent work by evolutionary game theorists (Blume and Easley 2002) demonstrates that, particularly in the case of external firm financing through capital markets, there is only a weak connection between profit maximization and firm survival.

I.B.3 Technical/Engineering Assessment of the Hypothesis

A different argument against the hypothesis is that the role of engineers and other technical staff within production facilities is to perpetually bring plants closer to the production possibilities frontier (PPF), as it moves outward by virtue of technological change. In this view, no firm ever manages to make it to the true profit maximizing point of the PPF, while always being in a state of at least partial renovation and improvement. The role of operations personnel is distinguished from product/process design personnel in this perspective. The mere existence of such personnel within the firm is *prima facie* evidence for violations of the profit maximization hypothesis. Of course, it becomes an empirical question as to just how from the profit maximizing position real firms actually lie.

Strong empirical evidence against literal profit maximization can be found in the industrial ecology literature. There the well-known program at Dow Chemical known as ‘Waste Reduction Always Pays’ (WRAP) involved a corporate-level commitment to energy savings investments (primarily) over several years, activities that would probably not have been the main activities of the engineering personnel in question. Over time many millions of dollars worth of cost savings were extracted from the myriad production facilities of this giant manufacturer. Indeed, today sizeable fractions of Dow’s profits are attributable to the return on investment from this program.

Further evidence of this kind for departures from profit maximization include the widespread use of business consultants by modern businesses. Presumably, if firms were already operating at their profit maximizing point they would have little need for such services.

I.B.4 Strategic Dimensions of the Hypothesis

In rational choice game theory an organization having a multi-divisional structure is considered to face various incentive problems associated with getting each division to pay an appropriate share of organizational overhead. This is the so-called *cost allocation* problem (Young 1985), about which a considerable amount is known. For instance, an important class of cost allocation problems have to do with the decision problem faced by a central manager when divisions strategically report their local cost functions. This problem is very similar to that faced by a corporate executive in charge of emissions compliance, who has to figure out how much abatement the divisions of the corporation will undertake. A variety of impossibility results are known in this domain, such that there exists no mechanism for the center to employ on its divisions that is incentive- compatable, individually-rational and efficient. Without going into the details, we simply assert that many of these issues will arise in the model described in section II below.

I.B.5 Empirical Plausibility of the Hypothesis

Today in the economics profession, when a behavioral specification is in doubt recourse can often be made to laboratory experiments for purposes of empirical testing of the specification. Such empirical approaches are well-known today as experimental or behavioral economics. Interestingly, it appears there has been essentially no work testing the profit maximization hypothesis in the lab, particularly from the multi-agent point of view. It is unclear whether this is a conceptual difficulty or simply lack of interest on the part of experimentalists. Lacking direct evidence, one might seek out case study evidence, but this too seems woefully lacking in quantity and, where it exists, in quality. Overall, today there is essentially *no* empirical evidence in favor of the hypothesis.

I.B.6 Other Dimensions of the Hypothesis

During the late 1990s, when many Internet businesses had yet to show any profit whatsoever, and one might have imagined replacing profit maximization behavioral specifications with loss minimization ones, it was common to hear investors suggest that profits did not matter, or at least were of subordinate importance. Rather it was firm share price that was the key indicator of firm performance. Profitless firms with large share prices were suitable investments, as long as prices kept rising.

Even after the burst of the Internet bubble, it is still possible to find investors who ostensibly care more about share price than profits *per se.* Now arguably, the latter are worked into the former, but this is in any case an imperfect process, and there is presumably weaker causalty from the former to the latter. 'What do firms maximize?' is a question a diehard neoclassical might ask. More relevant, it seems to us is simply 'What do firms do.

Relatedly, consider profit maximization from a macro-social perspective. Imagine an industry where there exist a variety of firms engaged in various non-profit maximizing behavior. Does it make sense for a fully rational firm competing with these firms to even want to maximize profits? Perhaps having fewer profits but more loyal customers is a way to compete against a very large, price cutting competitor. Or maybe losses in the short run are sufficient to drive a competitor out of business and make more money in the long run.

I.B.7 Falsifiability of the Hypothesis

These considerations bring up another aspect of the hypothesis: to what extent can it even be falsified. That is, are there definitive data on the behavior of firms that could be gathered to convincingly demonstrate the hypothesis is false? I believe that the existence on on-going plant improvements and technical personnel, combined with case studies such as Dow's WARP, are evidence against the hypothesis. But some economists will perhaps assert that these goal-directed activities are evidence of profit maximization objectives, even if actual maximization is never achieved, possibly even if real firms are a long way from operating at the profit maximal output level. For people holding such an extreme view the hypothesis ceases to be a scientific statement, since it cannot be falsified.

