Preliminary draft 5/17/22

"We Test": An Imagined Regulatory Future

Cass R. Sunstein*

Abstract

There can be a serious tension between the commitment to cost-benefit analysis and a realistic appreciation of the limits of official knowledge. Without significant efforts to reduce those limits, that analysis might be inadequately informed. Whenever regulators face significant informational deficits, or what is sometimes called "the knowledge problem," it is important to explore tools that take advantage of what the private sector knows; market-friendly tools, such as economic incentives, have important advantages on that count. An advanced regulatory system should also try to reduce the knowledge problem through three routes: (1) creative use of notice-and-comment rulemaking; (2) retrospective analysis of regulations and their costs and benefits; and (3) advance testing, as a way of informing ex ante analysis. For the future, the most promising approach is (3).

Two Ideas in Tension

Several years ago, I had occasion to be a member of a working group that contained former public officials and people in the business world, including Silicon Valley. Cheerfully, and seeking to make common ground, I emphasized the importance of cost-benefit analysis. A leader of a large company responded: "Cost-benefit analysis? We don't do that. We can't." I was flabbergasted. How could a company make business decisions if it did not do cost-benefit analysis?

His answer came in two words: "We test."

With respect to the past and future of regulation, there are two truly indispensable ideas. Unfortunately, they are in serious tension with one another. Potential solutions lie in three reforms, all connected with democracy itself – but perhaps not quite in the way that most

* Robert Walmsley University Professor, Harvard University. This essay was the basis for the keynote address at a conference on Agile Governance, held in Washington, DC on April 26, 2022. It draws on a chapter in Cass R. Sunstein, The Cost-Benefit Revolution (2018). Readers are asked for their indulgence with an essay originally intended for oral presentation. Many thanks to Susan Dudley and Michael Fitzpatrick for the invitation and for immensely valuable comments and suggestions.

people think. Of these, the most important and radical reform idea is the third. Let's give it a name: "We test." But I am getting ahead of the story.

The first indispensable idea is that it is immensely important to measure, both in advance and on a continuing basis, the effects of regulation of social welfare. As an empirical matter, what are the human consequences of regulatory requirements? That is the right question to ask, but inside and outside of government, it is tempting to focus on other things. These include the opinions of relevant officials and interests ("what does the business community think?" or "do environmentalists feel strongly?"), or purely symbolic or expressive considerations, as in the unhelpful and potentially damaging view that more stringent environmental regulation, an increase in the minimum wage, or strengthened protection of occupational safety is desirable because "it makes a statement."

At the present time, the standard way to answer the right question is to try to identify the costs and benefits of regulations in advance, in order to catalogue and to compare the various consequences, and to help make sensible tradeoffs. To be sure, cost-benefit analysis can create serious challenges, and at the present time, it is a highly imperfect tool. I shall spend some time here on those challenges and on that imperfection. Nonetheless, it has long been thought to be the best tool we have. (We don't test.)

The second idea, attributable above all to Friedrich Hayek, is that government's knowledge is modest, and what is necessary to know is widely dispersed in society.[1] As Hayek and his followers emphasize, government planners cannot possibly know what individuals know, simply because they lack that dispersed knowledge. The multiple failures of plans, and the omnipresence of unintended consequences, can be attributed, in large part, to the absence of relevant information. Are cost-benefit analysts planners? In a way, they certainly are. How can they possibly obtain the knowledge that would allow them to compare costs and benefits? Often they cannot.

Hayek was particularly concerned about socialist-style planning. He contended that even if socialist planners are well-motivated and if the public interest is their true concern, they will fail, because they will not know enough to succeed. Hayek celebrated the price system as a "marvel," not for any mystical reason, but because it can aggregate dispersed information, and do so in a way that permits rapid adjustment to changing circumstances, values, and tastes.

Hayek's arguments offer a serious cautionary note for planners of all kinds, including contemporary regulators who are committed, at least as a general rule, to free markets and freedom of contract. Even if they despise socialism and are simply correcting market failures (as, for example, in the domains of pollution, health care, or occupational safety), they might well lack indispensable information. Suppose that they are seeking to reduce levels of particulate matter in the ambient air. What, precisely, are the health benefits of a reduction of

[1] See Friedrich Hayek, The Uses of Knowledge in Society, 35 Am. Econ. Rev. 519 (1945).

existing levels to 12 parts per billion (ppb), or 11 ppb, or 8 ppb[2]? And what would be the costs, economic and otherwise, of mandating such reductions? When should reductions be required? How should they be obtained? Should small businesses receive exemptions? Of what kinds? What are the alternative approaches, and of these, which is best?

