As Prepared for Delivery:

Regulation in an Uncertain World
National Academy of Sciences
Government-University-Industry Roundtable Meeting
Washington, DC
June 20, 2012

Cass R. Sunstein

As you may have noticed, the national debate over regulation has become unusually politicized and polarized. One of our main efforts in recent years has been to go beyond soundbites and slogans and to focus on evidence and data – to ensure that regulation is empirically informed, that relevant uncertainties are acknowledged, and that sensible tradeoffs are made.

The plea for empirical foundations may seem obvious, a little like a plea for sunshine rather snow. But think for a moment about *Moneyball*, the best-selling book and Oscar-nominated film about Billy Beane, who worked with statistician Paul DePodesta to bring the Oakland Athletics into the top-tier of baseball teams, and in a short time transformed baseball itself, by substituting statistics and empirical data for longstanding dogmas, intuitions, and anecdote-driven judgments.

Consider this exchange:

> *"The guy's an athlete, Bill," the old scout says. "There's a lot of upside there."*
>
> *"He can't hit," says Billy.*
>
> *"He's not that bad a hitter," says the old scout.*
>
> *"Yeah, what happens when he doesn't know a fastball is coming?" says Billy,*
>
> *"He's a tools guy," says the old scout . . .*
>
> *"But can he hit?" asks Billy.*
>
> . . .

Paul reads the player's college batting statistics. They contain a conspicuous lack of extra base hits and walks."

"My only question," says Billy, "if he's that good a hitter why doesn't he hit better?" . . .

Over and over the old scouts will say, "The guy has a great body," or "This guy may be the best body in the draft." And every time they do, Billy will say, "We're not selling jeans here," and deposit yet another highly touted player, beloved by the scouts, onto his shit list.

For a long time, those thinking about regulation have been a bit like scouts in the era before Billy Beane. Scouts said that someone is "a tools guy," or that he "has a great body." Those favoring or opposing rules make analogous claims. "It won't cost a penny," or "it's completely uncertain," or "ten years ago, the following happened," or "here's the worst case," or "let's be as precautionary as possible," or "the industry will completely shut down." Regulatory systems need their own Billy Beanes and Paul DePodestas, carefully assessing what rules will do before the facts and testing them after the fact.

In these remarks, I will focus on two new directions in the last three years.

- **The regulatory lookback**, designing to promote retrospective analysis of rules on the books, to catalogue their effects, to streamline or improve them, and to eliminate those that are not justified. An important goal of the lookback is to use ex post analysis to inform and improve ex ante analysis.

- **New requirements for promoting public participation and requiring quantification,** designed to take advantage of the dispersed and specialized knowledge of the public, while also using state-of-the-art tools to deal with uncertainty, to promote accurate measurement, and thus to ensure that regulation is empirically justified, in advance, by reference to both costs and benefits.

Now for some details.

The Regulatory Lookback

The problem: Rules are placed on the books, often for good reasons. They accumulate. The original analysis of costs and benefits may turn out to be wrong; retrospective analysis may reveal opportunities for streamlining or instead

expansion. The original reason for some of them vanishes. What made sense once no longer makes sense, maybe because of technological change, maybe because of what has happened elsewhere in the system. Rules may produce unintended harm. A particular problem is the rise of cumulative burdens, stemming from the aggregation of rules that may make sense in individual cases, but that when taken as a whole, are not easy to justify.

Retrospective analysis has long been recommended by those interested in empirical assessment of regulations, including Michael Greenstone, former chief economist at the Council of Economic Advisers: "The single greatest problem with the current system is that most regulations are subject to a cost–benefit analysis only in advance of their implementation. This is the point when the least is known and any analysis must rest on many unverifiable and potentially controversial assumptions." By contrast, retrospective analysis can help show what works and what does not, and in the process can promote the repeal or streamlining of less effective rules and the strengthening or expansion of those that are working well. Greenstone thus urges a series of reforms, including randomized controlled trials, designed to "instill a culture of experimentation and evaluation."