I.B.8 Scientific Status of the Hypothesis

Indeed, I suggest that many practicing economists actually 'believe in' the hypothesis in this way, not as a scientific claim to be falsified, but rather as a tenet of faith and point of departure for modeling building. Insofar as this claim is true, profit maximization should be treated not as an hypothesis but rather as an axiom of a non-scientific field called pure mathematical economics.

However, for those interested in science, the status of the hypothesis—under the significant and growing weight of the above criticisms—is essentially untenable today and needs to be replaced with richer specifications that admit the multi-agent character of all real firms.

I.C Related Literatures

There are also a number of related literatures—beyond economics—that are directly relevant to models of firms, although they typically make no use of any profit maximization hypothesis. Specifically, the field of so-called 'organization science,' in both its computational and mathematical incarnations, builds models of organizations in which individual agents are represented. These models study the effect of distinct organizational forms, the flow on information within firms, and the resulting organizational efficiencies. Good overviews of work in this includes Carley and Prietula (1994; 1998), Prietula, Carley and Gasser (1998) and Lomi and Larsen (2001).

II Team Pollution

In the more or less standard team production environment, agents get together to work in an economic environment that offers some advantage to cooperation, for example, increasing returns to scale. But each agent behaves non-cooperatively, in the usual sense that there are not contracts in place that dictate the exact quantity and quality of her input to production. Rather, each agent seeks to maximize her utility by contributing to production and receiving compensation, but also from non-productive activities performed outside of the team.

II.A Team Abatement With Market-Based Regulations

When, during the production process, externalities are generated then team production can be thought of as 'team pollution.' Stipulating that pollution abatement is costly, it is of course the case that, in the absence of regulations, the individuals composing the team do not abate.

Consider the case of emissions costing τ/unit, either a tax or price of a permit. Each member of a multi-agent team receives reward for the total amount abated, according to some compensation system. In a team of size N, each agent contributes individually to abatement activities, with a_i being the contribution of the i^{th} agent.[3] The total amount of abatement is

$$A = \sum_{i=1}^{N} a_i$$

and we use the following notation to denote the contributions of all agents except i

$$A_{\sim i} \equiv A - a_i$$

The value of this abatement is τA, a sum that is divided up among the agents; the simplest case is equal sharing and each agent receives $\tau A/N$. Let us say that agents have Cobb-Douglas utility for this income due to abatement, but also value the time they spend not engaged in abatement activities. This is a kind of labor-leisure specification and looks like

$$U_i\left(a_i; A_{\sim i}, \theta_i, N, \tau\right) = \left(\frac{\tau\left(a_i + A_{\sim i}\right)}{N}\right)^{\theta_i} \left(1 - a_i\right)^{1-\theta_i} \qquad (1)$$

For reasonable values of parameters, this utility function is single-peaked in a_i—too little abatement is costly in monetary terms, while too much abatement leaves little time for leisure.

---

[3] It is possible to render this model in completely equivalent terms using emission and their reduction, instead of abatement. We have found the formulation to follow notationally simpler and so use it exclusively herein.

A slightly different interpretation of this basic set-up is also possible. With a richer specification of the cost structure (see Appendix) one imagines a plant manager faced with the following internal governance problem. Many distinct facilities within his plant contribute to overall emissions. These distinct 'stove pipes' are managed in an integrated way, with some underlying connectivity between units, but overall each process manager has significant autonomy with respect to varying process inputs, manpower and staffing issues, production schedules and operating parameters, and emissions abatement. It is desired to have each manager contribute to abatement, but the plant manager has only partial knowledge of the technical feasibility of abatement from specific facilities. That is, the plant manager must rely heavily on his or her subordinates in order to accomplish a given level of abatement. In order to deal with this 'knowledge problem,' the manager provides incentives to the process managers, telling them all operations will have to share the cost of emissions taxes or permits, and will similarly share in the cost savings from abatement decisions. However, each manager has a limited amount of attention or effort to supply to the abatement project. Therefore, there is an internal incentive problem that the plant manager cannot completely solve.