The problem should not be overstated. With respect to costs, regulators often have a good sense of potential outcomes, in part because of information from the regulated sector. To be sure, that information might well be self-serving, but regulators often have sufficient experience to discount alarmist or excessive claims. With respect to benefits, quantification and monetization present separate issues. In some domains, existing knowledge is sufficient to permit the identification of sufficiently narrow ranges with respect to (say) mortalities averted or accidents prevented. Well-established (and continuously improving) tools are in place to convert various values into monetary equivalents. In the day-to-day life of cost-benefit analysis, regulators are hardly making a stab in the dark. Usually they have, or are able to accumulate, a great deal of relevant information.

But still. Modern followers of Hayek are correct to emphasize what they call the knowledge problem, which can be a problem for contemporary regulators of all kinds, working (for example) to implement the Clean Air Act, the Occupational Safety and Health Act, the Affordable Care Act, and the Dodd-Frank Wall Street Reform and Consumer Protection Act. If cost-benefit analysis is essential to sensible judgments, incomplete knowledge, when it exists, would appear to be a serious and potentially devastating problem. In some cases, agencies do face serious challenges in cataloguing costs and benefits. Retrospective analysis attests to those challenges, because it has identified a number of mistakes.

The tension, in short, is that regulators have to focus on costs and benefits (the first indispensable idea), but they will sometimes lack the information that would enable them to make accurate assessments (the second indispensable idea). In light of the knowledge problem, can they produce reliable cost-benefit analyses, or any other kind of projection of the human consequences of what they seek to do, and of potential alternatives? What I am urging, in short, is that regulators are in large part technocrats, charged with measuring and assessing consequences, but their technocratic enterprise runs into a serious objection. Of course the force of the objection will depend on the context, but in some situations, the effort to assess the likely effects of a regulatory intervention (involving pollution, health care, energy, transportation safety, communications, or homeland security) might go badly wrong.

Three reforms can help. I am not suggesting that Hayek himself would be satisfied. Consider this remarkable passage:

[2] On some of the challenges here, see Francesca Dominici et al., Particulate Matter Matters, 244 Science 257 (2014).

> This is, perhaps, also the point where I should briefly mention the fact that the sort of knowledge with which I have been concerned is knowledge of the kind which by its nature cannot enter into statistics and therefore cannot be conveyed to any central authority in statistical form. The statistics which such a central authority would have to use would have to be arrived at precisely by abstracting from minor differences between the things, by lumping together, as resources of one kind, items which differ as regards location, quality, and other particulars, in a way which may be very significant for the specific decision. It follows from this that central planning based on statistical information by its nature cannot take direct account of these circumstances of time and place and that the central planner will have to find some way or other in which the decisions depending on them can be left to the "man on the spot."'[3]

In my view, Hayek's claim here is a mystification, at least as applied to the regulatory context. Statistical information "by its nature" can indeed "take direct account of these circumstances of time and place." How many refrigerator companies are affected by an energy efficiency requirement? How much would they have to pay, in total? Of course it is true that companies differ from one another, and the cost-benefit analyst may not know that for some of them, the cost will be far higher or far lower than anticipated. And of course it is true that for some purposes and activities, statistical knowledge is inadequate. But if the goal is to understand the damage done by air pollution, or the cost of reducing it, statistical knowledge can bring us much of the way home. On, then, to the three reforms.

The first involves the process of notice-and-comment rulemaking, updated for the current era, and with a clear sense of the underlying substantive goal, which is to obtain dispersed information about the likely consequences of regulations (including costs and benefits).

The second involves retrospective analysis of rules, which can both produce changes in those rules and lead to significant improvements in prospective analysis. Sometimes retrospective analysis can be performed a few years after a regulation has been issued, to see if it is having the anticipated effects. Much more ambitiously, new technologies should enable regulators to learn essentially immediately, and to see, in weeks or months, whether regulations are achieving their intended goals or have adverse side-effects. A pressing need is to enlist those technologies to provide that learning. (We're going to get there.) To realize its potential, retrospective analysis should be undertaken with public comment.