In January 2011, President Obama called for a government-wide review of regulations on the books to reduce costs, to eliminate unnecessary burdens, and to get rid of what the President has called "absurd and unnecessary paperwork requirements that waste time and money." The lookback process, formally undertaken under a process known as "retrospective analysis," has catalyzed a large and growing number of significant initiatives.

In November 2011, and after careful consideration of public comments, about two dozen agencies produced final plans, spanning over 800 pages and offering more than 500 proposals. A large number of independent agencies followed suit, responding to a historic request from the President to eliminate unjustified costs on their own.

A great deal has happened in a short time. In the next five years, well over $10 billion in savings are anticipated from just a very small fraction of the more than 500 initiatives. In the future, the lookback process is expected to deliver a great deal more.

In addition to economic savings, that process is creating an improved empirical culture, in which rules are carefully assessed after they are on the books, and in which the assessment sheds light on what was right, and what was wrong, in

prospective analysis. It is no accident that the title of the relevant section of Executive Order 13563 is "retrospective analysis." Agencies are already (1) using retrospective analysis to improve prospective analysis, (2) developing rules in a way that will facilitate ongoing evaluation of their effects, and (3) showing an interest in randomized controlled trials, which can help to reduce uncertainty, and to discipline the range of benefits and costs, before rules are proposed.

The lookback effort, designed above all to ensure retrospective analysis and improve rules accordingly, is no one-time endeavor. Very recently, the President issued a historic Executive Order to institutionalize the process of retrospective review. Executive Order 13610 states that "further steps should be taken . . . to promote public participation in retrospective review, to modernize our regulatory system, and to institutionalize regular assessment of significant regulations."

To promote participation, the Order directs agencies to seek public comments on rules in need to review. To promote priority-setting, the Order directs agencies to emphasize reforms that produce significant quantifiable savings. To promote accountability, the Order requires agencies to provide the public with regular reports on their past efforts and their future plans -- with details and deadlines.

With these steps, the process of retrospective analysis, no less than prospective analysis, is becoming a standard part of the assessment of federal regulations. This is a truly fundamental shift – very possibly the most important development in regulatory review in decades.

New Requirements for the Issuance of Rules

Executive Order 13563 is not merely about retrospective analysis. It also provides a series of directives to govern future rulemaking. Many of these directives are designed to promote careful attention to tradeoffs, to promote the use of the best available evidence, and to ensure quantification while also acknowledging its limits.

Let me emphasize four key points.

1. **Public participation**. Executive Order 13563 makes an unprecedented commitment to promoting public participation in the rulemaking process – with a central goal of ensuring that rules will be informed, and improved, by the dispersed knowledge of the public. Agencies are not merely required to provide the public with an opportunity to comment on their rules; they must

also – and this is a critical point -- provide timely online access to relevant scientific and technical findings, thus allowing them to be scrutinized.

2. **Simplification and harmonization**. Executive Order 13563 specifically directs agencies to take steps to harmonize, simplify, and coordinate rules. It emphasizes that some sectors and industries face redundant, inconsistent, or overlapping requirements. In order to reduce costs and to promote simplicity, it requires greater coordination. The order also explicitly connects the goal of harmonization with the interest in innovation, directing agencies to achieve regulatory goals in ways that promote that interest.

3. **Quantification**. The Executive Order firmly stresses the importance of quantification. It directs agencies "to use the best available techniques to quantify anticipated present and future benefits as accurately as possible" – and to proceed only on the basis of a reasoned determination that the benefits justify the costs. Importantly, it also recognizes that some values cannot be quantified or monetized at the present time.

4. **Flexibility**. The Executive Order directs agencies to identify and to consider flexible approaches that reduce burdens and maintain freedom of choice for the public. Such approaches may include, for example, public warnings, appropriate default rules, or provision of information "in a form that is clear and intelligible." We know that simplification of existing requirements can promote compliance and participation and that complexity can have serious unintended consequences. We also know that flexible performance objectives are often better than rigid design standards, because performance objectives allow the private sector to use its own creativity to identify the best means of achieving social goals. To promote flexibility, we have recently several calls to all executive agencies to reduce reporting burdens on small business and to eliminate unjustified complexity. We have received many important initiatives in response.