II.B Agents Undersupply Abatement at Nash Equilibrium

Each agent must decide how much abatement to contribute to the team. There are two specifications of agent behavior that yield identical results, one conventional and requiring significant cognitive abilities on the part of the agents, while the other is myopic and inductive. We treat these in turn.

First, assume the parameters of the problem are known to each agent, including the abatement activities of the other agents. The i^{th} agent chooses its optimal abatement level, according to

$$a_i^* \equiv \arg\max_{a_i}\left[U_i\left(a_i; A_{\sim i}, \theta_i, N, \tau\right)\right]$$

This can be solved in closed form as

$$a_i^*\left(A_{\sim i}, \theta_i\right) = \max\left[0, \theta_i - A_{\sim i}\left(1 - \theta_i\right)\right]$$

Note that the optimal amount of abatement does not depend explicitly on either the size of the firm or the cost of abatement.

The essence of this team production view of abatement is that an individual's level of abatement is decreasing in the amount of abatement undertaken by his teammates. That is, each member of the team is a 'social shirker,' someone who puts in less abatement effort as her teammates put in more. For each agent in the team there is a sufficiently large value of team effort beyond which the agent expends no effort in abatement activities at all.

In the second interpretation of individual behavior, agents have limited knowledge of the functional form of their utility function or its parameters, and do not know the contributions to abatement of their teammates. Rather, given some current abatement level, they simply grope around for utility-improving abatement levels. Once all this mutual groping stops when all agents are at the abatement levels given by () above.

II.B.1 Nash Equilibrium

When each agent behaves in this way it yields a Nash equilibrium in abatement levels, i.e.

$$a_i^*\left(A_{\sim i}^*, \theta_i\right) = \max\left[0, \theta_i - A_{\sim i}^*\left(1 - \theta_i\right)\right] \tag{2}$$

where $A^*_{\sim i}$ is the sum of Nash abatements.

II.B.2 Identical Agents: Symmetric Nash

This can be seen more clearly in the case of homogeneous agents. Substituting $(N\text{-}1)a^*_i$ for $A_{\sim i}^*{}_{\sim i}$ in (2), one then solves directly to get

$$a^{Nash} = \frac{\theta}{\theta + N\left(1 - \theta\right)}$$

The Pareto optimal level of abatement for identical agents is determined by making the previous substitution, $A = Na$, before optimizing, and then solving for the a that extremizes utility. Doing so yields.

$$a^{Pareto} = \theta$$

Comparing these two abatement levels we see that agents abate less at the Nash equilibrium, i.e.,

$$a^{Pareto} \geq a^{Nash} \Rightarrow \theta \geq \frac{\theta}{\theta + N\left(1 - \theta\right)}$$

which is true as long as

$$\begin{aligned}\theta + N(1-\theta) &\geq 1\\ N(1-\theta) &\geq (1-\theta)\\ N &\geq 1\end{aligned}$$

which is true in general.

It can be shown that utility is also greater in the Pareto case. We can see this by proceeding step by step:

$$\begin{aligned}U^{Pareto} &\geq U^{Nash}\\ (c\theta)^{\theta}(1-\theta)^{1-\theta} &\geq \left(\frac{c\theta}{\theta + N(1-\theta)}\right)^{\theta}\left(\frac{N(1-\theta)}{\theta + N(1-\theta)}\right)^{1-\theta}\\ (c\theta)^{\theta}(1-\theta)^{1-\theta}\left(\theta + N(1-\theta)\right) &\geq (c\theta)^{\theta}\left(N(1-\theta)\right)^{1-\theta}\\ \theta + N(1-\theta) - N^{1-\theta} &\geq 0.\end{aligned}$$

It is analytically difficult to prove that this is true in general, but a plot of the LHS as a function of N and θ, $f(\theta, N)$, demonstrates this clearly (Figure 2).

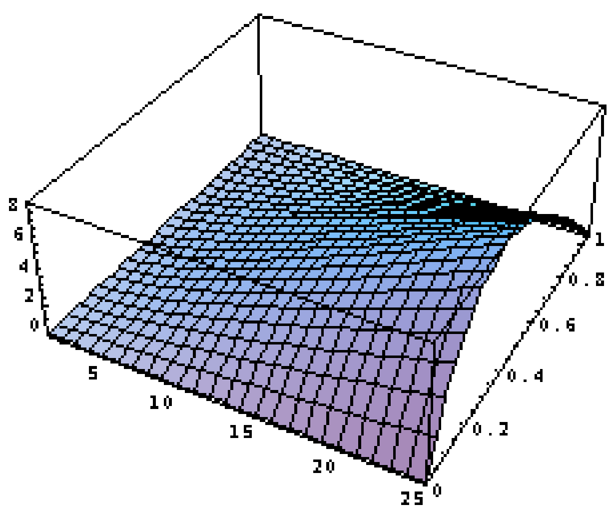

Figure 1: $f(\theta, N) \geq 0$

What this means is that there exist higher abatement levels that make all agents better off; they make each process manager agent better off, and of course the plant manager, but also because more abatement means less pollution, society is better off as well.