The third reform involves testing. It includes careful experiments, above all randomized controlled trials, which can provide far better information than expert judgments.[4] The "we

---

[3] Hayek, supra note.

[4] Duncan Watts, Everything is Obvious (2011); Jim Manzi, Uncontrolled: The Surprising Payoff of Trial-and-Error for Business, Politics, and Society (2012); Michael Greenstone, Toward a Culture of Persistent Regulatory Experimentation and Evaluation, in **N**EW **P**ERSPECTIVES ON **R**EGULATION 113, 113 (David Moss & John Cisterno eds., 2009). A relevant discussion, involving

test" reform is the most promising, but it does present challenges in terms of feasibility. Here as well, there is reason for optimism. Increasingly, those challenges can be, and are being, surmounted.

An Old Debate

During and after Franklin Delano Roosevelt's New Deal, the United States saw an intense debate about government regulation.[5] The competing sides were the New Deal enthusiasts, receptive to the larger regulatory state,[6] and the New Deal critics, insisting that the new administrative institutions were a betrayal of constitutional ideals. One of the enduring products of that debate was the Administrative Procedure Act (APA), enacted in 1946. The APA contained a genuine innovation, now called "notice-and-comment rulemaking." The basic idea is that regulators must provide the public with advance notice of what they are planning to do, and why, and then they solicit comments. It was expected that when agencies finalized rules, they would incorporate what they learned from the public.

Why did Congress call for notice-and-comment rulemaking? The historical record does not give an unambiguous answer, but we can isolate two quite different factors. The first involves the idea of self-government itself. During and after the New Deal, some people have been greatly concerned that regulators are not directly accountable to the people, and have contended that they may suffer from some kind of "democracy deficit." In an extreme view, the result is some kind of crisis of legitimacy. For such critics, notice-and-comment rulemaking may not be sufficient, but it is an important way to help legitimate the administrative process, by increasing accountability and responsiveness. A kind of democratic participation is built into the very idea of notice-and-comment rulemaking. That admittedly technical idea is designed to help to ensure ultimate rule, or at least access, by We the People. If administrators are not directly accountable through elections, at least it can be said that they must listen and respond to the public insofar as its members are willing to raise objections and concerns.

The second idea is less abstract and high-flown, and it is, I think, even more important. It involves information, not legitimation, and it has roots in Hayek's concerns. It is also closely connected with the ideals of the cost-benefit state. We have seen that if government is attempting to make air cleaner or food safer, to reduce deaths in the workplace or on the highways, or to increase homeland security, it might well have incomplete information about the effects of its plans.

---

quasi-experimental techniques in an especially important area, is Francesca Dominici et al., Particulate Matter Matters, 244 Science 257 (2014). It is true that randomized controlled experiments have been subject to some concerns. See Angus Deaton, Instruments of Development (2009), available at http://www.nber.org/papers/w14690

[5] See Daniel Ernst, Toqueville's Nightmare (2013), for a valuable discussion of the background.

[6] The classic discussion is James Landis, The Administrative Process (1935).

Some nonhypothetical examples: If regulators are concerned to eliminate emissions of ozone-depleting chemicals, they might end up banning asthma inhalers, and such bans might have adverse effects on human health. If regulators take steps to make the food supply safer, they might impose high costs on farmers, including small farmers, and potentially create serious economic dislocations. If government imposes high costs on electricity producers, it might produce spikes in the cost of electricity, which would be particularly harmful for the poor. If regulators require motor vehicles to be more fuel-efficient, they might make motor vehicles more dangerous, and thus cause losses of life.

To make sensible decisions, regulators need to obtain a great deal of information about questions of this kind. As hard as they might try, they will not know everything, and they may have significant gaps in their knowledge. Within government, those with technical expertise will try to fill those gaps, but their efforts might be insufficient. Here is the potential of the notice-and-comment process. If regulators have made mistakes or been too optimistic, there is a real chance that members of the public will tell them about it. Regulators' assessment of the costs of a proposed rule might depend on unrealistic assumptions. If so, someone might well object on that ground. Regulators might not have seen how a well-intended rule would affect small business. Their scientific projections might not be consistent with recent scientific findings. They might have neglected local circumstances, falling to understand that what makes sense in Los Angeles and New York is unnecessary or even harmful in Carson City and Boise. They might have missed the potential effects of a low-cost technology. They might not have appreciated the possibility that another approach would have higher net benefits.

