Regulation and Rationality

Richard L. Revesz

It is a great pleasure for me to be delivering the opening lecture of the academic year at the Universitat Pompeu Fabra. I am grateful to the Rector for his invitation and to Professor Pablo Salvador, who has been a wonderful academic colleague for many years, for having encouraged me to be here today. The last time I was here was to chair a dissertation defense and I have very fond memories of that time at your University and in Barcelona.

My talk will focus on four elements of regulatory policy that I believe are necessary for a rational regulatory system. In this talk, I will primarily discuss health and safety regulation. Examples of such regulation, which abound in the regulatory systems of advanced economies, include environmental regulation designed to reduce the concentration of contaminants that have an adverse impact on human health, worker safety regulation designed to avoid accidental deaths or injury in the workplace, and consumer product regulation designed to avoid such harms resulting from consumer products.

The four elements of a rational regulatory policy that I will discuss in today's talk are the following. First, how to determine the appropriate stringency of regulatory standards designed to reduce adverse impacts on human health, including premature deaths. Second, once we have decided the level of protection we want as a society, how to operationalize it into requirements that the actors imposing risks have to meet, and what trade-offs does that choice present? Third, how to set the relative stringency of standards for new and existing sources, respectively, to

avoid the pernicious consequences of excessive grandfathering. Fourth, how to appropriately allocate regulatory authority in federal or quasi-federal systems. These four elements are not necessarily exhaustive, but they are very important and regulatory systems often make the wrong choices.

I. Cost-Benefit Analysis

As to the first element, when the government seeks to reduce the risk that people face from exposure to substances that have adverse impacts on their health and safety, how stringently should it regulate? I will argue that in a rational regulatory system, this determination needs to be guided by cost-benefit analysis. Let me start by discussing the regulation of carcinogens because a lot of risks that regulation seeks to reduce are risks of cancer, often from environmental exposure. For regulatory purposes, carcinogens are treated as no-threshold contaminants. This means that they are assumed to impose health risks, which translate into premature deaths, at every concentration. The lower the concentration, the lower this risk is, but the risk does not go to zero until concentration goes to zero. Why is each regulatory standard not set at zero? The reason is obvious. We cannot have an industrial society that operates in this manner. As a result, concentrations of zero are not what regulatory systems seek to accomplish across the board. Instead, regulatory regimes have typically indicated that they seek to achieve "safe" levels of exposure. What is safe in this context? It is quite simply deciding whether it is appropriate for an individual to be exposed to a lifetime of cancer of 1 in 10,000, 1 in 100,000, 1 in a million, or another number of your choice. The choice of these numbers is simply arbitrary. It is not a scientific inquiry. And, the reason I know that is because no leading university such as

yours would award a PhD for a dissertation saying that safe means a 1 in 100,000 probability of death as opposed to a 1 in 10,000 probability.

Translated into a population level, if we are talking about an environmental contaminant that affects 10 million people, we are asked to decide whether we want to have 1000 deaths from lifetime exposure to this carcinogenic contaminant, which we can accomplish by setting the individual probability to 1 in 10,000. Or, instead, should we choose 100 deaths (by setting the individual probability to 1 in 100,000)? Or should we choose 10 deaths by setting the individual probability to 1 in a million? And, if you think that 10 is too many, would you prefer only one death. That, too, is possible, if the individual probability is set to one in 10 million. And, of course, because of the no-threshold nature of carcinogens, there is no stopping point.

You probably figured out by now that there are consequences to these decisions outside of the domain of public health. We can have any level of protection we want. But there will be costs attached. The more stringent the protection, the higher the costs. And, typically, costs rise rapidly as the level of protection increases. The function is not linear but convex.

There is nothing wrong with imposing costs of the sources of carcinogenic pollutants to give them incentives to reduce their emissions. Quite to the contrary: polluters are imposing costs on other people that they are not internalizing into their production decisions. So, for example, a steel factory optimizes its use of resources, such as iron, labor, technology, and electricity, so that it can produce steel as cheaply as possible. If the price of labor goes up, it might use more technology and so on. Otherwise, it will be run out of market by more effective competitors. But the factory does not pay the costs of using the clean air. Economists refer to this pathology by reference to the divergence between the private costs (those such as iron, labor, and technology, which the company has to pay) and social costs (the clean air that it can take for

free). As a result, the factory will use a suboptimally large amount of clean air, which is to say that it will pollute at a suboptimally high level. The goal of regulation should be to internalize the externality so that the factory sees the harm it imposes on the breathers of air and so that it can therefore make appropriate tradeoffs among all the resources (private and social) that it consumes.