II.C Command-and-Control Regulations

Non-market mechanisms have a long history in environmental policy (Portney and Stavins 2000). While the common perception today is that these regulatory instruments are typically costlier than market-based approaches, we shall demonstrate here that this is not the case in general. Indeed, we shall argue that for a wide range of 'team pollution' problems quite the contrary is true.

Commonly, two types of command-and-control regulations are distinguished, so-called technology standards and performance standards. In the former, regulations prescribe specific technologies, e.g., scrubbers for sulphur dioxide emissions, and thereby attempt to guarantee that industry is using known technology with proven reliability in order to reduce pollution levels. In the latter, regulations simply stipulate the maximum level of emissions and the technology employed to accomplish this is left up to industry. We investigate these in turn.

II.C.1 Technology Standards

In the context of the present model, it is clear, conceptually, the way that technology standards can be used to effectively control emissions. Given that processes with a plant are under-abating, the role of a specific technology can be to increase abatement up to some 'typical' level. If this level exceeds the 'under-supply' within the firm then welfare advantages accrue to all agents.

In this circumstance the incentive problems within the firm are 'solved' by technology. The technological standard is more efficient than the incentives facing polluting agents. An analogous situation occurs whenever a new computer networking standard, for instance, supplants an old one. Agents who have coordinated on an archaic standard cannot unilaterally depart from it, but rather need guidance as to what the new standard will be. Standards setting bodies do this for electronic components, but no analogous institution does this for environmental problem. That is, in the absence of a third party that says, in effect, 'all scrubbers at coal-fired power plants will emit no more than xx ppm of sulphur,' a technology prescription from industry can produce the same effect.

II.C.2 Performance Standards

Performance standards seek pollution reductions while permitting firms to choose the technological they will use. While somewhat more flexible than technology standards, they are quite similar in spirit. For our purposes, performance standards permit us to clearly demonstrate how market-based instruments in the context of strategic intra-firm behavior can be dominated by other instruments.

Consider the simplest case, a two agent team in which each agent contributes to production while also polluting. Agent utilities are as in (1) above. For the moment consider these agents to be identical, and then compute their optimal Nash and Pareto abatement levels. These are shown by 'N' and 'P' in Figure 2 below.

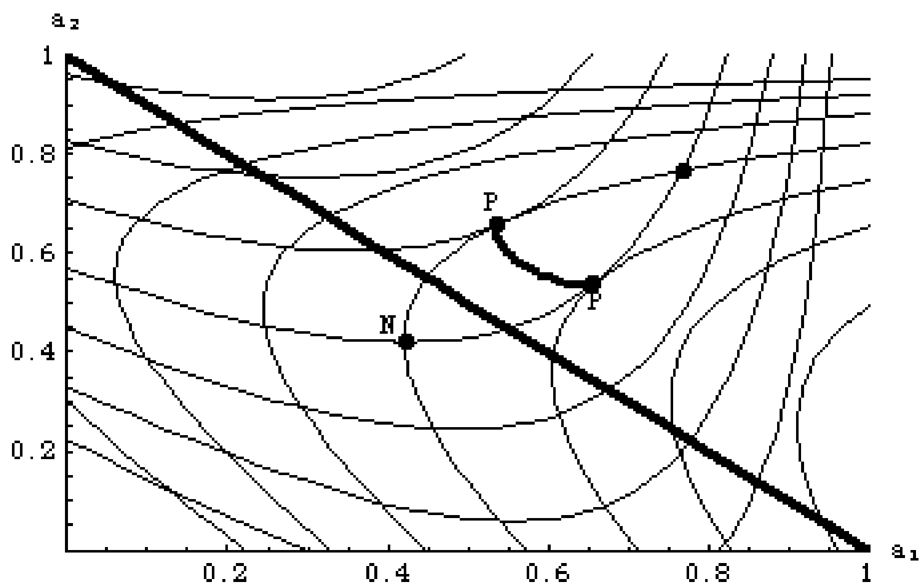

Figure 2: Iso-utility lines (light) for a two agent firm as a function of abatement levels; also shown is an iso-abatement line (heavy, straight) along with the Pareto optimal levels of abatement (heavy, curved).