On this view, the most important goal of notice-and-comment rulemaking is to increase the likelihood that agencies will obtain relevant information. Some of that information might come from technical specialists outside of government, who can correct agency errors (about, for example, the carcinogenic properties of silica, the social cost of carbon, or the likely costs of emissions controls). Some of it might come from private associations with distinctive knowledge of particular sectors. Some of it might come from people with highly localized knowledge, which might escape the regulators' attention. Some of it might show, consistent with Hayek's cautionary notes, that efforts to lump people together might go badly wrong. If the knowledge problem cannot always be eliminated – and it would be foolish to think that it can be – at least it can be reduced, in part through institutions that increase the likelihood that public officials will learn from what members of the public know.

In the modern era, regulators are in a far better position to collect the dispersed information of the public. On this view, the goal of notice-and-comment rulemaking is emphatically not to take an opinion poll, to take some kind of political temperature, to see how much applause a proposal is able to attract, to defuse public opposition, to engage in some communications strategy, or to collect the digital equivalent of postcards (even though a number of those are sometimes sent in). Instead the goal is overwhelmingly substantive – to fill gaps in knowledge and to see what might have been overlooked. In particular, the agency's assessment of the likely consequences is subject to close scrutiny. If the agency has inaccurately assessed costs and benefits, public participation can and often will supply a corrective.

Democratization of the regulatory process, through the comment process, has an *epistemic* value. It helps to collect dispersed knowledge and to bring it to bear on official choices.

It is true that many rules continue to be lengthy and complex. For the public comment process to work, they must be comprehensible rather than opaque, and technical language and sheer length can reduce comprehensibility. Executive summaries, now required for long or complex rules, can help, but they are not sufficient. Experiments with the idea of a "regulation room,"[7] a website offering informal, plain language versions of regulatory proposals, are designed to promote broader understanding, but those experiments have not been an unambiguous success. In my view, they have been a failure. Some of the relevant comments, in those experiments, look like simple thumbs-up or thumbs-down, and because they fail to provide new information, they do not genuinely inform regulatory judgments.

To come to terms with the challenge of complex rules, it is important to distinguish between two conceptions of the purpose of the public comment process. On one conception, connected with the old idea of legitimation, the goal is to allow We the People – anyone, really – to participate in the process. If so, intelligibility and clarity are indispensable, and complexity is a genuinely serious problem. On another conception, connected with the epistemic idea, the goal is to allow participation by those who have relevant information to contribute. If so, intelligibility and clarity are important, but because the process typically has a large technical component, the absence of genuine "plain language" need not be a fatal flaw, so long as those who have relevant information to add are in a position to do so.

Of Hubris and Risks

It is important to acknowledge that even in its most ambitious forms, and even if we emphasize the overriding importance of substantive contributions, the public comment process might fail to solve the knowledge problem, and there are imaginable risks as well. Centralized planning, of the old-style socialist sort, could hardly be redeemed by public notice-and-comment. It is true that five-year plans would likely be improved by efforts to receive public comments, but we should not, for that reason, favor five-year plans.

Dedicated followers of Hayek would urge that the process of aggregating information will inevitably be imperfect and (in their view) probably worse than that. Markets encode the emphatically local knowledge and values of everyone who produces or purchases relevant products. By contrast, most citizens, including many with substantive contributions to make, are unlikely to know about notice-and-comment rulemaking, or to have the commitment and background that would enable them to participate. There are strong reasons to demand a convincing demonstration of some kind of market failure before embarking on regulation. And even if a market failure is shown, market-friendly responses (such as disclosure of information or corrective taxes) deserve pride of place, in part because they reduce informational demands on regulators and can enlist rather than displace private knowledge.

[7] See http://regulationroom.org/

There is an independent point. It would be possible to fear that well-organized groups, of one of another kind, will inevitably play the most important part in notice-and-comment rulemaking, and even to dominate it. If so, there will be a kind of epistemic "skew" on the part of regulators. The supposed solution to the knowledge problem may make things even worse. The most frequent comments, and the most competent ones, will come from those whose self-interest is at stake, thus creating a distinctive form of capture, one that is epistemic in nature.