The Problem of Uncertainty

For many years, OIRA has offered guidance on the appropriate treatment of uncertainty (through OMB Circular A-4). The Appendix sets out the relevant discussion. For present purposes, let me simply quote, directly, several key points:

- In some cases, the level of scientific uncertainty may be so large that you can only present discrete alternative scenarios without assessing the relative

likelihood of each scenario quantitatively. For instance, in assessing the potential outcomes of an environmental effect, there may be a limited number of scientific studies with strongly divergent results. In such cases, you might present results from a range of plausible scenarios, together with any available information that might help in qualitatively determining which scenario is most likely to occur.

- For major rules involving annual economic effects of $1 billion or more, you should present a formal quantitative analysis of the relevant uncertainties about benefits and costs. In other words, you should try to provide some estimate of the probability distribution of regulatory benefits and costs.

- Your estimates cannot be more precise than their most uncertain component. Thus, your analysis should report estimates in a way that reflects the degree of uncertainty and not create a false sense of precision. Worst-case or conservative analyses are not usually adequate because they do not convey the complete probability distribution of outcomes, and they do not permit calculation of an expected value of net benefits. Whenever possible, you should use appropriate statistical techniques to determine a probability distribution of the relevant outcomes.

- For rules that exceed the $1 billion annual threshold, a formal quantitative analysis of uncertainty is required. For rules with annual benefits and/or costs in the range from 100 million to $1 billion, you should seek to use more rigorous approaches with higher consequence rules. This is especially the case where net benefits are close to zero.

Those more rigorous approaches include the following techniques:

- Use a numerical sensitivity analysis to examine how the results of your analysis vary with plausible changes in assumptions, choices of input data, and alternative analytical approaches.

- Apply a formal probabilistic analysis of the relevant uncertainties, possibly using simulation models and/or expert judgment as revealed, for example, through Delphi methods. Such a formal analytical approach is appropriate for complex rules where there are large, multiple uncertainties whose analysis raises technical challenges, or where the effects cascade; it is required for rules that exceed the $1 billion annual threshold.

Quantification and its Limits

A significant goal of Executive Order 13563 is to promote quantification of both benefits and costs and to ensure, to the extent permitted by law, that agencies proceed only after a reasoned determination that the benefits justify the costs. To that end, the Executive Order states "each agency is directed to use the best available techniques to quantify anticipated present and future benefits and costs as accurately as possible." But the same section adds, "Where appropriate and permitted by law, each agency may consider (and discuss qualitatively) values that are difficult or impossible to quantify, including equity, human dignity, fairness, and distributive impacts."

It should be clear that such values, where relevant, come in different (if sometimes overlapping) categories. For example:

- The issue may involve *specifying the magnitude of relevant effects*. An agency may know, for example, that a rule will reduce the risk of a terrorist attack, but it may not be able to quantify the reduction and thus be unable to convey the extent of the reduced risk. The direction of an effect may be clear, but the magnitude may be difficult or impossible to specify.

- The issue may involve *monetization*. An agency may know, for example, that a particular rule will have a beneficial effect on ecosystems and likely will preserve a known number of a certain species of fish, but it might not be able to translate that effect into monetary equivalents.

- A rule may have significant beneficial or adverse *distributional effects on lower income groups.* Those effects may or may not themselves be quantifiable; perhaps the agency is aware that poor people will be particularly affected, but perhaps it is unable to say to what degree. Even when quantification of the relevant effects is possible, they are not easily used as part of a standard analysis of costs and benefits.

- A rule might be designed to protect *human dignity.* It might, for example, reduce the incidence of rape, or allow wheelchair-bound employees to have easier access to bathrooms. Alternatively, an adverse effect on human dignity may be an unintended cost of a rule. For example, a security rule might involve body searches or scans that some might

consider to be an invasion of dignity or privacy.

It is not always easy to decide how to treat variables of this kind as part of regulatory impact analysis. The appropriate initial step is to promote transparency. OMB has recommended that the best practice is to accompany all significant regulations with (1) a tabular presentation, placed prominently and offering a clear statement of qualitative and quantitative benefits and costs of the proposed or planned action, together with (2) a presentation of uncertainties and (3) similar information for reasonable alternatives to the proposed or planned action.