Lines of iso-utility have been drawn through these points, along with other iso-utility lines. For agent 1, whose abatement level is shown along the horizontal axis, its iso-utility lines are 'u' shaped and increase up the page. For agent 2, whose abatement level is shown along the vertical axis, its iso-utility lines are 'c' shaped and increase to the right. The lens shaped region that starts at 'N' and moves up and to the right and which includes the 'P' to 'P' locus of points, represents an area in the abatement space in which both agents are made better off—each receives higher welfare and society receives windfall abatement. Any performance specification that deposits the two agents into the lens-shaped (core) region yields welfare gains in comparison to the market-based outcome at 'N'. A line of constant total abatement is shown in the figure as a heavy black line running from NW to SE. All point along this line that are also in the core represent regulatory standards that dominate market-based instruments.

III Estimating the Social Welfare of Regulation

Having demonstrated now non-market regulatory mechanisms can out-perform market-based ones, we now turn our attention to determinations of the corresponding cost savings. From Figure 2 it is clear that the performance standard (of the black line, say) provides higher welfare for all participants. It requires higher levels of abatement, which is costly, but each polluting agent is happy to pay this cost since each is better off in the new regime. In this sense, it is not abatement costs which are saved but rather agents are willing to spend more on abatement if they are, in essence, guaranteed that their partner in production will do the same.[4]

In conventional assessments of regulatory performance, it is the explicit costs of compliance or abatement that are compared. Often, such comparisons are not made on the basis of 'before-and-after' studies, but rather by resort to either engineering or economics calculations, and therefore such assessments are indirect. For instance, it is common in estimating smooth cost of abatement functions from engineering data to assert that certain coefficients have certain values because firms are profit maximizers or, what is the same thing, price-takers in a perfectly competitive market (e.g., Newell and Stavins 2003). Alternatively, very often emissions are not measured at the 'end-of-the-pipe,' as it were, but are the result of materials balance, chemical kinetics, or other technical calculations whose input data originate, for instance, from samples of feedstock, such as when sample sulphur levels in coal are used to computer total sulphur dioxide loadings.

In each of these cases the cost saving accruing to market-based instruments are biased insofar as they assume firm behavior that is merely hypothetical. That is, if firms are not profit maximizing in practice then estimates of cost savings based on such suppositions are biased. This is an important lacuna in extant cost-benefit methodology that has apparently gone unnoticed or whose significance is at least is under-appreciated.

Given profit-maximizing behavior, heterogeneous costs of abatement *require* that there be cost savings for market-based instruments. The floor of cost savings is $0, since firms could always abate at the command-and-control level if it were optimal to do so. But as soon as any alternative to profit-maximization is substituted as a behavioral specification, these results no longer hold. This is the sense in

[4] There is a sense in which we have turned the original non-cooperative team production problem into a cooperative game, by introducing the welfare-improving non-market regulation.

which we say that the optimality of market-based instruments is sensitive to firm behavior specifications.

IV Summary and Conclusion

The practice of environmental economics lags developments in economic theory. The conception of firm behavior as profit maximizing, which pre-dates modern mathematical economics, began to give way approximately a generation ago, at about the same time, coincidently, that market-based approaches to regulation began to actively considered. Although this view of the firm is still taught to beginning graduate students in economics, it is not one that a student of industrial organization will believe in having successfully completed her studies.

Today, on a variety of fronts, this conventional conception of the firm is under attack. Within economics a variety of game theoretic versions of the firm are available, in both cooperative and non-cooperative flavors. In none of these does profit-maximization withstand the gales of strategic interactions intra-firm. The field of organization science also yields models of organizational behavior that depart systematically from profit-maximization.

However, this paper has demonstrated that the desirability of market-based regulations depends sensitively on the profit maximization specification, and to relax this, as was done above, renders the market-based instruments incapable of producing a first-best solution, leaving the door open for non-market mechanisms to generate welfare improvements.