In the coming decades, the ability to elicit and compile information will inevitably expand at an extraordinarily rapid rate. A pressing question is how to use that information to reduce the risks of regulatory error, not least with respect to cost-benefit analysis itself. We could easily imagine large-scale improvements on this count as massive data sets become increasingly simple to compile and analyze. In the fullness of time, notice-and-comment rulemaking will become, far more than it is today, a major contributor to forms of data acquisition that are indispensable to accurate cost-benefit analysis.

Retrospective Analysis

Agency rulemaking occurs before the fact, when information-gathering is highly likely to be imperfect. An agency issues a rule involving food safety, occupational health, or energy efficiency, and then it moves on. But even with public participation, the original rule will be based on projections that have a high degree of speculation. The costs are expected to be $X and the benefits are expected to be $Y – or often, both costs and benefits are anticipated to fall within wide ranges, and the agency cannot produce a point estimate. The agency might also be aware of potential unintended bad consequences, but it might think that they are unlikely in the extreme. It might be wrong. Or it might not anticipate some terrible consequence, which occurs two months, two years, or five years after it has acted.

A sensible regulatory system gives continuing scrutiny to regulatory requirements to test whether they are working as anticipated. A central question is whether the ex ante estimates square with what is known ex post. If they do not, regulations can be changed. There is another advantage to retrospective analysis: It can help to inform and improve prospective analysis, as agencies learn about their own mistakes, and can become less likely to make them in the future. In this respect, retrospective analysis can help to reform prospective analysis.

On these counts, existing knowledge remains incomplete; we remain in early stages. But some valuable research can be found. Consider, for example, Winston Harrington's careful study.[8] Building on previous work, Harrington explored sixty-one rules for which benefit-cost ratios could be compared before and after the fact. He found significant errors – but no

8 Winston Harrington, *Grading Estimates of the Benefits and Costs of Federal Regulation* (Res. for the Future, Paper No. 06-39, 2006), *available at* http://papers.ssrn.com/sol3/papers.cfm?abstract_id=937357.

systematic bias. In his account, agencies overestimated both benefits and costs with about equal frequency. Specifically, in sixteen of the sixty-one cases, the ratios were found to be essentially accurate. In twenty-four cases, the ratio was better, not worse, than the agency had anticipated. In twenty-one cases, the ratio was worse than anticipated. Harrington's general conclusion is that while both costs and benefits tend to be lower than estimated, no bias can be found in estimates of benefit-cost ratios.

While highly illuminating, Harrington's study leaves many questions unanswered. The sample size is exceedingly small. Harrington focuses on benefit-cost ratios, which is certainly a relevant question but not the central one. What most matters is *net benefits* and whether agencies have accurately calculated them. Nor does Harrington specify the degree to which benefits and costs were underestimated or overestimated. Other studies do explore the question of underestimation or overestimation. One such study analyzed twenty-one environmental and occupational safety regulations for which retrospective estimates could be found.[9] The basic conclusion is that agencies display a modest tendency to overestimate costs. For thirteen rules, agencies overestimated costs; they estimated costs accurately for four; they underestimated for three; and the costs were indeterminate for one.

In 2005, the Office of Management and Budget, and in particular the Office of Information and Regulatory Affairs, provided an overview of many retrospective analyses based on an examination of forty-seven case studies.[10] A particular concern was the risk that ex ante estimates might be inadequately informed and therefore erroneous. The overview offers three key conclusions. First, agencies were far more likely to overestimate benefits than to underestimate them. More particularly, agencies overestimated benefits forty percent of the time, whereas they underestimated benefits only two percent of the time. Second, agencies tended to overestimate the benefit-cost ratio, and in that sense to be a bit too optimistic about the consequences of their rules. Agency estimates were accurate twenty-three percent of the time, while the ratio was overestimated forty-seven percent of the time and underestimated thirty percent of the time. Third, agencies were slightly more likely to overestimate than to underestimate costs. Agencies were accurate twenty-six percent of the time, overestimated costs thirty-four percent of the time, and underestimated costs twenty-six percent of the time.