A key advantage of this approach is that it promotes transparency for the public. If, for example, it is possible to quantify certain benefits (such as protection of water quality) but not to monetize them, then the public should be made aware of that fact. At the same time, qualitative discussion of nonquantifiable benefits and costs should help the public, and relevant decisionmakers, to understand the goal of the regulation and how it might achieve that goal.

Transparency is important, but even when it exists, agencies face serious challenges in resolving the question when and how to proceed when important effects cannot be quantified. For agencies to know how to weigh the relevant values, the governing law is crucial, and it is important for agencies to try to evaluate and rank such values to give a sense of their relative significance. In terms of those values, rule may have significant benefits or relatively small ones. For example, a water pollution rule might cover a large or small number of ecosystems, and those systems may or may not be ecologically significant. To the extent feasible, and with close reference to governing law, agencies should give a sense of the relative magnitude and importance of nonquantifiable variables.

When quantification and monetization are not possible, many agencies have found it both useful and informative to engage in "breakeven analysis." Under this approach, agencies specify how high the unquantified or unmonetized benefits would have to be in order for the benefits to justify the costs. Suppose, for example, that a regulation that protects water quality costs $105 million annually, and that it also has significant effects in reducing pollution in rivers and streams. It is clear that the regulation would be justified if and only if those effects could reasonably be valued at $105 million or more. Once the nature and extent of the water quality benefits are understood, it might well be easy to see whether or not the benefits plausibly justify the costs – and if the question is difficult, at least it would be clear why it is difficult. Breakeven analysis is an important tool, and it has analytical value when quantification is speculative or impossible.

In such cases, retrospective review may turn out to be doubly important. First, it might produce helpful numbers after the fact. Perhaps the agency had a wide range of benefits figures in advance – from moderate benefits to large benefits – and perhaps retrospective review can narrow the range or even produce a point estimate. Second, it can show, in particular contexts, that prospective estimates are in fact possible. We might learn that quantification is possible even if monetization is not. We might learn that monetization is less challenging that it appeared.

In the last thirty years, there have been truly extraordinary advances with respect to analysis of regulatory options; these advances have helped us to save both money and lives. Here as elsewhere, predictions may go wrong, simply because we live in an uncertain world. But I think that it is safe to say that we will see at least equally extraordinary advances in the coming decades. Let's get to work.

Appendix: Excerpt from OMB Circular A-4

Treatment of Uncertainty

The precise consequences (benefits and costs) of regulatory options are not always known for certain, but the probability of their occurrence can often be developed. The important uncertainties connected with your regulatory decisions need to be analyzed and presented as part of the overall regulatory analysis. You should begin your analysis of uncertainty at the earliest possible stage in developing your analysis. You should consider both the statistical variability of key elements underlying the estimates of benefits and costs (for example, the expected change in the distribution of automobile accidents that might result from a change in automobile safety standards) and the incomplete knowledge about the relevant relationships (for example, the uncertain knowledge of how some economic activities might affect future climate change). By assessing the sources of uncertainty and the way in which benefit and cost estimates may be affected under plausible assumptions, you can shape your analysis to inform decision makers and the public about the effects and the uncertainties of alternative regulatory actions.

The treatment of uncertainty must be guided by the same principles of full disclosure and transparency that apply to other elements of your regulatory analysis. Your analysis should be credible, objective, realistic, and scientifically balanced. Any data and models that you use to analyze uncertainty should be fully identified. You should also discuss the quality of the available data used. Inferences and assumptions used in your analysis should be identified, and your analytical choices should be explicitly evaluated and adequately justified. In your presentation, you should delineate the strengths of your analysis along with any uncertainties about its conclusions. Your presentation should also explain how your analytical choices have affected your results.

In some cases, the level of scientific uncertainty may be so large that you can only present discrete alternative scenarios without assessing the relative likelihood of each scenario quantitatively. For instance, in assessing the potential outcomes of an environmental effect, there may be a limited number of scientific studies with strongly divergent results. In such cases, you might present results from a range of plausible scenarios, together with any available information that might help in qualitatively determining which scenario is most likely to occur.