IV.1 The Death of Theoretical Justifications for Command-and-Control Regulations Has Been Greatly Exaggerated

The moniker of ‘command-and-control’ applied to the class of regulatory mechanisms that dicatate technology or performance standards clearly has a pejorative connotation. This is unfortunate, for in many cases it would seem completely obvious that regulators working in specific domains would have sufficiently advanced knowledge of pollution abatement equipment that they would be well-positioned to recommend technologies to polluting facilities that might outperform the solutions arrived at by the engineers and managers in the firm itself, none of whom, *a priori*, would have comparable knowledge about abatement technology.

Furthermore, the argument for performance standards is similar. A set of regulators would develop intimate knowledge of typical abatement technologies so that in prescribing a standard they would, either implicitly or explicitly, be not only be raising the standards of all facilities—taking them all closer to the production possibilities frontier and being a source of technology transfer—but also providing a device to move firms out of low-performance 'traps' as well, subverting the incentive problems that exist within all real firms.

When firms are viewed as perfect profit maximizers then it is completely obvious that market-based regulation will always out-perform other schemes, since the market approaches add degrees of freedom to the firms' behaviors. But as soon as one admits that firms are unlikely to do anything like profit maximization in practice then the dominance of market-based mechanisms in general is lost. Now it may be that in particular cases market approaches will outperform non-market regulations. But at the end of the day this is an empirical question, not one that can be answered with once-and-for-all with grand theorizing.

IV.2 The Decline of the 'Neoclassical Synthesis' and the End of Universal Economic Principles

This argument is actually a special case of a more general one that goes as follows: As economic models are enriched behaviorally, through the demise of rational choice specifications and replacement with empirically-credible behavioral specifications, and as such models are taken from the micro-level to the macro via computational (agent-based) techniques—given the intractability of purely analytical methods in the face of such specifications—the notion of universal economic principles will become increasingly archaic, until not a single one is left standing. Ideas like 'free trade should always be promoted,' or 'the law of one price,' or 'rent controls are bad,' or 'minimum wage laws cause unemployment,' which are today seen by many as always and everywhere true, will become more contingent, more dependent on institutional behavior, and less forceful as policy proposals. In the future the only way to compare policies will be by recourse to models that depict alternative worlds, and the dominant features of these alternatives—equity, efficiency—will determine policy.

For at the end of the day the only universal economic principle is behavioral, not economic—that agents are purposive, self-interested, boundedly rational. From this foundation arise the fortresses and cathedrals, the ghettos and kingdoms, the institutions and unintended consequences of economic reality.

Appendix: Richer Cost and Behavioral Specifications

Here we show that more detailed cost and utility functions yield results that are qualitatively identical to those described in the main text, albeit at a cost of significant analytical complexity.

Imagine that each agent has a financial budget, *fi*, and an effort budget, *bi*. The financial costs of abatement are quadratically increasing with abatement level, while the effort costs are linear. This leads to the following specification of utility:

$$U_i\left(a_i; A_{\sim i}, \theta_i, N, \tau, f_i, e_i, k_i, c_i\right) = \left(f_i - c_i a_i^2 + \frac{\tau\left(a_i + A_{\sim i}\right)}{N}\right)^{\theta}\left(e_i - k_i a_i\right)^{1-\theta}$$

Note that one recovers the earlier specification by setting and to 0 and and to unity. First order conditions are now messy and there are several branches of the solution, depending on parameters. In the symmetric case (identical agents) one can develop an expression for the optimal level of abatement, once again with several branches, depending on parameters. It is very cumbersome and is displayed here merely for purposes of completeness

$$a_i^* = \frac{k\tau\left(\theta + N\left(1-\theta\right)\right) + 2ckN\theta \pm \sqrt{\left(k\tau\left(\theta + N\left(1-\theta\right)\right) + 2ceN\theta\right)^2 - 4ckN\left(1+\theta\right)\left(ef\theta\tau - fkN\left(1-\theta\right)\right)}}{2ckN\left(1+\theta\right)}$$

where all subscripts denoting heterogeneity have been omitted. Comparative static analysis of this expression reveals that it behaves in accord with intuition, i.e.,

$$\frac{\partial a_i^*}{\partial \tau} \geq 0$$
$$\frac{\partial a_i^*}{\partial \theta} \geq 0$$
$$\frac{\partial a_i^*}{\partial c} \leq 0$$
$$\frac{\partial a_i^*}{\partial k} \leq 0$$

at least within those parts of the parameter space investigated.[5] The first expression simply means that as the cost of polluting goes up, each agent abates more. The second inequality implies that agents who value income more do more abatement. The third line indicates that less abatement is optimal as the marginal cost of abatement rises. The final result indicates that as the effort cost of abatement rises, less abatement is accomplished.