From existing work, the most sensible general conclusion is that agencies do make many mistakes (attesting to the reality of the knowledge problem), but there does not appear to be a

---

9 Winston Harrington et al., *On the Accuracy of Regulatory Cost Estimates*, 19 J. POL'Y ANALYSIS & MGMT. 297 (2000).

10 *See See* OFFICE OF MGMT. & BUDGET, VALIDATING REGULATORY ANALYSIS: 2005 REPORT TO CONGRESS ON THE COSTS AND BENEFITS OF FEDERAL REGULATIONS AND UNFUNDED MANDATES ON STATE, LOCAL, AND TRIBAL ENTITIES 41-46 (2005), *available at* http://www.whitehouse.gov/sites/default/files/omb/assets/omb/inforeg/2005_cb/final_2005_cb_report.pdf (collecting studies comparing ex ante and ex post analyses of regulations' costs and benefits, including examples where cost and benefit estimates were off by more than a factor of ten).

systematic bias in any one direction. That is useful and important to know. But it is even more important to acknowledge that we need to know a great deal more than we now do. The existing studies cover only a trivially small fraction of rules on the books. Much more can and should be done to compare prospective estimates to what actually happens in the world. And the real goal should not be the give the government an aggregate grade or to test for systematic bias. It is to see when particular rules are working as anticipated and when they are not – and to fix them when they have gone wrong.

Beyond Expert Projections

My colleague from Silicon Valley was unenthusiastic about the idea of cost-benefit analysis because he did not trust expert projections. Why guess? Why not learn?

To get the facts right, we need to be much more ambitious. We need more rigorous learning and experimentation. The central goal is not to rely even on expert judgments about likely effects, but instead to compile evidence from the real world, not retrospectively but in advance or in real time. In the past decade, there has been growing interest in the use of randomized controlled trials as a means of learning the effects of policy initiatives.

In medicine, of course, it is standard to rely on such trials to see if a drug is safe and effective. For drugs, it would not make a great deal of sense simply to guess, to rely on informed hunches, or even to make simple "before and after" assessments. Suppose that we learn that people who use a certain asthma medicine do better after taking the medicine than before. If so, we know something important—but we do not know nearly enough. The risk with before-and-after assessments is that they may not control for confounding variables. Perhaps people are doing better because of some change in the environment that is not adequately understood by those who are making the assessment. In the medical domain, the value of randomized controlled experiments is that they have the potential to provide a clear sense of the actual effects of the intervention.

Esther Duflo, a Nobel Prize-winning economist at MIT, has pioneered, along with many others, the use of randomized controlled trials for purposes of policy evaluation.[11] In principle, such trials are the best way to solve the knowledge problem; they provide the closest thing to the gold standard, and in at least in some context, they should adequately respond to the Hayekian concern of regulators not having enough information. Duflo has shown that in many cases, small measures can have significant effects. In the regulatory area, the use of such trials remains in a preliminary state. Analysis of costs and benefits is rarely informed by them. But it is easy to imagine serious evaluations. Consider a few examples:

1. Would state regulators save lives by banning the use of cell phones while driving? This is a disputed question. Laboratory experiments, showing that people's reaction times slow

[11] See generally ABHIJIT BANERJEE & ESTHER DUFLO, POOR ECONOMICS (2011).

down when they are distracted, strongly suggest that the answer is affirmative, and indeed that driving while talking on a phone is not unlike driving while inebriated, producing a fourfold increase in relative crash risk.[12] But perhaps those experiments are an unreliable guide to the real world. We could test whether a ban on cell phone use would have major effects on safety by comparing similarly situated localities, one with such a ban and one without. Or we could test whether accidents increase in periods in which cell phone use goes up – for example, when rates decrease after 9 p.m. (In fact precisely that question has been studied, with a surprising finding of no such increase.[13])

2. What are the effects of different methods of increasing rear visibility in cars? If monitors are placed in the dashboard, do accidents drop? How much, and compared to what? Do improved mirrors have a significant effect? What about sonar devices, making beeping noises? Do they work as well as cameras? Randomized trials might help (assuming that sufficiently large sample sizes could be obtained)

3. It is important to evaluate different disclosure requirements.[14] We might test whether different fuel economy labels have different effects on similarly situated consumers. Does one label produce different choices? How different? If labels draw attention to annual fuel costs, are people affected? Do people care about environmental factors? How much? The same kinds of questions might be asked about disclosure requirements for credit cards, mortgages, cell phones, and school loans.