When uncertainty has significant effects on the final conclusion about net benefits, your agency should consider additional research prior to rulemaking. The costs of being wrong may outweigh the benefits of a faster decision. This is true especially for cases with irreversible or large upfront investments. If your agency decides to proceed with rulemaking, you should explain why the costs of developing additional information—including any harm from delay in public protection—exceed the value of that information.

For example, when the uncertainty is due to a lack of data, you might consider deferring the decision, as an explicit regulatory alternative, pending further study to obtain sufficient data. Delaying a decision will also have costs, as will further efforts at data gathering and analysis. You will need to weigh the benefits of delay against these costs in making your decision. Formal tools for assessing the value of additional information are now well developed in the applied decision sciences and can be used to help resolve this type of complex regulatory question.

"Real options" methods have also formalized the valuation of the added flexibility inherent in delaying a decision. As long as taking time will lower uncertainty, either passively or actively through an investment in information gathering, and some costs are irreversible, such as the potential costs of a sunk investment, a benefit can be assigned to the option to delay a decision. That benefit should be considered a cost of taking immediate action versus the alternative of delaying that action pending more information. However, the burdens of delay—including any harm to public health, safety, and the environment—need to be analyzed carefully.

1. Quantitative Analysis of Uncertainty

Examples of quantitative analysis, broadly defined, would include formal estimates of the probabilities of environmental damage to soil or water, the possible loss of habitat, or risks to endangered species as well as probabilities of harm to human health and safety. There are also uncertainties associated with estimates of economic benefits and costs, such as the cost savings associated with increased energy efficiency. Thus, your analysis should include two fundamental components: a quantitative analysis characterizing the probabilities of the relevant outcomes and an assignment of economic value to the projected outcomes. It is essential that both parts be conceptually consistent. In particular, the quantitative analysis should be conducted in a way that permits it to be applied within a more general analytical framework, such as benefit-cost analysis. Similarly, the general framework needs to be flexible enough to incorporate the quantitative analysis without oversimplifying the results. For example, you should address explicitly the implications for benefits and costs of any probability distributions developed in your analysis.

As with other elements of regulatory analysis, you will need to balance thoroughness with the practical limits on your analytical capabilities. Your analysis does not have to be exhaustive, nor is it necessary to evaluate each alternative at every step. Attention should be devoted to first resolving or studying the uncertainties that have the largest potential effect on decision making. Many times these will be the largest sources of uncertainties. In the absence of adequate data, you will need to make assumptions. These should be clearly identified and consistent with the relevant science. Your analysis should provide sufficient information for decision makers to grasp the degree of scientific uncertainty and the robustness of estimated probabilities, benefits, and costs to changes in key assumptions.

For major rules involving annual economic effects of $1 billion or more, you should present a formal quantitative analysis of the relevant uncertainties about benefits and costs. In other words, you should try to provide some estimate of the probability distribution of regulatory benefits and costs. In summarizing the probability distributions, you should provide some estimates of the central tendency (e.g., mean and median) along with any other information you think will be useful such as ranges, variances, specified low-end and high-end percentile estimates, and other characteristics of the distribution.

Your estimates cannot be more precise than their most uncertain component. Thus, your analysis should report estimates in a way that reflects the degree of uncertainty and not create a false sense of precision. Worst-case or conservative analyses are not usually adequate because they do not convey the complete probability distribution of outcomes, and they do not permit calculation of an expected value of net benefits. In many health and safety rules, economists conducting benefit-cost analyses must rely on formal risk assessments that address a variety of risk management questions such as the baseline risk for the affected population, the safe level of exposure or, the amount of risk to be reduced by various interventions. Because the answers to some of these questions are directly used in benefits analyses, the risk assessment methodology must allow for the determination of expected benefits in order to be comparable to expected costs. This means that conservative assumptions and defaults (whether motivated by science policy or by precautionary instincts), will be incompatible with benefit analyses as they will result in benefit estimates that exceed the expected value. Whenever it is possible to characterize quantitatively the probability distributions, some estimates of expected value (e.g., mean and median) must be provided in addition to ranges, variances, specified low-end and high-end percentile estimates, and other characteristics of the distribution. Whenever possible, you should use appropriate statistical techniques to determine a probability distribution of the relevant outcomes. For rules that exceed the $1 billion annual threshold, a formal quantitative analysis of uncertainty is required. For rules with annual benefits and/or costs in the range from 100 million to $1 billion, you should seek to use more rigorous approaches with higher consequence rules. This is especially the case where net benefits are close to zero. More rigorous uncertainty analysis may not be necessary for rules in this category if simpler techniques are sufficient to show robustness. You may consider the following analytical approaches that entail increasing levels of complexity:

- Disclose qualitatively the main uncertainties in each important input to the calculation of benefits and costs. These disclosures should address the uncertainties in the data as well as in the analytical results. However, major rules above the $1 billion annual threshold require a formal treatment.
- Use a numerical sensitivity analysis to examine how the results of your analysis vary with plausible changes in assumptions, choices of input data, and alternative analytical approaches. Sensitivity analysis is especially valuable when the information is lacking to carry out a formal probabilistic simulation. Sensitivity analysis can be used to find "switch points" -- critical parameter values at which estimated net benefits change sign or the low cost alternative switches. Sensitivity analysis usually proceeds by changing one variable or assumption at a time, but it can also be done by varying a combination of variables simultaneously to learn more about the robustness of your results to widespread changes. Again, however, major rules above the $1 billion annual threshold require a formal treatment.
- Apply a formal probabilistic analysis of the relevant uncertainties B possibly using simulation models and/or expert judgment as revealed, for example, through Delphi methods. Such a formal analytical approach is appropriate for complex rules where there are large, multiple uncertainties whose analysis raises technical challenges, or where the effects cascade; it is required for rules that exceed the $1 billion annual threshold. For example, in the analysis of regulations addressing air pollution, there is uncertainty about the effects of the rule on future emissions, uncertainty about how the change in emissions will affect air quality, uncertainty about how changes in air quality will affect health, and finally uncertainty about the economic and social value of the change in health outcomes. In formal probabilistic assessments, expert solicitation is a useful way to fill key gaps in your ability to assess uncertainty. In general, experts can be used to quantify the probability distributions of key parameters and relationships. These solicitations, combined with other sources of data, can be combined in Monte Carlo simulations to derive a probability distribution of benefits and costs. You should pay attention to correlated inputs. Often times, the standard defaults in Monte Carlo and other similar simulation packages assume independence across distributions. Failing to correctly account for correlated distributions of inputs can cause the resultant output uncertainty intervals to be too large, although in many cases the overall effect is ambiguous. You should make a special effort to portray the probabilistic results—in graphs and/or tables—clearly and meaningfully.

New methods may become available in the future. This document is not intended to discourage or inhibit their use, but rather to encourage and stimulate their development.

2. Economic Values of Uncertain Outcomes

In developing benefit and cost estimates, you may find that there are probability distributions of values as well for each of the outcomes. Where this is the case, you will need to combine these probability distributions to provide estimated benefits and costs.

Where there is a distribution of outcomes, you will often find it useful to emphasize summary statistics or figures that can be readily understood and compared to achieve the broadest public understanding of your findings. It is a

common practice to compare the "best estimates" of both benefits and costs with those of competing alternatives. These "best estimates" are usually the average or the expected value of benefits and costs. Emphasis on these expected values is appropriate as long as society is "risk neutral" with respect to the regulatory alternatives. While this may not always be the case, you should in general assume "risk neutrality" in your analysis. If you adopt a different assumption on risk preference, you should explain your reasons for doing so.

3. Alternative Assumptions

If benefit or cost estimates depend heavily on certain assumptions, you should make those assumptions explicit and carry out sensitivity analyses using plausible alternative assumptions. If the value of net benefits changes from positive to negative (or vice versa) or if the relative ranking of regulatory options changes with alternative plausible assumptions, you should conduct further analysis to determine which of the alternative assumptions is more appropriate. Because different estimation methods may have hidden assumptions, you should analyze estimation methods carefully to make any hidden assumptions explicit.