That is good social policy. But should the factory reduce its pollution to a level greater than could be justified in light of the harms that it produces? Some might think that we should always privilege individuals over economic entities. But the steel factory is not a faceless entity. People will bear the costs of regulation. Workers might be paid less. Or the factory might close and cause people to lose their jobs. Or the factory will be less profitable and its shareholders (perhaps pension plans for retirees) will be adversely affected.

Once one looks at all the consequences for a no-threshold contaminant, one is inevitably drawn to choose the level of protection that has the highest net benefits: the level that maximizes the differences between benefits and costs. That is what cost-benefit analysis does. And, in the process of doing so it leads to the internalization of externalities.

I drive a car that I believe to be safe. It is heavy and has reinforced sides that enable it to withstand relatively well a side impact caused for example by a negligent vehicle crossing an intersection with a red light. I have paid a premium to buy such a car. But I would not have bought an armored tank even if one was available in the market and was safer. I traded off my preferences for safety against my preferences for cost in deciding how to control a risk, such as driving, that I undertake voluntarily. The government should pay attention to similar tradeoffs in deciding how to regulate risks, such as pollution from factories, to which people are exposed

to involuntarily. There is no compelling reason to suggest that it should act without considering the full consequences of regulation.

You might be thinking that we are likely to know the costs of regulation, how do we know what the benefits are and how do we value things like reducing premature death. Economists have figured out ways of doing this by observing decisions people make in market settings. For example, workers who take riskier jobs as opposed to less risky jobs that are similar with respect to other characteristics, get higher wages. From this premium, economists calculate a willingness to pay to be free of the additional risk caused by risky job. And, then, they perform an extrapolation to calculate what is known in this field as the value of a statistical life.

I am not going not going to make any argument here about whether these studies are all great studies or whether they give rise to conceptual problems. It is a complicated area worthy of many lectures. But I do want to say that, in the case of no-threshold contaminants, one cannot defensibly set regulatory standards without considering the tradeoff between the benefits of additional protection and the negative consequences associated with providing such protection. This is what cost-benefit analysis seeks to do.

You might worry that embracing cost-benefit analysis would lead to lax regulatory standards. In the United States, every President since 1981 has had an Executive Order requiring that major federal rules be justified through cost-benefit analysis. There are exceptions to this rule, for instances in which a particular statute, or a judicial interpretation of that statute, says otherwise. So, we have a natural experiment that allows us to make the comparison. In a recently published article, a colleague and I show that under the principal provision for which cost-benefit

analysis is not allowed, the resulting regulatory standards were less stringent than those that would have resulted from the use of cost-benefit analysis.

So far, I have talked about no-threshold contaminants. You might be thinking that the inquiry will be less complex for threshold contaminants—that is for contaminants that have a level below which there are no adverse health consequences. Non-carcinogens are typically treated as threshold contaminants. For such contaminants, it might be desirable to set the level of protection under the threshold. Even in this situation, the level of costs necessary to achieve this goal might be higher than could be justified by the health benefit. But let us leave that problem aside for a moment and focus on the threshold.

The threshold models are problematic for three reasons. First, in order to get there, scientists make indefensible assumptions by creating a sharp discontinuity in their assessment of the strength of the scientific evidence above and below the threshold, respectively. The probability of an adverse effect is treated as 100% above the threshold and as 0% below the threshold. But, typically, the science does not reveal a step function of this sort. Instead, it is more consistent with a continuous function. The probability of an adverse impact below the threshold is not zero, but a positive level that is arbitrarily ignored. Economists, in contrast, have a standard technique for dealing with scientific uncertainty of this sort, which is the concept of expected value. The expected adverse impact of a lower probability is lower than that of a higher probability, but it is not 0% below a certain point and 100% above that point.