It is tedious but possible to demonstrate (numerically) that this Nash equilibrium level of abatement is everywhere less than the corresponding Pareto optimal level, and that utility is also reduced at the Nash equilibrium.

References

Axtell, R. L. (1999). The Emergence of Firms in a Population of Agents: Local Increasing Returns, Unstable Nash Equilibria, and Power Law Size Distributions. Working paper. Santa Fe, N.M., Santa Fe Institute.

Blume, L. E. and D. Easley (2002). "Optimality and Natural Selection in Markets." Journal of Economic Theory **107**(1): 95-135.

Carley, K. M. and M. J. Prietula (1994). Computational Organization Theory. Hillsdale, N.J., Lawrence Erlbaum Associates.

Cyert, R. M. and J. G. March (1963). A Behavioral Theory of the Firm. Englewood Cliffs, N.J., Prentice-Hall.

EPA (2000). Guidelines for Preparing Economic Analyses. Washington, D.C., U.S. EPA.

Friedman, M. (1953). Essays in Positive Economics. Chicago, Ill., University of Chicago Press.

Hahn, R. W. (1984). "Market Power and Transferable Property Rights." Quarterly Journal of Economics **99**: 753-765.

Hahn, R. W. and R. L. Axtell (1995). "Reevaluating the Relationship between Transferable Property Rights and Command-and-Control Regulation." Journal of Regulatory Economics **8**(2): 125-148.

Hahn, R. W. and R. G. Noll (1982). Designing a Market for Tradable Emissions Permits. Reform of Environmental Regulation. W. A. Magat. Cambridge, Mass., Balinger.

Holmstrom, B. (1982). "Moral Hazard in Teams." Bell Journal of Economics **13**: 324-340.

[5] The actual expressions for these quantities are not reproduced here due to their long length. They are sufficiently complicated that it seems only feasible to work with them using symbolic mathematics software, e.g., Mathematica.

Lomi, A. and E. R. Larsen, Eds. (2001). Dynamics of Organizations: Computational Modeling and Organization Theories. Cambridge, Mass., MIT Press.
Maleug, D. A. (1990). "Emission Credit Trading and the Incentive to Adopt New Pollution Abatement Technology." Journal of Environmental Economics and Management **16**: 52-57.
Montero, J.-P. (forthcoming). "Pollution Markets with Imperfectly Observed Emissions." RAND Journal of Economics.
Montgomery, D. W. (1972). "Markets in Licenses and Efficient Pollution Control." Journal of Economic Theory **5**: 395.
Newell, R. G. and R. N. Stavins (2003). "Cost Heterogeneity and Potential Savings from Market-Based Policies." Journal of Regulatory Economics **23**(1): 43-59.
Portney, P. R. and R. N. Stavins, Eds. (2000). Policies for Protecting the Environment. Washington, D.C., Resources for the Future.
Prietula, M. J., K. M. Carley and L. Gasser, Eds. (1998). Simulating Organizations: Computational Models of Institutions and Groups. Cambridge, Mass., MIT Press.
Simon, H. A. (1947). Administrative Behavior. New York, N.Y., Macmillan.
Simon, H. A. and C. Bonini (1958). "The Size Distribution of Business Firms." American Economic Review **48**(4): 607-617.
Stavins, R. (forthcoming). Environmental Economics. The New Palgrave Dictionary of Economics. L. Blume and S. Durlauf.
Stavins, R. N. (2003). Experience with Market-Based Enviornmental Policy Instruments. The Handbook of Environmental Economics. K. G. Maler and J. Vincent. Amsterdam, North Holland/Elsevier Science.
Tschirhart, J. T. (1984). Transferable Discharge Permits and the Control of Stationary Source Air Pollution. Economic Perspectives on Acid Deposition Control. T. D. Crocker. Boston, Mass., Butterworth.
Varian, H. (1992). Microeconomic Analysis. New York, N.Y., W.W. Norton & Co.
Winter, S. G. (1993). On Coase, Competence, and the Corporation. The Nature of the Firm: Origins, Evolution, and Development. O. E. Williamson and S. G. Winter. New York, N.Y., Oxford University Press.
Young, H. P., Ed. (1985). Cost Allocation Methods, Principles and Applications. New York, N.Y., North Holland Publishing Company.