In important areas, experimentation might take the form of advance testing of regulatory alternatives through randomized controlled trials. A movement in this direction would have major advantages over current approaches, such as focus groups, which are often highly artificial and which sometimes test what people like rather than what they would actually do. A presentation might be pleasing without having much of an effect on what people understand and do. And we can supplement randomized controlled trials *by continuing to test after regulations have been adopted – to see where they might be failing and what they are actually achieving*.

In the United Kingdom, there has been a great deal of interest in using randomized

[12] CHARLOTTE L. BRACE ET AL., ANALYSIS OF THE LITERATURE: THE USE OF MOBILE PHONES WHILE DRIVING (2007), available at http://www.nsc.org/news_resources/Resources/Documents/Analysis%20of%20the%20Literature,%20The%20Use%20of%20Mobile%20Phones%20While%20Driving.pdf.

[13] Saurabh Bhargava & Vikram Pathania, Driving Under the (Cellular) Influence: The Link Between Cell Phone Use and Vehicle Crashes 5 AM. ECON. J.: ECON. POLICY 92, available at http://pubs.aeaweb.org/doi/pdfplus/10.1257/pol.5.3.92.

[14] See George Lowenstein et al., Disclosure: Psychology Changes Everything, 6 Annual Review of Economics 391 (2014).

controlled trials, above all through the work of the Behavioral Insights Team (sometimes called the Nudge Unit). What are the effects of different kinds of communications in reducing tax delinquency? In producing higher rates of organ donation? In reducing overprescription of antibiotics? All of these questions can be and are being tested. Related efforts have been made in the United States and elsewhere. If randomized trials are not feasible, we might be able to design experiments that replicate actual behavior by asking people concrete questions about what they would do if provided with certain information or if given a range of options.

Of course there are constraints—involving not merely law but also resources and feasibility (and perhaps equity as well)—in using randomized controlled trials in the regulatory context. Among other things, sufficient sample sizes might be difficult to obtain. But in some cases, they would be both appropriate and useful. We should expect far more progress in the future.

Duncan Watts has offered a variation on this theme. Sounding a lot like my colleague from Silicon Valley, he calls it the "measure-and-react strategy."[15] By way of explanation, he points to the Spanish clothing retailer Zara, which acknowledges its limited foresight and its inability to predict what consumers are going to buy. Instead of trying, it hires people to go to shopping malls to see, in real time, what people are buying and wearing. On the basis of that information, it produces a big portfolio of the styles, fabrics, and colors that seems popular, and and sends that portfolio to stores, where consumer behavior can be tested. Zara has a fast and flexible system for manufacturing and distribution, which makes it possible to design, produce, and ship new clothing all over the world with tremendous speed.

Measure-and-react is not a randomized controlled trial, but it serves the same functions, and in some ways it is better. It is increasingly used by private sector actors, who know what they do not know, and try to adjust, on the fly, to what people are doing. Can government do the same thing? In many contexts, it certainly can. With respect to security lines at airports, for example, it can make rapid adjustments as the number of travellers varies over time. With respect to traffic fatalities, it can test interventions and monitor diverse situations to test what works and what does not. The sky is the limit here.

Less Tension

But let us underline the more general point. There can be a serious tension between the commitment to cost-benefit analysis and a realistic appreciation of the limits of official knowledge. Without significant efforts to reduce those limits, that analysis might be inadequately informed. It might move governments in the wrong directions. Whenever regulators face significant informational deficits, it is important to explore tools that take advantage of what the private sector knows; market-friendly tools, such as economic incentives, have important advantages on that count. In some cases, regulators might

[15] See Duncan Watts, Everything Is Obvious (2010).

appropriately decide to abstain entirely on the ground that the market failure is not clear and any cure might be worse than the disease.

In other cases, however, Congress has required agencies to act, or the argument for action is too powerful to be ignored. In such cases, cost-benefit analysis is often indispensable. It would be extravagant to contend that notice-and-comment rulemaking can eliminate the knowledge problem, even in the modern era, but it can produce a great deal of help. Retrospective analysis of regulations is good. Learning, in advance and in real time, is even better. When regulators are asked, a few years from now, what they do, here is a hoped-for answer: “We test.”