Second, the determination of the threshold itself often involves making unsupportable value judgments of what counts as adverse impact. For example, mercury is a very harmful substance. One of its consequences is that is has a very bad impact on the brain development of young children, which translates into the loss of IQ points. When the U.S. Environmental

Protection Agency recently regulated mercury, it determined the threshold to be an average loss of 2 IQ points in the affected population. But there is nothing magic about 2 points. There is just a continuum of harm and the agency picked an arbitrary point and called it a threshold. The reason the agency did that was because it was forced to make its decision under a regulatory straightjacket, caused by the convention that noncarcinogenic effects should be assumed to have thresholds. This is not science, but arbitrary line drawing.

Third, thresholds are determined by reference to a particular type of individual. With respect to smog, for example, the level that it will make it difficult for me to breathe comfortably if I run outdoors is different than the level that an asthmatic will be able to tolerate, which is much lower. And an average asthmatic will be able to tolerate a higher concentration than a particularly sensitive asthmatic. So, even if each individual has a threshold, the population as a whole will not have one as long as there are sufficiently sensitive individuals, which typically there are. Because populations have different levels of sensitivity, even contaminants that are threshold contaminants for an individual are not threshold contaminants for the population.

In summary, contaminants that are treated for regulatory purposes as no-threshold and threshold contaminants, respectively, actually exhibit similar characteristics. For both, the decisionmaker must decide how many premature deaths and how many serious adverse health effects to avoid. There is no intellectually defensible way to make this decision without considering the resources that need to be expended to achieve standards of different stringencies. In broad outlines, that it what cost benefit analysis does.

* * *

I will now turn from the determination of the stringency of the regulatory standard to the choice of regulatory tools to achieve this standard and, here, I will focus on three different elements.

II. Cost-Minimizing Regulatory Tools

Once the regulator has chosen the level of regulatory stringency, it must select the regulatory tool that will impose the necessary obligations on the regulated community. Command-and-control regulation imposes specified obligations on each member of the regulated community. In the environmental context, for example, the regulator might require each polluter to meet the emission standard that results from the use of the best available technology.

Command-and-control standards of this sort are typically not the least cost way of meeting the regulatory goal. The reason is that the regulator will not have sufficiently detailed information to be able to figure out how the regulatory goal can be met at least cost. And the difference between the cost of such command-and-control standards and least cost standards can be considerable.

More flexible approaches provide a solution to this problem. For example, under marketable permit schemes, the regulatory goal determines the total number of emission permits that will be allocated to a region. These permits are then distributed among polluters, generally either by means of an initial auction or through some sort of grandfathering. Subsequently, permits are traded in an open market. Assuming that a robust market for permits arises, a tradable emission permit regime reduces aggregate emissions to the chosen aggregate level for the least cost.

Marketable permit schemes have many design complications and there is a robust academic literature exploring their properties. But for every possible objection to their use, there is a design solution that preserves its attractive least-cost properties. For example, typical marketable permit schemes control the total amount of permissible pollution but not the distribution of this pollution. That is fine for global pollutants, such as carbon dioxide, where all that matters is the total atmospheric loadings. But for local pollutants, such as sulfur dioxide, local concentrations do matter. Nonetheless, it is possible to construct more complex trading schemes that respond to this complication, constraining, for example, trades that would violate ambient standards. The useful inquiries are on how to deal with these matters, rather than perpetuate regulatory tools, like command and control regulation, that have serious inefficiencies.

III. Grandfathering

I will now turn to a second design element. A key feature of the regulatory policies of many jurisdictions, including the United States, is the extensive grandfathering of existing sources from standards that apply to new sources. Grandfathering of this sort has bad incentive effects because it distorts the economic analysis that existing plant owners undertake when deciding whether to modernize or replace a plant. Stricter standards for new sources make building a new plant more expensive than it would otherwise be. As a result, existing sources, often dirty and obsolete ones, remain in operation longer than would otherwise be the case—a phenomenon known as the "old plant effect." This effect is both economically undesirable and may worsen environmental quality by delaying the replacement

of a dirty existing source with a new source, which would be more efficient, and therefore cleaner, even absent a regulatory requirement.

When I became interested in this area and started writing about it, I discovered that the law and economics literature committed the same error as government practice. In the first step, it determined the optimal level of controls on new sources, not taking into account the impact of these standards on existing sources. Then, it determined the optimal transition rule for existing sources, in light of the standards for new sources. Because the costs of retrofitting existing sources to meet a standards are typically much higher than the costs of building new sources with the standard in mind, these transition rules are typically very permissive.

This two-step process has a pernicious effect. Unless one is in an era of great economic growth with additional demand for the products of the regulated entities, existing sources will continue operating with no additional costs rather than building new sources. Sources that would otherwise have been obsolete and closed down because they could not produce a product sufficiently efficiently, would now stay in operation because of the additional large cost of building new sources. The result could be very stringent standards on the books that would not be applied widely.

There is a solution to this problem, which involves understanding that the *relative* stringency of the standards for new and existing sources needs to be considered. The mistake is to optimize the respective standards sequentially: first setting the optimal standard for new sources as if existing sources did not exist, and then setting the optimal transition rules for existing sources. Instead, they need to optimized jointly, so that the difference between the standards does not undesirably stand in the way of technological innovation.

It makes sense for new and existing sources to be subject to different standards for a

period of time, in light of the higher compliance costs of existing sources. But the impact of differential standards on the transition from existing to new sources needs to be considered.

This problem would not arise if regulatory standards were replaced by marketable permit schemes, which I advocated for earlier in this lecture. Then, new and existing sources would compete for permits in the same market and existing sources would stay in operation only if it was economically desirable for them to do so.

IV. Allocating Regulatory Authority in Federal and Quasi-Federal Systems

Let me now address the third and last design element that I'll focus on today. Interjurisdictional impacts provide the strongest argument for allocating regulatory responsibility at the federal (or quasi-federal) level. For example, a state externalizing its pollution to other states can capture economic benefits in the form of jobs and tax revenues, but imposes costs in the form of adverse health effects on other states. As a result, the upwind state is not affected by the full costs of its actions. Here, too, there is an externality: in this case, a divergence between the private costs borne by the state and the social costs that are imposed on downwind states. In the absence of bargaining among states, which is difficult to accomplish, the amount of pollution crossing state lines will be greater than is optimal.

Another prominent justification posits that the harmonization of regulatory standards promotes the establishment of a common market by putting different states at an equal footing in their competition for markets for their products. This justification is prevalent in the European Union and prevalent is a somewhat analogous form in the United States. It is far less compelling than the justification focusing on the presence of interjurisdictional externalities.

The harmonization rationale does have force in the case of product standards. Indeed, a

product cannot trade freely throughout a common market if states within the market can exclude it on environmental or health and safety grounds. Harmonization arguments, however, have also been invoked to justify the vesting of centralized responsibility over process standards, such as environmental ambient and emissions standards. There are several serious problems with extending the argument in this manner.

First, as long as product standards are harmonized, there can be a well-functioning common market regardless of the stringency of the process standards governing the products' manufacture. Thus, more accurately, the argument must call for the harmonization of the products' production costs, so as to deny a comparative advantage to states with lax environmental standards.

The second problem is that the costs of complying with environmental regulation, or, for that matter, the costs of complying with any regulation, are only one component of the total costs of production. Other components include a state's investments in infrastructure, health care, and education, as well as its wages, labor productivity, and access to raw materials. These factors, which can have a significant effect on production costs, are unlikely to be (or are incapable of being) the subject of the European Union's harmonization efforts. Thus, rather than eliminating cost differences, the harmonization of environmental standards has the effect of conferring a competitive advantage on states with lower non-harmonizable components of costs.

Third, the harmonization argument cannot be used, as it has been in the European Union, to justify both uniform ambient standards and uniform emissions standards. A centralized regulatory regime consisting only of uniform ambient standards, which permits the states to allocate the pollution control burden among existing and new sources in any way they see fit, would confer a competitive advantage on the states with smaller industrial bases. Indeed, states

with lower pollution output could offer their sources less stringent emissions standards without violating the ambient standards. The addition of centralized emissions standards moderates this comparative advantage, but it does not wholly eliminate it. Highly industrialized states, where the centralized ambient standards constrain further growth, would be unable to attract new sources without imposing additional costs on existing sources.

If regulated activity does not have interjurisdictional effects, then centralizated regulation means that the local preferences for the level of regulatory stringency are trumped. Typically, then, the resulting regulatory standard reduces social welfare.

* * *

In summary, cost-benefit analysis, marketable permit schemes and proper attention to grandfathering and federalism issues are necessary components of a rational regulatory policy.

I am very grateful to have been invited to give this lecture and thank you very much for your attention.